



"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

Question: 1

Refer to the exhibit.

PE-A	PE-B
vrf definition Customer-A	vrf definition Customer-A
rd 65000:1111	rd 65000:1111
route-target export 65000 1111	route-target export 65000:1111 route-target
route-target import 65000 1111	import 65000 1111 I
address-family ipv4	address-family ipv4 mdt default 233.0 0 1
mdt default 233 00 1	mdt data 233.0.0.3 0.0.0.0 threshold 100
mdt data 233 0 0 2 0.0.0.0 threshold 100	exit-address-family
exit-address-family	

Which tree does multicast traffic follow?

- A. shared tree
- B. MDT default
- C. source tree
- D. MDT voice

Answer: B

Explanation:

Question: 2

Refer to the exhibit.

```
R1
interface g0/0
 ip address 192.168.1.1 255.255.255.0
 ip router isis
 router isis
 net 49.0022.1111.1111.1111.00
 area-password ciSco
```

```
R2
interface g0/1
 ip address 192.168.1.2 255.255.255.0
 ip router isis
 router isis
 net 49.0022.1111.1111.1111.00
 area-password ciSco
```

After you applied these configurations to routers R1 and R2, the two devices could not form a neighbor relationship. Which reason for the problem is the most likely?

- A. The two routers cannot authenticate with one another.
- B. The two routers have the same area ID.
- C. The two routers have the same network ID.
- D. The two routers have different IS-types.

Answer: C

Explanation:

For those asking about the password, area authentication doesn't prevent neighboring to come up because it is carried only in LSP, CSNP and PSNP messages and not in IIH messages.

<https://www.cisco.com/c/en/us/support/docs/ip/integrated-intermediate-system-to-intermediate-system-is-is/13792-isis-authent.html>

Question: 3

Refer to the exhibit.

```
router bgp 65520
 timers bgp 30 240
```

Which effect of this configuration is true?

- A. It sets the keepalive timer to 30 seconds and the hold timer to 240 seconds.
- B. It sets the keepalive timer to 30 milliseconds and the hold timer to 240 milliseconds
- C. It sets the hold timer to 30 milliseconds and the keepalive timer to 240 milliseconds

D. It sets the hold timer to 30 seconds and the keepalive timer to 240 seconds

Answer: A

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgps1.html#wp1552800140

Question: 4

Refer to the exhibit.

```
RP/0/0/CPU0:XR1#show run

route-policy AGGRO
  if destination in (10.0.0.0/8 ge 8 le 25) then
    set community (10:825)
  endif
  if destination in (10.2.0.0/24) then
    drop
  endif
  if destination in (10.1.0.0/24) then
    suppress-route
  endif
end-policy
!
!
router bgp 1
  bgp router-id 192.168.0.7
  address-family ipv4 unicast
    aggregate-address 10.0.0.0/8 route-policy AGGRO
  
```

RP/0/0/CPU0:XR1#

A network operator is working to filter routes from being advertised that are covered under an aggregate announcement. The receiving router of the aggregate announcement block is still getting some of the more specific routes plus the aggregate. Which configuration change ensures that only the aggregate is announced now and in the future if other networks are to be added?

- A. Configure the summary-only keyword on the aggregate command
- B. Set each specific route in the AGGRO policy to drop instead of suppress-route

- C. Filter the routes on the receiving router
- D. Set each specific route in the AGGRO policy to remove instead of suppress-route

Answer: A

Explanation:

Question: 5

Refer to the exhibit.

```
RP/0/0/CPU0:XR3#show bgp 10.11 11.0
Thu Jun 20 20:44 15 749 UTC
BGP routing table entry for 10.11 11.0/24
Versions:
  Process          bRIB/RIB      SendTblVer
  Speaker          9             9
Paths: (2 available, best #2)
  Advertised to update-groups (with more than one peer):
    0.1
  Path #1 Received by speaker 0
  Not advertised to any peer
    1
    10 .0.0.9 from 10.0.0.9 (192 168 0 1)
      Origin IGP, metric 0, localpref 100, valid, external
      Received Path ID 0, Local Path ID 0, version 0
      Origin-AS validity not-found
  Path #2: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0 1
    1
    11 .0 0 13 from 10.0.0.13 (192 168 0 2)
      Origin IGP, metric 0, localpref 100, weight 651, valid, external, best, group-best
      Received Path ID 0, Local Path ID 0, version 9
```

A network operator is getting the route for 10.11.11 0/24 from two upstream providers on #XR3. The network operator must configure #XR3 to force the 10.11.11.0/24 prefix to route via next hop of 10.0.0.9 as primary when available. Which of these can the operator use the routing policy language for, to enforce this traffic forwarding path?

- A. weight of 0 on the prefix coming from 192.168.0.2
- B. lower local preference on the prefix coming from 192.168.0.2
- C. higher local preference on the prefix coming from 192.168.0.1
- D. weight of 100 on the prefix coming from 192.168.0.1

Answer: A


Explanation:

Because Weight is the first attribute in path selection algorithm and is 0 for eBGP routes by default. If we set it to 0, router will continue from weight to the bottom, and when comparing neighbor router

IDs it will find 192.168.0.1 less than 192.168.0.2 and st path will be selected as the best

Question: 6

Refer to the exhibit.



After troubleshooting an OSPF adjacency issue, routers 1, 2, and 3 have formed OSPF neighbor relationships. Which statement about the configuration is true?

- A. Router 2 receives a Type 5 LSAs from router 1 for its connected subnets
- B. Router 2 uses router 3 as the next hop for 192.168.0.0/24
- C. Router 2 uses router 1 as the next hop for 192.168.0.0/24
- D. Router 2 receives a Type 7 LSAs from router 3 for its connected subnets

Answer: C

Explanation:

From IOS 15.1(2)S onwards, Cisco uses the OSPF path selection order is O > O IA > N1 > E1 > N2 > E2


<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/212608-ospf-external-path-selection-external-t.html>

Question: 7

DRAG DROP


Compare different features between OSPFv2 and OSPFv3. Drag and drop the descriptions of OSPF from the left onto the correct OSPF versions on the right.

- introduced IPv6 support
- introduced MD5 authentication
- process network information on a per-link basis
- processes network information on a per-subnet basis
- uses a locally-significant instance ID



Answer:

Explanation:



Question: 8

A network consultant is troubleshooting IS-IS instances to identify why a routing domains is having communication problems between the two instances. Which description of the possible cause of issues in the routing domain is true?

- A. The same interface cannot be advertised in two different IS-IS instances
- B. The IS-IS "ISP" and "ISP2" instances are unrelated and unable to intercommunicate
- C. The configured IS-IS NSEL value is not allowing the routing systems to establish a neighborhood
- D. The interface mode ip router is-is command was not included in the script

Answer: A

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/xr-3s/irs-xr-3s-book/irs-instance-vrf.html

Question: 9

What is used by SR-TE to steer traffic through the network?

- A. shortest path calculated by IGP
- B. dynamic rules
- C. path policy
- D. explicit maps

Answer: C

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/segment-routing/configuration/guide/b-seg-routing-cg-asr9k/b-seg-routing-cg-asr9k_chapter_0100.html

Question: 10

For which reason can two devices fail to establish an OSPF neighbor relationship?

- A. The two devices have different process IDs
- B. The two devices have different network types
- C. The two devices have different router IDs
- D. The two devices have the same area ID

Answer: B

Explanation:

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13699-29.html>

Question: 11

Refer to the exhibit.

```
OSPF Router with ID (192.168.1.1) (Process ID 1)
Router Link States (Area 1234)
LS age: 691
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 192.168.1.1
```

Which LSA type is indicated by this router output?

- A. type 3 LSA
- B. type 4 LSA
- C. type 1 LSA
- D. type 2 LSA

Answer: C

Explanation:

Question: 12

Which statement about enabling segment routing for IGPs is true?

- A. Segment routing must first be enabled under then routing process and then globally
- B. Segment routing must first be enabled globally and then under the routing process
- C. Segment routing can be enabled only under the routing process
- D. Segment routing can be enabled only globally

Answer: B

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/seg_routing/configuration/xr-16-6/segrrt-xr-16-6-book/sr-ospfv2-node-sid.html

Question: 13

Which task is performed when troubleshooting LDP?


- A. Execute the ping utility to generate information about the MAC addresses used along the path
- B. Verify that MPLS is disabled globally and enabled on the necessary interfaces in a per-interface basis
- C. Execute the traceroute utility to generate information about the labels used along the path
- D. Verify that Cisco Express Forwarding has been disabled on the network

Answer: C

Explanation:

Question: 14

Refer to the exhibit.



XR1 and XR2 are sending the prefix 10.11.11.0/24 to XR3. A configured policy on XR1 is incorrectly prepending AS path 11 11 12 12 onto this prefix. A network operator wants to add a policy onto XR3 that will not allow the falsely prepending prefix from being installed. Which policy configuration applied to the XR3 neighbor configuration for XR1 can accomplish this requirement without impact to other or future received routes?

^A route-policy NO_PREPEND if as-path passes-through '11' then pass else drop endif end-policy

^B route-policy NO_PREPEND if as-path prepends drop else pass endif end-policy

^C route-policy NO_PREPEND if as-path passes-through 'T' then pass else drop endif end-policy

^D route-policy NO_PREPEND if as-path passes-through '11' then drop else pass endif end-policy

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

Answer: D

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-1/routing/command/reference/b_routing_cr41crs/b_routing_cr41crs_chapter_01000.html#wp3850885229

Question: 15

Refer to the exhibit.

Router 1:

```
interface TenGigE0/1
  point-to-point
  address-family ipv4 unicast
    fast-reroute per-prefix
    Fast-reroute per-prefix ti-lfa
```

R1#show isis fast-reroute 172.16.200.9/32

```
L2 172.16.200.9/32 [30/115]
  via 192.168.20.1, TenGigE0/1, R2, SRGB Base: 16000, Weight: 0
  FRR backup via 192.168.30.1, TenGigE0/2, R3, SRGB Base: 16000,
  Weight 0, Metric 40
```

Router 1 is connected to router 2 on interface TenGigE0/1. Which interface provides the alternate path to 172.16.200.9/32 when the link between router 1 and router 2 goes down?

- A. TenGigE0/1 interface provides the alternate path

- B. A backup path must be statically installed
- C. TenGigE0/2 interface provides the alternate path
- D. A primary path must be manually installed


Answer: C

Explanation:

Question: 16

Refer to the exhibit.

OS PF Area 0



Current SRGB 19000 19999
Expanded SRGB 18000 24999

A network operator wants to expand the segment routing global block in upcoming maintenance. The operator must ensure that the changes to the segment routing global block have no adverse impacts on the prefix-sid associated with the loopback0 interface used within the OSPF domain.

Which command can the operator use to enforce R2 to have a strict prefix-sid assignment to loopback0?

A- router ospf 1 area 0

```
interface Loopback0 prefix-sid index 19002 explicit-  
null
```

B. router ospf 1 area 0 interface Loopback0 prefix-sid
absolute 13002

C- router ospf 1 area 0 interface Loopback0 prefix-sid
absolute 19002 0- router ospf 1 area 0
interface Loopback0 prefix-sid index 19002

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

Answer: C

Explanation:

Question: 17

Refer to the exhibit.

Router 1 is a core ABR in a Cisco Unified MPLS environment. All of the router 1 BGP peers are established, but traffic between customers is failing. Which BGP configuration must be added to the configuration?

```
router bgp 65525  
ibgp policy out enforce-modifications  
bgp router-id 192.168.1.1  
address-family ipv4 unicast
```

- A. It must be configured for graceful restart
- B. It must be configured with a route reflector
- C. It must be configured with send labels
- D. It must be configured with PIC edge

Answer: C

Explanation:

Question: 18

What can be used to determine a path from the head-end to a tail-end router when implementing SR-TE with a head-end, with little information on the network topology?

- A. traffic controller
- B. path computation engine

- C. tail-end router
- D. SNMP server

Explanation:

Question: 19

Refer to the exhibit.

Answer: B

R1#sh ip int bri

Interface	IP-Address	OK?	Method	Status	Protocol
FastEthernet0/0	10.1.12.1	YES	manual	up	UP
FastEthernet0/1	10.1.13.1	YES	manual	up	

R1#sh run | s router bgp 1

```
router bgp 123
bgp log-neighbor-changes
neighbor TEST peer-group
neighbor TEST remote-as 2 alternate-as 3
neighbor 10.1.12.2 peer-group TEST neighbor
10.1.13.3 peer-group TEST
```

R2#sh ip int bri

Interface	IP-Address	OK?	Method	Status	Protocol
FastEthernet0/0	10.1.12.2	YES	manual	up	up

R2#sh run | s router bgp

```
router bgp 2
bgp log-neighbor-changes
neighbor 10.1.12.1 remote-as 123
```

R3#sh ip int bri

Interface	IP-Address	OK?	Method	Status	Protocol
FastEthernet0/1	10.1.13.3	YES	manual	up	up

R3#sh run | s router bgp

```
router bgp 3
bgp log-neighbor-changes
neighbor 10.1.13.1 remote-as 123
```

R1 is directly connected to R2 and R3. R1 is in BGP AS 123, R2 is in BGP AS 2, and R3 is in BGP AS 3. Assume that there is no connectivity issue between R1, R2 and R1, R3. Which result between BGP peers R1, R2 and R1, R3 is true?

- A. The BGP session does not come up between R1 and R2 and between R1 and R3.
- B. The BGP session comes up between R1 and R2 and between R1 and R3.
- C. The BGP session comes up between R1 and R3, but not between R1 and R2.
- D. The BGP session comes up between R1 and R2, but not between R1 and R3.

Answer: B

Explanation:

Question: 20

Refer to the exhibit.

Router 1:

```
interface tunnel-te12
ipv4 unnumbered loopback0
autoroute announce
destination 192.168.1.2
path-option 12 dynamic segment-routing
path-protection
```

Router 1 has established an SR-TE tunnel with router 2. Which statement describes this configuration?

- A. Router 1 has a list of labels used to explicitly lay out a path to router 2.
- B. Router 1 and router 2 have a bidirectional tunnel set up with dynamic path selection.
- C. Router 1 is the head-end tunnel and has dynamically chosen a path to router 2.
- D. Router 2 is the head-end tunnel and has explicitly set a path to router 1.

Answer: C

Explanation:

Question: 21

Which two conditions must be met before separate ISPs can provide interdomain multicast routing?
(Choose two.)

- A. Each ISP must configure MSDP to connect its individual multicast administrative domain to the domains at other ISPs.
- B. Each ISP must dedicate a single router to handle multicast traffic between providers.
- C. Each ISP must replace its RP assignment with a global RP.
- D. Each ISP must configure its network to use PIM-DM.
- E. Each ISP must support intradomain multicast routing.

Answer: AE

Explanation:

Question: 22

An engineer is troubleshooting a connectivity issue across the MPLS network and is verifying the forwarding

behavior of packets. Which table does the engineer look at to verify the forwarding

behavior of an IP packet as it enters the MPLS network at the ingress LSR?

- A. LFIB
- B. LIB
- C. RIB
- D. FIB

Answer: A

Explanation:

<https://community.cisco.com/t5/image/serverpage/image-id/44142iC3B9032F415B1395?v=v2>

Question: 23

DRAG DROP

An engineer is troubleshooting end-to-end customer traffic across an MPLS VPN service provider network. Which tasks should the engineer use to solve the routing issues? Drag and drop the table types from the left onto the most useful troubleshooting tasks/router types on the right. (Not all options are used.)

on the CE router to check for routing errors	
LFIB	on the P router to see LDP functionality
LIB	on PE and P router to verify expected forwarding
RIB	on VRF of the PE-CE connection
FIB	
adjacency table	
RIB	
LIB	
LFIB	
adjacency table	

Answer:

Explanation:

Question: 24

You have configured MSDP peering between two autonomous systems that pass traffic between two sites, but the peering has failed to come up.

Which task do you perform to begin troubleshooting the problem?

- A. Verify that multicast has been disabled globally
- B. Verify that PIM-DM is configured on the source interface
- C. Verify that both source interfaces are reachable from both peers
- D. Verify that the two MSDP peers allow asymmetric routing

Answer: C

Explanation:

Question: 25

Refer to the exhibit.

```
R2#sh ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.1.3.3	1	FULL/BDR	00:00:37	10.1.234.3	EthernetO/O.234
10.1.3.4	1	FULUDR	00:00:35	10.1.234.4	Ethernet0/0.234
10.1.3.5	1	2WAY/DROTHER	00:00:35	10.1.234.5	Ethernet0/0.234

Why is neighbor 10.1.5.5 stuck in “2WAY” state?

- A. Router ID 10.1.5.5 is not reachable from R2
- B. OSPF authentication has failed between R2 and 10.1.5.5
- C. It is an expected behavior when OSPF network type is broadcast
- D. OSPF parameters (Area ID or hello interval) are mismatched between R2 and 10.1.5.5

Answer: C

Explanation:

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13683-11.html>

Question: 26

For which reason do you deploy BGP confederations within a BGP transit backbone?

- A. to support a larger number of eBGP peer sessions
- B. to increase the number of routes that can be redistributed between the running IGP and BGP
- C. to reduce the number of eBGP routes that must be shared between autonomous systems
- D. to reduce the number of iBGP peering sessions

Answer: D

Explanation:

Question: 27

Which feature is used in multicast routing to prevent loops?


- A. STP
- B. inverse ARP
- C. RPF
- D. split horizon

Answer: C

Explanation:

Question: 28

Refer to the exhibit.



How are packets directed through the data plane when SRv6 is implemented?

- A. An ordered list of segments is encoded in a routing extension header
- B. The MPLS data plane is used to push labels onto IGP routes
- C. A stack of labels represents an ordered list of segments
- D. The packet is encapsulated with a header and trailer encoding the ordered list of segments


Answer: A

Explanation:

Reference: <https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2019/pdf/BRKIPM-2249.pdf>

Question: 29

Refer to the exhibit.



An engineer is addressing an IS-IS design issue which is running within the topology. All links are running on FastEthernet, except the link between R5 and R4, which is Gigabit Ethernet. Which statement about the design is true?

- A. R4 prefer to reach R5 using R1 as the next hop
- B. All links have equal cost if the default metric is used
- C. R5 prefers to use R4 as the next hop for all routes
- D. R1 prefer to use R5 as the next hop to reach R4

Answer: B

Explanation:

Question: 30

Refer to the exhibit.

```
"PEftshow ip msdp peer
```

```
MSDP Peer 10.10 10.10 (?), AS?
```

```
Connection status:
```

```
State: Listen. Resets: 0. Connection source none configured Uptime (Downtime):
```

00:00:07, Messages sent/received 0/0
Output messages discarded: 0
Connection and counters cleared 00:00:7 ago
SA Filtering:
Input (S, G) filter: none, route-map: none
Input RP filter: none, route-map: none
Output (S, G) filter: none, route-map: none
Output RP filter: none, route-map: none
SA-Requests:
Input filter: none
Peer ttl threshold: 0
SAs learned from this peer: 0
Input queue size: 0, Output queue size: 0"

A service provider technician is working on a multicast issue for a customer. While checking the multicast table, the technician notices that no flags are present for the (1.1.1.1, 239.1.1.1) entry, yet flags are present for the (1.1.1.1, 232.1.1.1) entry.

Which factor might explain this issue?

- A. Only the administratively scoped range is permitted
- B. Only ASM is permitted
- C. Only the default SSM range is permitted
- D. Only GLOP is permitted

Answer: C

Explanation:

Question: 31

After you change the IP address on an IOS XR router, you cannot ping the new address. Which step did you forget to complete?

- A. commit the configuration
- B. roll back the configuration
- C. merge the configuration
- D. save the running configuration

Answer: A

Explanation:

Question: 32

Which two statements about route reflectors are true? (Choose two.)

- A. Routes received from nonclient peers are reflected to route reflector clients as well as nonclient peers.
- B. Routes received from nonclient peers are reflected to route reflector cluster as well as OSPF peers.

- C. If a router received an iBGP route with the originator-ID attribute set to its own router ID, the route is discarded.
- D. Routes received from a route reflector client is reflected to other clients and nonclient peers.
- E. If a route reflector receives a route with a cluster-list attribute containing a different cluster ID, the route is discarded.

Answer: CD

Explanation:

Reference: <http://www.networkers-online.com/blog/2009/02/bgp-route-reflector-basics/>

Question: 33

Which two statements about mapping multicast IP addresses to MAC addresses are true? (Choose two.)

- A. All mapped multicast MAC addresses begin with 0x0100.5E
- B. The router performs the mapping before it hands the packet off to a switch
- C. All multicast MAC addresses end with 0x0100.5E
- D. The mapping process may generate overlapping addresses, which can cause receivers to receive unwanted packets
- E. All destination MAC addresses begin with an octet of binary 1s

Answer: AD

Explanation:

Question: 34

You have configured routing policies on a Cisco IOS XR device with routing policy language. Which two statements about the routing policies are true? (Choose two.)

- A. The routing policies affect BGP-related routes only.
- B. If you make edits to an existing routing policy without pasting the full policy into the CLI, the previous policy is overwritten.
- C. You can change an existing routing policy by editing individual statements.
- D. The routing policies are implemented in a sequential manner.
- E. The routing policies are implemented using route maps.

Answer: CD

Explanation:

Question: 35

Refer to the exhibit.

```
ip pim ssm interface g 1/0/0 ip pim sparse-mode
```

Which task must you perform on interface g1/0/0 to complete the SSM implementation?

- A. configure OSPFv3
- B. enable CDP
- C. disable IGMP
- D. configure IGMPv3

Answer: D

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti_pim/configuration/xe-16/imc-pim-xe-16-book/imc-ssm.html

Question: 36

Refer to the exhibit.

Router 1:

```
router bgp 65530
 address-family ipv4 unicast bgp additional-paths select all neighbor 192.168.1.1
 additional-paths send neighbor 192.168.1.1 advertise additional-paths all
```

Which statement about this configuration is true?

- A. Router 1 sends and receives multiple best paths from neighbor 192.168.1.1
- B. Router 1 sends up to two paths to neighbor 192.168.1.1 for all routes
- C. Router 1 receives up to two paths from neighbor 192.168.1.1 for all routes in the same AS
- D. Router 1 receives only the best path from neighbor 192.168.1.1

Answer: A

Explanation:

Question: 37

Which cost is the default when redistributing routes from BGP to OSPF?

- A. 20
- B. 1
- C. infinite
- D. automatic

Answer: B

Explanation:

Reference: <http://www.ciscopress.com/articles/article.asp?p=27573>

Question: 38

Refer to the exhibit.

Router 1:

```
router ospf 20  
redistribute eigrp 1  
network 192.168.0.0 0.0.0.255 area 0
```

An engineer is troubleshooting an OSPF issue. Router 1 has a neighbor relationship with router 2. Only router 1 classful EIGRP routes can be seen on router 2. In order for all EIGRP routes to be redistributed correctly, which action should be taken?


- A. Router 1 must have the keyword subnets included in the redistribution command for all EIGRP routes to be redistributed.
- B. Router 1 must remove the AS number 1 from the redistribution command for all EIGRP routes to be redistributed.
- C. Router 1 must have the keyword ospf-metric included in the redistribution command for all EIGRP routes to be redistributed.
- D. Router 1 must have the keyword metric-type 1 included in the redistribution command for all EIGRP routes to be redistributed.

Answer: A

Explanation:

Question: 39

Refer to the exhibit.



A network operator must inject a Level 1 route from XR2 (10.16.16.0/24) into the ISIS topology. Which configuration allows the injection in a way that XR3 and XR1 have a valid and working route for 10.16.16.0/24?

A- #XR3

```
route-policy ISIS_PROPO
  if destination in (10.0.0.0/8 ge 8 le 22) then
    pass endif
end-policy I
router isis 1
  net 49.1921.6800.0003.00
  address-family ipv4 unicast

propagate level 1 into level 2 route-policy ISIS_PROPO
```

^B #XR2

```
route-policy ISIS_PROPO
  if destination in (10.0.0.0/8 ge 8 le 32) then
    pass endif
end-policy
router isis 1
  net 49.1921.6800.0003.00
  address-family ipv4 unicast

propagate level 2 into level 1 route-policy ISIS_PROPO
```

^C #XR2

```
route-policy ISIS_PROPO
  if destination in(10.0.0.0/8 ge 8 le 32) then
    pass
  endif
end-policy
```

```
router isis 1
  net 49.1921.6800.0003.00
  address-family ipv4 unicast
```

```
  propagate level 1 into level 2 route-policy ISIS_PROPO
```

^D #XR3

```
route-policy ISIS_PROPO
  if destination in(10.0.0.0/8 ge 8 le 32) then
    pass
  endif
end-policy
```

```
router isis 1
  net 49.1921.6800.0003.00
  address-family ipv4 unicast
  propagate level 2 into level 1 route-policy ISIS_PROPO
```

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


Answer: D

Explanation:

L1 route injected by XR2 won't be advertised to XR1, since L2 routes are NOT leaked to L1 by default.
An [L2-L1] route leak policy should be configured on L1/L2 router (XR3).

Question: 40

Refer to the exhibit.



CE1 and CE2 cannot communicate through the service provider BGP peering is established between PE1 and PE2. IS-IS is the only routing protocol running in the service provider CORE.

What step can be done to troubleshoot the issue?

- A. Switch the IGPs running in the core from IS-IS to OSPF to support a Cisco MPLS TE tunnel from PE1 to PE2.
- B. Configure BGP between CE and PE routers.
- C. Confirm that IS-IS is running with metric-style narrow.
- D. Verify the MPLS LSPs.

Answer: D

Explanation:

Question: 41

Refer to the exhibit.

Router 1:	Router 2:
mpls traffic-eng tunnels	mpls traffic-eng tunnels
router ospf 2	router ospf 2
mpls traffic-eng router-id loopback 0	mpls traffic-eng router-id loopback 0
mpls traffic-eng area 0	mpls traffic-eng area 0
interface gigabitethernet0/1	interface gigabitethernet0/1
ip rsvp bandwidth	mpls traffic-eng tunnels
	ip rsvp bandwidth
interface tunnel 1	
ip unnumbered loopback 0	
tunnel destination 192168.4 1	
tunnel mode mpls traffic-eng	
tunnel mpls traffic-eng bandwidth 150	
tunnel mpls traffic-eng autoroute announce	

Router 1 has attempted to establish a Cisco MPLS TE tunnel to router 2, but the tunnel has failed. Which statement about this configuration is true?


- A. Router 1 must define an explicit path to router 2
- B. Router 1 and router 2 must define the RSVP bandwidth reserved on the physical interfaces
- C. Router 2 must have a tunnel interface created with router 1 as the destination
- D. Router 1 must have Cisco MPLS TE enabled on interface gigabitethernet0/1

Answer: D

Explanation:

Question: 42

Refer to the exhibit



Which router does R1 install as an alternate next hop when trying to reach R3 if LFA is enabled?

- A. R5
- B. R3
- C. R4
- D. R2

Answer: D

Explanation:

Question: 43

Refer to the exhibit.

