

VCEplus.com

VCEplus.com

642-883

VCEplus.com

Number: 642-883
Passing Score: 800
Time Limit: 120 min
File Version: 5.0

Cisco 642-883 Questions & Answers

Deploying Cisco Service Provider Network Routing

(SPROUTE)

Version: 5.0
Cisco 642-883 Exam

Exam A

QUESTION 1

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 2

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 gi0/0/0/0
ip address 10.1.1.1 255.255.255.0
!
router ospf 1
network 10.1.1.1 0.0.0.0 area 0
```
- B.

```
interface gi0/0/0/0
ip address 10.1.1.1 255.255.255.0
!
router ospf 1
network 10.1.1.1 255.255.255.255 area 0
```
- C.

```
router ospf 1
area 0
interface GigabitEthernet0/0/0/1
```
- D.

```
interface gi0/0/0/0
ip address 10.1.1.1 255.255.255.0
ip ospf 1 area 0
```
- E.

```
router ospf 1
address-family ipv4 unicast
interface GigabitEthernet0/0/0/1 area 0
```
- F.

```
router ospf 1
```

```
address-family ipv4 unicast
interface GigabitEthernet0/0/0/1
area 0
```

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 3

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 4

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 5

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#wp_1148617

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 6

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 7

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

- A. point-to-point networks
- B. multi-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 8

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 9

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:

QUESTION 10

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 11

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 EIGRP, 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 12

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 13

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.s.html

QUESTION 14

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:

QUESTION 15

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 16

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 17

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 18

When using the Cisco IOS-XR show bgp command to examine the BGP table, the Metric value being displayed is used to represent which BGP attribute?

- A. Weight
- B. Local Preference
- C. MED
- D. Cost Community

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 19

In Cisco IOS and Cisco IOS XE Software images, when redistributing routes from other routing protocols into OSPF, what is a common reason why some of the routes might not be redistributed into OSPF?

- A. The OSPF external metric type (E1 or E2) is not defined.
- B. The OSPF seed metric is not defined.
- C. The OSPF level (Level 1, Level 2, or Level-1-2) to which the routes will be redistributed into is not defined.
- D. The subnets option in the redistribute command is missing.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 20

When configuring Cisco IOS route maps, which command allows the route map processing to jump to another statement instead of exiting?

- A. jump
- B. next
- C. set
- D. continue
- E. goto

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Match Operations With Continue Clauses

If a match clause does not exist in the route-map entry but a continue clause does, the continue clause will be automatically executed and go to the specified route-map entry. If a match clause exists in a route-map entry, the continue clause is executed only when a successful match occurs. When a successful match occurs and a continue clause exists, the route map executes the set clauses and then goes to the specified route-map entry. If the next route map contains a continue clause, the route map will execute the continue clause if a successful match occurs. If a continue clause does not exist in the next route map, the route map will be evaluated normally. If a continue clause exists in the next route map but a match does not occur, the route map will not continue and will "fall through" to the next sequence number if one exists.

Set Operations With Continue Clauses

Set clauses are saved during the match clause evaluation process and executed after the route-map evaluation is completed. The set clauses are evaluated and executed in the order in which they were configured. Set clauses are only executed after a successful match occurs, unless the route map does not contain a match clause. The continue statement proceeds to the specified route-map entry only after configured set actions are performed. If a set action occurs in the first route map and then the same set action occurs again, with a different value, in a subsequent route map entry, the last set action may override any previous set actions that were configured with the same set command unless the set command permits more than one value. For example, the set as-path pretend command permits more than one autonomous system number to be configured.

QUESTION 21

When configuring Cisco IOS XR route policy nesting, which command is used within a route policy to call another route policy?

- A. Apply
- B. Continue
- C. Jump
- D. Go to
- E. Call

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.7/routing/command/reference/rr37plcy.html

Examples

In the following example, the policy CustomerIn applies the route-policy SetLocalPref to conditionally set the local preference on a route. The parameters 20, 30, 40, and 50 are passed to the parameterized policy SetLocalPref, where the local preference is set to:

- 20, if the community 217:20 is present in the route
- 30, if the community 217:30 is present in the route
- 40, if the community 217:40 is present in the route
- 50, if the community 217:50 is present in the route

