642-883.examcollection.premium.exam.130q

642-883

Deploying Cisco Service Provider Network Routing (SPROUTE)

Version 5.0

Exam A

QUESTION 1
Refer to the OSPF command exhibit.

```
RP/0/RSP0/CPU0:P1(config-ospf)#area 1 stub no-summary
```

Which effect does the no-summary command option have?

A. It will cause area 1 to be able to receive non-summarized inter-area routes.
B. It will cause area 1 to not receive any inter-area routes and will use a default route to reach networks in other areas.
C. It will cause area 1 to not receive any external routes and will use a default route to reach the external networks.
D. It will convert the NSSA area into a NSSA totally stubby area.
E. It will convert the stubby area into a NSSA.
F. It will disable OSPF auto-summary.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

To define an area as a stub area, use the **area stub** command in router configuration mode. To disable this function, use the **no** form of this command.

> **area** *area-id* **stub** [**no-summary**]
>
> **no area** *area-id* **stub**
>
> **no area** *area-id*

Syntax Description

area-id	Identifier for the stub area; either a decimal value or an IP address.
no-summary	(Optional) Prevents an ABR from sending summary link advertisements into the stub area.

Defaults

No stub area is defined.

Command Modes

Router configuration

Command History

Release	Modification
10.0	This command was introduced.

Usage Guidelines

You must configure the **area stub** command on all routers and access servers in the stub area. Use the **area** router configuration command with the **default-cost** option to specify the cost of a default internal router sent into a stub area by an area border router.

There are two stub area router configuration commands: the **stub** and **default-cost** options of the **area** router configuration command. In all routers attached to the stub area, the area should be configured as a stub area using the **stub** option of the **area** command. Use the **default-cost** option only on an ABR attached to the stub area. The **default-cost** option provides the metric for the summary default route generated by the area border router into the stub area.

To further reduce the number of link state advertisements (LSAs) sent into a stub area, you can configure **no-summary** on the ABR to prevent it from sending summary LSAs (LSA type 3) into the stub area.

QUESTION 2

When troubleshooting OSPF neighbor errors, which three verification steps should be considered? (Choose three.)

A. Verify if neighboring OSPF interfaces are configured in the same area.
B. Verify if neighboring OSPF interfaces are configured with the same OSPF process ID.
C. Verify if neighboring OSPF interfaces are configured with the same OSPF priority.
D. Verify if neighboring OSPF interfaces are configured with the same hello and dead intervals.
E. Verify if neighboring OSPF interfaces are configured with the same area type.

Correct Answer: ADE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 3
On Cisco IOS XR Software, which set of commands is used to enable the gi0/0/0/1 interface for OSPF in area 0?

A. interface gi0/0/0/0ip address 10.1.1.1 255.255.255.0!router ospf 1network 10.1.1.1 0.0.0.0 area 0
B. interface gi0/0/0/0ip address 10.1.1.1 255.255.255.0!router ospf 1network 10.1.1.1 255.255.255.255 area 0
C. router ospf 1area 0interface GigabitEthernet0/0/0/1
D. interface gi0/0/0/0ip address 10.1.1.1 255.255.255.0ip ospf 1 area 0
E. router ospf 1address-family ipv4 unicastinterface GigabitEthernet0/0/0/1 area 0
F. router ospf 1address-family ipv4 unicastinterface GigabitEthernet0/0/0/1area 0

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 4
Which three statements are true regarding the OSPF router ID? (Choose three.)

A. The OSPF routing process chooses a router ID for itself when it starts up.
B. The router-id command is the preferred procedure to set the router ID.
C. If a loopback interface is configured, its address will always be preferred as the router ID over any other methods.
D. After the router ID is set, it does not change, even if the interface that the router is using for the router ID goes down. The router ID changes only if

the router reloads or if the OSPF routing process restarts.
E. In OSPF version 3, the OSPF router ID uses a 128-bit number.

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 5
Which two OSPF network scenarios require OSPF virtual link configuration? (Choose two.)

A. to connect an OSPF non-backbone area to area 0 through another non-backbone area
B. to connect an NSSA area to an external routing domain
C. to connect two parts of a partitioned backbone area through a non-backbone area
D. to enable route leaking from Level 2 into Level 1
E. to enable route leaking from Level 1 into Level 2
F. to enable OSPF traffic engineering

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 6
What is function of the RP/0/RSP0/CPU0:PE1(config-ospf)#distance Cisco IOS-XR command?