```
R1
interface gigabitethernet0/0
 ip address 192.168.2.1 255.255.255.0
 ip router isis
 router isis
  net 49.0022.1111.1111.1111.00
  is-type level-1

R2
interface gigabitethernet0/1
 ip address 192.168.1.2 255.255.255.0
 ip router isis
 router isis
  net 49.0021.1111.1111.1112.00
  is-type level-1
```

Routers R1 and R2 cannot form a neighbor relationship, but the network is otherwise configured correctly and operating normally. Which two statements describe the problem? (Choose two.)

- A. The two routers are in the same area
- B. The two routers are in different subnets
- C. The two routers have password mismatch issues
- D. The two routers have the same network ID
- E. The two routers are in different areas

Answer: BE

Explanation:

Question: 44

Refer to the exhibit.

RP/0/0/CPU/0:P1#

!

key chain BGP

key 1

accept-lifetime 13:14:06 february 14 1993 infinitive

send-lifetime 13:14:06 february 14 1993 infinitive

key-string password cisco123

cryptographic-algorithm MD5

!

!

router bgp 1

address-family ipv4 unicast

!

neighbor 192.168.13.3

remote-as 1

keychain BGP

address-family ipv4 unicast

RP/0/0/CPU/0:PE3#

!

key chain BGP

key 1

accept-lifetime 13:14:06 february 14 1993 infinitive

send-lifetime 13:14:06 february 14 1993 infinitive

key-string password cisco123

cryptographic-algorithm MD5

!

!

router bgp 1

address-family ipv4 unicast

!

neighbor 192.168.13.1

remote-as 1

keychain BGP

address-family ipv4 unicast

P1 and PE3 Cisco IOS XR routers are directly connected and have this configuration

applied. The BGP session is not coming up. Assume that there is no IP reachability problem and both routers can open tcp port 179 to each other. Which two actions fix the issue? (Choose two.)

- A. Change MD5 to HMAC-SHA1-12
- B. Change MD5 to HMAC-ESP
- C. Change MD5 to SHA-1
- D. Change MD5 to HMAC-MD5
- E. Remove the send and accept lifetime under key 1

Answer: AD

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-0/security/configuration/guide/sc40crsbook_chapter5.html

Question: 45

Which command is used to enable BIDIR-PIM under global configuration mode for Cisco IOS XE Software?

- A. ip pim bidir-enable
- B. ipv4 pim bidir-enable
- C. ip multicast-routing
- D. ip pim bidir

Answer: A

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti_pim/configuration/xe-3s/imc-pim-xe-3sbook/imc_basic_cfg.html

Question: 46

Which output from the show isis interface command helps an engineer troubleshoot an IS-IS adjacency problem on a Cisco IOS-XR platform?

- A. metric
- B. priority
- C. circuit type
- D. hello interval

Answer: C

Explanation:

RP/0/0/CPU0:XR1#show isis interface | i Cir
Tue Jun 29 17:56:42.369 UTC
Circuit Type: level-1-2

Circuit Number: 0

Circuit Type: level-1

Circuit Number: 1

RP/0/0/CPU0:XR1#show isis interface | i Hel

Tue Jun 29 17:56:46.609 UTC

Question: 47

A network engineer is troubleshooting OSPF multiare

A. Which Cisco IOS XR feature should the engineer use in order to streamline OSPF issue?

A. hierarchical CLI

B. DR support for topology management

C. routing process enabled by default on all interfaces

D. show ip ospf topology command

Answer: A

Explanation:

OSPF Hierarchical CLI and CLI Inheritance

Hierarchical CLI is the grouping of related network component information at defined hierarchical levels such as at the router, area, and interface levels. Hierarchical CLI allows for easier configuration, maintenance, and troubleshooting of OSPF configurations. When configuration commands are displayed together in their hierarchical context, visual inspections are simplified. Hierarchical CLI is intrinsic for CLI inheritance to be supported.

With CLI inheritance support, you need not explicitly configure a parameter for an area or interface. In the software, the parameters of interfaces in the same area can be exclusively configured with a single command, or parameter values can be inherited from a higher hierarchical level—such as from the area

configuration level or the router ospf configuration levels.

More information: https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/routing/71x/b-routing-cg-ncs5500-71x/b-routing-cg-ncs5500-71x_chapter_011.html#con_1059437

Reference: https://www.cisco.com/c/en/us/td/docs/routers/xr12000/software/xr12k_r4-0/routing/configuration/guide/rc40xr12k_chapter4.html#con_1059437

Question: 48

Refer to the exhibit.

```
router bgp 65515
  neighbor 192.168.1.1 route-map ciscotest in
  neighbor 192.168.1.1 remote-as 65516

ip as-path access-list 1 permit _65517_

route-map ciscotest permit 10
  match as-path 1
  set local-preference 150
```

After troubleshooting BGP traffic steering issue, which action did the network operator take to achieve the correct effect of this configuration?

- A. Routes that have passed through AS 65517 have the local preference set to 150.
- B. Routes that have originated through AS 65517 have the local preference set to 150.
- C. Routes directly attached to AS 65517 have the local preference set to 150.
- D. Routes that have passed through AS 65517 have the local preference set to 150 and the traffic is denied.

Answer: A

Explanation:

Question: 49

For which reason can two BGP peers fail to establish a neighbor relationship?

- A. Their BGP send-community strings are misconfigured
- B. Their BGP timers are mismatched
- C. Their remote-as numbers are misconfigured
- D. They are both activated under an IPv4 address family

Answer: C

Explanation:

Question: 50

In a PIM-SM environment, which mechanism determines the traffic that a receiver receives?

- A. The receiver explicitly requests its desired traffic from the RP on the shared tree.
- B. The receiver explicitly requests traffic from a single source, which responds by forwarding all traffic.
- C. The RP on the shared tree floods traffic out of all PIM configured interfaces.

D. The receiver explicitly requests traffic from each desired source, which responds by sending all traffic.

Answer: A

Explanation:

A shared tree is built first between receiver and the RP. The receiver is then able to switch to a Source tree as needed.

Question: 51

Which statement about BFD on Cisco IOS XR Software is true?

- A. Cisco IOS XR router must use LDP to route back to the Cisco IOS router to establish the peer relationship.
- B. Cisco IOS XR Software does not support BFD multihop for IPv4.
- C. Cisco IOS XR router must use dynamic routing or a static route back to the Cisco IOS router to establish the peer relationship.
- D. BFD is not compatible between Cisco IOS XR and Cisco IOS Software.

Answer: C

Explanation:

A router running BFD in Cisco IOS software can designate a router running BFD in Cisco IOS XR software as its peer using the bfd neighbor command; the Cisco IOS XR router must use dynamic routing or a static route back to the Cisco IOS router to establish the peer relationship.

https://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r6-1/routing/configuration/guide/b-routing-cg-crs-61x/b-routing-cg-crs-61x_chapter_0100.html

Reference: https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k-r6-3/routing/configuration/guide/b-routing-cg-asr9000-63x/b-routing-cg-asr9000-63x_chapter_0100.html

Question: 52

Which two routing protocols have extensions capable of running SRv6? (Choose two.)

- A. OSPF
- B. BGP
- C. RIP
- D. IGRP
- E. EIGRP

Answer: AB

Explanation:

Question:

53


DRAG DROP

Drag and drop the attributes for the BGP route selection on the left into the correct order on the right.
Not all options are used.

lowest router ID	Step 1
highest router ID	Step 2
lowest local preference	Step 3
if both paths are external, prefer the newest path	Step 4
highest local preference	Step 5
highest weight	Step 6
eBGP over iBGP paths	
if both paths are external, prefer the oldest path	
lowest MED	

Answer:

Explanation:



Reference: <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html>

Question: 54

Refer to the exhibit.

```
RP/0/0/CPU0:iosxr# show run segment-routing
```

```
segment-routing
  global-block 18000 24999
```

```
!
```

```
RP/0/0/CPU0:iosxr#
```

A network engineer implemented this segment routing configuration. Which statement about the output is true?

- A. This range conflicts with the segment routing local block range.

- B. The device must be reloaded for these ranges to be allocated and used.
- C. The default segment routing global block range is being used on this device.
- D. A nondefault segment routing global block range is being used on this device.


Answer: D

Explanation:

Reference: <https://www.cisco.com/c/en/us/td/docs/routers/asr920/configuration/guide/segment-routing/segment-routing-book/seg-routing-global-block.html>

Question: 55

Refer to the exhibit.



CE1 is the gateway router into the provider network via PE1. A network operator must inject a default route into OSPF area 0. All devices inside area 0 must be able to reach PE1. Which configuration achieves this goal?

A #CE1

```
router ospf 1 default-information originate always
```

B #CE1

```
ip route 0.0.0.0 0.0.0.0 GigabitEthernetO/2 10.0.0.14 ( router ospf 1 redistribute static
```

C- #CE1

```
ip route 0.0.0.0 0.0.0.0 GigabitEthernetO/2 10.0.0.14 j
```

```
router ospf 1 default-information originate
```

D #CE1

```
ip route 0.0.0.0 0.0.0.0 GigabitEthernetO/2 10.0.0.14 l router ospf 1 redistribute static
subnets
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/47868-ospfdb9.html>

Question: 56

Which two characteristics unique to SSM when compared to ASM are true? (Choose two.)

- A. It uses SPT switchover
- B. It uses (*,G) exclusively
- C. It uses IGMPv3
- D. It uses RP
- E. It uses (S,G) exclusively

Answer: CE

Explanation:

Question: 57

Refer to the exhibit.

```
RP/0/0/CPU/0:P1#
```

```
key chain BGP
key 1
key-string password cisco123
cryptographic-algorithm HMAC-MD5
1
router bgp 1
address-family ipv4 unicast
neighbor 192.168.13.3
remote-as 1
keychain BGP
address-family ipv4 unicast
```

```
RP/0/0/CPU/0:PE3#j
```

```
key chain BGP
key 1
key-string password cisco123
cryptographic-algorithm HMAC-MD5
1
router bgp 1
address-family ipv4 unicast
neighbor 192.168.13.1
remote-as 1
keychain BGP
address-family ipv4 unicast
```

P1 and PE3 Cisco IOS XR routers are directly connected and have this configuration applied. The BGP session is not coming up. Assume that there is no IP reachability problem and both routers can open tcp port 179 to each other. Which action fixes the issue?


- A. Change HMAC-MD5 to HMAC-SHA1-20
- B. Configure the send and accept lifetime under key 1
- C. Change HMAC-MD5 to MD5
- D. Change HMAC-MD5 to HMAC-SHA1-12

Answer: B

Explanation:

Question: 58

Refer to the exhibit.



An engineer has successfully fixed BGP peering issue. R1 has an established eBGP

peering with R2 and R3. Which mechanism should the engineer apply in order to steer the traffic correctly?

- A. The MED attribute can be applied on R2 to influence R1 to use it as the primary path.
- B. The local preference attribute can be applied on R3 to influence AS 65513 to use AS 65515 as the secondary path.
- C. The weight attribute can be applied on R2 to influence AS 65513 to use AS 65515 as the primary path.
- D. The IGP metric can be manipulated on R1 to allow traffic to be load balanced between R2 and R3.

Answer: D

Explanation:

Some level of load balancing is achieved by default as BGP prefers shortest AS_PATH for certain prefix (Cisco BGP path selection algorithm - rule 4) . If 1-7 rules gives no preference for certain prefix then with rule 8.

Prefer lowest IGP metric to BGP next hop we can force traffic to R2 or R3 influencing load distribution.

Question: 59

Refer to the exhibit.

```
Router1# show ip bgp
```

BGP table version is 4, 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

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

	Network	Next Hop	Metric	LocPrf	Weight	Path
>	192.168.10.0/24	192.168.1.2	0		0	65525
>	192.168.3.0/24	192.168.2.2	0		0	65535
	192.168.3.0/24	192.168.4.2	0		0	65545
>	192.168.20.0/24	0.0.0.0	0		32768	i

Which attribute can router 1 alter so that only other iBGP peers prefer to use 192.168.4.2 as the next hop for route 192.168.3.0/24?

- A. MED
- B. local preference
- C. origin
- D. weight


Answer: B

Explanation:

Weight is local to the router while LP allows routers in an AS with multiple exit points to choose which exit point is used to reach a particular NLRI.

Question: 60

Refer to the exhibit.



A network operator must stop 10.33.33.33/32 from being redistributed into Level 1 router XR1. Which configuration meets this need?

```

A #XR2
  prefix-set NO_33
    10.33.33.33/32
  end-set

  route-policy ISIS_NO_33
    if destination in NO_33 then drop
    else
      pass
    endif
  end-policy

router isis 1
  address-family ipv4 unicast
    propagate level 2 into level 1 route-policy ISIS_NO_33
  
```

```

B #XR3
  prefix-set NO_33
    10.33.33.33/32
  end-set
  
```

```
route-policy ISIS_NO_33
  if destination in NO_33 then drop
endif
end-policy !
router isis 1
  address-family ipv4 unicast
  propagate level 2 into level 1 route-policy ISIS_NO_33
```

```
#XR3
prefix-set NO_33
  10.33.33.33/32
end-set
```

```
i
route-policy ISIS_NO_33
  if destination in NO_33 then drop
else
  pass endif
end-policy i
router isis 1
  address-family ipv4 unicast
  propagate level 2 into level 1 route-policy ISIS_NO_33
```

```
D #XR3
  prefix-set NO_33
    10.33.33.33/23
  end-set

  route-policy ISIS_NO_33
    if destination in NO_33 then drop
    else
      pass
    endif
  end-policy ;
router isis 1
  address-family ipv4 unicast
  propagate level 2 into level 1 route-policy ISIS_NO_33
```


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

Answer: C

Explanation:

Question: 61

Refer to the exhibits.



All links inside the network are configured at a default cost of one inside the fully converged OSPF domain. Given the configuration from XR1, which interface does traffic from XR1 that is destined to the loopback interface of XR6 select for the exiting interface?

- A. Interface GigabitEthernet 0/0. The tie breaker of the node index priority is lower and trumps the path cost.
- B. Interface GigabitEthernet 0/1. The tie breaker of the path cost being lower. The node index priority does not impact this selection process.
- C. Interface GigabitEthernet 0/1. The tie breaker of the node index priority is lower and trumps the path cost.
- D. Interface GigabitEthernet 0/0. The tie breaker of the path cost being lower. The node index priority does not impact this selection process.

Answer: A

Explanation:

Question: 62

After an engineer configures BGP in R1, it starts receiving this message

*Jun 29 13:30:50.122: %BGP-5-ADJCHANGE: neighbor 192.168.10.1 Down User reset

Jun 29 13:30:52.341: %BGP-3-NOTIFICATION: sent to neighbor 192.168.10.1 2/6 (unacceptable hold time) 0 bytes

Which action makes the peering come back up again?

- A. Make a soft reset to the peer.
- B. Set up a minimum hold-down timer higher.
- C. Set up a hello timer higher.
- D. Set up a hold-down timer higher.

Answer: D

Explanation:

Question: 63

Refer to the exhibits.

24000	Pop	192.168.0.2/32	Gi0/0/0/3	10.0.0.5	1644
24001	24000	192.168.0.4/32	Gi0/0/0/2	10.0.0.30	24647
	24000	192.168.0.4/32	Gi0/0/0/3	10.0.0.5	0
24002	Pop	192.168.0.6/32	Gi0/0/0/2	10.0.0.30	12412
24003	24001	192.168.0.7/32	Gi0/0/0/2	10.0.0.30	22359
	24001	192.168.0.7/32	Gi0/0/0/3	10.0.0.5	1473
24004	Pop	10.0.0.20/30	Gi0/0/0/3	10.0.0.5	0
24005	Pop	10.0.0.16/30	Gi0/0/0/2	10.0.0.30	0
	Pop	10.0.0.16/30	Gi0/0/0/3	10.0.0.5	0
24006	Pop	10.0.0.40/30	Gi0/0/0/2	10.0.0.30	0
24007	24002	10.0.0.32/30	Gi0/0/0/2	10.0.0.30	0
	24002	10.0.0.32/30	Gi0/0/0/3	10.0.0.5	7045024
24009	Unlabelled	10.1.1.1/32	Gi0/0/0/0	10.0.0.9	7037648
RP/0/0/CPU9:PE1#					

A network operator is troubleshooting packet loss seen from the R1 loopback interface to the R2 loopback interface over the core network. The operator is attempting to identify the next leg in the path from PE1.

Which interface and label path should the operator investigate next?


- A. PE1 – Gi0/0/0/3 - forwarding label 24002
- B. PE1 – Gi0/0/0/2 - forwarding label 24002
- C. PE1 – Gi0/0/0/3 - forwarding label 24001
- D. PE1 - Gi0/0/0/2 - forwarding label 24001

Answer: C

Explanation:

Question: 64

Refer to the exhibit.



Routers R1, R2, R3, and R4 have been configured to run OSPF, and router R5 is running EIGRP. Traffic from R1 to R5 is expected to pass via R4, but OSPF routing has calculated the best path via R3. Which action corrects the problem?

- A. Configure R3 to use metric-type 1 with a higher metric than R4.
- B. Configure R1 with a static route to the R5 network and set R5 as the next hop.
- C. Reconfigure R1 in Area 0.
- D. Change the metric-type value on R3 to 2.

Answer: D

Explanation:

Question: 65

Refer to the exhibit.

```
RP/0/0/CPU0:PE1#show run
mpls traffic-eng
  soft-preemption
    timeout 10
!
interface tunnel-te 1
  ipv4 unnumbered Loopback0
  signalled-bandwidth 1000
  autoroute announce
  destination 192.168.0.7
  fast-reroute
```

A network operator is troubleshooting TE tunnels and discovers that the soft preemption is not working as desired. Where should the soft preemption be applied to solve the issue?

- A. on the tail-end device for tunnel-te1 192.168.0.7
- B. under the tunnel-te1 interface on PE1
- C. explicitly for tunnel-te 1 under the Cisco MPLS TE configuration mode after the timeout setting

A. Under the interface Loopback0 necessary for this deployment type.

Answer: B

Explanation:

Question: 66

Refer to the exhibit.

Router 1:

peer

address ipv4 192.168.1.1

Router 2:

router isis
distribute link-state instance-id 40

router bgp 65515
address-family link-state link-state neighbor 192.168.1.1 remote-as
65515 update-source LoopbackO
address-family link-state link-state

What is the relationship between Router 1 and Router 2?

- A. Router 1 centrally learns the topology of the network to aid in SR-TE path selection, and Router 2 is a node that feeds Router 1 topology information.
- B. Router 1 and Router 2 are participating in SR-TE tunnels and are both head-end routers.
- C. Router 1 and Router 2 centrally learn the topology of the network to aid in SR-TE path selection for peers.
- D. Router 2 is the head-end router in an SR-TE tunnel, and it is learning topology of the network from the PCE enabled on Router 1.

Answer: A

Explanation:

Question: 67

Which command is used to enable BIDIR-PIM under global configuration mode for Cisco IOS XE Software?

- A. ip pim bidir ip
- B. multicast-routing

- C. ipv4 pim bidir-enable
- D. Ip pim bidlr-enable

Answer: D

Explanation:

Question: 68

What is the purpose of a BGP confederation?

- A. It limits the number of routes a device receives from its peers, which reduces CPU load.
- B. It improves service by increasing the number of simultaneous iBGP peering sessions.
- C. It redirects traffic away from route reflectors, which reduces their operating load.
- D. It reduces the number of iBGP peers and increases stability.

Answer: D

Explanation:

Question: 69

What is the characteristic of enabling segment routing for IGPs?

- A. Segment routing must first be enabled under the routing process and then globally.
- B. Segment routing must first be enabled globally and then under the routing process.
- C. Segment routing must be enabled only globally.
- D. Segment routing must be enabled only under the routing process.

Answer: B

Explanation:

Question: 70

Refer to the exhibit.



R1# show ip ospf interface serial1/0

(output limited)

Serial 1/0 is up, line protocol is up

Internet Address 172.16.1.0/32 Area 0

Process ID 1 Router ID 11.11 Network Type BROADCAST, Cost 64

Transmit Delay is 1 sec State DR, Priority 0

Designated Router (ID) 172.16.1.0. Interface address 172.16.1.0

While configuring router 2 with all the default values, a network engineer cannot see any route received in router 1. How should the engineer solve the issue?


- A. Set up a priority different than 0 in the interface.
- B. Modify the router ID to be the interface IP on the serial.
- C. Modify the IP address or mask of the interface to a valid one.
- D. Set the network type in S1/0 to point-to-point.

Answer: C

Explanation:

Question: 71

Refer to the exhibit.



An engineer has configured all routers in the environment to run IS-IS Level 1 and Level 2 routing.

The engineer wants traffic from R1 to R5 to pass via R2, but IS-IS routing has calculated the best path via R4.

Which action corrects the problem?

- A. Configure routers R1, R4, and R5 for Level 2 routing only.
- B. Set the link metric for the link from router R1 to router R4 to 30 or more.
- C. Set the link metric on R2 for the links from router R2 to routers R3 and R4 to 30 or more.
- D. Configure routers R1, R2, and R5 for Level 1 routing only.

Answer: B

Explanation:

Question: 72

What are three requirements for a static IPv6-in-IPv4 tunnel? (Choose three.)

- A. A static IPv6 address must be configured.
- B. A dynamic IPv6 address must be configured.
- C. Cisco Express Forwarding must be enabled.
- D. Each tunnel endpoint must support IPv4 and IPv6.
- E. A static IPv4 address must be configured.
- F. A dynamic IPv4 address must be configured.

Answer: A, D, E

Explanation:

An IPv6 address is manually configured on a tunnel interface, and manually configured IPv4 addresses are assigned to the tunnel source and the tunnel destination. The host or device at each end of a configured tunnel must support both the IPv4 and IPv6 protocol stacks. Manually configured tunnels can be configured between border devices or between a border device and a host. Cisco Express Forwarding switching can be used for IPv6 manually configured tunnels, or Cisco Express Forwarding switching can be disabled if process switching is needed.

Question: 73

Which expression, when used with the as-path-set command, matches the AS path attribute of the route?

- A. neighbor-is
- B. dfa-regtex
- C. ios-regex
- D. Length

Answer: C

Explanation:

Question: 74

What are two differences between OSPF and IS-IS? (Choose two.)

- A. OSPF is a link-state routing protocol, and IS-IS is a distance-vector routing protocol.

- B. OSPF uses a router ID to identify a router, and IS-IS uses a system ID.
- C. OSPF elects a DR and a BDR, and IS-IS elects a DIS.
- D. Unlike OSPF, IS-IS supports virtual links.
- E. Unlike IS-IS routers, an OSPF router belongs to only one area in addition to the backbone area.

Answer: B, C

Explanation:

Question: 75

A network engineer is troubleshooting OSPF multiarea

a. Which Cisco IOS XR feature should the engineer use in order to streamline OSPF issue?

- A. hierarchical CLI
- B. routing process enabled by default on all interfaces
- C. DR support for topology management
- D. show ip ospf topology command

Answer: A

Explanation:

Question: 76

Refer to the exhibit.

Routed config-router)#no bgp client-to-client reflection intra-cluster cluster-id 192.168.11

Routers within the cluster are not receiving the desired prefixes. What must be done to fix the issue?

- A. Clients in that cluster must have full mesh connectivity between eBGP peers.
- B. No client-to-client must be disabled.
- C. Clients in that cluster must have full mesh connectivity between iBGP peers.
- D. No client-to-client reflection must be enabled.

Answer: C

Explanation:

Question: 77

Refer to the exhibit.

```
RP/0/0/CPU0:XR1#
route-policy ALL_LABELS
  drop
end-policy
!
router bgp 1
  bgp router-id 192.168.0.3
  address-family ipv4 unicast
    allocate-label route-policy ALL_LABELS
```

A network operator is working to deploy a Unified BGP design and allow it to be available only in selected markets and services in the future. The label allocation is not functioning as desired. Which action will fix the issue?

- A. Remove the allocate-labels command.
- B. Reconfigure the route-policy ALL_LABELS to match all after the drop.
- C. Reconfigure the route-policy ALL_LABELS to pass.
- D. Remove the route policy and only configure "allocate-label".

Answer: C

Explanation:

Question: 78

An engineer is working to implement segment routing protocol on the customer's core network. Which step should the engineer take before the segment routing is enabled and is running with BGP?

- A. Segment routing must be configured with ISIS.
- B. Segment routing must be configured with EIGRP.
- C. Explicit-null must be configured for all neighbors.
- D. MPLS must be configured.

Answer: D

Explanation:

Question: 79

After changing the IP address on an IOS XR router, an engineer cannot ping the new address. Which step did the engineer forget to complete?


- A. commit the configuration
- B. roll back the configuration
- C. save the running configuration
- D. merge the configuration

Answer: A

Explanation:

Question: 80

Refer to the exhibit.



The green policy is designed with low delay, and the pink policy is designed with low cost. What is the next hop for node A on its way to node D, with the objective of achieving a low delay?


- A. 192.168.3.2
- B. 192.168.1.6
- C. 192.168.3.6
- D. 192.168.1.2

Answer: D

Explanation:

Question: 81

Refer to the exhibit.



A network operators configuring BGP PIC on CE1 on already established neighborships with PE1 and PE2 inside the fully converged MPLS network Which element needs to be implemented to make this feature function effectively?

- A. Bidirectional Forwarding Detection must be applied to the upstream facing BGP interfaces.
- B. The operator must ensure that all prefixes have the same next-hop from PE1 and PE2 for BGP PICA reserved BGP
- C. community of 1 10 must be used to denote the PIC feature set to the routing protocol
- D. BGP import export policies must be applied on all devices for the routes needing BGP for PIC

Answer: A

Explanation:

Question: 82

Refer to the exhibit.

```
router bgp 65520
  timers bgp 30 240
```

What is the effect of this configuration?

- A. It sets the keepalive timer to 30 milliseconds and the hold timer to 240 milliseconds.

B. It sets the keepalive timer to 30 seconds and the hold timer to 240 seconds.

C. It sets the hold timer to 30 milliseconds and the keepalive timer to 240 milliseconds.

D. It sets the hold timer to 30 seconds and the keepalive timer to 240 seconds.

Answer: B

Explanation:

Question: 83

DRAG DROP

Drag and drop the features from the left into the order of operations for SRv6 SSH field creation and forwarding on the right.

segments left		first
last entry		second
packet is forwarded		third
next header		fourth
routing type		fifth
segment list		sixth

Answer:

Explanation:

Features as




Table 1. Example of a Segment Routing Header

Next Header	Hdr Ext Len	Routing Type	Segments Left
Last Entry	Flags		Tag

Segment List[0]
(128-bit IPv6 address)

Segment List[n]

(128-bit IPv6 address)

Optional Type Length Value objects (variable)

Question: 84

Which two differences should be considered when deciding whether to implement routing policies or route maps? (Choose two.)