```
RP/0/RP0/CPU0:router(config)# route-policy SetLocalPref ($lp0, $lp1, $lp2, $lp3, $lp4)
RP/0/RP0/CPU0:router(config-rpl)# if community matches-any ($lp0:$lp1)then RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp1 RP/0/
RP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp2) then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp2 RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp3)
then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp3 RP/0/RP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp4)
then
RP/0/RP0/CPU0:router(config-rpl-elseif)# set local-preference $lp4

RP/0/RP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
RP/0/RP0/CPU0:router(config)# route-policy CustomerIn($cust) RP/0/RP0/CPU0:router(config-rpl)# apply SetLocalPref ($cust, 20, 30, 40, 50) RP/0/
RP0/CPU0:router(config-rpl)# end-policy
RP/0/RP0/CPU0:router(config)# route-policy Cust_217
RP/0/RP0/CPU0:router(config-rpl)# apply CustomerIn(217) RP/0/RP0/CPU0:router(config-rpl)# end-policy
```

QUESTION 22

What is defined by using the Cisco IOS XR policy-global configuration command?

- A. the default BGP route policy
- B. global variables that can be referenced by any route policy
- C. the global default route policy
- D. hierarchical route policy
- E. nested route policy

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Global Configuration Mode

Prompt: (config)

Enter global configuration mode from executive (EXEC) mode by using the configure command. Global configuration commands generally apply to the whole system rather than just one protocol or interface. You can enter all other configuration submodes listed in this section from global configuration mode.

RP/0/RP0/CPU0:router# configure

RP/0/RP0/CPU0:router(config)#

Global Parameter Configuration Mode

Prompt: (config-rpl-gl)

Enter global parameter configuration mode by using the policy-global command in global configuration mode.

In global parameter configuration mode, you can create or modify a global policy by entering successive commands and then terminating the configuration by entering the end-global command.

For example, to configure global parameters:

RP/0/RP0/CPU0:router(config)# policy-global

RP/0/RP0/CPU0:router(config-rp-gl)# glbpathtype ebgp

RP/0/RP0/CPU0:router(config-rp-gl)# glbtag 100

RP/0/RP0/CPU0:router(config-rp-gl)# end-global

QUESTION 23

What is the default OSPF seed metric and type?

- A. 10 and E1
- B. 10 and E2
- C. 20 and E1
- D. 20 and E2
- E. 0 and E1
- F. 0 and E2

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 24

When redistributing routes into OSPF, which kind of routes will be redistributed by default on Cisco IOS XR Software but will not be automatically redistributed by default on Cisco IOS Software and Cisco IOS XE Software?

- A. Type 1 external routes
- B. Type 2 external routes
- C. subnetted routes
- D. /32 host routes
- E. tagged static routes

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Table 12 Redistribution Configurations for OSPF in Cisco IOS and Cisco IOS XR

Cisco IOS	Cisco IOS XR
<pre>router ospf 1 redistribute connected subnets redistribute static subnets route-map map</pre>	<pre>router ospf lab redistribute connected redistribute static policy</pre>

You no longer need the **subnets** keyword in Cisco IOS XR, and you declare an RPL policy rather than a route map.

If you want to redistribute only classful routes into OSPF in Cisco IOS XR, you can use the **redistribute** command with the **classful** keyword. The use of policies rather than route maps is covered in the ["Border Gateway Protocol \(BGP\)"](#) section and also in the existing Cisco IOS XR documentation on RPL. The mainpiece of information to get from this discussion is that redistribution into OSPF on Cisco IOS XR is a policy attach point.

QUESTION 25

When redistributing EIGRP routes into OSPF as type E2 external OSPF routes, what is the default OSPF seed metric?

- A. Infinite
- B. 0
- C. 1

- D. 10
- E. 20

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

QUESTION 26

Which high-availability routing feature requires the neighbor router to support the graceful restart capability?

- A. BFD
- B. NSR
- C. NSF
- D. MTR

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

On Cisco IOS XR software, NSF minimizes the amount of time a network is unavailable to its users following a route processor (RP) failover. The main objective of NSF is to continue forwarding IP packets and perform a graceful restart following an RP failover.

When a router restarts, all routing peers of that device usually detect that the device went down and then came back up. This transition results in what is called a routing flap, which could spread across multiple routing domains. Routing flaps caused by routing restarts create routing instabilities, which are detrimental to the overall network performance. NSF helps to suppress routing flaps in NSF-aware devices, thus reducing network instability.