A. To modify the administrative distance of the OSPF routes
B. To modify the default seed metric of the OSPF external routes
C. To modify the OSPF default reference bandwidth
D. To modify the OSPF cost

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

Explanation:

distance (OSPF)

To define an administrative distance, use the **distance** command in the appropriate mode. To remove the **distance** command from the configuration file and restore the system to its default condition in which the software removes a distance definition, use the **no** form of this command.

distance *weight* [*ip-address wildcard-mask* [*access-list-name*]]

no distance *weight ip-address wildcard-mask* [*access-list-name*]

Syntax Description		
weight		Administrative distance. Range is 10 to 255. Used alone, the *weight* argument specifies a default administrative distance that the software uses when no other specification exists for a routing information source. Routes with a distance of 255 are not installed in the routing table. Table 1 lists the default administrative distances.
ip-address		(Optional) IP address in four-part, dotted-decimal notation.
wildcard-mask		(Optional) Wildcard mask in four-part, dotted decimal format. A bit set to 1 in the *mask* argument instructs the software to ignore the corresponding bit in the address value.
access-list-name		(Optional) Name of an IP access list to be applied to incoming routing updates.

QUESTION 7
Which four statements are correct regarding IS-IS operations? (Choose four.)

A. By default, Level 1 routers within an IS-IS area do not carry any routing information external to the area to which they belong. They use a default route to exit the area.
B. Summarization should be configured on the Level 2 routers, which injects the Level 2 routes into Level 1.

C. IS-IS supports "route leaking" in which selected Level 2 routes can be advertised by a Level 1/Level 2 router into Level 1.

D. The IS-IS backbone is a contiguous collection of Level 1 capable routers, each of which can be in a different area.

E. With IS-IS, an individual router is in only one area, and the border between areas is on the link that connects two routers that are in different areas.

F. Cisco IOS XR Software supports multitopology for IPv6 IS-IS unless single topology is explicitly configured in IPv6 address-family configuration mode.

Correct Answer: ACEF
Section: (none)
Explanation

Explanation/Reference:
Explanation:

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.0/routing/configuration/guide/rc3isis.html#wp1148617

Single-Topology IPv6 Support
Single-topology IPv6 support on Cisco IOS XR software allows IS-IS for IPv6 to be configured on interfaces along with an IPv4 network protocol. All interfaces must be configured with the identical set of network protocols and all routers in the IS-IS area (for Level 1 routing) or the domain (for Level 2 routing) must support the identical set of network layer protocols on all interfaces.

When single-topology support for IPv6 is used, only narrow link metrics, also known as old-style type, length, value (TLV) arguments, may be employed. During single-topology operation, one shortest path first (SPF) computation per level is used to compute both IPv4 and IPv6 routes. Using a single SPF is possible because both IPv4 IS-IS and IPv6 IS-IS routing protocols share a common link topology.
Because multitopology is the default behavior in the software, you must explicitly configure IPv6 to use the same topology as IPv4 in order to enable single-topology IPv6. Configure the single-topology command in ipv6 address family configuration submode of the IS-IS router stanza.

Multitopology IPv6 Support
Multitopology IPv6 support on Cisco IOS XR software differs from Cisco IOS software in that IS-IS assumes that multitopology support is required as soon as it detects interfaces configured for both IPv6 and IPv4 within the IS-IS stanza.
You must use the metric-style wide command to configure IS-IS to wide link metrics as multitopology link advertisements.

QUESTION 8
When configuring IPv4 and IPv6 IS-IS routing on Cisco IOS XR routers, which three statements are correct? (Choose three.)

A. By default, a single SPF is used for both IPv4 and IPv6, so the IPv4 and IPv6 topology should be the same.

B. By default, the IS-IS router type is Level 1 and Level 2.

C. All IS-IS routers within the same IS-IS area must be configured with the same IS-IS routing process instance ID.

D. By default, metric-style narrow is used.

E. By default, the IS-IS interface circuit type is Level 1 and Level 2.

F. The area IS-IS address-family configuration command is used to specify the IS-IS area address.

Correct Answer: BDE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

is-type {level-1 | level-1-2 | level-2-only}	(Optional) Configures the system type (area or backbone router).
Example: RP/0/RP0/CPU0:router(config-isis)# is-type level-2-only	• By default, every IS-IS instance acts as a **level-1-2** router. • The **level-1** keyword configures the software to perform Level 1 (intra-area) routing only. Only Level 1 adjacencies are established. The software learns about destinations inside its area only. Any packets containing destinations outside the area are sent to the nearest **level-1-2** router in the area. • The **level-2-only** keyword configures the software to perform Level 2 (backbone) routing only and the router establishes only Level 2 adjacencies, either with other Level 2-only routers or with **level-1-2** routers. • The **level-1-2** keyword configures the software to perform both Level 1 and Level 2 routing. Both Level 1 and Level 2 adjacencies are established. The router acts as a border router between the Level 2 backbone and its Level 1 area.

The default metric style for single topology is narrow metrics. However, you can use either wide metrics or narrow metrics. How to configure them depends upon how single-topology is configured. If both IPv4 and IPv6 are enabled and single-topology is configured, the metric style is configured in the address-family ipv4 stanza. You may configure the metric style in the address-family ipv6 stanza, but it will be ignored in this case. If IPv6 only is enabled and single topology is configured, then the metric style is configured in the address-family ipv6 stanza.