- A. Sequences are added after implementing a route map, but routing policies must be reconfigured when change is needed.
- B. Route maps are implemented using hierarchical policies, but routing policies must be implemented sequentially.
- C. Route maps require an explicit deny at the end of the sequence, but routing policies have an implicit deny at the end of the program.
- D. Route policies require sequence numbers, but route maps are implemented without sequencing.


E. Route maps are supported in Cisco IOS Software and routing policies are supported in Cisco IOS XR Software.

Answer: A, E

Explanation:

Question: 85

Refer to the exhibit.



All routers on this network have been configured with PIM-SM and R1 is the rendezvous point. However, when asymmetric routing is implemented to modify link usage, the network begins to drop certain multicast traffic. Which action corrects the problem?

- A. Place the routes affected by asymmetric routing in a VRF.
- B. Remove the asymmetric routing and use spanning tree to manage link usage.
- C. Add a static mroute for routes that are failing.
- D. Configure the routers to use PIM-DM instead of PIM-SM.

Answer: A

Explanation:

Question: 86

Refer to the exhibit.

RI

```
ip prefix-list ciacotast 6a<y 1*2.1a*.1.0/24 1* 22 rowtar Mpf 1  
arc* 5 tliter>llat praflil cucotatt out
```

An engineer applied the configuration on R1 to prevent network 192.168.1.0/24 from being propagated outside of area 5. After the change, users have reported they are not able to access any of the application servers that were working before. While checking the routing table of the peer router, the engineer notices R1 stopped propagating any routes outside area 5. Which action must be

taken to fix the problem?

- A. Configure an additional filter-list on R1 for inbound traffic to allow external router into area 5.

- B. Reconfigure the filter-list statement to apply in the inbound direction
- C. Change the prefix-list action to permit and add an explicit deny statement for network 192.168.1.0/24
- D. Add a permit statement for 0.0.0.0/0 to the end of the prefix-list 10 to override the implicit deny.

Answer: D

Explanation:

Question: 87

What is the role of segment routing mapping server?

- A. It advertises a local SID mapping policy to all the mapping clients.
- B. It works with IGP instances to calculate the prefix-SIDs in the absence of a mapping policy.
- C. It selects multiple mapping entries to create overlapping active mapping policies.
- D. It reads and translates remotely received SIDs from other mapping servers to create SID mapping entries.

Answer: A

Explanation:

E. A router that acts as a mapping server allows the user to configure SID mapping entries to specify the prefix-SIDs for some or all prefixes. This creates the local SID-mapping policy. The local SID-mapping policy contains non-overlapping SID-mapping entries. **The mapping server advertises the local SID-mapping policy to the mapping clients.**

Question: 88

Which difference should be considered when intradomain or interdomain multicast routing is implemented?


- A. Interdomain multicast routing requires an IS-IS or OSPF neighbor relationship between the domains but intradomain multicast routing requires only BGP
- B. Interdomain multicast routing relies on PIM-DM and intradomain multicast routing relies on PIM-SM.
- C. interdomain multicast routing uses BIDR-PIM to establish neighbor relationships between AS and interdomain multicast routing uses MSDP
- D. A network uses intradomain multicast routing without interdomain routing but networks that use interdomain multicast routing must also apply intradomain routing

Answer: D

Explanation:

Question: 89

Refer to the exhibit.



Routers RA and RB are IS-IS peers configured for NSF but router RC is an IS-IS peer without NSF capability. If RA undergoes processor switchover, what is the effect on the network environment?

- A. If RC is operating without the Cisco configuration option, all three routers tear down their peering relationships and re-establish peering.
- B. All peer relationships remain up and the link-state database is unchanged.
- C. All peer relationships remain up, but the link-state database is rebuilt on each device.
- D. If RC is operating without the Cisco configuration option, only 2 routers tear down their peering relationships and re-establish peering.

Answer: C

Explanation:

Question: 90

Refer to the exhibit.

■ IK* ip runtT Dipt

```

a LJI.114.1.0/24 U-O-LII VU i7J.lt.14.J, KiLMJD, Lilmi.^iU J
Q LA E1J-L48.J .4/14 [1LOFJ1] <ti iTJ.il.14.1< ttiUUI. CthainatIIJO
(I LA E12. L11.5.0U4 IJLD/I1J Vie :u.!».n.lj 00147117, L'.ter mill':
CI LA ktJ.IU.20.0FH [ilfl/21] via ITj.it.K.I, api4t[Di, EtbametOfO
a FA E1J. L4>.2L.D/H [110/21] Vil IU.it.14.1, nj e L1E3J, El rwnwtfl/(1
a ii in.118.22.0/14 [in/21] ?!■ in.it.it.I, gmidij r:ht™tP/Q
a u 1*2. LH.25.fl/2l [110/211 via 171.lt.£4.E, girlJtDI, I-Ji* 1 wtO/Q
a LA Ltl. L11.JJ.0/24 [114/2EH via 1U.H.E4.I. CnilFtU, OUaraatO/Q
a LA Hi.141.35.4/24 1110/21) fl* 171.11.14.1, DIIOItl', tll»EMI«/tl
a u E1J LM.54.0/14 (110/2l| via 1T1.II. 14.1, ODjQJrJI, IEharwtO/n
D 14 Eli. Lit. 13,0/24 [110/11] via L13.II.i4.il noiQJtH, EtbarnclO/D
a u 112. IM. 54.0/24 (L10/21J via 112.10.14.1, QaiOJiDI, ELbarutO/O
a TA EW. LM. 37.0/24 [LIIViEJ via 113.16. £4.1, MiftliM, E^barnetO/C
(I LA 1»».114.51.0/24 [110/21) via 172.11.i4.1, COiHill, EinaroaEO/D
a u LJJ 148.39.0/24 'JOUi; via 172.10.14.1, COsMtlI, E*.ba:wtO/D
  
```

An engineer applied the summarization configuration on R1 for four networks (192.168.20.0/24 to

192.168.23.0/24) in area 1 and eight networks (192.168.32.0/24 to 192.168.39.0/24) in area 2 to stop the flooding of all the customer routes. While checking the routing table of R2, the engineer noticed that R1 is still sending only specific routes to R2. Which configuration should the engineer apply on R1 to summarize routes?

R1 (config F router ospf 1

R1 (config-routerF area 1 range 192,168.20.0 255.255 252.0

R1 (config-router F area 2 range 192.168.32.0 255 255 240 0

R1 (config F router ospf 1

R1 (config-routerF area 1 range 192.168.20.0 255.255.248.0

R 1 (config-r outer F area 2 range 192.168.32.0 255 255 240 0

R1 (config F router ospf 1

R1 (config-router F area 1 range 192.168.20.0 255 255 252.0

R1 (config-routerF area 2 range 192.168.32.0 255 255 248 0

R1 (config F router ospf 1

R1 (config-routerF area 1 range 192.168.20.0 255.255.252.0

R1 (config-routerF area 2 range 192.168.32.0 255.255.253.0


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

Answer: C

Explanation:

Question: 91

Refer to the exhibit.



Traffic flow from router R1 to router R6 is delay-sensitive. It must consider potential link-failure and node-failure conditions. Which configuration must an engineer apply to router R1 to route traffic to router R6 if router R2 fails?

```
router ospf 1
area 1
interface GigabitEthernet0/0/1 fast-reroute per-prefix
fast-reroute per-prefix tiebreaker node-protecting index 100
fast-reroute per-prefix tiebreaker srlg-disjoint index 200
```

```
router ospf 1
area 1
interface GigabitEthernet0/0/1
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa
fast-reroute per-prefix tiebreaker node-protecting index 100
```

```
router ospf 1
area 1
interface GigabitEthernet0/0/1
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa
```

```
router ospf 1
area 1
interface GigabitEthernet0/0/1
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa
fast-reroute per-prefix tiebreaker srlg-disjoint index 100
```

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

Answer: B

Explanation:

Question: 92

An engineer is troubleshooting slow performance issues on a customer's network after the last multicast configuration change was applied on it. While checking the running configuration on the router, the engineer notices there are many ip igmp join-group commands applied on several interfaces of the router which caused the high CPU utilization usage. What action must the engineer take to solve this issue?


- A. Configure ip igmp static-group command on all interfaces
- B. Remove ip igmp join-group command on all unnecessary interfaces
- C. Configure all router interfaces to be process-switched by increasing the query interval
- D. Remove unnecessary members from the IGMP group

Answer: B

Explanation:

Question: 93

Refer to me exhibit.



A network engineer sets up a multihoming eBGP topology where multiple Autonomous Systems connect to ASN 64501. The engineer wants to block all the routers coming from ASN 64502, but allow all the others.

For that purpose, the following AS Path prefix 1st s being used:

```
(config) # ip as-path access-list 10 deny _64502$
```

What must be fixed to achieve this result?

- A. The AS-PATH filter must be defined inside the route-map mode
- B. At the end ip as-path access-list 10 permit .* must be included
- C. The statement must be modified with ip as-path access-list 1 deny _64502_
- D. The statement must CM modified with ip as path access list 1 deny A64502\$

Answer: B

Explanation:

Question: 94

Refer to the exhibit.

```
router isis
 net 49.0012.0000.0000.0002.00
 log-adjacency-changes
 summary-address 1.0.0.0 255.0.0.0
```

A network engineer applied configuration on R1 to summarize all ISIS routes, but R2 is still receiving specific routes from R1. The engineer has confirmed that both routers are configured with the correct summarization configuration, but R1 is not sending the correct summary routes. Which configuration must be applied to router R1 to summarize routes within Level 1?

- A. R1 (config-router)#summary-address 1.0.0.0 255.0.0.0 level-1-2
- B. R1(config-router)#no summary-address 1.0.0.0 255.0.0.0 level-1
- C. R1(config-router)#summary-address 1.0.0.0 255.0.0.0 level-1
- D. R1(config-router)#summary-address 1.0.0.0 255.0.0.0 level-2

Answer: C

Explanation:

Question: 95

Refer to the exhibit.

```
R1
interface g0/0
 ip address 192.168.1.1 255.255.255.0
 ip router isis
router isis
 net 49.0022.1111.1111.1111.00
 area-password ci5co

R2
interface g0/1
 ip address 192.168.1.2 255.255.255.0
 ip router isis
router isis
 net 49.0022.1111.1111.1112.00
 area-password ci5co
```

After these configurations were applied to routers R1 and R2, the two devices cannot form a

neighbor relationship What is the reason for the problem?

- A. The two routers have the same area ID.
- B. The two routers have different iS-types
- C. The two routers have the same network ID
- D. The two routers cannot authenticate with one another

Explanation:

Answer:

D

Question:

96

Refer to the exhibit.

Router 1:

```
router bgp 65515
 address-family ipv4 unicast
  bgp additional-paths receive
  bgp additional-paths select group-best
  neighbor 192.168.1.1 activate
  neighbor 192.168.1.1 additional-paths send receive
  neighbor 192.168.1.1 advertise additional-paths group-best
```

An engineer working for a private telecommunication company with an employee id 3977 74 814 implemented the configuration on Router 1. what is the effect of it?

- A. Router 1 sends and receives multiple best paths from neighbor 192.168.1.1
- B. Router 1 sends up to three paths to neighbor 192.168.1.1 for all routes
- C. Router 1 sends only one best path to neighbor 192.168.1.1.
- D. Router 1 receives only one best path from neighbor 192.168.1.1

Answer:

A

Explanation:

Question:

97

Refer to the exhibit.

```
R4#show ip pim interface
```

Address	Interface	Ver/ Mode	Nbr Count	Query Intvl	DR Prior	DR
10.1.1.1	FastEthernet0/0	v2/S	1	30	1	10.1.1.2

Which two commands must the engineer configure for the company's PIM-PIM network to enable Auto-RP mappings to be sent over the FastEthernet0/0 interface without affecting normal operation? (Choose two)


- A. enable auto-rp listener
- B. enable sparse-dense mode
- C. enable sparse-mode
- D. enable Auto-RP announcements
- E. enable dense mode

Answer: A B

Explanation:

Question: 98

Refer to the exhibit.



A network engineer applied configuration on R5 to summarize all OSPF routes, but R4 is still receiving specific routes from R5. The engineer has confirmed that both R5 and R4 routers are configured with correct summarization configuration but R5 is not sending the summary routes. What action must the engineer take to fix the problem?

- A. Remove summarization configuration on R5 and configure it on R4
- B. Configure a sham link between R4 and R5 to support summarization within Area 2
- C. Move R4 and R5 in separate areas as now they maintain the same link-state database

D. Clear link-state database on both R4 and R5 routers for summarization to work

Answer: C

Explanation:

Question: 99

Refer to the exhibit.

```
router1#show ip ospf interface serial 2
Serial1/0 is up, line protocol is up
  Internet Address 192.168.2.1/24, Area 0
  Process ID 1, Router ID 192.168.2.1, Network Type BROADCAST, Cost: 4
  Transit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 192.168.2.1, Interface address 192.168.2.1
  Backup Designated router (ID) 192.168.2.2, Interface address 192.168.2.2
  Timer Intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 9 Hello due in 00:00:07
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 192.168.2.2 (Backup Designated Router)
  Suppress hello for 0 neighbor(s)

router2#show ip ospf interface serial 1/0
Serial1/0 is up, line protocol is up
  Internet Address 192.168.2.2/24, Area 0
  Process ID 1, Router ID 192.168.2.2, Network Type POINTTOPOINT, Cost: 44
  Transit Delay is 1 sec, State POINT TO POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 9 Hello due in 00:00:03
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 192.168.2.1
  Suppress hello for 0 neighbor(s)
```

Router 1 and Router2 have shared routes in the OSPF database but the routes are missing from their routing tables. Checking the prefix-list configuration on both routers, the engineer confirmed all networks are allowed. What action should the engineer take to fix the problem?


- A. Configure the two routers with different process IDs
- B. Configure the two routers with different hello and dead timer values
- C. Switch the DR and BDR roles between the two routers
- D. Configure interface Serial1/0 on Router1 as a point-to-point interface

Answer: D

Explanation:

Question: 100

Refer to the exhibit.



A customer reports that Host-1 is failing to receive streaming traffic from the IPTV source. The engineer has confirmed that hosts on router R2 are receiving traffic normally and that Host-1 is correctly sending subscription messages to join the IPTV stream. Which action must the engineer take to correct the problem?

- A. Configure IP PIM SSM and IGMP version 2 under interface GigabitEthernet 1/0/1 on R3
- B. Configure IGMP version 3 under interface GigabitEthernet 1/0/1 on R3
- C. Remove IP PIM SSM and IGMP from interface GigabitEthernet 1/0/1 on R3 and configure under global configuration
- D. Remove IP PIM SSM from the global configuration on R3 and configure it under the GigabitEthernet 1/0/1 interface

Answer: B

Explanation:

Question: 101

Refer to the exhibit.

<pre>RI interface FastEthernet0/0 ip address 192.168.1.1 255.255.255.0 ip ospf authentication ip ospf authentication-key Cisco ip ospf 1 area 0 speed auto duplex auto</pre>	<pre>R2 interface FastEthernet0/0 ip address 192.168.1.2 255.255.255.0 ip ospf authentication-key Cisco ip ospf 1 area 0 speed auto duplex auto router ospf 1</pre>
--	--

```
router ospf 1
  log-adjacency-changes
  area 0 authentication
  message-digest
```

While applying the configurations on two routers an engineer notices that OSPF adjacency Between them remains down Through the ping test the engineer confirmed that both me routers have Layer 3 reachability between them Which action should me engineer take to make the adjacencies to full?

- A. Enter the command ip ospf authentication in R2 interface
- B. Enable OSPF just inside the router OSPF process not in the interfaces of any router
- C. Delete the area 0 authentication message-digest command from the OSPF process in R1
- D. Delete the area 0 authentication message-digest command from the OSPF process in R2

Answer: A

Explanation:

Question: 102

Refer to the exhibit.

```
ip route 0.0.0.0 0.0.0.0 192.168.0.1
router isis
 redistribute static
```

An administrator is troubleshooting Internet access issues on a customer's network. After applying this ISIS configuration to R1, the administrator notices that it fails to redistribute the default route into IS-IS. After checking the connectivity between the ISIS router and the ISP router the engineer confirmed there is Layer 3 connectivity between them Which action should be taken to correct the problem?

- A. Associate the default route with a VRF
- B. Add the default-information originate command to the configuration
- C. Configure the default route under any routing protocol other than IS-IS
- D. Configure R1 as a Layer 1 router

Answer:

B

Explanation:

Question: 103

Refer to the exhibit.

Router R2 is configured with the following commands:

```
router bgp 100
  neighbor 10.10.10.10 remote-as 200
  neighbor 10.10.10.10 weight 100
  neighbor 10.10.10.10 label 100
  neighbor 10.10.10.10 route-policy 100 out
  neighbor 10.10.10.10 route-policy 100 in
```

Router R3 is configured with the following commands:

```
router bgp 100
  neighbor 10.10.10.10 remote-as 200
  neighbor 10.10.10.10 weight 100
  neighbor 10.10.10.10 label 100
  neighbor 10.10.10.10 route-policy 100 out
  neighbor 10.10.10.10 route-policy 100 in
```

Router R4 is configured with the following commands:

```
router bgp 100
  neighbor 10.10.10.10 remote-as 200
  neighbor 10.10.10.10 weight 100
  neighbor 10.10.10.10 label 100
  neighbor 10.10.10.10 route-policy 100 out
  neighbor 10.10.10.10 route-policy 100 in
```

There is a connectivity issue between Customer-1 and Customer-2. File servers between the customers cannot send critical data. R3 routes are missing from the routing table on the Customer-1 router. All interfaces on Customer-1 are up. Which configuration must be applied to router R2 to correct the problem?

Router R2 is configured with the following commands:


```
router bgp 100
  neighbor 10.10.10.10 remote-as 200
  neighbor 10.10.10.10 weight 100
  neighbor 10.10.10.10 label 100
  neighbor 10.10.10.10 route-policy 100 out
  neighbor 10.10.10.10 route-policy 100 in
```

Router R3 is configured with the following commands:

```
router bgp 100
  neighbor 10.10.10.10 remote-as 200
  neighbor 10.10.10.10 weight 100
  neighbor 10.10.10.10 label 100
  neighbor 10.10.10.10 route-policy 100 out
  neighbor 10.10.10.10 route-policy 100 in
```

Router R4 is configured with the following commands:

```
router bgp 100
  neighbor 10.10.10.10 remote-as 200
  neighbor 10.10.10.10 weight 100
  neighbor 10.10.10.10 label 100
  neighbor 10.10.10.10 route-policy 100 out
  neighbor 10.10.10.10 route-policy 100 in
```



```

router bgp 1
^jdm>4iPTtfty vpvvt UH icon
allwatt-UbtE all

```

```

routtr bgp 1
vrf on#
rd 1:1
address-family =pv4 unicast
Jllouw-lab<I all

```

```

iwter bgp 1
neighbor
rvmcte-as 1
updatc-MUte# LoapbachO
a^dnn-family ipv^ mh^lad^nieim dEocaM-labtl ail

```

```

rawter h0p i
J(ldr><-I>m lv Ip¥4 uric Jit jHwatt-lab#: iN

```


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

Answer: D

Explanation:

Question: 104

Refer to the exhibit.



A network engineer configured routers R1 and R5 to run in IS-IS Level 1 mode and router R6 to run in

IS-IS Level 2 mode. All other routers are running as Level 1 / Level 2 routers. An engineer expects traffic from R1 to R6 to pass via R2, but IS-IS routing has calculated the best path via R4. Which action corrects the problem?

- A. Configure all routers as Level 1 routers.
- B. Remove the link metric for the link from router R1 to router R2.
- C. Change the link metric for the link from router R1 to router R2 to 1.
- D. Configure all routers as Level 1 / Level 2 routers.

Answer: C

Explanation:

The default link metric for IS-IS is 10, so changing the link from R1 to R2 to 1 will change the overall metric to R6 to 9 (8+1).

Question: 105

Refer to the exhibit.



The extension header of the IPv6 header is ignored when which value is equal to zero?

- A. Segments Left

- B. Hdr Ext Len
- C. Routing Type
- D. Next Header

Answer: D

Explanation:

Reference: <https://www.ciscopress.com/articles/article.asp?p=31948>

Question: 106

What are two features of dual-stack? (Choose two.)

- A. It is implemented without excessive processing power.
- B. It supports IPv4 and IPv6 traffic on the same infrastructure.
- C. It allows an engineer to stack multiple Layer 2 switches to serve as a single switching unit.
- D. It requires tunneling IPv6 traffic over an IPv4 infrastructure.
- E. It securely routes and tunnels traffic between separate sites over the Internet.

Answer: AB

Explanation:

Question: 107

Refer to the exhibit.

```
Router1# show bgp 1.1.1.1
BGP routing table entry for 1.1.1.1/32, version 10
Community: 64502 64503
```

```
Router2# show bgp 1.1.1.1
BGP routing table entry for 1.1.1.1/32, version 10
Community: 100:100 100:64102
```

An engineer working on a client's network is trying to solve the BGP community issue on R1 and R2 routers. After displaying the BGP entries, the engineer notices that both routers still have different outputs. Which action must the engineer perform to any of the routers to correct the problem and get the same output?

- A. Execute the `ip bgp-community new-format` command in Router1 .
- B. Remove the `ip bgp-community new-format` command in Router1 .
- C. Advertise the route to Router2 as standard community.
- D. Advertise the route to Router2 as extended community.

Answer: A

Explanation:

Reference: <https://www.noction.com/blog/understanding-bgp-communities>

Question: 108

When deploying a nationwide network of routers, what is the benefit of using BGP confederations?

- A. availability
- B. scalability
- C. security
- D. automatability

Answer: B

Explanation:

Reference: <https://www.routerfreak.com/bgp-network-design-bgp-confederation/>

Question: 109

A network operator working for a telecommunication company with an employee id: 4074:92:707 is planning to implement the Nonstop Forwarding (NSR) feature on the customer's core network. After getting the configuration ready for NSR, on which router should the operator implement NSR changes?

- A. on the CE router
- B. on the ASBR router
- C. on the ABR router
- D. on the PE router


Answer: D

Explanation:

Reference: https://www.cisco.com/en/US/technologies/tk869/tk769/technologies_white_paper0900aecd801dc5e2.html

Question: 110

Refer to the exhibit.



A network engineer has divided AS into confederations. Due to repeated ASN, when the 10.0.0.0/8 prefix from R1 arrives to R2, BGP automatically rejects it. What should the engineer do to fix the problem so that BGP allows that prefix on R2?

- A. Configure the command as-override on R1.
- B. Configure the command allowas-in on R2.
- C. Configure the command allowas-in on all the PE routers.
- D. Configure the command as-override on R2.

Answer: B

Explanation:

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/112236-allowas-in-bgp-config-example.html>

Question: 111

Refer to the exhibit.

```
R1#
interface LoopBack0 no ip address ipv6 address 2000:1::1/96
interface Loopback2 no ip address ipv6 address 2000:1::1/86 ipv6 router isis areal
```

```
R2#show isis ipv6 unicast rib
IS-IS IPv6 process areal, local RIB
```

* 2000:1::/96

via FE80::C801:4CFF:FE54:8/FastEthernet0/0, type L1 metric 10 tag 0 LSP [5/6]

via FE80::C801:4CFF:FE54:8/FastEthernet0/0, type L2 metric 10 tag 0 LSP [4/6]

2001:688:1001:1000::2/128

via FE80::C801:4CFF:FE54:8/FastEthernet0/0, type L2 metric 20 tag 0 LSP [4/6]

After configuring IS-IS on routers R1 and R2, an engineer notices that only the loopback interface at 2000:1::1/96 is known to router R2. Which change must be made so that only Loopback2 is advertised from R1 to R2?


- A. Configure the router isis areal command under the Loopback0 interface on R1.
- B. Remove the advertise passive-only command under the IS-IS address family ipv6 configuration.
- C. Remove the ipv6 router isis areal command under the Loopback2 interface on R1.
- D. Remove the passive-interface Loopback0 command under the router isis areal configuration.

Answer: B

Explanation:

Question: 112

Refer to the exhibit.



```

Cisco PE1P ip erf VPN Bank
rd 100:1 route-target
export 100:1 route-target
import 100:1

network 2.2.24 0.0.0.0 area 0
network 192.163.10 0 0.0 0.255
area 0
router bgp 100 no bgp default
ipva-unicast bgp log-neighbor-
changes neighbor 2.2X2 remote-
as 100 neighbor 2.2 2.2 update-
source loopback0

router ospf 2 erf VPN-Bank router-
id 24.2.2 log-adjacency-changes
redistribute bgp 100 subnets
network 10.0.0.0
0.255.255.255 area 0

router ospf 1 mpls ldp autoconfig
router-id 0 0.0.2 log-adjacency
changes

address-family vpnv4 neighbor
2,2.2.2 activate neighbor 2.2.2.2
send-community extended
exit-address-family

address-family ipv4 vrf VPN Bank
no synchronization eut-address-
family
  
```

Customer traffic from the branch site to the hub is experiencing packet drops. The engineer verified that:
 Customer traffic from the hub is able to reach the branch site.
 The connection between PE1 and PE2 is working normally.
 Routers P1 and P2 are able to ping devices on the hub site.

Which action resolves the issue?

- A. Redistribute the CE1 OSPF routes in the PE1 VRF.
- B. Redistribute the CE2 OSPF routes in the PE2 VRF.
- C. Advertise the VPNv4 routes of the branch site in CE2.
- D. Advertise the VPNv4 routes of the hub in CE1.

Answer: A

Explanation:

Question: 113

Which difference must an engineer consider when choosing whether to implement OSPF or IS-IS as the routing protocol on the network?

- A. The OSPF DR election process is deterministic, and the IS-IS DS election process is not deterministic.
- B. To support fast convergence. OSPF uses the LSA arrival timer and IS-IS uses the PRC-interval timer.
- C. IS-IS supports ordinary, stub, totally-stub, and NSSA areas, but OSPF supports only stub areas.
- D. IS-IS links are associated with either a Level 1 area or a Level 2 area at one time, but OSPF links are associated with multiple areas at one time.

Answer: B

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/15-s/irs-15-s-book/irs-fscnt.html

Question: 114

While configuring Cisco NSF awareness, a network engineer enters the bgp graceful-restart command after the BGP session is established in a router that runs IOS XE Software. Graceful restart capabilities are not exchanged. Which two actions should be taken? (Choose two.)

- A. Reload the router
- B. Verify that BGP route dampening is configured
- C. Reduce BGP convergence time
- D. Issue the clear ip bgp * command
- E. Issue the show ip bgp neighbors command.