NSF allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following an RP failover. When the NSF feature is configured, peer networking devices do not experience routing flaps. Data traffic is forwarded through intelligent line cards while the standby RP assumes control from the failed active RP during a failover. The ability of line cards to remain up through a failover and to be kept current with the Forwarding Information Base (FIB) on the active RP is key to NSF operation.

When the Cisco IOS XR router running IS-IS routing performs an RP failover, the router must perform two tasks to resynchronize its link-state database with its IS-IS neighbors. First, it must relearn the available IS-IS neighbors on the network without causing a reset of the neighbor relationship. Second, it must reacquire the contents of the link-state database for the network. The IS-IS NSF feature offers two options when configuring NSF:

- IETF NSF
- Cisco NSF

If neighbor routers on a network segment are NSF aware, meaning that neighbor routers are running a software version that supports the IETF Internet draft for router restartability, they assist an IETF NSF router that is restarting. With IETF NSF, neighbor routers provide adjacency and link-state information to help rebuild the routing information following a failover. In Cisco IOS XR software, Cisco NSF checkpoints (stores persistently) all the state necessary to recover from a restart without requiring any special cooperation from neighboring routers. The state is recovered from the neighboring routers, but only using the standard features of the IS-IS routing protocol. This capability makes Cisco NSF suitable for use in networks in which other routers have not used the IETF standard implementation of NSF

QUESTION 27

Which high-availability mechanism is a detection protocol that is enabled at the interface and at the routing protocol levels?

- A. NSF
- B. SSO
- C. NSR
- D. BFD
- E. SDR

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/fs_bfd.html

QUESTION 28

Which OSPF feature allows a router with redundant route processors to maintain its OSPF state and adjacencies across planned and unplanned RP switchovers and does this by checkpointing state information from OSPF on the active RP to the standby RP? This feature does not require the OSPF neighbor to support graceful restart.

- A. NSR
- B. NSF
- C. BFD
- D. MTR
- E. SDR

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Continuous Forwarding

An important aspect of high availability is maintenance of traffic forwarding, even in the case of control-plane switchovers. Cisco IOS XR Software has several built-in features that can provide continuous forwarding, including RSP stateful switchover (SSO), Nonstop Forwarding (NSF), Graceful Restart, and NSR.

NSF: Cisco IOS XR Software supports forwarding without traffic loss during a brief outage of the control plane through signaling and routing protocol implementations for Graceful Restart extensions as standardized by the IETF. In addition to standards compliance, this implementation has been compatibility tested with Cisco IOS Software and third-party operating systems.

Graceful Restart: This control-plane mechanism ensures high availability by allowing detection and recovery from failure conditions while preserving NSF services. Graceful Restart is a way to recover from signaling and control-plane failures without affecting the forwarding plane. Cisco IOS XR Software uses this feature and a combination of check pointing, mirroring, RSP redundancy, and other system resiliency features to recover prior to timeout and avoid service downtime as a result of network reconvergence.

NSR: This feature allows for the forwarding of data packets to continue along known routes while the routing protocol information is being refreshed following a processor switchover.

NSR maintains protocol sessions and state information across SSO functions for services such as Multiprotocol Label Switching (MPLS) VPN. TCP connections and the routing protocol sessions are migrated from the active RSP to the standby RSP after the RSP failover without letting the peers know about the failover. The sessions terminate locally on the failed RSP, and the protocols running on the standby RSP reestablish the sessions after the standby RSP goes active, without the peer detecting the change. You can also use NSR with Graceful Restart to protect the routing control plane during switchovers. The Cisco IOS XR Operating System provides system resiliency through a comprehensive set of high-availability features including modularity, process restart, fault handling, continuous forwarding, and upgradability.

QUESTION 29

Which two statements regarding OSPFv2 or OSPFv3 authentication are correct? (Choose two.)

- A. OSPFv2 supports MD5 authentication.
- B. OSPFv2 supports MD5 or SHA authentication.
- C. OSPFv2 relies on the native security stack that uses IPsec.
- D. OSPFv3 supports MD5 authentication.
- E. OSPFv3 supports MD5 or SHA authentication.
- F. OSPFv3 relies on the native security stack that uses IPsec.

Correct Answer: AF

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 30