`circuit-type {level-1 | level-2-only | level-1-2}`
 	
Example:
`RP/0/RP0/CPU0:router(config-isis-if)# circuit-type level-1-2`
	(Optional) Configures the type of adjacency.
	
	• The default circuit type is the configured system type (configured through the is-**type** command).
	
	• Typically, circuit type needs to be configured when the router is configured only **level-1-2** and you want to constrain an interface to form only **level-1** or level-2-only adjacencies.

QUESTION 9
Refer to the PE1 router routing table output exhibit.

```
RP/0/RSP0/CPU0:PE1#show route ipv4 isis

<output omitted>
i su 10.1.10.0/24 [115/30] via 0.0.0.0, 00:40:34, Null0
i L1 10.1.10.1/32 [115/30] via 192.168.101.11, 00:42:39, GigabitEthernet0/0/0/0
i L1 10.1.10.2/32 [115/24] via 192.168.112.21, 00:44:40, GigabitEthernet0/0/0/1
i L1 10.1.10.3/32 [115/32] via 192.168.113.22, 00:38:23, GigabitEthernet0/0/0/2
i L1 10.1.10.4/32 [115/22] via 192.168.114.23, 00:14:10, GigabitEthernet0/0/0/3
<output omitted>
```

What is causing the i su 10.1.10.0/24 [115/30] via 0.0.0.0, 00:40:34, Null0 entry on the PE1 router routing table?