Answer: AD

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-16/irg-xe-16-book/bgp-nonstop-forwarding-awareness.html

Question: 115

What is the difference between basic IS-IS and OSPF packet types?

- A. IS-IS and OSPF use area packets, but only IS-IS uses sequence number packets.
- B. IS-IS and OSPF use link-state update packets, but only OSPF uses link-state ACK packets.
- C. IS-IS and OSPF use Hello packets, but only OSPF uses DBD packets.
- D. IS-IS and OSPF use link-state update packets, but only IS-IS uses DBD packets.


Answer: C

Explanation:

Reference: <https://www.ciscopress.com/articles/article.asp?p=26850&seqNum=4>

Question: 116

Refer to the exhibit



Router R1 is configured to advertise outgoing SA messages to routers R2 and R3, but to receive incoming SA messages only from R2. Which additional configuration must an engineer apply to router R1 so it filters all MSDP SA messages from Domain-C?

- A. R1# ip msdp sa-filter in 192.168.3.1
- B. R3# ip msdp sa-filter out 192.168.3.1
- C. R3# ip msdp password peer 192.168.1.1
- D. R1# ip msdp password peer 192.168.3.1

Answer: A

Explanation:

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/ip-multicast/13717-49.html>

Question: 117

What is the difference between a source tree and a shared tree in a multicast environment?

- A. To route traffic from source to receiver a source tree uses a link-state routing protocol and a shared tree uses a distance-vector routing protocol.
- B. A source tree has its root at the source, and a shared tree has its root at a designated rendezvous point.
- C. To stream multicast from source to receiver, a source tree uses PIM-SM and a shared tree uses PIM-DM.
- D. Source trees are the default type for bidirectional PIM. and PIM-DM uses shared trees by default.


Answer: B

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti_pim/configuration/xr-16-5/imc-pim-xr-16-5-book/imc-tech-overview.html

Question: 118

Refer to the exhibit.



R3 is:

- failing to accept multicast RP information from Domain-A
- advertising MSDP SA messages to R1 and R4
- receiving SA messages only from R4

Which command must the engineer implement to resolve the issue?

- A R3# ip msdp sa-filter in 192.168.3.1
- B R3# no ip msdp sa-filter in 192.168.1.1
- C R3# no ip msdp peer 192.168.1.1
- D R3# ip msdp sa-filter out 192.168.1.1

Answer: B

Explanation:

Reference: <https://allthingsnetworking.wordpress.com/2014/06/27/pim-msdp-sa-filter/>

Question: 119

Refer to the exhibit.

```
R1

ip as-path access-list 10 permit ^65516$

router bgp 65515
  neighbor 192.168.1.2 remote-as 65516
  neighbor 192.168.1.2 route-map ciscotest in

route-map ciscotest permit 10
  match as-path 10
```

R1 is expected to receive routes originating from AS 65516 and from any ASs that are directly attached to it. However, R1 is receiving routes only from AS 65516. Which action corrects the configuration?


- A. Change the regular expression in the AS-path permit filter to .*.
- B. Change the regular expression in the AS-path permit filter to ^65516_[0-9]*\$.
- C. Add the regular expression ^\$. in the AS-path filter to permit the traffic from R2.
- D. Change the regular expression in the AS-path permit filter to _65516_.

Answer: B

Explanation:

Question: 120

Refer to the exhibit.



An engineer is troubleshooting IS-IS configuration between two areas. IS-IS Area-A network 30.1.1.0/24 is

leaked into IS-IS Area-B. R2 is failing to filter the route updates from network 10.1.1.0/24. Which configuration must the engineer apply to resolve the issue?

- A. R2(config)# ip prefix-list List2 seq 5 deny 10.1.1.0/24
R2(config)# interface fastethernet 0/0
R2(config-if)# ip router isis 100
R2(config-if)# router isis 100
R2(config-router)# distribute-list gateway List2 in
- B. R2(config)# ip prefix-list List1 seq 3 deny 10.1.1.0/24
R2(config)# ip prefix-list List1 seq 5 permit 30.1.1.0/24 ge 25 le
R2(config)# ip prefix-list List1 seq 10 permit 0.0.0.0/le 32
R2(config)# interface fastethernet 0/0
R2(config-if)# ip router isis 122
R2(config-if)# router isis 122
R2(config-router)# distribute-list prefix List1 in
- C. R1(config)# ip prefix-list List2 seq 5 deny 10.1.1.0/24
R1(config)# interface fastethernet 0/0
R1(config-if)# ip router isis 100
R1(config-if)# router isis 100
R1(config-router)# distribute-list gateway List2 in
R (config-if)# router isis 150
R1(config-router)# distribute-list route-map Map1 in
- D. R2(config)# access-list 101 deny ip any 10.1.1.0 0.0.0.127
R2(config)# access-list 101 permit ip any 30.1.1.0 0.0.0.63
R2(config)# access-list 101 deny ip any 0.0.0.0 0.0.0.0
R2(config)# interface fastethernet 0/0
R2(config-if)# ip router isis 121
R2(config-if)# router isis 121
R2(config-router)# distribute-list 101 in

Answer: C

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/15-mt/irs-15-mt-book/isis-inbound-filtering.html

Question: 121

Which keyword is used with the match route-type command to redistribute the external BGP and IGP routes using route map?

- A. match route-type type-1
B. match route-type nssa-external
C. match route-type type-2
D. match route-type external


Answer: D

Explanation:

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/49111-route-map-bestp.html>

Question: 122

Refer to the exhibit.



CE1 and CE2 are iBGP neighbors in AS 65516. All traffic that exits AS 65516 must use the link from CE1 to PE1. CE1 is advertising a higher local preference to CE2, but traffic from CE2 still prefers the PE2 link. Which action corrects the problem?

- A. Add the lower local-preference value on PE2 towards CE2.
- B. Configure CE1 to send routes to CE2 with a higher MED.
- C. Configure CE1 to send routes to CE2 with a higher weight.
- D. Add the next-hop self command to the CE1 neighbor statement for CE2.

Answer: D

Explanation:

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753->

Question: 123

Refer to the exhibit.

```
R1
router bgp 65530
  neighbor 192.168.1.2 remote-as 65531
  address-family ipv4
    neighbor 192.168.1.2 activate
    neighbor 192.168.1.2 route-map ciscotest in
  network 192.168.2.0 mask 255.255.255.0
  exit-address-family

route-map ciscotest permit 10
  set as-path prepend 65531 65531
```

Routers R1 and R2 reside in AS 65530, which is multihomed to the Internet. A network engineer expects devices in the AS to use R2 to access the Internet, but they are using R1 as the exit point from the AS. Which action corrects the problem?

- A. Add a sequence number to the route map to remove the implicit deny.
- B. Change the route map direction in the neighbor statement to out
- C. Configure the route map with a local preference of 200 or higher.
- D. Remove the neighbor statements from the address family configuration and activate the neighbor globally on R1.

Answer: B

Explanation:

Reference: <https://community.cisco.com/t5/networking-blogs/bgp-as-path-prepend-configuration/ba-p/3819334>

Question: 124

What is the role of a segment routing mapping server?

- A. It advertises a local SID mapping policy to all of the mapping clients
- B. It works with IGP instances to calculate the prefix-SIDs in the absence of a mapping policy
- C. It selects multiple mapping entries to create overlapping active mapping policies
- D. It reads and translates remotely received SIDs from other mapping servers to create SID mapping entries

Answer: A

Explanation:

Question: 125

Refer to the exhibit.


```

R1#ip explicit path name FRR_EXP_PATH enable
  next-address 10.10.23.3
  next-address 10.10.34.4
  next-address 10.10.45.5
  next-address 10.10.57.7
  next-address 7.7.7.7

interface Tunnel127
 ip unnumbered Loopback0
 tunnel destination 7.7.7.7
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autroute announce
 tunnel mpls traffic-eng path-option 10 explicit name
 FRR_EXP_PATH
 tunnel mpls traffic-eng path-option 20 dynamic

R1#tracert 8.8.8.8 source 100
Type escape sequence to abort.
Tracing the route to 8.8.8.8
 0 10.10.12.2 36 msec 24 msec 28 msec
 1 10.10.23.3 [MPLS: Labels 3011/7012 Exp 0] 80 msec 48 msec 40 msec
 2 10.10.34.4 [MPLS: Labels 4009/7012 Exp 0] 40 msec 32 msec 40 msec
 3 10.10.45.5 [MPLS: Labels 5010/7012 Exp 0] 48 msec 60 msec 28 msec
 4 10.10.78.7 [MPLS: Label 7012 Exp 0] 64 msec 60 msec 48 msec
 5 10.10.78.8 64 msec * 56 msec

```



An MPLS core network has connectivity issues R4 has failed. It impacts traffic loss between R1 and R8. Customers report no access to their file servers, which delays their transformation work. Which quick action resolves the issue until R4 recovers?

- A. Implement Link and Node protection on routers R2 and R7.
- B. Disable traffic engineering so that traffic prefers the IGP path
- C. Enable MPLS TE fast reroute on router R1 and Link and Node protection on router R2.
- D. Configure IBGP full mesh for faster convergence.

Answer: C

Explanation:

Question: 126

Refer to the exhibit.

```
PE1#show mpls ldp discovery
Local LDP Identifier:
  192.168.12.1:0
Discovery Sources:
Interfaces:
  Gi1/0 (ldp): xmit/recv
LDP Id: 192.168.1.1:0; no route
```

An administrator is troubleshooting network issues on a customer's network. PE1 cannot form an LDP relationship with PE2 using LDP ID 192.168.1.1. Due to this connectivity issue, the customer's routes behind both PE routers cannot be exchanged. Which action must the administrator take to correct the problem?

- A. Configure PE2 to automatically select its own LDP router ID with 192.168.1.1.
- B. Create a targeted LDP session on PE2 to override the missing routing table entry on PE1.
- C. Configure PE1 with a static route to 192.168.1.1.
- D. Enable LDP session protection on the link that PE1 uses as the next hop to PE2.

Answer: C


Explanation:

Question: 127

Refer to the exhibit.

```
R1#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, JU - 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
Gateway of last resort is not set
10.0.0.0/32 is subnetted, 1 subnets
C 10.0.0.0/32 is directly connected, Loopback0
2.0.0.0/32 is subnetted, 1 subnets
O 2.2.2.2 [110/11] via 10.0.0.2, 01:38:48, FastEthernet 0/0
3.0.0.0/32 is subnetted, 1 subnets
O 3.3.3.3 [110/21] via 10.0.0.2, 01:02:29, FastEthernet 0/0
10.0.0.0/24 is subnetted, 2 subnets
C 10.0.0.0 is directly connected, FastEthernet 0/0
O 10.0.1.0 [110/20] via 10.0.0.2, 01:02:39, FastEthernet 0/0
R1#sh ip bgp vrf REO
BGP table version is 9, local router id is 1.1.1.1
Status codes: s suppressed, d damped, h history, v valid,
> best, r RIB-failure, S Stale Origin codes: i - IGP, e - EGP, ? - incomplete
Network Next Hop Metric LocPrf Weight Path Route Distinguisher: 4:4 (default for vrf REO)
10.0.0.0/32 15.5.5.5/32 3.3.3.3 11
100 0 ? > 1192.168.2.0 3.3.3.3 0 100 0 ?
R4#sh ip route
4.0.0.0/32 is subnetted, 1 subnets
C 4.4.4.4 is directly connected, LoopbackB
```

C 192.168.1.0/24 is directly connected, FastEthernet 0/



An engineer is troubleshooting connectivity issues on the MPLS core network. A customer connected through R4 cannot reach the OSPF domain on R5. While checking the routing table of R1, the engineer cannot see all the routes from R3 and R5. Which task must the engineer perform so that R4 is able to reach R5?


- A. Enable OSPF peering and configure route redistribution between routers R4 and R1.
- B. Enable route filtering between routers R1 and R3.
- C. Enable MP-BGP peering on routers R1, R3, R4, and R5.
- D. Enable OSPF on the Area-0 routers and configure MP-BGP between routers R1 and R3.

Answer: A

Explanation:

Question: 128

Refer to the exhibit.




- A. 120
B. 230
C. 110
D. 10

Explanation:

Question: 129

Refer to the exhibit



- ```
A. Router2(config)# segment-routing
Router2(config-sr)# ldp mapping-server
Router2(config-sr-ms)# prefix-sid-map
```

```

Router2(config-sr-ms-map)#
Router2(config-sr-ms-map-af)# 2.2.2.2/32 500 range 4
B. Router2(config)# segment-routing
Router2(config-sr)# sr mapping-server
Router2(config-sr-ms)# ldp-sid-map
Router2(config-sr-ms-map)# address-family ipv4
Router2(config-sr-ms-map-af)# 10.1.1.1/32 500 range 50
C. Router2(config)# segment-routing
Router2(config-sr)# mapping-server
Router2(config-sr-ms)# prefix-sid-map
Router2(config-sr-ms-map)# address-family ipv4
Router2(config-sr-ms-map-af)# 10.1.1.1/32 500 range 50
D. Router2(config)# segment-routing
Router2(config-sr)# ldp mapping-server
Router2(config-sr-ms)# prefix-sid-map
Router2(config-sr-ms-map)# address-family ipv4
Router2(config-sr-ms-map-af)# 2.2.2.2/32 500 range 40

```


**Answer: C**

Explanation:

Reference: [https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/segment-routing/configuration/guide/b-seg-routing-cg-asr9k/b-seg-routing-cg-asr9k\\_chapter\\_01001.html](https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/segment-routing/configuration/guide/b-seg-routing-cg-asr9k/b-seg-routing-cg-asr9k_chapter_01001.html)

### Question: 130

Refer to the exhibit.



PE1 | show mpls forwarding-table

| Local Outgoing Prefix |       |       | Bytes Label Outgoing Next Hop |   |           |
|-----------------------|-------|-------|-------------------------------|---|-----------|
| Label                 | Label |       | or Tunnel Id Switched         |   | interface |
| 16                    | No    | Label | 172.16.1.1/32                 | 0 | drop      |
| 17                    | No    | Label | 192.168.12.12/32              | 0 | drop      |
| 20                    | No    | Label | 192.168.2.2/32                | 0 | drop      |
| 21                    | No    | Label | 10.1.212.0/24                 | 0 | drop      |
| 22                    | No    | Label | 10.1.211.0/24                 | 0 | drop      |
| 23                    | No    | Label | 192.168.11.11/32              | 0 | drop      |
| 24                    | No    | Label | 172.16.11.0/24                | 0 | drop      |
| 25                    | No    | Label | 172.16.14.0/24                | 0 | drop      |

```

PE2#show ip route 192.168.1.1
Routing entry for 192.168.1.0/24
Known via "bgp 100", distance 200, metric 0
Tag 1, type internal
Last update from 192.168.1.12 20:10:38 ago
Routing Descriptor Blocks:
 • 192.168.1.12, from 192.168.12.12, 20:10:38 ago
Route metric is 0, traffic share count is 1
AS Hops 5

```

```

PE1#show ip route 192.168.11.11
Routing entry for 192.168.11.11/32
Known via "ospf 100", distance 110, metric 2, type intra area
Last update from 10.1.111.11 on Gi0/1 00:04:34 ago
Routing Descriptor Blocks:
 • 10.1.111.11, from 192.168.11.11, 00:04:34 ago via GigabitEthernet0/1
Route metric is 2, traffic share count is 1

```

VPN users that are connected to PE routers are facing network issues. Traffic that originates from CE1 drops before reaching CE2. An engineer finds no outgoing traffic statistics on PE1 and PE2 routers toward CE devices and finds that the PE1 router is running the older software image. Which action must be implemented to resolve the issues?

- A. Enable LDP protocol on PE1 and PE2 routers.
- B. Advertise P1 router loopback on PE1 in OSPF.
- C. Enable CEF-based forwarding on PE1 router.
- D. Advertise PE2 router loopback on PE1 in OSPF.

**Answer: C**

Explanation:

Reference: [https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp\\_basic/configuration/xr-3s/mp\\_basic-xr-3s-book/mp-mpls-cisco-rtrs.html](https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_basic/configuration/xr-3s/mp_basic-xr-3s-book/mp-mpls-cisco-rtrs.html)

## Question: 131

What are the two characteristics of route reflectors? (Choose two.)


- A. If a router received an iBGP route with the originator-ID attribute set to its own router ID, the route is discarded.
- B. Routes received from nonclient peers are reflected to route reflector clients as well as nonclient peers.
- C. Routes received from nonclient peers are reflected to route reflector cluster as well as OSPF peers.
- D. If a route reflector receives a route with a cluster-list attribute containing a different cluster ID, the route is discarded.
- E. Routes received from a route reflector Client are reflected to other clients and nonclient peers.

**Answer: AE**

Explanation:

### Question: 132

Refer to the exhibit.



Refer to the exhibit After recent configuration changes to a customer's network, a network engineer notices that R2 cannot communicate with R3 Both FastEthernet interfaces on R2 and R3 are up and configured with the correct IP addresses MD5 password configured on R2 and R3 match with no issues What is the minimum change the engineer must make to enable R2 and R3 to communicate and fix the problem?


- A. Configure a loopback interface on R2 and assign it to area 0.
- B. Define area 2 as a NSSA on R2 and R3
- C. Configure virtual links between R1 and R3.
- D. Configure interface F0/0 on R1 and R2 to be in area 0.

**Answer: A**

Explanation:

### Question: 133

Refer to the exhibit.



Refer to the exhibit. Customers A and C are experiencing packet drops when connecting to the application server. While troubleshooting the problem, the network engineer confirms that the IS-IS Level-1/2 adjacency is up between routers A, B, C, and D, and both customers can communicate with the database server without packet loss. Which action must the engineer take to resolve the issue?


- A. Advertise the application server subnet in the router D IS-IS database.
- B. Advertise a static default route to the router B IS-IS database.
- C. Leak the 172.17.10/24 route in the IS-IS databases on routers A and C.
- D. Leak the customer A and customer C subnets in the router A IS-IS database.

**Answer: A**

Explanation:

### Question: 134

Refer to the exhibit.



TZ4 show ip Mdp sa-cache© rejected-SA det read-only <snip> 86854209.328, (10.80.148.101, 239.192.0.1), RP: 10.2.255.241, Peer: 10.1.255.241, Reason: rpf-fail -> learned from central site RT1 but not accepted (originated from site A RT1) 86854209.328, (10.88.84.101, 239.192.0.2), RP: 10.3.255.241, Peer: 10.1.255.241, Reason: rpf-fail -> learned from central site RT1 but not accepted (originated from site B RT1)

TZ» show ip rpf 10.1.255.241  
 RP: information for ? (10.1.255.241)  
 RPF interface: Vlan10  
 RPF neighbor: ? (10.111.254.9)  
 RPF route/aask: 10.1.255.241/32  
 RPF type: unicast (ospf 15)  
 Doing distance-preferred lookups across tables  
 RPF topology: ipv4< multicast base, originated from ipv4 unicast base

TZ» show ip route 10.1.255.241  
 Routing Table CEKT1  
 Routing entry for 10.1.255.241/32  
 Known via "ospf 15", distance 110, metric 3, type intra area Last update from 10.111.254.9 on Vlan10, ld22h ago  
 Routing Descriptor Blocks:  
 • 10.111.254.9, from 10.205.0.197, ld22h ago, via Vlan10 Route metric is 3, traffic share count is 1

TY4 sh ip msdp sa-cache  
 MSDP Source-Active Cache - 2 entries  
 (10.80.148.101, 239.192.0.1>, HP 10.2.255.241, AS ?,ld23h/00:05:42, Peer  
 10.2.255.241 -> learned from RT1 at site A (which xs 10.2.255.241)  
 (10.88.84.101, 239.192.0.2), RP 10.3.255.241, AS ?,ld21h/00:05x31, Peer  
 10.3.255.241 -> learned from RT1 at site B (which is 10.3.255.241)

TY# sh ip rpf 10.2.255.241  
 RPF information for ? (10.2.255.241)  
 RPF interfaces Ft>9/1.1035  
 RPF neighbors ? (10.111.252.86)  
 RPF route/mask: 10.2.255.241/32  
 RPF type: unicast (ospf 15)  
 Doing distance-preferred lookups across tables  
 RPF topology: ipv4 multicast base, originated from ipv4 unicast base

TY4 sh ip route 10.2.255.241  
 Routing Table: CU4T1  
 Routing entry for 10.2.255.241/32  
 Known via "ospf 15", distance 110, metric 150, type extern 2, forward metric 2  
 Last update free. 10.111.252.86 on FortyGigabltBthernet0/1.IOSS, 84:06:26 ago  
 Routing Descriptor Blocks:

\* 10.111.252.86, from 10.205.120.195, 04:06:26 ago, via

FortyGigabitEthernet9/1.1035  
Route metric is 150, traffic share count is 1

Refer to the exhibit. Multicast traffic destined from T1 and T2 routers to RP routers works well. A network engineer observes problems with multicast traffic flows between Site-A and Site-B. Site-A users fail to receive multicast stream on Site-B via RPTY site, while Site-B users fail to receive multicast stream on Site-A via RPTZ site. Which action must be implemented to resolve the issues?


- A. Establish MDSP peering with interface IP subnet.
- B. Configure Site-A and Site-B in 10.80.14804
- C. Allow the OSPF and MSDP packets on the firewall.
- D. Configure direct OSPF peering between Site-A and Ste-B

**Answer: C**

Explanation:

### Question: 135

Refer to the exhibit.



```

CE1
router bgp 65530
neighbor 192.168.1.2 remote-as 65531
neighbor 192.168.2.2 remote-as 65530
address-family ipv4
neighbor 192.168.1.2 activate
neighbor 192.168.1.2 route-map ciscotest in
neighbor 192.168.2.2 activate
network 192.168.2.0 mask 255.255.255.0 exit-address-family
route-map ciscotest permit 10
set local-preference 50

```

Refer to the exhibit. Routers CE1 and CE2 are in AS 65530, which is multihomed for Internet access.

An engineer expects inbound traffic to AS 65530 to arrive from PE1, but it is coming from PE2 instead. PE1 and PE2 routers are connected with CE routers through the same bandwidth. Which action must be taken to correct the problem?


- A. On router CE2, configure inbound routes from PE2 to CE2 with a local-preference value of 50 or greater.
- B. Configure router CE1 to prepend the AS path to routes it receives from PE1.
- C. Set the local-preference value on router CE1 to 100 or greater.
- D. On router PE1, change the origin for routes that are redistributed from CE1 to CE2.

**Answer: C**

Explanation:

### Question: 136

Refer to the exhibit.



Refer to the exhibit. Customer A is a small media company with two offices connected by a 512 Kbps line. Their NY office is connected to several external partners by static routes on router R3. VoIP services use VoIP codec G729 Users reported poor voice quality and slow data transfer between the offices A network engineer configured ip tcp header-compression iphc-format on R2 and R3 routers Which additional action must the engineer take to fix the issue?


- A. Configure the ip ospf area 0 command under Serial 1/1 interfaces on R2 and R3 to avoid routing loops
- B. Change the OSPF router ID on either router so that the router IDs are unique.
- C. Configure the summary-address 10.10.0.0 255.255.240.0 command on R3 to optimize OSPF communication
- D. Configure the BGP routing protocol between R2 and R3 to control route propagation.

**Answer: C**

Explanation:

### Question: 137

Refer to the exhibit.



Refer to the exhibit While troubleshooting a networking issue an engineer identified a suboptimal communication issue on route reflector RR2 In the current environment

Router A is a non-route-reflector client for RR1 and RR2

Routers D and E are directly connected iBGP peers.

Router F is not an iBGP peer of routers D and E

Which action resolves the issue?

- A. Disable BGP Client-to-Client reflection on router RR2.
- B. Enable next-hop-self for BGP peering on router C.
- C. Remove the route-reflector configuration on router RR2.
- D. Enable next-hop-self for BGP peering on router D.

**Answer: B**

Explanation:

### Question: 138

Refer to the exhibit.



Refer to the exhibit. An engineer working for a private telecommunication company with an employee id 4115 46 881 is enabling a segment routing solution with these requirements. A service provider is using the default range for prefix SID.

PE1 must allocate the first SID from the default range for the loopback address

PE1 and PE2 loopback SID allocation should have a minimum difference of 500.

Which configuration must be implemented to meet the requirements?

```
PE1(config-iss-if-af)# adjacency-sid absolute 16201
PE2(config-ists-if-ai> adjacency-sid absolute 16710
PE1(config-ists-if-al)# prefix-sid absolute 16001
PE2(C0nfig it-rf-3f)S prefix-sid index 610
PE1(config-ist$jf-af)# prefix-sid absolute 16201
PE2(config4Sis-if-af)# prefix-sid absolute 16710
PE2(Config-ii-if-a0» adjacency-sid absolute 16001
PE1(ConfigHSts-rf-af)# adjacency-sid index 610
```

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

**Answer: B**

Explanation:

**Question: 139**

What is a requirement of PIM-SM?

- A. It requires Cisco Express Forwarding to be enabled.
- B. It must be enabled on loopback interfaces only
- C. It requires OSPF to be configured on the network.
- D. It must use an RP

**Answer: D**

Explanation:

**Question: 140**

Refer to the exhibit.



Refer to the exhibit. CE1 and CE2 use connectivity over the service provider cloud to reach CE3. In the event of a link failure in the service provider cloud, which BGP feature relies on IGP convergence to quickly assist in the installation of a backup path?


- A. BGP graceful restart
- B. BGP route dampening
- C. BGP PIC core
- D. BGP confederations

**Answer: C**

Explanation:

**Question: 141**

Refer to the exhibit.



Refer to the exhibit Networks 172.16.31.0/24 and 172.16.30.0/24 are advertised in area 34, and network 172.22.10.0/24 is advertised in area 0. A recent security review discovered that users connected to routers R1 and R2 have been making unauthorized access to an application running on network 172.16.31.0/24. An engineer determined that routers R1 and R2 are receiving updates for network 172.16.31.0/24. Which action resolves the issue?


- A. Apply route filtering on routers R3 and R4.
- B. Apply route filtering on router R3 only.
- C. Apply route filtering on routers R1 and R2.
- D. Apply route filtering on router R4 only.

**Answer: D**

Explanation:

**Question: 142**

Refer to the exhibit.