What are three common problems that can cause a BGP neighbor state to toggle between the idle state and the active state? (Choose three.)

- A. BGP network command misconfiguration
- B. route policy misconfiguration
- C. AS number misconfiguration
- D. route map misconfiguration
- E. BGP neighbor peering to wrong IP address
- F. IGP routing problem: not able to reach the source IP address of the BGP open packet

Correct Answer: CEF

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 31

BGP peerings can be secured using which protection mechanism?

- A. SHA authentication
- B. MD5 authentication
- C. SSH
- D. SSL
- E. AAA
- F. digital certificates

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080b52107.s.html

QUESTION 32

In an AS with 10 routers running IBGP, how many IBGP sessions will be required to establish fully meshed IBGP peerings?

- A. 10
- B. 20

- C. 45
- D. 50
- E. 99
- F. 100

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 33

You are implementing OSPF as the IGP using a single OSPF area design. The router memory usage for OSPF is too high. Which two methods can lower the OSPF memory usage? (Choose two.)

- A. Enable OSPF synchronization
- B. Implement multi-area OSPF
- C. Enable OSPF sham links to reduce the number of LSAs
- D. Implement route summarization on the ABRs
- E. Enable route leaking between Level 1 and Level 2 areas

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 34

Routes that are received from an IBGP peer will be propagated to which other routers by default?

- A. to the EBGP peers only
- B. to the IBGP peers only
- C. to both EBGP and IBGP peers
- D. to no other peers

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 35

What are two consequences of having constant link flaps, resulting in the OSPF neighbor adjacencies going up and down repeatedly? (Choose two.)

- A. routes getting into the "Stuck In Active" state
- B. constant flooding of LSAs
- C. OSPF route dampening to occur
- D. many SPF recalculations
- E. routing loops may temporarily be introduced into the network

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 36

When monitoring the OSPF operations in the network, you notice that the OSPF process is not load balancing traffic across multiple OSPF links. Which configuration adjustment could be made to enable the OSPF process to load balance traffic across multiple OSPF links?

- A. Enable BFD on the OSPF interfaces and on the OSPF routing process.
- B. Enable NSR on the OSPF routing process.
- C. Enable NSF on the OSPF routing process.
- D. Adjust the OSPF cost on the OSPF interfaces.
- E. Adjust the OSPF variance parameter.
- F. Use the wide-style OSPF metric instead of the narrow-style OSPF metric.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 37

Which of the following is used by an IS-IS router to detect other IS-IS neighbor routers and to form adjacencies?

- A. ESH
- B. ISH
- C. IIH
- D. PSNP
- E. CSNP

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

ISIS actually features three different hello types

An ES Hello (ESH) is sent by all End Systems, and all IS devices listen for this Hello. An IS Hello (ISH) announces the presence of an IS - An IS Hello is sent by all IS devices, and End Systems listen for these hellos. Finally, an IS-to-IS Hello (IIH) is used by an IS to discover other ISes and to form adjacencies with them. A router will send an IIH to another router on the link to form or maintain an adjacency, but it will still send an ISH as well in case there are end systems located on that segment.

QUESTION 38

Instructions

- Enter the proper CLI commands and analysis the outputs on the Cisco routers to answer the multiple-choice questions.
- From the network topology diagram, click on each of the router icon to gain access to the console of each router.
- No console or enable passwords are required.
- There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.
- Not all the CLI commands or commands options are supported or required for this simulation.
- For example, the show running-config and the ping commands are **NOT** supported in this simulation.
- All the devices in this simulation have been pre-configured and you are not required to enter in any configurations.

Scenario


Referring to the network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers and interpret the supported CLI commands outputs to answer the four multiple choice questions.

Note: The CE5 router is an IOS router, the PE5 router is an IOS-XR router, and the PE6 router is an IOS-XE router.


Click on the CE5/PE5/PE6 icons to access the respective router console



IGP = IS-IS







There are how many IS-IS area(s) defined in this network Scenario?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation: #show clns

QUESTION 39

Which two of the following are true regarding the BGP Prefix-Based outbound route filtering feature? (Choose two.)

- A. IP multicast routes are not supported.
- B. Outbound route filtering is configured only on a per-address family basis.
- C. Outbound route filtering can be configured for either iBGP or eBGP sessions.