A. The PE1 router is receiving the 10.1.10.0/24 summary route from the upstream L1/L2 IS-IS router.
B. The PE1 router has been configured to summarize the 10.1.10.x/32 IS-IS routes to 10.1.10.0/24.
C. The 10.1.10.0/24 has been suppressed because IS-IS auto-summary has been disabled on the PE1 router.
D. The 10.1.10.0/24 has been suppressed because of a route policy configuration on the PE1 router.
E. The 10.1.10.0/24 has been suppressed because the more specific 10.1.10.x/32 IS-IS routes have been configured to leak into the IS-IS non-backbone area.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       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, E - EGP, i - ISIS, L1 - IS-IS level-1
       L2 - IS-IS level-2, ia - IS-IS inter area
       su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local
```

QUESTION 10
In comparing IS-IS with OSPF, a Level-1-2 IS-IS router is similar to which kind of OSPF router?

A. ASBR on a normal OSPF area
B. ASBR on NSSA
C. ABR on totally stubby OSPF area
D. ABR on stubby OSPF area
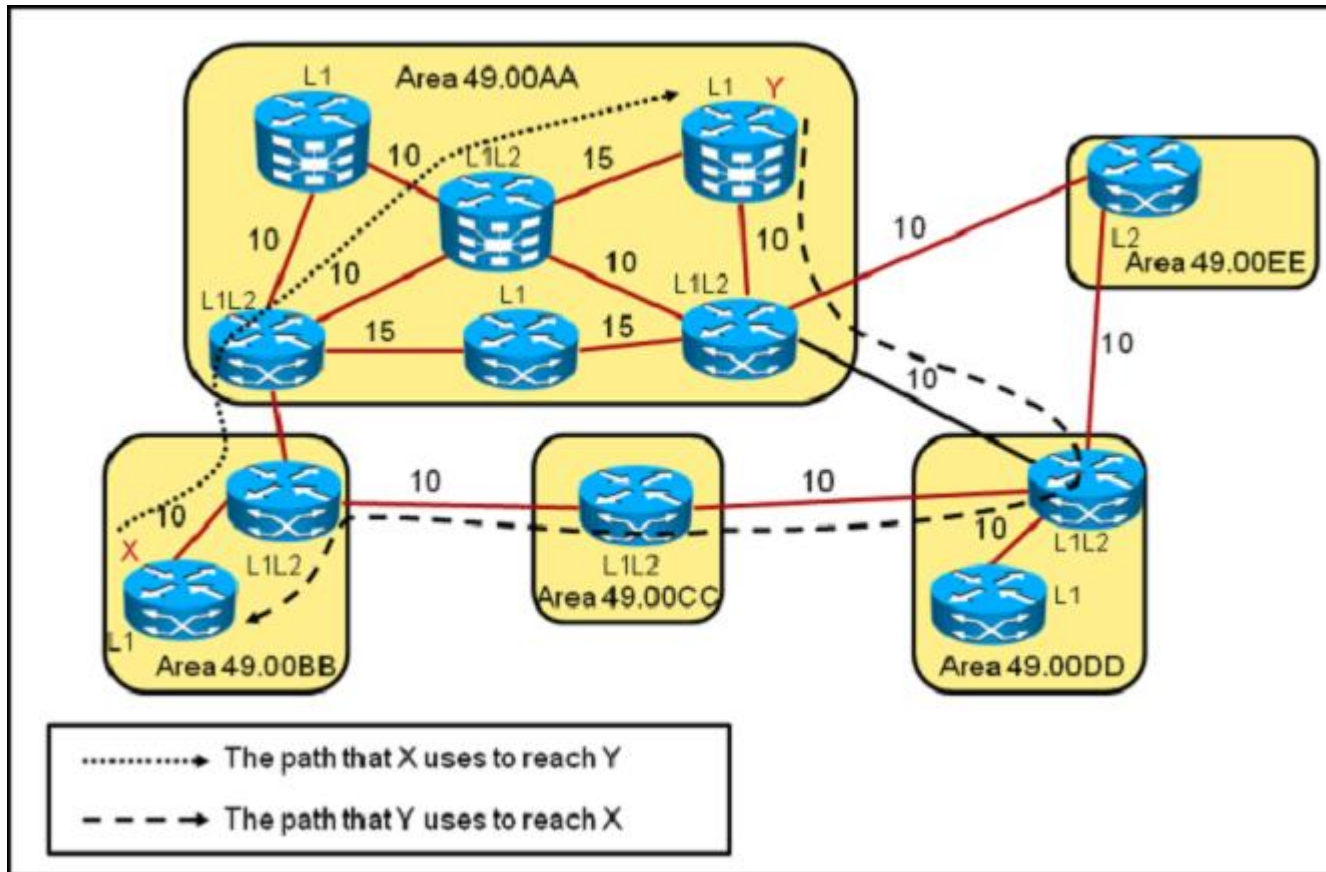E. ABR on a normal OSPF area

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 11
Refer to the topology diagram in the exhibit.

Which IS-IS feature could be implemented so that the return path for the packets from router Y in area 49.00AA to router X in area 49.00BB will use the more optimal path?

A. Enable route leaking to pass Level 2 information into the Level 1 routers.
B. Change the area 49.00AA type from a stub area to a regular area.
C. Change the IS-IS administrative distance on router Y in area 49.00AA.
D. Change the IS-IS metric type from narrow to wide on all IS-IS routers.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

http://www.cisco.com/application/pdf/paws/13796/route-leak.pdf

The IS–IS routing protocol allows for a two–level hierarchy of routing information. There can be multiple Level 1 areas interconnected by a contiguous Level 2 backbone. A router can belong to Level 1, Level 2, or both. The Level 1 link–state database contains information about that area only. The Level 2 link–state database contains information about that level as well as each of the Level 1 areas. An L1/L2 router contains both Level 1 and Level 2 databases. It advertises information about the L1 area to which it belongs into L2. Each L1 area is essentially a stub area. Packets destined for an address that is outside of the L1 area are routed to the closest L1/L2 router to be forwarded on to the destination area. Routing to the closest L1/L2 router can lead to sub–optimal routing when the shortest path to the destination is through a different L1/L2 router.