Refer to the exhibit. Excessive routes are flooding from network 150.0.0.0 into AS100. Internet traffic between AS400 and AS300 is working normally. No route controlling mechanism is applied on incoming and outgoing traffic Which configuration resolves the issue?

```

R2#router bgp 200
neighbor 12.12.12.2 remote-as 100
neighbor 23.23.23.2 remote-as 300
neighbor 12.12.12.12 filter-list 1 out
ip as-path access-list 1 deny *400$
ip as-path access-list 1 permit .*

R2#router bgp 200
address-family ipv4 unicast
neighbor 12.12.12.2 remote-as 100
neighbor 12.12.12.2 activate
neighbor 12.12.12.2 route-map PREPEND out exit-address-family
exit
route-map PREPEND permit 10
set as-path prepend 100 100

R2#router bgp 200
neighbor 12.12.12.2 route-map FLOODING out
ip as-path access-list 1 permit M00.
route-map FLOODING permit 10
match as-path 1 set metric 50000

R1#router bgp 100
neighbor 12.12.12.1 remote-as 200
neighbor 12.12.12.1 route-map SET-LOCAL-PREF in
route-map SET-LOCAL-PREF permit 10
match ip address 2
set local-preference 700
route-map SET-LOCAL-PREF permit 20
access-list 2 permit 150.0.0.0 0.255.255.255
access-list 2 deny any

```

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

**Answer: A**

Explanation:

### Question: 143

An engineer wants to map a multicast IP address to a multicast MAC. How many bits are used to make the conversion?


- A. high-order 24 bits
- B. higher-order 23 bits
- C. low order 23 bits
- D. lower-order 24 bits

**Answer: C**

Explanation:

## Question: 144

Refer to the exhibit.



Refer to the exhibit. An engineer is troubleshooting an issue with this network and notices that prefixes from R3 are missing on the R1 routing table. Due to repeated ASN when the 10.0.0.0/8 prefix from R3 arrives at R1, BGP automatically rejects it. There is no prefix-list on R1 which blocks the traffic from R3. What should the engineer do to fix the problem so that BGP allows that prefix on R1?

- A. Configure R2 as a route reflector client of R1.
- B. Configure the allowas-in command on R1.
- C. Configure the next-hop-self command on R2.
- D. Configure identical confederation ASNs on R1 and R2.

**Answer: B**

Explanation:

## Question: 145

Refer to the exhibit.

```
R1#show route-map route-map filtering, permit, sequence 10 Match clauses:
ip address (access-lists): 1 Set clauses:
Policy routing catches: 0 packets, 0 bytes
route-map filtering, deny, sequence 20 Match clauses:
ip address (access-lists): 2 Set clauses:
Policy routing catches: 0 packets, 0 bytes route-map filtering, permit, sequence 30
Match clauses: Set clauses: Policy routing catches: 0 packets, 0 bytes
```

```
R1#show access-lists
Standard IP access list 1
10 permit 10.0.0.0, wildcard bits 0.0.0.255 (0 matches)
Standard IP access list 2
10 deny 10.0.1.0, wildcard bits 0.0.0.255 (1 match)
```

Refer to the exhibit. A network engineer configured the redistribute connected subnets route-map filtering command on R1 to redistribute connected interfaces to the OSPF process. The engineer also wants to filter out IP address 10.0.1.0/24, but the prefix still appears in the routing tables of the other routers on the network. Which action corrects the problem?

- A. Remove route-map sequence 30.
- B. Add a set statement to route-map sequence 20.

- C. Change the deny statement in access list 2 to permit
- D. Remove the subnets keyword from the redistribute connected subnets route-map filtering command.

**Answer: C**

Explanation:


### Question: 146

Refer to the exhibit.

```
PE1#sh run | sec router bgp router bgp 65000
no synchronization
bgp log-neighbor-changes
network 10.1.0.0 mask 255.255.255.0 neighbor 4.4.4.4 remote-as 65000 neighbor 4.4.4.4
update-source Loopback0 no auto-summary

PE2#sh run | sec router bgp router bgp 65000 no synchronization bgp log-neighbor-
changes network 10.2.0.0 mask 255.255.255.0 neighbor 1.1.1.1 remote-as 65000 neighbor
1.1.1.1 update-source Loopback0 no auto-summary

PE1#sh ip cef exact-route 10.1.0.1 10.2.0.1
10.1.0.1 -> 10.2.0.1 : FastEthernet0/0 (next hop 172.16.0.2)
```



Refer to the exhibit. Network connectivity between bank A and bank B has been lost. Users at bank A and bank B are able to successfully reach their directly connected PE routers. All routers in OSPF area 0 are correctly advertising and learning routing updates. Which action resolves the issue?


- A. Enable next-hop-self under the iBGP peering configuration on routers PE1 and PE2
- B. Configure the P routers to redistribute BGP routes within OSPF area 0.
- C. Configure router P1 to advertise the IP prefix of PE1.
- D. Configure MPLS with an end-to-end label-switched path on each router.

**Answer: D**

Explanation:

### Question: 147

Refer to the exhibit.



Refer to the exhibit A network engineer installed a new router (router 3) at the regional hub running MPLS services for scalability Router 3 is connected to the 10.44.4.0/24, 10.44.5.0/24, 10.44.6.0/24, and 10.44.7.0/24 subnets The new router has been configured for OSPF area 2, and it is advertising the four connected networks. The engineer noticed that the same networks are listed as interarea summary routes, and they are being flooded into each area on the area borders Which action resolves the issue?

- A. On router 3, configure an access list to filter the networks.
- B. On router 2, configure a route map to filter the networks.
- C. Under the OSPF configuration on router 3, add area 2 range 10.44.4.0 255.255.252.0.
- D. Under the OSPF configuration on router 2, add area 2 range 10.44.4.0 255.255.252.0.

**Answer: D**

Explanation:

### Question: 148

How does SRv6 function on the control plane?


- A. It enables SRH-capable nodes to terminate IPv6 packets at the network egress to carry the SRv6 locator.
- B. The Ingress node of the SR domain adds a uSID format IPv6 header to carry the SRv6 locator.
- C. The ingress node of the SR domain swaps the SRv6 header for the IPv6 header
- D. The egress node of SR domain imposes a new outer header

**Answer: B**

Explanation:

## Question: 149

Refer to the exhibit.



Refer to the exhibit An engineer is troubleshooting a networking issue with several symptoms

The shortest path from router R1 to R8 is underused and the longest path is overused

Traffic from 10.1.1 1 to 10 8 8 8 is routed on the longest path

Traffic between the R1 and the application server is experiencing packet drops and latency problems.

Which action resolves the issue?


- A. Configure route leaking for the IP address Of the application server on router R1.
- B. Increase the R1 to R2 link metric to 20.
- C. Configure a Level 2 IS-IS domain on router R1.
- D. Block the advertisement of the application server IP address to router R6.

**Answer: C**

Explanation:

## Question: 150

Refer to the exhibit.



Refer to the exhibit The branch office in area 10 is connected to HQ via Frame Relay uplinks with bandwidth constraints After a recent implementation of QoS on the R2 and R3 networks the system has been logging %SYS-2-MALLOCFAIL: Memory allocation of 65536 bytes failed from 0x224E321, alignment 0 messages To reduce traffic load and memory utilization on R2 and R3 the network engineer configured R1 to announce only one user subnet per location by issuing the summary address 192.168.0.0 255.255.248.0 command en R1 However, the engineer noticed that router R2 still has two routes and a summary address from HQ and R3 also has two routes from HQ Which two actions must the engineer take on R1 to fix the issue so that only one route is announced' (Choose two.)


- A. Configure the subnet range with the network command under the IS-IS process.
- B. Configure R1 as a Level 1 device.
- C. Redistribute both routes into the ISIS process.
- D. Configure a summary route for Level 1-Level 2 devices.
- E. Configure a summary route for Level 1 devices.

**Answer: D, E**

Explanation:

## Question: 151

Refer to the exhibit.



Refer to the exhibit. Routers R2, R3, R4, and R5 all reside in the same area, with R1 in a different area. R3 is overutilized, and the engineer wants to reduce its CPU load. The engineer configured router R4 to summarize routes that it receives from R5, but R3 is still receiving all of the R5 routes. Which action resolves the issue?


- A. Configure R3 in a new area
- B. Configure R2 as a Level 1 router
- C. Configure the summary routes on R5
- D. Configure R4 as a Level 1-Level 2 router

**Answer: A**

Explanation:

### Question: 152

Refer to the exhibit.



Refer to the exhibit There is a BGP traffic path issue between Customer-1 and Customer-2 Users from Customer-2 have reported file transfer issues High utilization on the path between both customers CAUSES many packet drops. Which configuration resolves the issue?

```

R1#neighbor 10.0.24.4 routemap LOCAL-PREF-1 BO In route-map LOCAL-PREF-
150 set local-preference 150
ip prefix-list 5-5-5-5 seq 5 permit 5.5.5.5/32
route-map LOCAL-PREF-150 permit 10
match ip address prefix-list 5-5-5-5
set local-preference 150

R4#router bgp 65504
neighbor 10.0.23.3 remote-as 65501
neighbor 10.0.23.3 filter-list 1 out
ip as-path access-list 1 deny *65505$
ip as-path access-list 1 permit .*

R4#router bgp 65504
address-family Ipv4 unicast
neighbor 10.0.23.3 remote-as 65501
neighbor 10.0.23.3 activate
neighbor 10.0.23.3 route-map PREPEND in exit-address-family
exit
route-map PREPEND permit 10
set as-path prepend 65506 65507

R1#neighbor 10.0.23.3 route-map LOCAL-PREF-150 out route-map LOCAL-PREF-
150 set local-preference 150

```

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

**Answer: B**

Explanation:

### Question: 153

What is a characteristic of a segment routing mapping server?

- A. It must be placed in the core of the network.
- B. It serves multiple VRFs.
- C. It must have an IGP adjacency.
- D. It applies SID mappings from one IGP instance to another IGP instance.

**Answer: C**

Explanation:

### Question: 154

Refer to the exhibit.

```

R1#sh run int tu222 interface Tunnel222
description R1 to R2 bandwidth 33000 ip
unnumbered Loopback0 load-interval 30 tunnel
destination 10.10.11.1 tunnel mode mpls traffic-
eng tunnel mpls traffic-eng autor tunnel mpls
traffic-eng priority 112 tunnel mpls traffic-eng
path-option 10 dynamic tunnel mpls traffic-eng
record-route

```

```

End
R1#show rsvp reservation

```

| ip To      | From       | Pro | DPort | Sport | Next Hop    | 1/F   | Fi Serv | BPS |
|------------|------------|-----|-------|-------|-------------|-------|---------|-----|
| 10.0.1.4   | 10.10.11.1 | 0   | 5542  | 203   | 10.0.1.4    |       | SE LOAD | 33M |
| 10.0.1.4   | 10.10.11.1 | 0   | 5543  | 35    | 10.0.1.4    |       | SE LOAD | 33M |
| 10.10.11.1 | 10.0.1.4   | 0   | 5543  | 1154  | 10.0.252.18 | Ge0/0 | SE LOAD | 33M |

```

R1#sh rsvp reservation

```

(TUMUID DeatinaUon! 10.10.11.1)

Otatui:  
 Maur up open down path- Wild,signalling- MW nguUtag proceeding path option 10  
 type dynamo (Bull for Setup, path night )

config Pamterr  
 Bandvldth: 33000 thpt (Olabal) Priority: 1 1 anility OiO/Oarm  
 Metric Type T1 (default)  
 Autotouter Mibled taekhow ill tabled Lotdahare- 1)000 bwbiaed auto-bw di tab led

RSVP Signalling info:  
 Ore 10.0.1.0, bit 10.10.11.1, Tun Id 122, Tua Initaataa 73

Dhorteit Umaatralnad Path Into:  
 path night: 3 (TSI  
 Explicit Hute 10.0.1.250 10.10.11.1

Hit lory;  
 Turjjel  
 Tine since created: 0 homa. 10 ainuUr  
 Tin BlBtt path change 1 kliwla\*. 33 lewnda

Current UP  
 Setup TiM 3 Minutes, 31 lecondo twining

HU>r UP  
 ID: path optian 10 |72( Mncval Trigger letup lined out

Refer to the exhibit. A network engineer is investigating a report of packet drops between the branch site and the central site.

The two sites are connected via OSPF and RSVP-TE tunnels.

Traffic from the central site to the branch site is passing normally.

Technicians at both sites successfully ping the loopback IP addresses on routers R1 and R2. Which configuration corrects the packet-drop problem?

- R1(Config-if)# interface Tunnel333
- R1(Config-if)# tunnel mpls traffic-eng bandwidth 33000
- R2(Config)# interface Tunnel333
- R2(Config-if)# tunnel mpls traffic-eng bandwidth 33000
- R2(Config)# interface Ge0/0
- R2(Config-if)# ip rsvp bandwidth 33000 3300

R1(Config)# interface Ge0/0  
R1(Config-ltr)# ip rsvp bandwidth 33000 3300

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

**Answer: D**

Explanation:

### Question: 155

Which difference must an engineer consider when Implementing Inter-domain and Intra-domain multicast routing on the network?

- A. Intra-domain routing allows the service provider to control incoming and outgoing multicast data streams on its network, but inter-domain routing limits the service provider's control.
- B. Intra-domain routing uses the PIM and MBGP protocols for multicast routing, but inter-domain routing must use PIM.SSM or MSDP.
- C. Intra-domain routing is dependent on the RP router within the same SP network, but inter-domain routing reduces the dependency on the other SP network.
- D. Inter-domain routing supports policy routing to connect different multicast domains using PIM.SM, but intra-domain routing supports policy routing using PIM-SM only within a single domain.

**Answer: C**

Explanation:

### Question: 156


DRAG DROP

Drag and drop the BGP attributes from the left into the order of route selection preference on the right.




**Answer:**

**Explanation:**



**Question:**

157 to the exhibit.



Refer to the exhibit. A network engineer configured three new PE routers to expand the network. The new routers run in the IS-IS routing protocol and reside in the data center in the same exchange as the existing routers. However, the network is now experiencing suboptimal routing. The Layer 2 configuration and VLANs are configured correctly to provide segregation between networks, but the Level 1 routes are not being converted to Level 2 routes. Which action resolves the issue?


- A. On Router 1, redistribute the routes into IGP.
- B. On Router 1, summarize internal routes between areas.
- C. On Router 2, redistribute the routes into IGP.
- D. On Router 2, summarize internal routes between areas.

**Answer: D**

Explanation:

### Question: 158

Refer to the exhibit.



Refer to the exhibit. A network engineer just replaced five routers on this OSPF network. When the routing protocol is brought up, R5 cannot reach routes that originate on R1. The engineer verified that all connected links have established neighbor relationships. R5 reaches routes originating on R3 and R4. Which action resolves the issue?

- A. Configure an OSPF virtual link to bridge Area 0 on routers R3 and R4.

B. Configure automatic neighbor discovery on R1 and R5.

C. Configure OSPF to have a contiguous Area 0.


D. Configure each link to be point-to-point.

**Answer: A**

Explanation:

## Question: 159

Refer to the exhibit.



Refer to the exhibit. A network engineer is configuring router R3 to handle multicast streams, but Host-2 cannot send subscriptions messages to the IPTV source. Which configuration must the engineer apply to router R3 so it passes the IPTV stream to Host-2?

```
R3# configure terminal
R3ir.onfg * ip multi«st-routing
R3iconf«g «interface gigabitetbamat 1/0/0
R3iconfig-it.* ip pim sparse-mode
R3iconCg-rf> Ip Igmp version 3
RJlclwif$ * interface grgabitetbemat 110/1
RSicontgutp ip pim sparse-mode
RSjConTig-iri# ip igmp version 3
RSrconfig-JW ip pim «m default
```

```
R3J# configure terminal
R3|conf>g l* no ip pim asm de fault
RSiconfg rinterface gigabitethemet 1/0/0
R3|confg-ll^ Ip pim sparse-mode
R3lconfg-lf)# Ip Igmp version 3
R3(contg-rt^ ip pint asm default
R3(conf>g f interface gigabitethomet 1/0/1
R^onfig-if A Ip pim spatse-mode
RJlconro-if^ Ip igmp version 3
RJlc-Mq-i'^ip pim atm default
```

```
R3(cMgi# interface gigs Met hornet 1/0/0
R3(conf>g-if^ ip pim sparse-mode
R3(config4fpi ip igmp version 3
R3(config)# interface gige bitether not 1/0/1
R3(wnfig-if jr ip pim sparse-mode
R3(config-<W Ip igmp version 3
```

```
R3(config)# Interface gigabitethernet 1/0/0
R3(config)# ip pim sparse-mode
R3(config)# Interface gigabitethernet 1/0/1
R3(config)# ip pim sparse-mode
R3(config)# ip igmp version 3
```

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

**Answer: D**

Explanation:

### Question: 160

What is the purpose of ACL type prefix set entries in RPL prefix sets?

- A. They hold IPv4 or IPv6 prefixes that do not match specifications.
- B. They hold IPv4 or IPv6 prefix match specifications.
- C. They hold IPv6 prefix match specifications.
- D. They hold IPv4 prefixes that do not match specifications.

**Answer: B**

Explanation:

### Question: 161

An ISP has an MPLS VPN-based network with 12 PE routers. How many peerings are required between the 12 routers if the engineer has not configured route reflectors?

- A. 60
- B. 66
- C. 78
- D. 84

**Answer: B**

Explanation:

### Question: 162

Refer to the exhibit.

```

UMhM ip hh>
B4I' UbM MMB106 LI IHtMl. iMBI NU«M ID 11 >19 ll.Ill.lt
Cutui coM« i ivvpM«iMi e eu^«e. n Micorj * *1111. > twit i inMtMl
0*1*11 rwM* 1- IS*. *- MP. 1 - Ir.rne*>#1»

 0*t*Uk IM** BOf M»tii« Loe Pi f Might Pith
*> u.n io.e/11 ila 11411 4 0 o Mti mu ua i
*> 11.22 11 0/74 171111154 0 0 1421 MtO 12 41 35473 154 7 3 4
*> 11.1) 15 0/24 101 lf 71 11 0 0 Mil IWO M41 W>) IN 1
«> ix.i4.ilt/14 17.1.112 TV 0 0 Mil 1)41 410 O1H »O1 11411 4
*> n n n »/M 10 C 11 * uo o MN iMi in two »en n««t i
*> 11.14.10.0/2J 11.14.111.7 sis 0 MN Mil lilt mt 1

```

Refer to the exhibit. Company A established BGP sessions with several ISPs. A network engineer at the company must filter out all traffic except for routes that transit AS 152. The engineer configured the filtering policy “permit \_152S (\_[0.9])” on R1, but after applying the configuration, the engineer notices that other routes are still visible. Which action resolves the issue?

Change the filtering policy to ip as-path access-list 1 permit \_152\_.

Add the following filtering policy in the configuration: ip access-list 1 permit \*152[M]\*)

Change the filtering policy to ip as-path access-list 1 permit \$152\*

Add a second filtering policy in the configuration: ip as-path access-list 1 permit \*152\*


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

**Answer: A**

Explanation:

## Question: 163

Refer to the exhibit.



```
R4 (config)# mpls label
protocol Ldp
R4 (config)# mpls Ldp router-
id loopback 0
R4 (config)# interface serial
0/0/1
R4 (config-if) # mpls-lp
R4 (config)# interface serial
0/0/2
```

```
R2 (config)# mpls label
protocol Ldp
R2 (config)# mpls Ldp router-
id loopback 0
R2 (config)# interface serial
0/0/1
R2 (config-if) # mpls-ip
```

```
T^Hcontigj^mpIsTa^Tourer^a'iiop
Dacir
```

```
R4 (config)# interface serial
0/0/1
R4 (config if) # mpls-ip
R4 (config)# interface serial
0/0/2
```

```
R2 (config)# mpls label
protocol Ldp
R2 (config)# mpls Ldp router-
id loopback 0
R2 (config)# Interface serial
0/0/1
R2 (config if) # mpls-ip
R2 (config)# interface serial
0/0/3
```

Refer to the exhibit. MPLS traffic from 192.168.4.0/24 to 192.168.5.0/24 is failing to pass over the link from R4 to R2. The engineer verified that:

Cisco Express Forwarding Is enabled on all routers.

All routers reach all networks via OSPF.

MPLS traffic from 192.168.1.0/24 to 192.168.3.0/24 is passing normally over the link from R1 to R2. Which action resolves the issue?

A. On router R4, remove the mpls ip command on the s/0/0/2 interface.


- B. On router R2, configure the mpls ip command on the SIO/O/4 interface.
- C. On router R4, configure the mpls label protocol ldp command on all serial interfaces.
- D. On router R2, configure the mpls label protocol ldp command on the interface.

**Answer: B**

Explanation:

### Question: 164

Refer to the exhibit.



Refer to the exhibit. A service provider has LDP and segment routing running in the network. Mobility traffic is carried through LDP and enterprise traffic is carried through segment routing. Which configuration must be implemented to connect the bank branch with the HUB site on routers?


- A. Configure segment-routing Sr-prefer prefix-list on AG1 and AG2 router for 10.10.10.10/24.
- B. Enable segment-routing Mpls Sr-prefer on CR1 and CR2 routers for 10.0.0.0/8.
- C. Enable segment-routing Mpls Sr-prefer on AG1 and AG2 routers for 10.0.0.0/8.
- D. Configure segment-routing sr-prefer prefix-list on CR1 and CR2 routers for 10.20.20.10/24.

**Answer: D**

Explanation:

### Question: 165

Refer to the exhibit.



```

R22#
*Apr 5 02:38:03.313: AS65111-R11: Received RP-Discovery, from 192.168.221.11, RP: 224.1.1.1
*Apr 5 02:38:03.319: AS65111-R11: Update C224.0.0.0/4, RP: 172.16.0.1, RPF: 172.16.0.1
R22(config-if)#
Apr 5 02:38:05.413: RPF: 172.16.0.1: RPF: 172.16.0.1: Received (, 224.1.1.1)
 Join from 0.0.0.0 for invalid RP 172.16.0.1

```

Refer to the exhibit. R21 is a multicast source sending multicast traffic 224.1.1.1 to R23, with R22 serving as the rendezvous point inside AS65222. A network engineer noticed that when R21 goes down, R12 in AS65111 starts to send the same multicast group 224.1.1.1 through R11. Which action resolves the issue?

- A. Block service groups 224.0.1.39 and 224.0.1.40 between the two autonomous systems
- B. Disable PIM parse mode between RP11 and RP22 in the two autonomous systems.
- C. Advertise RP2 with a high local preference in AS65222.
- D. Enable passive interface under EIGRP between the two autonomous systems.

**Answer: A**

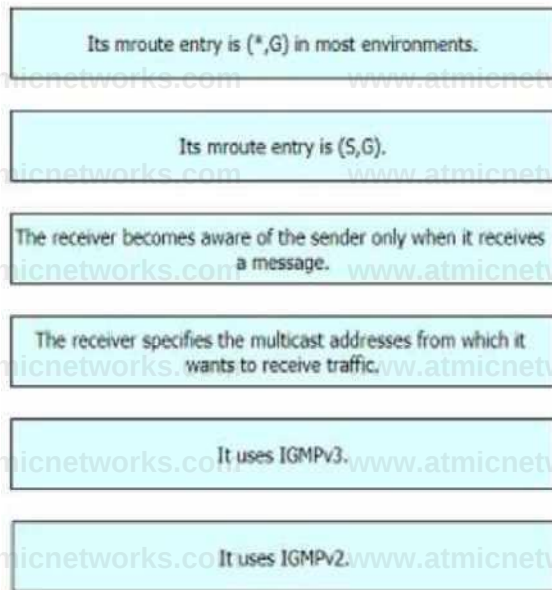
Explanation:

## Question: 166

DRAG DROP

Drag and drop the features about multicast from the left onto the multicast protocols on the right.


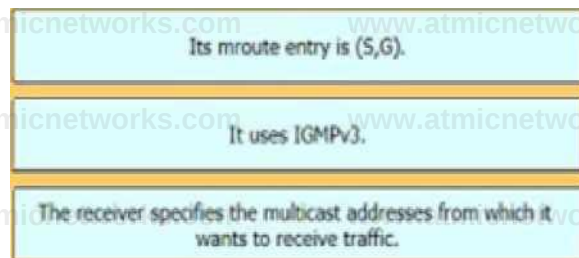
Not all options are used.



**Answer:**

**Explanation:**

SM



**Question: 167**

Refer to the exhibit.

PIMw rm | 9 iMUilliliitaM ntwlia  
aifibltlihirMtOJ \*wiUUw I\* M IF aiiitn » IC  
II JU III ),, 4 if rater LFU Con 1  
UtiilK, Mt\*tiid1Mrn»tO?)  
iaicilHla\* U uuct ir tttiH nr it it 1  
JU JU 115 t it h\*m ittt I mui tiki cm  
nt 4t 0113 wo: mi.mi.w u-qp\* M-|  
JogMtf)«ancr\*c\*iAgi 1  
L t AM run paata; :41s IOMMI  
Ulf CtW

U-lyps laval-I  
MO tl 0i3l.H4l OKI MU M  
Io# irt^aran^f tAingai  
UT\*r1»M uijAutTwrMtO,;  
addins'tanli? :jvi util cast

1  
Atarla^ aijshittisiM to, I »MsMi-fx Aily i>\*«  
nauut



Refer to the exhibit. Users at the branch office on R1 reported issue with an application at the home office on R4. While troubleshooting the issue, a network engineer determined that The branch-office users can connect to the home office.

The IS-IS adjacencies between R1 and R2 and R1 and the branch office are up.

Traffic from R1 to the R2 10.20.1.0/24 network is moving normally.

The application at the home office is experiencing packet drops on the connection to the Branch, and R3 cannot reach the R1 172.16.10.0/24 network.

Which action resolves the issues?


- A. Redistribute static connected routes in IS-IS on router R1.
- B. Configure the IS-IS core instance on the R1 GigabitEthernet0/3 interface.
- C. Redistribute static connected routes in IS-IS on router R4.
- D. Configure the IS-IS core instance on the R2 GigabitEthernet0/1 interface.

**Answer: B**

Explanation:

## Question: 168

Refer to the exhibit.



Refer to the exhibit. After a recent network implementation project, customer A is performing stress testing to verify network redundancy at the branch office connected to R4. When the link from R2 is shut down as shown, the SLA tracking object fails and the cost of the link between R2 and R4 increases to 100. However, a traceroute operation from a PC in the Branch office shows that traffic to HQ is still routed via R2. Which solution corrects the problem and optimizes traffic flow via R3 without creating operational overhead?

- A. Configure two OSPF processes on R2 and R3 and redistribute traffic between them.
- B. Redistribute routes from BGP to OSPF as type E1.
- C. Use multiarea adjacency to extend Area 10 to the link between R2 and R3.
- D. Create a virtual channel from R3 to R4.