- D. The outbound route filter can be defined in a Prefix list, Distribute list or Access lists.
- E. Outbound route filtering is more effective when a distance vector IGP is used.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fsbgporf.pdf

Restrictions for BGP Prefix-Based Outbound Route Filtering · The BGP Prefix-Based Outbound Route Filtering feature does not support IP multicast routes. · IP addresses that are used for outbound route filtering must be defined in an IP prefix-list. BGP distribute lists and IP access lists are not supported. · Outbound route filtering is configured on only a per-address family basis and cannot be configured under the general session or BGP routing process (Router(config-router)#). · Outbound route filtering is configured for only external peering sessions.

QUESTION 40

How can you prevent multihomed customers with connections to two service providers from acting as a transit AS?

- A. Enable BGP synchronization on all the customer routers
- B. Use MED to influence the inbound traffic from the ISPs
- C. Use static routing to the ISPs
- D. Use an AS-path access-list to filter the BGP updates to the ISPs
- E. Use conditional advertisements when sending BGP updates to the ISPs

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 41

When implementing LDP, what is liberal label retention mode?

- A. To reduce the convergence time, the LSR will retain all the received labels in its LIB even if all the neighbor LSRs go down temporary
- B. The LSR can assign it's own label for each destination network even though it has not been assigned a next hop label from the neighbor LSR
- C. The LSR will store the label received from the downstream LSRs in its LIB even if the downstream LSRs are not the next hop for the destination
- D. The LSR will not perform PHP if it is operating in liberal label retention mode

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

By default, LDP accepts labels (as remote bindings) for all prefixes from all peers. LDP operates in liberal label retention mode, which instructs LDP to keep remote bindings from all peers for a given prefix. For security reasons, or to conserve memory, you can override this behavior by configuring label binding acceptance for set of prefixes from a given peer.

The ability to filter remote bindings for a defined set of prefixes is also referred to as LDP inbound label filtering

QUESTION 42

Which two statements about a transit AS are correct? (Choose two.)

- A. A transit AS has eBGP connection(s) to only one external AS.
- B. Routes between ASs are always exchanged via eBGP.
- C. A transit AS uses an IGP like OSPF or ISIS to propagate the external networks within the transit AS.
- D. Core routers within a transit AS normally use default routing to reach the external networks.
- E. iBGP sessions can be established between non directly connected routers.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 43

Which two statements are true about an EBGP session or an IBGP session? (Choose two.)

- A. IBGP uses AS-Path to detect routing information loops within the AS.
- B. EBGP routes have a default Admin Distance of 20 and IBGP routes have a default Admin Distance of 200.
- C. No BGP attributes are changed in EBGP updates except for the next-hop attribute if next-hop- self is configured.

- D. Routes learned from an EBGp peer not advertised to another EBGp peer to prevent routing information loops.
- E. IBGP uses split horizon to prevent routing information loops; routes learned from an IBGP peer are not advertised to another IBGP peer.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 44

Which option describes the configuration for a Cisco IOS XR OSPF router to give it the highest chances of becoming DR?

- A. ip ospf priority 0 under the OSPF interface statement
- B. ip ospf priority 0 under the OSPF area statement
- C. ip ospf priority 255 under the OSPF interface statement
- D. ip ospf priority 255 under the OSPF area statement
- E. Leave priority unchanged, but assign a high-numbered IP address on a loopback interface.
- F. Leave priority unchanged, but assign a high-numbered IP address on any physical interface.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 45

Which option describes how an engineer must configure route summarization on a Cisco IOS XE OSPF-enabled network of routers?

- A. area range command under the router OSPF process of the ABR
- B. area range command under the router OSPF process of the router wanting to originate the summary route from
- C. summary-address command under the OSPF process of the ABR
- D. summary-address command under the OSPF process of the router wanting to originate the summary route from
- E. network statement for the aggregate address under the router OSPF process on the ABR
- F. network statement for the aggregate address under the router OSPF process of the router wanting to originate the summary route from

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 46

Which series of commands configures area 1 as an OSPF totally stubby area on a Cisco IOS XR router?

- A.