Route leaking helps reduce sub–optimal routing by providing a mechanism for leaking, or redistributing, L2 information into L1 areas. By having more detail about inter area routes, an L1 router is able to make a better choice with regard to which L1/L2 router to forward the packet.

QUESTION 12
Refer to the exhibit.
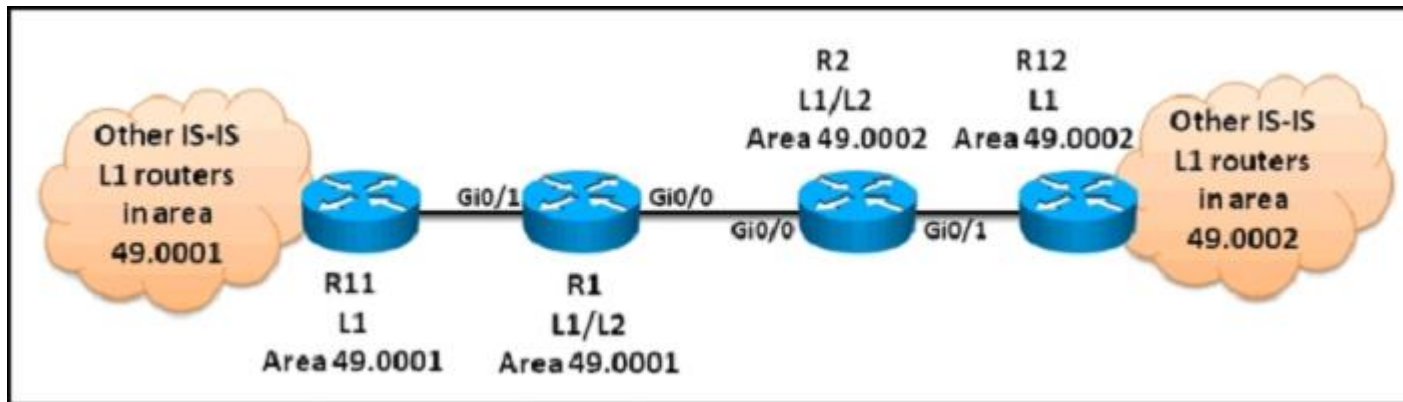
Which two configuration options can be used to optimize the IS-IS network scenario? (Choose two.)

A. Change the R1 and R2 IS type to Level 2.
B. Change the R1 and R2 IS type to Level 1.
C. Change the gi0/0 interface IS-IS circuit type on R1 and R2 to Level 2 only.
D. Change the gi0/1 interface IS-IS circuit type on R1and R2 to Level 1.
E. Change the IS type for all the routers to Level-1-2.
F. Change the IS-IS network type for all the routers to point-to-point.

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 13
The Cisco IOS XE Software summary-address router IS-IS configuration command can be used to send a summarized route into which IS-IS hierarchy?

A. Level 1 only
B. Level 2 only
C. Level-1-2 only
D. Level 1 or Level 2 or Level-1-2

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

summary-address (IS-IS)

To create aggregate addresses for IS-IS, use the **summary-address** command in router configuration mode. To restore the default, use the **no** form of this command.

summary-address *address mask* {level-1 | level-1-2 | level-2}

no summary-address *address mask* {level-1 | level-1-2 | level-2}

Syntax Description

address	Summary address designated for a range of addresses.
mask	IP subnet mask used for the summary route.
level-1	Only routes redistributed into Level 1 are summarized with the configured address and mask value.
level-1-2	Summary routes are applied when redistributing routes into Level 1 and Level 2 IS-IS, and when Level 2 IS-IS advertises Level 1 routes as reachable in its area.
level-2	Routes learned by Level 1 routing are summarized into the Level 2 backbone with the configured address and mask value. Redistributed routes into Level 2 IS-IS will be summarized also.

QUESTION 14
Refer to the Cisco IOS XE IS-IS configuration exhibit.