**Answer: B**

Explanation:

### Question: 169

Refer to the exhibit.

```

Cisco(config)# extcommunity-set opaque overlay-color
Cisco(config-ext)# 1 co-flag 01
Cisco(config-ext)# end-set
Cisco(config)#
Cisco(config)# route-policy color
Cisco(config-rpl)# if destination in (10.10.10.1/32) then
Cisco(config-rpl-if)# set extcommunity color overlay-color
Cisco(config-rpl-if)# endif
Cisco(config-rpl)# pass
Cisco(config-rpl)# end-policy

```

Refer to the exhibit. An engineer is troubleshooting an issue with traffic steering using the color-only automated steering mechanism. BGP is failing to automatically steer traffic into an SR policy with the given color of a route, regardless of the next hop. The layer 2 configuration is correct, and the physical connection between the devices is working normally. Which additional command sequence must the engineer add to correct the issue?

```

Cisco?? configure
Cisco(config)# segment-routing
O>>>(config-si j# traffic-* ng
CSeconfig-sr-le)# policy FI
Ciscoconfig-er-te-poltyiF color 1 end-point

```

```

Ciscoconfigj# segment-routing
Ciscoconfig-st in traffic-eng
dsco(conng4r-le)^ policy PI
Oicoconfig-sMo-r-Noyii- color 1 end ipw 1,1.1 1
Cjscoicunng-st-le-t'dicy)# autoroule Include all

```

```

Ctecoi config)# segment-routing traffic-ong
Ctscoi config-3Mc)4 policy PI
Cisco(CDnng-sr-le-poiky)<< color 1 end-point <pv4 1.1.1.1
Cisco(conhg-5r-le-pciicy f# automate
Ctscoiconfig^r.le-poiicy-airtoroute# include ell

```

```

Ci>>>(CDnlig)d segment-routing
Ciscoi ainlig-51 in traffic-eng
Cbtmconfig-sr-le)# policy PI
Osco(crinng-sr-te-pci!cy)P color 1 end-point pv4 0.04.0

```

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

**Answer: D**

Explanation:

## Question: 170

Refer to the exhibit.



Refer to the exhibit. This network is deployed with all connected links configured to run IS-IS. The routing protocol is enacted globally on each router, and the network engineer expects full routing Information to be shared among all routers. R5 is receiving routes from R4 but is missing routes from R1. Which action corrects the issue so that all routes are shared among the routers?

- A. Configure all routers to reside in the same area.
- B. Configure R1 and R5 as Level 1 and Level 2.
- C. Configure R3 as a Level 2 router.
- D. Configure R5 in the same area as R1.

**Answer: A**

Explanation:

### Question: 171

Refer to the exhibit.


```

*Jan 26 12:18:15.035: %BGP-3-NOTIFICATION: sent to neighbor 10.0.0.8 2/2
(peer in wrong AS) 2 bytes 000A
RMS FFFF FFFF FFFF FFFF FFFF FFFF FFFF FFFF 0035 0104 000A 02B4 0A0A 0A03
1802 0601 0400 0100 0102 0280 0002 0202 0002 0641 0400 0000 0A

RX#
*Jan 26 12:17:57.423: %BGP-3-NOTIFICATION: received from neighbor 10.0.0.8
2/2 (peer in wrong AS) 2 bytes 000A
RX#
*Jan 26 12:18:00.679: %BGP-3-NOTIFICATION: received from neighbor 10.0.0.8
2/2 (peer in wrong AS) 2 bytes 000A

R#show ip bgp
Network Next Hop Metric LocPrf Weight Path
*> 10.0.0.8/30 20.0.0.2 0 0 10 ?
*> 10.0.0.28/30 20.0.0.2 2 0 10 ?
*> 10.0.0.29/30 20.0.0.2 0 0 10 ?
*> 10.10.10.3/32 20.0.0.2 0 0 10 ?
*> 10.10.10.4/32 20.0.0.2 2 0 10 ?
*> 192.168.30.1/32 20.0.0.2 2 0 10 ?
R#

```



Refer to the exhibit. An engineer working for a private telecommunication company with an employee id: 4233:46:364 notices that the customer network going through AS30-AS65001-AS65000 is experiencing packet drops when it accesses an application at 172.16.20.1/32 in the DC cloud. The BGP link between AS20 and AS30 is inaccessible because of a fiber cut. Routers RL, RN, and RZ are configured with confederation identifier 10. Which action resolves this issue?

- RX#config sprouter bgp 65001
- RX#config-fObtep no synchronization
- RX#ftcnfig^ouiarp bgp confederation identifier 10
- RB#corrtilgprouter bgp 30
- RB: config +lOLlerp neighbor 10,10.10.3 remote-as 65001
- RX#conflg pouter bgp 65001
- RX(CQofig-router)4i bgp confederation peers 65000
- RM(wnflg prowler bgp 65000
- RMiConfig-iouterp bgp confederation peers 65001

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

**Answer: A**

Explanation:

### Question: 172

Which type of BGP attribute does a route reflector attach to routes learned from iBGP peers that allows them to be accepted by other iBGP peers, thereby eliminating the need for a full-mesh BGP topology?

- A. well-known mandatory
- B. optional transitive
- C. well-known discretionary
- D. optional non transitive

**Answer: D**

Explanation:

### Question: 173

Refer to the exhibit.


Guidelines Topology Tasks

R1 R2

Configure and verify an OSPF neighbor adjacency between R1 and R2 in OSPF area 0 according to the topology to achieve these goals:

1. R1 pings the LoopbackO interface of R2. Use interface-level configuration to complete this task,
2. R2 pings the LoopbackO interface of R1. Use interface-level configuration to complete this task.
3. R2 receives a single summary route 172.16.100.0/22 for networks 172.16.100.0/24, 172.16.101.0/24, and 172.16.103.0/24.

BI>0



R1 R2

```
Current configuration : 1302 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
```

```
!
no aaa new-model
```

```
!
clock timezone PST -8 0 mmi polling-interval 60 no mmi auto-configure no mmi pvc
—More—
```

```
interface Loopback0
 ip address 10.1.1.1 255.255.255.255

interface Loopback10
 ip address 172.16.100.1 255.255.255.0
 ip ospf 10 area 1 r

interface Loopback20
 ip address 172.16.101.1 255.255.255.0
 ip ospf 10 area 1 t

interface Loopback30
 ip address 172.16.103.1 255.255.255.0
 ip ospf 10 area 1
```

R1 R2

```
interface*Loopback10
 ip address 172.16.100.1 255.255.255.0 ip ospf 10 area 1
t
interface Loopback20
 ip address 172.16.101.1 255.255.255.0 ip ospf 10 area 1
```

```
interface Loopback30
 ip address 172.16.103.1 255.255.255.0 ip ospf 10 area 1
t
```

```
interface Ethernet0/0
 ip address 172.16.0.1 255.255.255.0
 ip ospf 10 area 0 duplex auto i
```

```
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 i
```

```
router ospf 10 router-id 10.1.1.1
```

R2

R1 R

```
R2> R2> R2> R2> R2>en R2lsh run
Building configuration...
```

```
Current configuration : 1059 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 no mmi pvc
--More--
```

```
1
!
t
interface Loopback0
ip address 10.2.2.2 255.255.255.255
```

```
interface Ethernet0/0
ip address 172.16.0.2 255.255.255.0
ip ospf 10 area 0
duplex auto
```

```
interface Ethernet0/1
no ip address shutdown duplex auto
```

```
interface Ethernet0/2 no ip address
```

R1 ^

```
interface Ethernet0/0
ip address 172.16.0.2 255.255.255.0 ip ospf 10 area 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 ospf 10 router-id 10.2.2.2
```

```
ip forward-protocol nd l
```

```
no ip http server
no ip http secure-server
```

```
ipv6 loam timestamp t
t j
control-plane
```

**Answer: See  
the solution  
below in  
Explanation.**

Explanation:

Solution:-

R1

```
Int loopback0
Ip ospf 10 area 0
```

```
Int loopback10
Ip ospf network point-to-point
```

```
Int loopback20
Ip ospf network point-to-point
```

```
Int loopback30
Ip ospf network point-to-point
```

```
Router ospf 10
Area 1 range 172.16.100.0 255.255.252.0
```

Copy run start

Verification: -

```
R2lsh ip route ospf
```

```
Codes: L - local, 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, 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, I - next hop override, p - overrides from P
```

```
ER
```

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

```
10.0.0.0/32 is subnetted, 2 subnets
```

```
0 10.1.1.1 (110/111 via 172.16.0.1, 00:02:25, Ethernet0/0
```

```
172.16.0.0/16 is variably subnetted, 3 subnets, 3 masks
```


```
0 IA 172.16.100.0/225(110/11] via 172.16.0.1, 00:00:50, Ethernet0/0
```

```
R2f|
```

## Question: 174

Refer to the exhibit.

| Guidelines | Topology | Tasks |
|------------|----------|-------|
|------------|----------|-------|



## Guidelines Topology Tasks

Troubleshoot and configure BGP according to the topology to achieve these goals:

1. R1 and R3 establishes IBGP connectivity using Loopback addresses.  
The updates should come from LoopbackO.
2. R3 should be able to ping loopbackO interface of R2. These changes must be accomplished through BGP.
3. R1 advertises only the summary route of 172.16.100.0/22 to R2 and R3.

- Submit feedback about this item

**Answer: See the solution below in Explanation.**

### Explanation:

Solution :-

R1

Router bgp 100

Neigh 10.3.3.3 remote-as 100

Neigh 10.3.3.3 update-source loopback0

Address-family ipv4

Neigh 10.3.3.3 next-hop-self

Aggregate-address 172.16.100.0 255.255.252.0 summary-only

Copy run start

R3

Router bgp 100

Neigh 10.1.1.1 remote-as 100

Neigh 10.1.1.1 update-source loopback 0

Copy run start

Verification:-

```
R3|ping 10.2.2.2
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 10.2.2.2, timeout is 2 seconds: 11111111
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1 /1/1 ms
```

```
R3i|
```

```
R3fshow ip route
```

```
Codes: L - local, 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, LI - IS-IS level-1, L2 - IS-IS level-2
```

```
ia - IS-IS inter area, * - candidate default, U - peruser static route
```

```
o - ODR, P - periodic downloaded static route, H - NHR
```

```
P, 1 - LISP
```

```
a - application route
```

```
+ - replicated route, % - next hop override, p - overrides from PfR
```

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

```
1.1.1.1/32 is subnetted, 3 subnets
```

```
S 10.1.1.1 [1/0] via 172.20.2.1
```

```
B 10.2.2.2 [200/0] via 10.1.1.1, 00:00:19
```

```
C 10.3.3.3 is directly connected, Loopback0
```

```
172.16.0.0/22 is subnetted, 1 subnets
```

```
B 172.16.100.0 [200/0] via 10.1.1.1, 00:00:02
```

```
172.20.0.0/16 is variably subnetted, 3 subnets, 2 masks
```

```
B 172.20.1.0/24 [200/0] via 10.1.1.1, 00:00:19
```

```
C 172.20.2.0/24 is directly connected, Ethernet0/1
```

```
L 172.20.2.3/32 is directly connected, Ethernet0/1
```

## Question: 175

### Guidelines

This is a lab item in which tasks will be performed on virtual devices.

Refer to the Tasks tab to view the tasks for this lab item.

Refer to the Topology tab to access the device console(s) and perform the tasks.

Console access is available for all required devices by clicking the device icon or using the tab(s) above the console window.

All necessary preconfigurations have been applied.

Do not change the enable password or hostname for any device.

Save your configurations to NVRAM before moving to the next item.

Click Next at the bottom of the screen to submit this lab and move to the next question.

When Next is clicked, the lab doses and cannot be reopened.

### Topology

## OSPF Process ID 10

### Area 0



Lo0:  
172.16.100.0/24  
Lo20:  
172.16.101.0/24  
Lo30:  
172.16.103.0/24

### Tasks

Configure and verify an OSPF neighbor adjacency between R1 and R2 in OSPF area 0 according to the topology to achieve these goals:

1. R1 pings the Loopback0 interface of R2. Use interface-level configuration to complete this task.
2. R2 pings the Loopback0 interface of R1. Use interface-level configuration to complete this task.
3. R2 receives a single summary route 172.16.100.0/22 for networks 172.16.100.0/24, 172.16.101.0/24, and 172.16.103.0/24.

**Answer: See  
explanation below.**

Explanation:

R1 R2

```
R1(config) t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#int lo
R1(config)#int lo0
R1(config-if)#ip ospf 10 area 0
R1(config-if)#exit
R1(config)#int lo 10
R1(config-if)#ip ospf nei
R1(config-if)#ip ospf net
R1(config-if)#ip ospf network point-to-po
R1(config-if)#ip ospf network point-to-point
R1(config-if)#int lo 20
R1(config-if)#ip ospf net po
R1(config-if)#ip ospf net point-to-po
R1(config-if)#ip ospf net point-to-point
R1(config-if)#exit
R1(config)#int lo
R1(config)#int lo30
R1(config-if)#ip ospf netwo poi
R1(config-if)#ip ospf netwo point-to-po
R1(config-if)#ip ospf netwo point-to-point
R1(config-if)#
R1(config-if)#exit
R1(config)#int et
R1(config)#int ethernet 0/0
R1(config-if)#ip ospf 10 area 0
R1 (conf ig-if) Hexit
R1(config)#router ospf 10
R1(config-router)#area 1 rang 172.16.100.0 255.255.252.0
R1(config-router)#exit
R1(config)#
R1#
R1#
R1#copy r
♦Aug 26 11:44:29.078: %SYS-5-CONFIG_I: Configured from console by console
R1#copy run start
Destination filename [startup-config]?
Building configuration...
[OK]
R1#
R1#
```

R1 R2

```
R2>
R2>en
```

```
R2config t
```

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

```
R2(config)#int lo
```

```
R2(config)#int lo0
```

```
R2(config-if)#ip ospf 10 area 0
```

```
R2(config-if)#^Z
```

```
R2#
```

```
R2#
```

```
R2# c
```

```
♦Aug 26 11:44:48.122: %SYS 5 CONFIG_I: Configured from console by console
```

```
R2#copy run start
```

```
R2#copy run startup-config
```

```
Destination filename (startup-config)?
```

```
Building configuration...
```

```
[OK]
```

```
R2#
```

```
R2#sh ip route ospf
```

```
Codes: L - local, 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, H - NHRP, 1 - LISP
```


```
a - application route
```

```
+ - replicated route, % - next hop override, p - overrides from PfR
```

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

## Question: 176

Refer to the exhibit.



Refer to the exhibit. Routers R1, R2, R3, and R4 are peer routers that reside in different administrative domains. PIM-SM is running in each autonomous system, and EBGP is configured between the peers. A network administrator has just implemented MSDP between the connected peers. When the administrator enabled the MSDP configuration, R1 and R2 failed to establish a peering relationship. All other connected routers successfully established peering sessions. Which action must the engineer take to resolve the issue between R1 and R2?

- A. Change the MSDP authentication method to clear text
- B. Configure the peers to be in the same autonomous system.
- C. Set the same MSDP password on both peers
- D. Implement BGP authentication between the peers.

**Answer: C**

Explanation:

### Question: 177

Which IPv6 prefix format defines the destination of a dynamic 6to4 tunnel?

- A. 2001:db8::/32
- B. 2002::/16
- C. fe80::/10
- D. ff00::/8

**Answer: B**

Explanation:

### Question: 178

Refer to the exhibit.



```

PE1
router bgp 100
 bgp log-neighbor-changes
 network 10.10.10.1
 mask 255.255.255.255 neighbor 10.10.20.1
 remote-as 100 neighbor 10.10.20.1 update-
 source Loopback0 neighbor 10.10.20.1 send-
 label

RR1
router bgp 100
 bgp log-neighbor-changes
 neighbor 10.10.30.1 remote-as 100
 neighbor 10.10.30.1 update-source Loopback0
 neighbor 10.10.30.1 next-hop-self all
 neighbor 10.10.10.1 remote-as 100
 neighbor 10.10.10.1 update-source Loopback0
 neighbor 10.10.10.1 route-reflector-client
 neighbor 10.10.10.1 next-hop-self all
 neighbor 10.10.10.1 send-label

```

```

RR2
router bgp 100
 bgp log-neighbor-changes
 neighbor 10.10.20.1 remote-as 100
 neighbor 10.10.20.1 update-source Loopback0
 neighbor 10.10.20.1 next-hop-self all neighbor
 10.10.20.1 send-label neighbor 10.10.40.1
 remote-as 100
 neighbor 10.10.40.1 update-source Loopback0
 neighbor 10.10.40.1 route-reflector-client
 neighbor 10.10.40.1 next-hop-self all neighbor
 10.10.40.1 send-label

PE2
router bgp 100
 bgp log-neighbor-changes
 network 10.10.40.1
 mask 255.255.255.255 neighbor 10.10.30.1
 remote-as 100 neighbor 10.10.30.1 update-
 source Loopback0 neighbor 10.10.30.1 send-
 label

```

Refer to the exhibit. A network engineer is investigating a report of packet drops in an application running on a server connected to PE2. The engineer determined that:

The OSPF adjacency in area 0 is up, and it is learning the loopback addresses of all routers in area 0.

Traffic from users connected to PE1 to the application is also passing normally.

Packets from PE2 back to PE1 are being dropped.

Which action resolves the issue?

☐ RR2(config# router bgp 100

RR2(config-router)# no neighbor 10.10.20.1 send-label

☐ RR1 (config# router bgp 100

RR1 (config-router)# redistribute ospf 100

☐ RR1 (config# router bgp 100

RR1 (config-router)# neighbor 10.10.30.1 send-label

☐ RR2 (config# router bgp 100

RR2 (config-router)# redistribute ospf 100


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

**Answer: C**

Explanation:

### Question: 179

Refer to the exhibit.



Refer to the exhibit. HostB receives an IPTV traffic stream from the multicast source at 192.168.10.1. When router RB is powered down, it impacts multicast traffic forwarding to downstream peers. The RA and RB routers reside in the same physical location. Which two configurations must be implemented so that HostB receives the multicast stream even when RB is powered down? (Choose two )

- A. Enable PIM sparse mode on routers RC and RZ with a static RP address of 10.10.10.1.
- B. Set the backup RP address to 10.10.10.3 on routers RY and RZ.
- C. Enable bi-directional PIM on all routers.
- D. Enable PIM SSM mode on routers RA and RY.
- E. Set the RP address to 10.10.10.4 on all routers

**Answer: A, D**

Explanation:

### Question: 180

Refer to the exhibit.

Refer to the exhibit. Users connected to an application on 1.1.1.1 via routers CE1 and CE2 are complaining of frequent disconnections. While troubleshooting the issue, a network engineer with

an employee ID: 4333:96:916 determined that the application itself is working normally, but the link connected to the 1.1.1.1 application is flapping. The users on CE1 and CE2 are accessing applications on 2.2.2.2 without difficulty. Which action resolves the issue?

- A. Disable split horizon on router R2.
- B. Redistribute CE1 routes on router R2.
- C. Enable passive interface on the router R3 Loopback1 interface.
- D. Enable route dampening on routers R1 and R2

**Answer: D**

Explanation:

### Question: 182

The operator must leverage automated steering and coloring with the design. Which two elements from an SR Policy tuple are used for automated steering inside of the network? (Choose two.)

- A. endpoint
- B. candidate labels
- C. color
- D. CM paths
- E. headend

**Answer: A, E**

Explanation:

### Question: 183

Which feature provides service independence in service provider networks?

- A. binding SID that is used to stitch different domains
- B. prefix SID that contains router hops toward the destination
- C. adjacency SID that represents a path between the routers
- D. node SID that identifies the network device

**Answer: D**

Explanation:

### Question: 184

Refer to the exhibit.

```
RP/0/0/CPU0:router# configure
RP/0/0/CPU0:router(config)# mpls ldp
RP/0/0/CPU0:router(config-ldp)# router-id 1.1.1.1
RP/0/0/CPU0:router(config-ldp)# discovery hello holdtime 30
RP/0/0/CPU0:router(config-ldp)# discovery targeted-hello holdtime 180
```

Refer to the exhibit. A network consultant is updating the company's MPLS solution and configuring several new PE devices to enable full communication among PE devices connected to the network. The updates must

ensure the redundancy and scalability of the network The egress LSR is configured with mpls ldp explicit-null While configuring the given PE router to be adjacent to the egress LSR. the consultant notices that MPLS forwarding is not working in the network Which action resolves the issue?

- A. Configure MVRF on all routers in the MPLS network
- B. Open TCP and UDP ports 464
- C. Open TCP and UDP ports 179
- D. Configure Cisco Express Forwarding on all routers in the MPLS network

**Answer: D**

Explanation:

### Question: 185

Which lag command is used within route-map to redistribute routes that match with routes tagged with the specified tag numbers?

- A. match routing tag
- B. match all tag
- C. match route tag
- D. match tag

**Answer: D**

Explanation:

### Question: 186

DRAG DROP

Drag and drop the characteristics from the left onto the corresponding routing protocols on the right.

|                                                                 |       |
|-----------------------------------------------------------------|-------|
| Backbone routers reside in any area.                            | OSPF  |
| Routing between areas is distributed through the backbone area. |       |
| It operates on the network layer.                               |       |
| It uses the system ID to identify routers on the network.       | IS-IS |
|                                                                 |       |

**Answer:**

Explanation:

OSPF

Routing between areas is distributed through the backbone area.

It operates on the network layer

IS-IS

Backbone routers reside in any area.

It uses the system ID to identify routers on the network

### Question: 187

A network engineer recently deployed a wireless network for users at a new office to join video conferences wirelessly. The existing network consists of routers that are configured for EIGRP and BGP, which are located at several data centers in different physical locations. The engineer configured PIM-SSM within a single domain from a single source. However, during the testing phase, the video conference failed to work for any of the users. The IPs are reachable and firewalls in the network are configured correctly to allow the traffic.

Which action resolves the issue?


- A. Configure PIM dense mode to replace PIM-SSM.
- B. Configure PIM dense mode with multidomain multicast.
- C. Configure the devices to send join messages toward a multicast SPT.
- D. Configure PIM sparse mode to replace PIM-SSM.

**Answer: C**

Explanation:

### Question: 188

Refer to the exhibit.



Refer to the exhibit The engineering team noticed route disruptions when DSL subscriber 172.16.20.10 goes offline. In this service provider environment:

The OSPF backbone area is configured to advertise loopback prefixes

The PE routers are running BGP-IPv4 address family in a BGP-free core topology.

The DSL subscriber IP subnet 172.16.20.10/32 is redistributed in BGP on PE1

Which configuration on PE1 resolves the issue?

```

PE1(config)# router bgp 100
PE 1 (config-router)# network 172.16.20.0 255.255.255.255
PE 1 (Config)# Router ospf 100
PE1(Config-router)# redistribute bgp 100

```

```

PE1(Config)# ip route 172.16.20.0 255.255.255.0 Null0
PE 1 (Config)# Router BGP 100
PE1(Config-router)# redistribute static

```

```

PE1(Config)# ip route 172.16.20.0 255.255.255.255 Null0
PE1 (Config) # Router ospf 100
PE 1 (Config-router)* redistribute static

PE1(config)# router ospf 100
PE 1 (config-router)# network 172.16.20.0 0.0.0.255 area 0
PE 1(Con fig)# Router BGP 100
PE 1(Config-router)# redistribute ospf 100

```

A. Option A

B. Option B C. Option C D. Option D

**Answer: B**

Explanation:

### Question: 189

Which technique should be permitted on ASA Firewall to allow MSDP Peering?


- A. PIMv2 messages carrying MSDP messages
- B. TCP socket(s) carrying MSDP messages
- C. Multicast BGP carrying MSDP messages
- D. UDP socket(s) carrying MSDP messages

**Answer: B**

Explanation:

### Question: 190

Refer to the exhibit.



Refer to the exhibit. An engineering team implemented IS-IS on the network, with several different areas. Recently, router R6 in area 3 has been experiencing excessive CPU usage. To reduce the load, a network engineer implemented route summarization on R2. However, R6 is still receiving the full routes from all routers. Which action must the engineer take to resolve the excessive CPU usage?


- A. Configure routers R3 and R5 to send summary routes to R6.
- B. Configure the Level 2 router R6 to manage summarization for the network.
- C. Configure router R1 to send only the default route to R2 and R4.
- D. Configure all routers as Level 1/2 routers so that route summarization functions throughout the network.

**Answer: A**

Explanation:

**Question: 191**

Refer to the exhibit.



Refer to the exhibit. Mid-sized company Z connected two branch offices via a multicast-enabled ISP using the BGP routing protocol. PIM was implemented to support multicast streaming between the branches via MSDP. Client A cannot connect to the multicast stream source of company Z. The network engineer ran a debug on the edge of the network as shown. Which action resolves the issue?

- A. Implement MBGP and enable a multicast address family on R2
- B. Configure the PIM protocol on the Serial 1/1 interface and enable Auto-RP
- C. Deploy BSR on routers R1 and R2 and enable multicast on both LAN segments.
- D. Enable the BGP routing protocol and advertise the 10.10.0.0/24 subnet on R1.

**Answer: C**

Explanation:

### Question: 192

What is the difference between RPL and route-maps?


- A. Unlike RPL, route-maps are supported with IPv4 address families only
- B. Unlike route-maps, only RPL supports extcommunity sets to group extended communities.
- C. Unlike RPL, route-maps are designed for BGP protocol only.
- D. Unlike route-maps, only RPL is designed for bulk configuration of multiple devices.

**Answer: B**

Explanation:

### Question: 193

Refer to the exhibit.



Refer to the exhibit. An engineer is troubleshooting a network routing issue on this IS-IS network situated at the customer's regional hub. Router 2 that runs Cisco IOS XR Software cannot see the network 10.1.1.0/24 on router 1. The Layer 2 encapsulations are correct. The ARP connectivity for the Ethernet interface that runs IS-IS is working. Which action resolves the issue?

- A. On router 1, remove the CISCOLIVE keyword after the Ip router Isis command under interface GigabitEthernet0/2.
- B. Under the router isis command on router 1. change is-type level-1 to is-type level-2
- C. On router 1. add the CISCOLIVE keyword after the Ip router isis command under interface GigabitEthernet0/3.
- D. Under the router isis command on router 2. change is-type level-1 to is-type level-2

**Answer: C**

Explanation:

### Question: 194

What is the function of BGP confederations?

- A. They require routers to use private AS numbers.
- B. They eliminate the need for a fully-meshed internal BGP design.
- C. They enable a single device to distribute a common configuration to its peers
- D. They assign an external BGP router to serve as the route reflector

**Answer: B**

Explanation:

### Question: 195

Refer to the exhibit.

```

Cisco(config)* segment-routing
Cisco(config-sr)I traffic-eng

```

```
Cisco(config-sr-te) ♦ policy P1
Cisco(config-sr-te-policy) I color 10 end ipv4 1.1.1.1
```

Refer to the exhibit. The Cisco router has a single uplink to the upstream network, which is critical to the business. The router is configured for traffic steering with the Autoroute mechanism configured to steer all OSPF and IS-IS prefixes over non-shortest paths and divert traffic for those prefixes to the SR-TE policy. However, the Autoroute Announce functionality for all prefixes is not working. Which additional command set must the engineer configure on the router to correct this issue?

```
Cisco(config-srM; traffic-eng
Cisco! config-sr-te># policy P1
Cisco(config-sr-te-policy)<s color 1 end-point ipv4 0.0.0.0
```

```
Cisco(con fig-sr)»traffic-eng
Cisco! config-sr-le Ms policy P1
Ciscofconhg-sr-le-policyM* color 1 end-point ipv4 255.255.255.255
```

```
C Cisco! config-srMs traffic-eng
Ciscoconfig-sr-te^olicy P1
Cisco(config-sr-te-p<^cy M* autoroute all
Cisco(config-sr-te-policy-path)# preference 10
```

```
Cisco! config-srM» traffic-eng
Cisco! config-sr-le Ms policy P1
Ciscofconfig-sr-te-policy M* autoroute Include all
Cisco(config-sr-te*policy MS candidate-paths
Ciscoiconfig-sr-te-policy-path MS preference 100
```


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

**Answer: D**

Explanation:

### Question: 196

Refer to the exhibit.