```
interface GigabitEthernet0/0/0
 ip address 192.168.104.40 255.255.255.0
 ip router isis
 ipv6 address 2001:DB8:192:168:104::40/80
 ipv6 enable
 ipv6 router isis
!
interface GigabitEthernet0/0/1
 ip address 192.168.134.40 255.255.255.0
 ip router isis
 ipv6 address 2001:DB8:192:168:134::40/80
 ipv6 enable
 ipv6 router isis
!
router isis 1
 net 49.0004.0100.0400.1001.00
 summary-address 10.4.10.0 255.255.255.0
!
```

What are two problems with the configuration that are causing the IPv4, or the IPv6, or the IPv4 and IPv6 IS-IS operations to fail? (Choose two.)

A. The configuration is missing the interface gi0/0 and interface gi0/0 commands under router isis 1.
B. The configuration is missing the address-family ipv4 unicast and address-family ipv6 unicast commands under router isis 1.
C. IPv6 unicast routing has not been enabled globally on the Cisco IOS XE router.
D. The gi0/0 and gi0/1 interfaces are not participating in the router isis 1 routing instance.
E. Multitopology IS-IS must be enabled to support both IPv4 and IPv6.
F. Another router isis instance must be enabled to support IPv6.

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

http://www.cisco.com/en/US/docs/switches/metro/me3400/software/release/12.2_50_se/configuration/guide/swipv6.pdf

• When using user-network interface (UNI) or enhanced network interface (ENI) ports for any
IPv6-related features, you must first globally enable IP routing and IPv6 routing on the switch by entering the ip routing ipv6 unicast-routing

QUESTION 15
In which network environment is IS-IS adjacency check important?

A. in a multitopology environment where there are different instances of IS-IS running on the same router
B. in an IPv4/IPv6 environment and running single-topology IS-IS
C. when a level L1/L2 IS-IS router is neighboring with a Level 1 only or Level 2 only router
D. when IS-IS neighbors are in an NBMA environment like over Frame Relay
E. when IS-IS neighbors are in a broadcast environment like an Ethernet LAN

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Disabling IPv6 Protocol-Support Consistency Checks

Perform this task to disable protocol-support consistency checks in IPv6 single-topology mode.

For single-topology IS-IS IPv6, routers must be configured to run the same set of address families. IS-IS performs consistency checks on hello packets and will reject hello packets that do not have the same set of configured address families. For example, a router running IS-IS for both IPv4 and IPv6 will not form an adjacency with a router running IS-IS for IPv4 or IPv6 only. In order to allow adjacency to be formed in mismatched address-families network, the adjacency-check command in IPv6 address family configuration mode must be disabled.

Entering the no adjacency-check command can adversely affect your network configuration. Enter the no adjacency-check command only when you are running IPv4 IS-IS on all your routers and you want to add IPv6 IS-IS to your network but you need to maintain all your adjacencies during the transition. When the IPv6 IS-IS configuration is complete, remove the no adjacency-check command from the configuration.

QUESTION 16

Refer to the IS-IS configuration exhibit.

```
router isis 1
 net 49.0001.0100.0100.1001.00
 address-family ipv4 unicast
 !
 address-family ipv6 unicast
  single-topology
 !
 interface gi0/0/0/0
  address-family ipv4 unicast
  !
  address-family ipv6 unicast
```

This is the typical IS-IS configuration of the routers in an AS using IS-IS as the IGP. This AS is in the transition phase of integrating IPv6 into the network. During this transition phase, some of the routers within the AS might be running IPv4 only, some might be running IPv6 only, and others might be running both IPv4 and IPv6. To avoid any black holes for the IPv6 traffic, which configuration change can be made?

A. Disable IS-IS adjacency checks.
B. Enable IPv6 adjacency over IPv4 IS-IS peering.
C. Enable multi-topology IS-IS.
D. Disable the IPv4 unicast address-family.
E. Enable IS-IS wide metric to support the single-topology mode.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 17
When implementing OSPF, which type of networks require DR/BDR election?

A. point-to-point networks

B. mutli-access broadcast networks

C. non-broadcast multi-access networks (Hub and Spoke Frame Relay) using point-to-multipoint OSPF network type

D. All networks type

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 18
The S bit in the MPLS header is used for what purpose?

A. To indicate the bottom level in the label stack

B. To indicate if LDP is sync to the IGP

C. To indicate if LDP is sync to the IGP

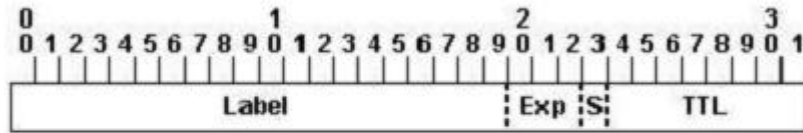D. To indicate the status of the LSP

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A label is a short, four-byte, fixed-length, locally-significant identifier which is used to identify a Forwarding Equivalence Class (FEC). The label which is put on a particular packet represents the FEC to which that packet is assigned.

```
0                   1                   2                   3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
|              Label              | Exp |S|     TTL           |
```

- Label—Label Value (Unstructured), 20 bits

- Exp—Experimental Use, 3 bits; currently used as a Class of Service (CoS) field.

- S—Bottom of Stack, 1 bit

- TTL—Time to Live, 8 bits

a 1-bit bottom of stack flag. If this is set, it signifies that the current label is the last in the stack.

QUESTION 19
What are two purposes of the BGP scan-time command? (Choose two.)

A. to tune the BGP process which walks the BGP table and confirms the reachability of next hops
B. to allow faster detection of downed BGP peers
C. to improve BGP convergence time
D. to tune the BGP update interval
E. to decrease the effects of unstable routes by increasing the route suppression time

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:

http://www.networkers-online.com/blog/2008/12/bgp-performance-tunning-convergence-stability-scalabilityand-nsf-part-2/

Background BGP scanner
It is responsible for BGP housekeeping by scanning both the BGP RIB and the IP RIB and cleaning and sorting things out.
BGP monitors the next hop of the installed routes to verify next-hop reachability and to select, install, and validate the BGP best path. By default, the

BGP scanner polls the RIB for this information every 60 seconds.

During the 60 second time period between scan cycles, IGP instabilities or other network failures can cause temporarily black holes and routing loops.

NOTE With Cisco IOS the default timer is 60 seconds for the IPv4 address family and 15 seconds for the

VPNv4 address family in order to optimize the VPNs routing table convergence.

This timer can be controlled via the following command:
Router(config-router)#bgp scan-time <5-60>

QUESTION 20
When using the show bgp ipv6 unicast summary command to verify the IPv6 BGP session status with the IPv6 BGP peers, you noticed the "St/PfxRcd" status for one of the IPv6 BGP peers is in the "Active" state. What does the "Active" state indicate?

A. The IPv6 BGP session has been established with the IPv6 BGP peer.
B. The router is in the process of sending BGP routing updates to the IPv6 BGP peer.
C. The router is in the process of establishing the IPv6 BGP session with the IPv6 BGP peer.
D. The router is exchanging BGP notification messages with its IPv6 BGP peer.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 21
Which three BGP configuration groupings are supported on Cisco IOS XR Software? (Choose three.)

A. peer-group
B. af-group
C. bgp-group
D. session-group
E. neighbor-group
F. as-group

Correct Answer: BDE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

•Commands relating to a peer group found in Cisco IOS Release 12.2 have been removed from Cisco IOS XR software. Instead, the af-group, session-group, and neighbor-group configuration commands are added to support the neighbor in Cisco IOS XR software:
–The af-group command is used to group address family-specific neighbor commands within an IPv4 or IPv6 address family. Neighbors that have the same address family configuration are able to use the address family group name for their address family-specific configuration. A neighbor inherits the configuration from an address family group by way of the use command. If a neighbor is configured to use an address family group, the neighbor will (by default) inherit the entire configuration from the address family group. However, a neighbor will not inherit all of the configuration from the address family group if items are explicitly configured for the neighbor.
–The session-group command allows you to create a session group from which neighbors can inherit address family-independent configuration. A neighbor inherits the configuration from a session group by way of the use command. If a neighbor is configured to use a session group, the neighbor (by default) inherits the session group's entire configuration. A neighbor does not inherit all the configuration from a session group if a configuration is done directly on that neighbor.

–The neighbor-group command helps you apply the same configuration to one or more neighbors. Neighbor groups can include session groups and address family groups. This additional flexibility can create a complete configuration for a neighbor. Once a neighbor group is configured, each neighbor can inherit the configuration through the use command. If a neighbor is configured to use a neighbor group, the neighbor (by default) inherits the neighbor group's entire BGP configuration.

–However, a neighbor will not inherit all of the configuration from the neighbor group if items are explicitly configured for the neighbor. In addition, some part of the neighbor group's configuration could be hidden if a session group or address family group was also being used.

QUESTION 22
Which AS path access list is used by a multihomed customer to only announce their own address space to their service providers to prevent the multihomed customer from becoming a transit AS?

A. ip as-path access-list permit.*
B. ip as-path access-list permit^$
C. ip as-path access-list permit _$
D. ip as-path access-list permit _^
E. ip as-path access-list permit local-asip as-path access-list deny *
F. ip as-path access-list deny.*ip as-path access-list permit *

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