**R1#show ip bgp 192.168.30.1**

**BGP routing table entry for 192.168.30.1/32, version 121**

**Paths: (1 available, best \*1, table Default-IP-Routing-Table) Not advertised to any peer 65201**

**10.10.10.5 (metric 12) from 10.10.10.5 (10.10.10.5)**

**Origin IGP, Mtric 0, localpref 100, valid, internal, best**

**R2#show ip bgp 192.168.30.1**

**BGP routing table entry for 192.168.30.1/32, version 9**

**Paths: (2 available, best 12, table DeFault-XP-Routing-Table) Advertised to update-groups:**

**65201**

**10.10.10.3 from 10.10.10.3 (10.10.10.3)**

**Origin IGP, localpref 100, valid, external 65201**

**10.10.10.5 (aetric 11) froa 10.10.10.5 (10.10.10.5)**

**Origin IGP, aetric 0, localpref 200, valid, internal, best**

Refer to the exhibit. The service provider operations team was alerted that hub site traffic from BGP AS 65101 to AS 65201 uses a non-primary path via the R5-R6 link. IBGP peering between R1 and R2 is up, and no fiber failure has been reported on the R2-R3 link. The team determined that the traffic flow between 10.10.10.1 and 192.168.30.1 is not considering the R1-R2-R3-R4 path. Which action resolves this issue?


- A. Change the metric attribute to 20 with neighbor 10.10.10.5 on R1.
- B. Change the local-preference attribute to 50 with neighbor 10.10.10.5 on R1.
- C. Set the local-preference to the default value with BGP neighbor 10.10.10.3 on R2 for receiving routes.
- D. Add metric 5 with BGP neighbor 10.10.10.6 on R5 for receiving routes.

**Answer: B**

Explanation:

### Question: 197

Refer to the exhibit.



```
*Oct 15 09:54:00.907: »UP-S MW«W LDP Neighbor 10.10.10.5:0 (?) is UP KX4
*Oct 15 09:55:21.247: tWe«S«MK90: LDP Neighbor 10.10.10.5:0 (?) is DQMN (Discovery Nello Bold Yiner expired)
```

```
«♦
*Oct 15 09:55:21.247: tUDP-S-NBICSU: LDP Neighbor 10.10.10.5:0 (?) is DOON (Discovery Italic Bold Tiner expired)
KLOshov npi# Ldp neighbor
Feer LDP I dent 10.10.10.20; Local WP I dent 192.17.1.1:0 TCP connection: 1010.102545 - 192.151.1.1.45472 State: <^er.* M*MJ« seet/revd: M/5; Dcenstteer Up Line: 00:20:5€ LDP
discovery sourceci
Gigabit/therat2/0, Bro IP addr: 10.0*0.? Addresses bound to peer LDP Ident*
10.10.10.2 10.0.0.5 10.0.0.2
interface FasethernetO/O ip address 10.0.0.14 255.255.255.252 (topics full f«pls 10.0.0.25
Ldp discovery transport a«Mress 10.10.10.4 flpls label protocol Ldp apls Ip end
```

s4i23SiS

Refer to the exhibit MPLS traffic from the packet core that is connected to router RX (10 10 10 1) to the radio network that is connected to router RB (10 10 10 6) uses the longer route RB R2 RY RX RX and R2 are configured with a default value of LDP hello and hold intervals The team verified that OSPF neighborship in Area 0 is established between RX and R2 Which action must be taken to choose the shortest path and resolve the issue?

A)

**RX(config)#mpls ldp holdtime 30**

B)

**RZ(config)#interface FastEthernet0/0  
R7(config-if)#no mpls ldp discovery transport-address 10.10.10.6**

C)

**R7(config-if)#interface FastEthernet0/0  
R2(config-if)#mpls ldp discovery transport-address 10.10.10.1**

D)

**RX(config)#mpls ldp neighbor 10.10.10.6 targeted ldp**

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


**Answer: B**

Explanation:

### Question: 198

#### DRAG DROP

Drag and drop the features from the left onto the corresponding routing protocols on the right




**Answer:**

Explanation:

OSPF: Fully supports NBMA and point to multipoint connections, entire routing table refreshed every 30 min  
ISIS: Provides limited support for NBMA and point to multipoint connections, elects designated router only

### Question: 199

Refer to the exhibit.



The UserA device fails to receive IPTV streaming traffic from the source. While debugging the issue, a network engineer finds that routers RB and RC receive the multicast traffic normally. RC receives subscription messages from UserA to join the IPTV stream. Which action must the engineer take on router RA to correct this problem?

- A. Configure ip pim ssm with ip igmp version 2 on interface GigabitEthernet 1/0.
- B. Remove the ip pim ssm default configuration and replace it with ip pim rp-address 10.10.1.3
- C. Change ip igmp version 2 to ip igmp version 3 on interface GigabitEthernet 1/0.
- D. Configure ip pim bidir-enable in global mode and add ip pim rp-address 10.10.10.3


**Answer: C**

Explanation:

## Question: 200

DRAG DROP

Drag and drop the route-manipulation methods from the left onto their features on the right.



**Answer:**

Explanation:

As-path filter – It uses regular expressions to filter traffic


Route-map – It matches packets with specified criteria and performs configured actions against them

Distribute list – It filters routes into and out of routing protocols

Offset-list – It converts incoming and outgoing metrics on RIP and EIGRP route

## Question: 201

Refer to the exhibit.



Refer to the exhibit. A network engineer commissioned Router 1, which connects to Router 2 and Router 5. Router 1 participates in IS-IS routing with other routers, but it is not used for transit traffic. Traffic from Router 2 to Router 5 is expected to travel via Router 1 but the traffic is traveling via Router 3 and Router 4. Which action resolves this problem?


- A. Add hello padding on Router 1
- B. Remove the hello padding on Router 5
- C. Add the overload bit on Router 5
- D. Remove the overload bit on Router 1.

**Answer: D**

Explanation:

## Question: 202

Refer to the exhibit.



Refer to the exhibit. Router 1 resides in a different data center than Router 2 and Router 3. Router 1 is

running the IS-IS routing protocol, and it is preconfigured with LDP and segment routing. Which configuration

must the engineer apply under router isis 111 on Router 1 to implement TI-LFA on the link to Router 2 without disrupting traffic to Router 3?

A)

**mttrfxt GigabitEthernetWOrt  
address-family Ipv4 unicast  
fast-reroute per-prefix**

B)

**interface GigabitEthernetftOd  
fast-reroute per-prefix tMfa  
enable**

C)

**interface GigabitEtherneKWI  
fast-reroute prefix enable**

D)

**interface GigabitEthernetO'Od  
address-family ipv4 unicast  
fast-reroute per-prefix  
fast-reroute per-prefix ti-lfa**

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

**Answer: D**

Explanation:

### Question: 203

Refer to the exhibit.



Refer to the exhibit. Unified MPLS is implemented on this network to provide end-to-end services A network engineer noticed that R3 is failing to receive labels from R1 However, R2 is receiving labels from R1 R3 can ping R1 normally Which action resolves the issue?

A. Add the cdp run command to the link configuration between R1 and R3.

- B. Add the send-label keyword in the R1 neighbor statement for R3.
- C. Configure R1 to run Unified MPLS on both interfaces.
- D. Disable mpls ip on the link between R1 and R3

**Answer: B**

Explanation:

## Question: 204

Refer to the exhibit.



Refer to the exhibit. Customer A and customer B are connecting through a service provider network with these requirements

R1 acts as the mapping server to connect 10.0.0.0/30 and 172.16.0.0/30 networks

R1 and R2 are using nondefault SRG8 range to allocate SID.

R2 acts as the mapping client for transient traffic originating from customer A Which configuration meets the requirements?

- A. Remove allocation of prefix SIDs for 10.0.0.0/30 and 172.16.0.0/30 on R1.
- B. Configure segment-routing prefix-sid-map advertise-local on R1
- C. Allocate prefix SIDs for 10.0.0.0-30 and 172.16.0.0/30 subnets on R2
- D. Configure segment-routing prefix-sid-map advertise-local on R2.

**Answer: B**

Explanation:

## Question: 205

What is a feature of SR-PCE?

- A. It learns the topology information via OSPF or IS-IS only
- B. It learns the topology information via OSPF or IS-IS or through BGP-LS.
- C. It uses the IGP metric in its path calculations to optimize reachability.
- D. It uses the TE metric in its calculations to optimize the cumulative TE metric

**Answer: B**

Explanation:

### Question: 206

Level 3 switch SW1 is part of an EIGRP enabled network located at the edge of an are

a. Cisco Express Forwarding and SSO are enabled on the router. A network engineer must minimize the downtime of data-transmit services in the network. HSRP group 2 has already been configured with standby preempt delay minimum 2 and standby timers 140. Which additional action must the engineer take on SW1 to suppress routing flaps for 2 minutes?

- A. Configure timers standby 2 wait 2 under the HSRP group 2 configuration.
- B. Configure timers nsf route-hold 120 under the global EIGRP configuration.
- C. Configure timers nsf signal 2 under the global EIGRP IPv4 address family.
- D. Configure timers standby 2 coverage 120 under the HSRP group 2 configuration.

**Answer: B**

Explanation:

### Question: 207

Which two BGP route reflector attributes should be implemented to prevent routing loops? (Choose two)

- A. CLUSTER\_ID
- B. ORIGINATOR\_ID
- C. CLUSTER\_LIST
- D. COMMUNITY
- E. AS\_PATH

**Answer: B, C**

Explanation:

### Question: 208

What is a feature of SRGB label values?

- A. They are used for adjacency segment identifiers.
- B. They allow IS-IS and OSPF to use overlapping ranges on a domain.
- C. They are locally significant.
- D. They are assigned to segment routing nodes and have global significance throughout the domain.

**Answer: D**

Explanation:

## Question: 209

Refer to the exhibit.

```
show ip bgp vpnv4 all neighbors
```

```
<omitted>
```

Local Policy Denied	Prefixes:	Outbound	Inbound
ORIGINATOR loop:		n/a	DWES

Refer to the exhibit. Which duplicate configuration causes this message during setup of an MPLS BGP network?


- A. ASNs
- B. originator ID
- C. router ID
- D. cluster ID

**Answer: B**

Explanation:

## Question: 210

Refer to the exhibit.



Refer to the exhibit. The initial BGP configuration of the network is complete, and all the routers have established BGP neighborship. The engineer configured Router 2 with additional BGP paths to install and advertise more paths to Router 1 and Router 3. Which configuration must the engineer apply to Router 1 and Router 3 so that they will use and install all additional paths?

```
router bgp 65005
```

address-family ipv4  
bgp additional-paths select all

☐ router bgp 65005

address-family ipv4  
neighbor 2.2.2.2 additional-paths receive  
bgp additional-paths select all  
bgp additional-paths all

☐ router bgp 65005

address-family ipv4  
bgp additional-paths select all  
bgp additional-paths install

router bgp 65005

address-family ipv4  
neighbor 2.2.2.2 additional-paths receive bgp additional-paths select all bgp additional-paths  
Install

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

**Answer: D**

Explanation:

### Question: 211

What is the difference between RPL and route map actions?

- A. RPL uses match, pass, set, and drop actions, and route maps use match, set, and drop actions.
- B. RPL uses pass, set, and drop actions only, and route maps use match, set, and drop actions.
- C. RPL uses pass, set, drop, and done actions, and route maps use match and set actions only.
- D. RPL uses match, set, drop, and stop actions, and route maps use match, set, and drop actions.

**Answer: C**

Explanation:

### Question: 212

Refer to the exhibit.



Refer to the exhibit. Company X is multihomed to a service provider. With the default settings, the link between PE2 and CE1 is the primary link, but company X has asked for the link between PE4 and CE2 to operate as the primary instead. A network engineer used a route map with local-preference to manipulate the routes from PE2 to CE1 to be less desirable and applied the route map on PE2 outbound to CE1. However, the path from PE2 to CE1 is still preferred. Which action must the engineer take so that the desired path is chosen?

- A. Apply local-preference on PE2 outbound to CE2 with a value less than 100.
- B. Configure the AS path prepend attribute to increase the number of hops between CE2 and PE4.
- C. Configure the weight attribute on PE2 in the outbound direction to CE2 with a value less than 65535.
- D. Apply MED in the outbound direction on PE2 to CE1 with a cost lower than the path from PE4 to CE2.

**Answer: B**

Explanation:

### Question: 213

The ISP\_A network is configured with PIMv2 with sparse mode. The engineering team at ISP\_A must implement multicast support following RFC 5059 in the network. The RP priority will be 0, and the solution must require limited supervision from the operations team. Which action must the team take to meet the requirements?

- A. Implement Auto-RP on the candidate RP and the mapping agent only.
- B. Implement Static RP across all routers in the network.
- C. Implement bootstrap router in the network core.
- D. Implement Anycast-RP on edge routers only.

**Answer: C**

Explanation:

### Question: 214

Refer to the exhibit.

**Area 1 Level 1 Routers: R1, R2, R4**

**Area 2 Level 1/2 Routers: R3, R5**

**Area 3 Level 2 Routers: R6**



IS-IS is the routing protocol on the network. After the team noticed excessive CPU usage on all routers, particularly routers R2 and R6, they configured route summarization on all routers. However, a network engineer noticed that summary routes and more specific routes are present in the routing table on R1. Which action must the team take so that R1 receives only the summary routes?

- A. Configure R2 and R4 to perform summarization for Level 1 routers.
- B. Configure R3 to perform summarization for Level 1 routers, and configure R6 to perform summarization for Level 2 routers.
- C. Configure R3 and R5 to perform summarization for Level 1 and Level 2 routers.
- D. Configure R6 only to perform summarization for all routers.

**Answer: B**

Explanation:

### Question: 215

Refer to the exhibit.

```
ip pim ssm
interface g1/0/0
ip pim sparse-mode
```

Which task must be performed on interface g1/0/0 to complete the SSM implementation?

- A. configure OSPFv3
- B. enable CDP
- C. disable IGMP
- D. configure IGMPv3

**Answer: D**

### Question: 216

An engineer configured OSPF on all routers in the data center in the same exchange as other routers Layer 2 connectivity was established. Now, the engineer plans to implement BFD to detect failures in the forwarding path. Which action must the engineer take?

- A. Configure BFD on any one of the routers
- B. Configure IS-IS routing before implementing BFD.
- C. Configure BFD between each adjacent router
- D. Configure CEF before implementing BFD

**Answer: C**

### Question: 217

An engineer with an employee ID: 4441 70:118 must implement PIM-SSM in the network to support a new

multicast messaging service. All the routers in the network run a distance vector routing protocol. Unicast routing is established on the network and is working normally. What must the engineer enable to continue the implementation process?

- A. PIM dense mode and IGMPv2 on all interfaces that participate in multicast
- B. IGMPv3 on all interfaces that participate in multicast
- C. PIM dense mode on all devices on the network
- D. PIM RP feature on the network

**Answer: D**

### Question: 218

Refer to the exhibit. R1 is sending data traffic to R6 using the same SRGB range as the other routers in the topology, but it is also using a non-default SRLB range. Which two configuration tasks must the engineer perform to enable normal SRGB and SRLB operations on this network? (Choose two.)


- A. Configure an adjacency SID from SRLB range 1002 to 6005.
- B. Configure router R2 with adjacency SID 4002 to enable it to reach R4.
- C. Configure router R1 with adjacency SID 1003 to enable it to reach R2.
- D. Set the default SRGB range to 16000 to 23999.

E. Set the SRGB range to 1002 to 6005

**Answer: A, D**

### Question: 219

Refer to the exhibit.



IS-IS is the IGP for the network, where routers R1, R2, R3, and R4 are Level 1 routers in Area 12 and routers R5 and R6 are Level 2 routers in area 16. According to the network design, router R1 must send traffic to R6 using R4 as its next hop. To meet the requirement, a network engineer configured route filtering on R2 to prevent the route to R6 from being advertised from R2. However, R1 is still using R2 as the next hop. Which action must the engineer take to resolve the issue?


- A. Disable wide area metrics and use only cost to perform the route manipulation.
- B. Manipulate link costs to provide the desired routing through R4.
- C. Move all routers into the same IS-IS area.

D. Configure routers R5 and R6 as Level 1/2 routers

**Answer: D**

### Question: 220

Refer to the exhibit.



OSPF is running in the network, which spans multiple areas as shown, and all OSPF neighbors are established. Hosts that reside on the LAN on R1 are trying to reach a server connected to R5, but R1 cannot reach the R5 routes. The server must maintain a TCP connection with the hosts. Which action must the engineer take to resolve the issue?


- A. Configure a static route on R1 pointing to R3 so that traffic reaches routes connected to R5.
- B. Configure a virtual link between R1 and R3 to allow area 1 to communicate to area 3.
- C. Configure a static route on R1 and R3 that uses R5 as the next hop for traffic destined to routes connected to R5.

D. Configure a virtual link between R1 and R2 to connect area 1 to area 0.

**Answer: D**

### Question: 221

Refer to the exhibit.




The four routers in this network have formed adjacencies. Router R2 is an IS-IS level 1-2 router that connects to other areas, and the attach bit is set in the R2 level 1 LSP advertisements to R1. According to the network design, R2 should push level 2 routes from its routing table to the R1 routing table. However, R1 installed a default route in its routing table. Which action resolves this problem?

- A. Use The is-type level-1-2 backdoor command to unset the attach bit on R1 and use a prefix list to filter prefixes from R2
- B. Set the attach bit on R1 and use a prefix list to filter prefixes from R2.
- C. Use the is-type level-1-2 backdoor command to unset the attach bit on R2 and configure a distribute-list to redistribute prefixes to R1
- D. Configure R2 as a level 1 router and use a prefix list to filter prefixes to R1.

**Answer: C**

## Question: 222

Refer to the exhibit.



ISP\_A in AS 635 has four BGP neighbors one-by-one, starting from AS 12 R1 has been configured with `bgp deterministic-med` and `bgp additional-paths` install to influence egress traffic After a recent reload of R10, all traffic was automatically re-routed via AS 87 A network engineer must correct the traffic flow and route the traffic via AS 101 Which configuration must the engineer apply to R1 to correct the problem without affecting the second BGP attribute?

- A. `bgp bestpath compare-routerid`
- B. `bgp bestpath med missing-as-worst`
- C. `bgp always-compare-med`
- D. `bgp default local-preference 200`

**Answer: A**

### Question: 223

How does the SRv6 IS-IS Flexible Algorithm function on the data plane?

- A. It uses the TE metric in its path calculations to optimize the cumulative TE metric.
- B. It uses the IGP metric in its path calculations to optimize reachability.
- C. It automatically delivers traffic to the destination over an engineered path computed by the IGP.
- D. It applies path computation algorithms to learn the topology and compute a pair of disjoint LSPs.

**Answer: C**

### Question: 224


A network administrator is planning an upgrade for a PE router that has BGP peering sessions with several CE routers. The SLA has stringent uptime requirements, so the administrator wants to implement BGP NSR with SSO to provide fast convergence during the upgrade. Which action must the engineer take as part of the NSR implementation?

- A. Configure a VRF on each CE router on the interface that connects to the PE router.
- B. Configure the `bgp graceful-restart` command on the PE router.
- C. Configure the `clear ip bgp 65514 soft in` command on the PE and CE routers.
- D. Configure MPLS on the link between the PE and CE routers.

**Answer: B**

**Question: 225**

Refer to the exhibit.




OSPF is an interior routing protocol on the network. Routers R4 and R6 are in area 0 and all tunnels are in area 2. R4 is connected to R6 with a 1 Gbps link and with 100 Mbps links to the rest of the LAN. After the most recent maintenance window, a network engineer noticed that R4 is using a suboptimal route toward 10.173.12.10 instead of the route via R6. Which action must the network engineer take to correct reachability toward 10.173.12.10/24?

- A. Configure auto-cost reference-bandwidth 1000 under the global OSPF configuration on R4
- B. Configure `ip ospf 0 area 0` under tunnel 2 on both R4 and R5
- C. Configure `ip ospf cost 120` under the Gi0/0 interface on R4
- D. Configure `ip ospf 2 area 0` under tunnel 3 on both R6 and R5

**Answer: A**

## Question: 226

Refer to the exhibit.



A network engineer must enable segment routing between the three autonomous systems with these requirements:

- OSPF link-state data for area 0 routers must be distributed to the controller via BGP LS.
- EPE must be enabled on the egress node in AS 100.
- Routers must advertise peer node SIDs to the controller R1 via BGP LS.

Which configuration must the engineer apply?

A)

```
R2(config)# router ospf 100
R2(config-ospf)# redistribute bgp 100
R2(config)# router bgp 100
R2(config-bgp)# neighbor 192.168.20.1
R2(config-bgp-nbr)# remote-as 100
R2(config-bgp-nbr)# egress-engineering
R2(config-bgp)# neighbor 10.10.10.1
R2(config-bgp-nbr)# remote-as 200
```

R2(config-bgp-nbr)# address-family link-state link-state

B)

R1(config)# router ospf 100

RI(config-ospf)# redistribute bgp 100

R1(config)# router bgp 100

RI(config-bgp)# neighbor 192.168.30.1

R1(config-bgp-nbr)# remote-as 100

RI(config-bgp-nbr)# egress-engineering

RI(config-bgp)# neighbor 10.10.10.1

RI(config-bgp-nbr)# remote-as 200

R1(config-bgp-nbr># address-family link-state link-state

C)

R1 (config)# router ospf 100

RI(config-ospf)# distribute link-state instance-id 100 throttle 10

R1 (config)# router bgp 100

RI(config-bgp)# neighbor 10.10.10.1

RI(config-bgp-nbr)# remote-as 200

RI(config-bgp-nbr)# egress-engineering

RI(config-bgp)# neighbor 192.168.30.1

RI(config-bgp-nbr)# remote-as 100

RI(config-bgp-nbr)# address-family link-state link-state

D)

R2(config)# router ospf 100

R2(config-ospf)# distribute link-state instance-id 100 throttle 10

R2(config)# router bgp 100

R2(config-bgp)# neighbor 10.10.20.1

R2(config-bgp-nbr)# remote-as 300

R2(config-bgp-nbr^ egress-engineering

R2(config-bgp)# neighbor 192.168.20.1

R2(config-bgp-nbr)# remote-as 100

R2(config-bgp-nbr)<< address-family link-state link-state

A. Option A

B. Option B

C. Option C

D. Option D

**Answer: B**

### Question: 227

Refer to the exhibit. After creating dynamic neighbors for the range specified, the network administrator with an employee ID: 4690:88:374 cannot see all the BGP neighbors established. What is the issue?

- A. The single activated neighbor cannot be within the listen range.
- B. A list of neighbors belongs only to a defined ASN.
- C. The peer groups must also be activated inside the address family.
- D. The listen limit is smaller than the prefix defined in the listen range.

**Answer: C**

### Question: 228

Refer to the exhibit.



OSPF is running in the core of a company. Router RA is running EIGRP. After the company acquired an additional office, an engineering team connected the two networks from router RA to router R1. OSPF and EIGRP have been mutually redistributed on R1. but the team cannot view the routes on R3. Which action must the team take to resolve the issue?

- A. Create a virtual link between R1 and R2.
- B. Configure the routers in Area 2 to be in Area 0.
- C. Disable stub routing in Area 1.
- D. Redistribute connected routes on R3.

**Answer: C**

### Question: 229


What is one of the purposes of a default Multicast Distribution Tree?

- A. to offload traffic from data MDT
- B. to define the RP for PIM sparse mode and the different leaves
- C. to support control traffic between different VRFs
- D. to handle high-bandwidth source streams

**Answer: C**

**Question: 230**

Refer to the exhibit.



A service provider is running IS-IS in the core, and all routers use the default IS-IS configuration and reside in the same area.


- a. All routers maintain a different 6 bytes for the system ID. but the attached bits in the NET are the same. After an engineer added router R6 to the network. R3 and R5 began to experience CPU processing issues. Which action must the engineer take to resolve the issue?
- A. Configure all routers to operate as IS-IS Level 1 routers by altering the attached bit to ensure that a default route is injected into the routing tables.
  - B. Configure R6 to operate as a Level 1 router only so that it maintains the Level 1 database for all routers.
  - C. Configure routers R5 and R3 to operate as Level 1 routers only so that they inject a default routing into the routing tables of the other routers

D. Configure all routers to operate as IS-IS Level 2 routers so that only a Level 2 database is maintained.

**Answer: D**

### Question: 231

Refer to the exhibit.



An engineer at a large ISP is troubleshooting suboptimal routing on a customer's IS-IS network with multiple Level 1-2 routers. The network consists of numerous routers that are configured for MPLS. When IS-IS picks the closest Level 1-2 router to exit the area, suboptimal routing occurs in Area 8 and Area 10. The other IGP protocols are running correctly, and the IS-IS adjacency between routers is established. Which action resolves the issue?

- A. Configure each interface as either Level 1 or Level 2 circuit type on router 1.
- B. Configure the router as DIS for each of the interfaces on router 1.
- C. Redistribute Level 2 prefixes into Level 1 on router 2.
- D. Disable advertisement of the default route on router 2.

**Answer: C**

**Question: 232**


Refer to the exhibit. A network engineer must correct the security policy on the EDGE\_R router, which is connected to two Tier 1 ISPs. After another engineer added ip as-path access-list 11 permit 11 to EDGE\_R, some routes from AS 11 and AS 4832 started to appear in the routing information base (RIB). Traffic control via Community options is disabled on both ISPs. How should the network engineer update the router configuration on EDGE\_R so that only routes for AS 11 are permitted?

- A. Configure ip as-path prefix-list 11 permit \_11\$.
- B. Configure ip as-path access-list 11 permit \_11\_.
- C. Configure ip as-path distribution-list 11 deny A[0-9]+\$.
- D. Configure ip as-path access-list 11 deny 4832\$.

**Answer: D**

**Question: 233**

Refer to the exhibit.




A network engineer must optimize link failures on the OSPF routing domain. The solution must follow RFC 5286, and downtime must be no more than 50 ms. All known subnets on R1 must be protected. Which solution must the engineer implement on R1?