```
+-----------------------------------------------------------+
| CHAR | USAGE                                              |
+------|----------------------------------------------------|
|  ^   | Start of string                                    |
|------|----------------------------------------------------|
|  $   | End of string                                      |
|------|----------------------------------------------------|
|  []  | Range of characters                                |
|------|----------------------------------------------------|
|  -   | Used to specify range ( i.e. [0-9] )               |
|------|----------------------------------------------------|
|  ( ) | Logical grouping                                   |
|------|----------------------------------------------------|
|  .   | Any single character                               |
|------|----------------------------------------------------|
|  *   | Zero or more instances                             |
|------|----------------------------------------------------|
|  +   | One or more instance                               |
|------|----------------------------------------------------|
|  ?   | Zero or one instance                               |
|------|----------------------------------------------------|
|      | Comma, open or close brace, open or close          |
|  _   | parentheses, start or end of string, or space      |
+-----------------------------------------------------------+
```

Some commonly used regular expressions include:

```
+-------------+------------------------------+
| Expression  | Meaning                      |
|-------------|------------------------------|
| .*          | Anything                     |
|-------------|------------------------------|
| ^$          | Locally originated routes    |
|-------------|------------------------------|
| ^100_       | Learned from AS 100          |
|-------------|------------------------------|
| _100$       | Originated in AS 100         |
|-------------|------------------------------|
```

QUESTION 23
What is recursive lookup in BGP and how does it work?

A. The router looks up the EBGP route and the EBGP next hop to reach a destination in the remote AS. Then the router looks up the route to reach the EBGP next hop using the IBGP.
B. The router looks up the IBGP route and the IBGP next hop to reach a destination in the remote AS. Then the router looks up the route to reach the IBGP next hop using the EBGP.
C. The router looks up the BGP route and the BGP next hop to reach a destination in the remote AS. Then the router looks up the route to reach the BGP next hop using the IGP.
D. The router looks up the route and the next hop to reach a destination in the remote AS using the IGP. Then the router looks up the route to reach the next hop using BGP.
E. The router perform three routing lookups to determine the route to reach a destination in the remote AS. The first lookup is done using EBGP, the second lookup is done using IBGP, and the third lookup is done using the IGP.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A few different approaches are available to deal with iBGP and synchronization. We may turn on the synchronization option on our routers and wait for the IGP to have a route for the destination before it's advertised to peers. Another option is to simply use a full mesh, so that iBGP convergence isn't an issue.
Clearly that isn't going to happen when a network's core needs to scale: it will implement something like reflectors that cause iBGP's full mesh to be broken.
The real alternative, if you don't enable synchronization, is to use route recursion. A recursive route lookup uses the BGP next-hop attribute to actually make a different route lookup. The IGP can use the destination network instead of the AS-path to determine where it gets sent. Even if the iBGP hasn't converged, the routers will still know how to get to that network, since it will exist in the router it was advertised from, who will know the next-hop.

QUESTION 24
Which reserved AS number or range of numbers is used for backward compatibility between old BGP peers using 16-bit AS number and new BGP peers using 32-bit AS number?

A. AS 65001 to 65535
B. AS 65512 to 65535
C. AS 12345
D. AS 23456
E. AS 64001

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6554/ps6599/4byte_asnios.pdf

New Reserved AS#
AS_TRANS = AS #23456
2-byte placeholder for a 4-byte AS number
Used for backward compatibility between OLD and NEW BGP speakers

QUESTION 25
Which BGP attribute is a set of generic tags that can be used to signal various routing policies between BGP routers?

A. AS path
B. MED
C. weight
D. communities
E. route tags

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a00801475b2.shtml

QUESTION 26
Which of the following is a characteristic of dual-multihomed connectivity between an enterprise network and the service provider network or networks?

A. An enterprise network that is connected to two or more different service providers with two or more links per service provider and using BGP to exchange routing updates with the service providers.
B. Each service provider announces a default route on each of the links that connect to the customer with a different metric.
C. An enterprise network announces a default route to each service provider.
D. Load balancing can be achieved using the maximum-paths command.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 27
What are two ways to advertise networks into BGP? (Choose two.)

A. using the neighbor router BGP command
B. using a route policy in Cisco IOS XR Software or using a route map in Cisco IOS Software or Cisco IOS XE Software
C. using route redistribution into BGP
D. using the network router BGP command
E. enabling an interface to run BGP using the interface router BGP command

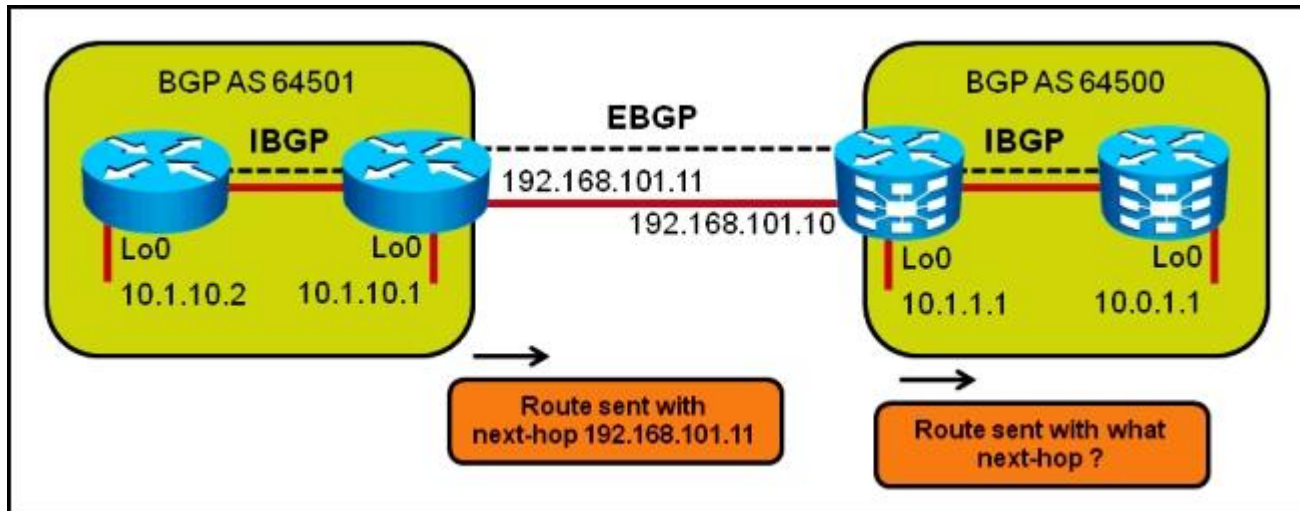**Correct Answer:** CD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 28
Refer to the network diagram in the exhibit.

Assuming the IBGP session within AS 64500 was established using the loopback 0 interface between the two routers, by default, what will be the next hop of the routes from AS 64501 when the routes appear on the router running IBGP only in AS 64500?

A. 192.168.101.11
B. 192.168.101.10
C. 10.1.1.1
D. 10.0.1.1
E. 10.1.10.1

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 29
When configuring BGP on Cisco IOS XR Software, which address-family is enabled by default?

A. IPv4 unicast
B. IPv6 unicast
C. VPNv4
D. IPv4 unicast and IPv6 unicast

E. IPv4 unicast and IPv6 unicast and VPNv4

F. No address-family is enabled by default.

Correct Answer: F
Section: (none)
Explanation

Explanation/Reference:
Explanation:

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.8/routing/command/reference/rr38bgp.pdf

An address family must be explicitly configured in the router configuration mode for the address family to be active in BGP. Similarly, an address family must be configured under the neighbor for the BGP session to be established for that address family. An address family must be configured in router configuration mode before it can be configured under a neighbor.

QUESTION 30
What are two characteristics of the multihomed customers to service providers connection option? (Choose two.)

A. Multihomed customers must use a private AS number.
B. The traffic load can be shared for different destination networks between service providers.
C. Multihomed customers must receive a full routing table from the service providers.
D. The routing methodology must be capable of reacting to dynamic changes. BGP is used to achieve this flexibility.
E. Multihomed customers must use a provider-assigned address space.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 31
Referring to the partial Cisco IOS-XR BGP configuration exhibit, when trying to commit this configuration, the following error is displayed:

% Failed to commit one or more configuration items during a pseudo-atomic operation. All changes made have been reverted. Please issue 'show configuration failed' from this session to view the errors.