- A. Configure `fp ospf frr per-prefix candidate enable high` under the interface configuration.
- B. Configure `ip ospf fast-reroute per-prefix candidate enable` under the interface configuration.
- C. Configure `fast-reroute per-prefix enable prefix-priority low` under the global OSPF configuration.
- D. Configure `frr per-prefix enable priority high` under the global OSPF configuration.

**Answer: C**

### Question: 234

Refer to the exhibit.



BGP is running on the network, and CE1 currently prefers routes to the internet via PE 1. A network engineer must


manipulate the BGP routes that CE1 receives from PE2 to be more desirable. The engineer configured a route map using MED and applied it on CE1 outbound from PE2. However, after a BGP soft reset, PE1 is still preferred. Which action must the engineer take so that PE2 is chosen as the next hop?

- A. Update the route map to use only weight to manipulate the path.
- B. Configure the route map with MED and local preference to manipulate the route outbound
- C. Apply the route map on CE1 In the inbound direction.
- D. Configure the route map with AS path prepend to manipulate the AS path to exit the autonomous system

**Answer: B**

### Question: 235

Refer to the exhibit.




Company\_A hosts multiple online VoD platforms. R3 is connected to MAN with a 100 Mbps link R2 has only a 10 Mbps connection, but it has the lowest number of hops to R1 A network engineer must allow users in area 49.421e.2331 5548.7624 00 to access online library resources, using the best possible route toward the online platforms to guarantee the highest connection quality. After the engineer implements the requirements, they notice an excessive number of routes in the R4 routing table and high memory consumption on the router. Which two actions must the engineer take to resolve the issue? (Choose two.)

- A. Create an access list on R2 and redistribute all routes from L1 into L2.
- B. Configure summary-address area level-1-2 64.45.52.0 255.255.252.0 on R3.
- C. Create a route map on R3 and redistribute all routes from L2 into L1.
- D. Configure summary-address level-2 64.45.52.0 0.0.3.255 on R2.
- E. Configure summary-address 64.45.52.0 255 255.252.0 level-1 on R3

**Answer: C, E**

### **Question: 236**

Refer to the exhibit.



OSPF is an interior routing protocol on the network. R2 is in the backbone area, and R4 is in a totally NSSA area.

a. R5 announces only the 0.0.0.0/0 route to R4 via BGP. R4 is configured with redistribute bgp 736 metric-type 2. After the most recent maintenance window, R2 no longer has the default route via R4. Which action must the network engineer take to correct the reachability problem?

A. Configure area 2 nssa no-summary under the global router configuration on router R2.

B. Configure default-information-originate under the OSPF configuration on router R4.

C. Configure ip route 0.0.0.0 0.0.0.0 null0 under the global router configuration on router R4.

D. Configure area 2 nssa translate type7 no-suppress-fa under the OSPF configuration on router R2.

**Question: 237**

**Answer: A**

What is the purpose of the confederation identifier in a BGP confederation?

- A. It identifies the autonomous system that belongs to the confederation.
- B. It advertises routes outside of the confederation's autonomous system.
- C. It reflects the route reflector cluster ID within the confederation.
- D. It carries the AS number within the confederation's local autonomous system

**Question: 238**

**Answer: A**

Refer to the exhibit.

## Route-Map

Route reconfigure terminal

Routeri(config)#ip prefix-list UM LP PREFIXES seq 10 permit 172.22.0.0/16

Routeri(config)#route-map LOCAL PREF permit 10

Routeri (config-route-map) #match ip address prefix-list UM LP PREFIXES Routeri(config-route-map)#set local-preference 50 Routeri(config-route-map)#exit

Routeri(config)#route-map LOCALPREF permit 20

Routeri(config-route-map)#match ip

Routeri(config-route-map)#match ip address prefix-list UM LP PREFIXES

Routeri(config-route-map)#set local-preference 50

Routeri(config-route-map)#set community (2:666) no-advertise

Routeri(config-route-map)#exit

Routeri(config)#end

Routeri#

Routeri#show run

ip prefix-list UM LP PREFIXES seq 10 permit 172.22.0.0/16

route-map LOCALPREF permit 20

match ip address prefix-list UM LP PREFIXES

set local-preference 50

set community (2:100) no-advertise

exit

router bgp 65500

address-family ipv4 unicast |

neighbor 192.1.1.5

remote-as 65510

address-family ipv6 unicast

An engineer is comparing these RPL and route map configurations. Which two conclusions will the engineer reach? (Choose two.)

- A. RPL is applying policy only on BGP IPv6 address families, but the route map is applying policy on BGP IPv4 and IPv6 address families.
- B. RPL and the route map are applying only local preference 50 for prefix 172.22.0.0/16.
- C. RPL is applying local preference 100 and community (2:100) for multiple prefixes, but the route map is applying default local preference for multiple prefixes
- D. RPL is applying local preference 50 and community (2:666) for prefix 172.22.0.0/16. but the route map is applying only local preference 50 for the same prefix
- E. RPL is applying local preference 50 and community (2:666) for all prefixes, but the route map is applying

only local preference 50 for all prefixes.

**Answer: C, D**

**Question: 239**


Which requirement must be met while using parameter with conditional matching in routing policy definition?

- A. Parameters must be used only in route maps.
- B. Parameters must start with the \$ symbol.
- C. Parameters must be applied only to BGP statements.
- D. Parameters must be used in a hierarchal structure.

**Answer: D**

**Question: 240**

Refer to the exhibit.



Router 1, Router 2, and Router 3 are co-located in an onsite data center of an office, but Router 4, Router 5, and Router 6 are installed in another data center 'eral miles away. A network engineer is troubleshooting an RSVP-TE issue on the network. On most of the devices, the engineer can move all LSPs away from a selected dace and take it down for maintenance without interrupting traffic.

However, when an interface on Router 3 is taken down, traffic on the network is disrupted. What must the engineer enable to resolve the issue on Router 3?

- A. FRR
- B. NSR
- C. soft-preemption feature
- D. Soft-disperse

**Answer: C**

### Question: 241

What is a restriction of 6PE?

- A. It is implemented only over eBGP peers.

- B. It is implemented only over iBGP peers.
- C. It adds a value of ::FFFF:: on the IPv4 next hop address.
- D. It cannot be implemented on VRFs.

**Answer: D**

### **Question: 242**

What does conditional matching operate with RPL?

- A. The matching identifies a group of routes that are removed from the routing table when the administrator defined tags match
- B. The matching identifies tags that are placed on routes from neighboring autonomous systems and drops the routes if they are the same.
- C. The matching identifies a set of criteria, and an action is taken based on whether the criteria are the same or different
- D. The matching identifies interfaces that run the same routing protocol and manipulates data based on the routing rules

**Answer: C**

### **Question: 243**

Which attribute is used to instruct the router that packets must be processed by an ASA that is associated with the segment?


- A. service SID with service instructions

- B. node SID that identifies a specific network device
- C. prefix SID that represents a specific route prefix
- D. adjacency SID with unidirectional adjacency instructions

**Answer: B**

### Question: 244

Refer to the exhibit.




Router 1, Router 2, and Router 3 are co-located in an onsite data center of an office, but Router 4, Router 5, and Router 6 are installed in another data center several miles away. A network engineer is troubleshooting an RSVP-TE issue on the network. On most of the devices, the engineer can move all LSPs away from a selected interface and take it down for maintenance without interrupting traffic. However, when an interface on Router 3 is taken down, traffic on the network is disrupted. What must the engineer enable to resolve the issue on Router 3?

- A. FRR
- B. NSR
- C. soft-preemption feature
- D. soft-disperse feature

Answer: C

### Question: 245

Refer to the exhibit.



ISP\_1 uses eBGP for connection to all clients LDP with OSPF and IS-IS is used in the ISP\_1 network Client\_A and Client\_B are in the same VRF\_VoD. After software upgrade on all routers on the ISP\_1 network, Client\_A has lost the connection to Client\_C. A network engineer must troubleshoot the connection from R1 to R4 via M. After initial checks, the engineer confirmed that all routes are installed correctly. Which two tasks must the engineer perform to correct the problem? (Choose two.)

- A. Configure next-hop-self under the BGP configuration on R42.
- B. Configure mpls label implicit-null under the global configuration on R31.
- C. Configure next-hop-unchanged under the BGP configuration on R31.
- D. Configure ip cef under the global configuration on R31.
- E. Configure send-label under the BGP configuration on R42

**Answer: A, E**

### Question: 246


network engineer must implement a segment routing policy rule that specifies that SIDs are bounded manually. The policy will be used for traffic steering based on incoming label ..with the binding SIDs, adjacency SIDs, and BGP peering SIDs. All labels will be in the local segment with the range from 40000 to 50000. How must the engineer implement ... label allocation?

- A. Implement the standard default SRGB range in the SR global segment.
- B. Implement a non-default SRGB range in the SR local segment.
- C. Implement a non-default SRLB range in the SR global segment.
- D. Implement the default SRLB range in the SR local segment

**Answer: B**

### Question: 247

Refer to the exhibit.



The network is situated at different locations on the customer site. The customer runs a mission-critical application behind CE2. which is connected over a 10 link to the local network of the customer After

recent configuration changes to the P and PE routers, an LSP failure has occurred, causing connectivity issues.


CE1 is unable to reach the CE2 internal IP address. Which action resolves this issue?

- A. Implement BGP in the core and configure an MPLS TE tunnel between PE1 and PE2
- B. Configure the ip ospf network point-to-multipoint command on PE2.
- C. Change the subnet mask on the PE2 loopback interface from /24 to 32
- D. Configure the ip ospf network point-to-point command on PE2

**Answer: C**

## Question: 248

Refer to the exhibit.



I>:ll 11 MV\*HUU ravy kalla\_y\*tk 4aMta yak tar kali\* UMI HI IS 10 1>1H II 10 1 a\*11\*4 KIWI atat\* tyak aat-1) 10:11 11 MV» MUO ravy.lMlH\*^\*t\*\_Mlai\* yak r\*t kaki\* Mat HI II 10 1>1Tt it io 1  
\*\*it\*« NMBT ot«M iy\*k a>i-o; 10:10 11 MV» BOU0 trvy kalla^yatk Mlau Mat ya\* d\*Mt\*4 kail\* Mat tat 111 10 10 I ain II 10 1 SCTIVO arSMlvo 10:10 11 MV\* KUO larybalM\_y\*M kalata yak fa\*  
hall\* Mat 111 10 10 I>111 li 10 1 aa.tad OOMI atat\*tyak cab»O> 10 IS 11 Mvr aauo rary Kall\*jaUTaaMt\* Laat yah dal\* tad kail\* tar HI.10 10 I ->H1 II 10.1 ACTIVS >MMIV« 10 IS St MV\* auo  
eaaBaauoatlaa Mat Mtk HI 11 10 1  
10 15 Si Mvy-aiUO ta\*y\_halla\_oaaBMkleatMa\_I\*at Bel\*kb»r HI II 10 1 W«a (\*«: (an\_iaat>

All backbone interfaces on the network are in VRF\_MPLS and ip rsvp signalling hello is preconfigured on all MPLS routers in the network. After router R2 is upgraded and its configuration restored from backup. LSP signaled via RSVP that the tunnel from R1 had been tom down The network engineer confirmed IP reachability line the backbone network Which action must the engineer take to fix the issue?

- A. Configure ip rsvp udp neighbor 172.16.10.1 vrf VRFMPLS under the global configuration on R2.
- B. Configure ip rsvp listener outbound reply 172.16.10.1 vrf VRF\_MPLS under the global configuration on R2.
- C. Configure ip rsvp signalling hello refresh interval 20 under the interface GiO/0 configuration on R1.
- D. Configure ip rsvp udp-multicasts 172.16.10.2 vrf VRFMPLS under the interface GiO/0 configuration on R1

**Answer: C**

**Question: 249**

A network administrator working for a private service provider with an employee ID 4784 41 907 wants to configure timers for the BGP process on their Cisco IOS XE router. Which configuration follows guideline timer recommendations and is a valid BGP timer deployment?

A)

```
router bgp 1.0
bgp router-id 1.1.1.1
no bgp default ipv4-unicast
bgp fast-external-fallover
timers keep-alive 40
timers hold 120
```

B)

```
router bgp 1.0 bgp router-id 1.1.1.1 no bgp default ipv4-
unicast bgp fast-external-fallover timers keep-alive 40
timers hold 60
```

C)

```
router bgp 1.0 bgp router-id 1.1.1.1
no bgp default ipv4-unicast bgp fast-external-fallover
timers bgp 40 60
```

D)

```
router bgp 1.0 bgp router-id 1.1.1.1 no bgp default ipv4-
unicast bgp fast-external-fallover
timers bgp 40 120
```

A. Option A

B. Option B


C. Option C

D. Option D

**Answer: D**

## Question: 250

Refer to the exhibit.



A network engineer must enable connections from customers on PE1 to the external network via the SR-MPLS forwarding plane with a range of 8000 SID. In the current environment

- The MPLS forwarding tables on PE1 and PE2 are populated with LDP-originated LDP labels
- The current ranges are 16100 - 16199 on PE1 and 16200 - 16299 on PE2

Which configuration must the engineer apply?

C PE1 (config)# router isis SR

PE1(config-isis)# segment-routing global-block 1042000 1050000

PE1 (config-isis)# segment-routing mpls sr-prefer

O PE 1 (config)# router isis SR

PE1(config-isis)# segment-routing global-block 90000 98000

PE1(config-isis)# address-family ipv4 unicast

PE1(config-isis-af)# segment-routing mpls sr-prefer

O PE1 (config)# router isis SR

PE1(config-isis)# segment-routing global-block 3000 11000

PE1(config-isis)# segment-routing mpls sr-prefer

CI PE 1 (config)# router isis SR

```
PE1(config)#isis S segment-routing global-block 16000 24000
PE1(config-isis)# address-family ipv4 unicast
PE1(config-isis-af)# segment-routing mpls sr-prefer
```

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

**Answer: B**

### Question: 251


Refer to the exhibit Dedicated MPLS Layer 3 services for VRF customer A are configured on the MPLS network for FTP access OSPF with area 0 configuration is used as the interior routing protocol between the routers, and LDP is up and running on the MPLS network. After recent maintenance work on the MPLS network. IP address 172.16.100.1 cannot access the FTP server with IP address 172.16.192.21 All loopbacks are reachable on the network Which configuration must the engineer apply to correct the problem?

- A. Configure passive-interface under the Gi0/2 interface on R22
- B. Configure ip ospf network point-to-point under the Lo0 interface on R33
- C. Configure network 33.33.10.1 0.0.0.255 area 0 under the OSPF configuration on R33
- D. Configure ip cef under the global configuration on R22

**Answer: B**

### Question: 252

Refer to the exhibit.




All routers in the network are in area 12. with IS-IS running in the default configuration as the IGP To avoid congestion on the link from R1 to R2 the network architect planned for traffic from R1 to R3 to pass through R4 To meet the requirement a network engineer set the metric from R1 to R4 to a cost of 15 but the traffic is still travelling through R2 Which action must the engineer take to resolve the issue?

- A. Change the default metric to 20 on all links and disable wide area metrics
- B. Change the cost on the link from R1 to R2 to 11 and enable narrow metrics
- C. Change the cost on the link from R1 to R4 to 5
- D. Change the IS-IS administrative distance to 10

**Answer: C**

### Question: 253

Refer to the exhibit.



```

router isis
 net 48.0000.0000.0001.00
 redistribute static ip route-map set-tag

route-map set-tag permit 5
 set tag 10

```

IS-IS runs in the network and all routers are in the same area. R1 statically routes to distribute traffic with R6 and the networking team has redistributed the static routes on R1 into IS-IS. Hosts on R3 need access to servers connected to R6. A tag has been applied to the routes to manipulate routes


along the path but routers running IS-IS are still unable to reach the redistributed routes. Which action must the team take to resolve the issue?

- A. Configure the routers as Level 1/Level 2 routers
- B. Configure R3 with a default route that uses R2 as the next hop
- C. Configure the IS-IS routers with metric-style wide
- D. Configure the route map with an additional route-map set-tag permit 10 sequence

**Answer: C**

**Question: 254**

Refer to the exhibit.



```

RP/0/0/CPU0:R2# show logging %BGP-3-MAXPFXEXCEED: No. of prefix
received from 10.0.0.1 (af1 0): 1560 exceed limit 1500 4ROUTING-
BGP-5-ADJCHANGE: neighbor 10.0.0.1 Down - BGP Notification sent,
prefix limit exceeded (VRF: default) %BGPSESSION-5-ADJCHANGE:
neighbor 10.0.0.1 VPNv4 Unicast topology base removed from session
BGP protocol initialization
%BGP_SESSION-5-ADJCHANGE: neighbor 10.0.0.1 IPv4 MDT topology
base removed from session BGP protocol initialization
%BGPSESSION-5-ADJCHANGE: neighbor 10.0.0.1 VPNv4 Unicast topology
base removed from session Unknown path error

```

After maintenance to activate BGP neighbors connecting a new data center attached to R1,

L3VPN customers attached to OSPF domain 100 lost connectivity to services in OSPF 200 and OSPF 300 However Unified MPLS services within OSPF 100 remain reachable Which action must the network engineer take to resolve the outage?

- A. Reset the BGP session between R1 and R2
- B. Activate BGP neighbors between the new data center and R2 and R3
- C. Establish new MPLS LSPs from R1 to R2, R3 and R4
- D. Raise the maximum prefix limit on R2 to 1560 or more

**Answer: A**

### Question: 255

Refer to the exhibit.



A company has been running OSPF and EIGRP on two separate portions of the network. To allow communication flow between these previously isolated parts of the business, the networking team just configured R2 to mutually redistribute routes between EIGRP and OSPF. R3 has the OSPF routes in its routing table, but R1 is failing to receive EIGRP routes for the link between R2 and R3. Which action must the team take to resolve the issue?

- A. Configure EIGRP on R1
- B. Configure the redistribution statement with the subnets keyword under the R2 OSPF configuration
- C. Configure a virtual link between R1 and R2 to handle traffic without an assigned area
- D. Configure R1 with a static route with the next hop set to R3 with the vrf keyword to avoid the global routing table

**Answer: B**

### Question: 256

Which type of policy does an IPv6 ISATAP tunnel support?


- A. policies that transmit Connectionless Network Service packets
- B. policies to transmit IPv6 traffic within a site

- C. policies to transmit IPv6 traffic over a point-to-point tunnel
- D. policies that require addresses to have a 2002 716 prefix

**Answer: B**

### Question: 257

Refer to the exhibit.




Routers R1 and R2 are RFC 5036-based MPLS-enabled core routers MPLS IP is enabled globally on both routers OSPF with area 0 configuration is used as an interior routing protocol between the routers After new services were enabled on the MPLS network, LDP connectivity failed on R1 TCP/UDP port 711 is confirmed to be open Which two actions must the engineer take to correct the problem? (Choose two.)

- A. Configure network 10.10.140.192 0.0.0.0 area 0 under the OSPF configuration on R2
- B. Configure mpls label protocol ldp under the Gi0/0 interface on R1
- C. Configure mpls ip under the Gi0/0 interface on R1
- D. Configure network 10.10.140.192 0.0.0.255 area 0 under the OSPF configuration on R2
- E. Configure labtl mpls ip protocol tdp under the Gi0 0 interface on R1

**Answer: A, C**

**Question: 258**

Refer to the exhibit.



A regional service provider provides connectivity for several midsize businesses that run mission-critical applications between multiple offices connected via PE 1 and PE2 Links throughout the service provider network are connected using shared resources and an SR-TE tunnel is in use to route traffic over different paths depending on congestion policies The networking team must implement SR updates to reduce downtime in the event of a link failure. Which action must the team take to meet the requirements?

- A. Implement SRLG protection to identify a post-convergence backup path
- B. Implement BGP timers to support immediate reconvergence when the BGP neighbor relationship status changes.
- C. Implement targeted LDP neighbors to prevent unwanted LDP relationships from forming
- D. Implement VRFs to segment the network and avoid large-scale network outages when a link fails

**Answer: A**

### Question: 259

Refer to the exhibit.

The service provider provides IPv6 connectivity to large enterprise customers with offices spanning several geographic regions. SRv6 is running in the service provider core. The service provider just signed a contract with a new customer that is still running IPv4. The new customer needs connectivity between its two sites connected via CE1 and CE2 and the internet. Which action must the service provider's network-engineering team take to provide connectivity across the SRv6 data plane?

- A. Encapsulate IPv4 traffic from the CEs on their respective PEs through a TE path between PE1 and PE2.
- B. Translate the IPv4 addresses to an IPv6 address on each CE.
- C. Redistribute the IPv4 routes from CE2 and CE1 on their respective PEs.
- D. Implement an NTP master and sync the timers on the P and PE routers.

**Answer: A**

### Question: 260

Refer to the exhibit.



All four routers on the network reside in the same data center. The four routers are running IS-IS between them in Level 1 mode within an IPv4 network. After a network engineer installed several new servers connected to R4, users connected via R1 reported they could not connect to applications on the servers. The engineer has determined that R4 is receiving routes from R3 but it is failing to receive routes from R1.


Which action must the engineer take to resolve the issue?

- A. Configure R3 only to be a Level 1/Level 2 router.
- B. Configure R2 and R3 to be Level 1/Level 2 routers
- C. Configure R2 only to be a Level 2 router
- D. Configure R1 and R4 to be Level 2 routers

**Answer: B**

## Question: 261

Refer to the exhibit.



EBGP peering is up between BGP AS 65101 and BGP AS 65201 The service provider expects traffic convergence within 300ms between the two autonomous systems but the network operations team sees that route and traffic convergence is too slow Which two actions improve the convergence time? (Choose two.)


- A. Change the bfd multiplier value to 3 on routers RN and RM
- B. Configure fall-over bfd under the BGP instance for BGP neighbors on RN

- C. Configure bfd interval 50 min\_rx 100 multiplier 6 on the connected interfaces on RN and RM
- D. Configure bfd interval 50 min\_rx 50 multiplier 6 on the connected interfaces on RL and RY
- E. Change the bfd interval value to 100ms on routers RN and RM

**Answer: A, E**

### **Question: 262**

Refer to the exhibit.



Early in the growth cycle of a small enterprise the network-engineering team implemented segment routing to support automation of the IPv4 network. As the network grew organically more devices were added and the company migrated to SRv6. The company is now planning to significantly expand the number of hosts connected to R2 and R6 which will require further path manipulation. Which action must the team take to customize the route that traffic takes between endpoints in the network?

- A. Implement the flexible algorithm to customize the shortest path through the network
- B. Implement MLAG to reduce traffic load on heavily congested links
- C. Implement spanning-tree protocol to avoid loops between excess devices in the network
- D. Implement priority queuing to drop non-mission-critical traffic.

**Answer: A**

### Question: 263

What are two major differences between OSPF and IS-IS routing protocols? (Choose two.)

- A. OSPF uses IP protocol 89 for encapsulation and IS-IS uses an IEEE 802.3 standard-based protocol for encapsulation
- B. The OSPF default administrative distance is 115 and the default administrative distance for IS-IS is 110
- C. The OSPF DR role can be pre-empted only when the active DR goes down, but the IS-IS DR role can be taken by a router with higher priority at any time
- D. OSPF partially refreshes the LSA database every 30 minutes, and IS-IS refreshes the complete database table every 60 minutes
- E. An OSPF-enabled router is part of only one area but an IS-IS-enabled router is part of multiple areas

**Answer: A, C**

### Question: 264

RPL has several fundamental capabilities that differ from those present in configurations oriented to traditional route maps. Routing policy language introduces the notion of sets. Which group of set types is valid in RPL?

- A. acl-path-set, community-set, pre-community-set, prefix-set, and bgp-set
- B. as-path-set, community-set, extcommunity-set, prefix-set, and rd-set
- C. access-control-set, community-set, post-community-set, pre-prefix-set, and routed-set
- D. route-path-set, community-set, in-community-set, prefix-set, and rd-set

**Answer: B**

### Question: 265

Which fact must a network architect consider when implementing IGMP in a multicast routing environment?

- A. If SSM is in use, IGMPv3 must be enabled

- B. If BISS support is required IGMPv2 must be enabled
- C. If leave latency is critical IGMPv3 must be enabled
- D. If PIM-DM is required, IGMPv2 must be enabled

**Answer: A**

### **Question: 266**


Which BGP well-known attributes are preserved between peers with a BGP confederation?

- A. origin and MED
- B. originator and origin
- C. next-hop and local-preference
- D. next-hop and weight

**Answer: C**

### **Question: 267**

Refer to the exhibit.



A network engineer at a large ISP with an employee ID 4876 94 052 started to configure OSPF on two routers. The engineer notices that router 1 learned about network 20.20.20.0/24 from router 2, but router 2 did not learn about 10.10.10.0/24 from router 1. Layer 2 and Layer 3 connectivity are working between the two routers, and the physical connection between the devices is functioning. Which action must the engineer take to resolve the issue?

- A. Replace the router 1 OSPF network configuration with network 10.10.10.0 0.0.0.255 area 0.
- B. On router 1, configure a route map with a distribution list for the 10.10.10.0/24 network.
- C. Replace the router 2 OSPF network configuration with network 10.10.10.0 0.0.0.255 area 0.
- D. On router 2, configure a route map with a distribution list for the 10.10.10.0/24 network.

**Answer: A**

## Question: 268

The engineering team at ISP\_1 must improve network stability time and minimize the impact of routing instabilities from other domains to unicast transit services in the backbone. R1 is part of an OSPF-enabled


network, and acts as an ASBR connected via interface Gi0/1 to the external domain. The OSPF exponential backoff algorithm RFC 3623 with helper mode is already implemented on R1. Which additional configuration must the ISP\_1 team apply to suppress routing instability for 3 minutes'?

- A. Configure ip ospf resync-timout 180 under the Gi0/1 interface configuration
- B. Configure ip ospf timer nsf interval 3 under the global Gi0/1 interface configuration
- C. Configure nsf ittf restart-inttrval 180 under the global OSPF configuration
- D. Configure nsf ospf wait 3 under the global OSPF configuration

**Answer: A**

### **Question: 269**

Refer to the exhibit.



The service provider deployed segment routing in its core. P2 is connected to P3 and P4 via a point-to-multipoint link as shown. The service provider provides connectivity for a large enterprise with three main offices in different geographic regions. Each office is connected via one of the three PEs. The enterprise wants to implement multicast traffic across the service provider to stream mission-critical applications between sites. The applications will be sourced from a server connected to the network through PE2. Which action must the network-engineering team take so that multicast traffic will flow between the enterprise sites over the point-to-multipoint link?

- A. Implement RSVP to reserve bandwidth on the link between PE2 and P2
- B. Implement PIM-DM to flood multicast traffic from the enterprise into the service provider core
- C. Implement MLDP to allow multicast traffic to flow within the single AS
- D. Implement SR Tree-SID to build a multicast replication tree over the point-to-multipoint link

**Answer: D**