```
router ospfv3 1
router-id 10.10.10.1
area 0
interface GigabitEthernet 0/0/0/1
area 1
stub no-summary
interface GigabitEthernet 0/0/0/2
```
- B.

```
router ospfv3 1
router-id 10.10.10.1
default-information originate
area 0
interface GigabitEthernet 0/0/0/1
area 1
stub
interface GigabitEthernet 0/0/0/2
```
- C.

```
ipv6 prefix-list default permit ::0/0
router ospfv3 1
router-id 10.10.10.1
default-information originate
area 0
interface GigabitEthernet 0/0/0/1
distribute-list prefix-list default in
area 1
interface GigabitEthernet 0/0/0/2
```
- D.

```
router ospfv3 1
router-id 10.10.10.1
area 0
interface GigabitEthernet 0/0/0/1
area 1
interface GigabitEthernet 0/0/0/2
no-summary
passive
```
- E.

```
router ospfv3 1
router-id 10.10.10.1
```

```
area 0
interface GigabitEthernet 0/0/0/1
area 1
stub
interface GigabitEthernet 0/0/0/2
default-cost 20
```

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 47

Refer to the exhibit.

Area 1 routers must have only a default route in their OSPF routing tables. Which configuration must be completed on XR2 to meet this requirement?

- A. router ospf 1
address-family ipv4
interface GigabitEthernet0/0/0/1
area 1 stub no-summary
- B. router ospf 1
area 1 stub no-summary
- C. router ospf 1
area 1 stub
interface GigabitEthernet0/0/0/1
- D. router ospf 1
area 1 stub no-summary
interface GigabitEthernet0/0/0/1

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 48

Which configuration implements the most secured OSPF authentication mode on Cisco IOS XE?

- A. interface Gi0/0/0/0
ip ospf authentication message-digest
ip ospf message-digest-key 1 md5 cisco
- B. router ospf 1
area 1
authentication message-digest
interface GigabitEthernet0/0/0/0
message-digest-key 1 md5 cisco
- C. router ospf 1
area 1
interface GigabitEthernet0/0/0/0
authentication message-digest
message-digest-key 1 md5 cisco
- D. router ospf 1
area 1 authentication
interface GigabitEthernet0/0/0/0
ip ospf authentication-key cisco
- E. interface Gi0/0/0/0
ip ospf authentication
ip ospf authentication-key cisco

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 49

Which three potential issues can cause an OSPF adjacency to not come up? (Choose three.)

- A. wrong area configured on OSPF peers
- B. wrong authentication key configured on OSPF peers
- C. different OSPF instance numbers configured on OSPF peers
- D. different OSPF instances and VRFs configured with the command: router ospf <number1> vrf <number2>
- E. having the no passive-interface statement configured for the interfaces on which OSPF adjacencies are expected to form
- F. Layer 2 connectivity problems on the shared medium of the participating OSPF-enabled routers

Correct Answer: ABF

Section: (none)

Explanation

Explanation/Reference:

QUESTION 50

Refer to the exhibit.

R1 and XR2 cannot establish OSPF adjacency. Which option is the root cause of this issue?

- A. stub flag mismatch
- B. mismatched passwords
- C. process ID mismatch
- D. mismatched authentication types
- E. MTU mismatch

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 51

A network engineer wants to test a path between any two points in the network for multicast traffic to identify congestion points in the network. Which command in a Cisco IOS XE router shows multicast traffic drops?

- A. mstat
- B. mtrace
- C. mrinfo
- D. ping

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 52

Which option is the minimum required configuration to enable PIM on a Cisco IOS XR router interface?

- A. router pim
address-family ipv4
interface GigabitEthernet0/0/0/0
- B. router pim
interface GigabitEthernet0/0/0/0
address-family ipv4
enable
- C. router pim
address-family ipv4
interface GigabitEthernet0/0/0/0
enable
- D. router pim
interface GigabitEthernet0/0/0/0
enable
- E. router pim
address-family ipv4
interface GigabitEthernet0/0/0/0
address-family ipv4
enable

Correct Answer: D

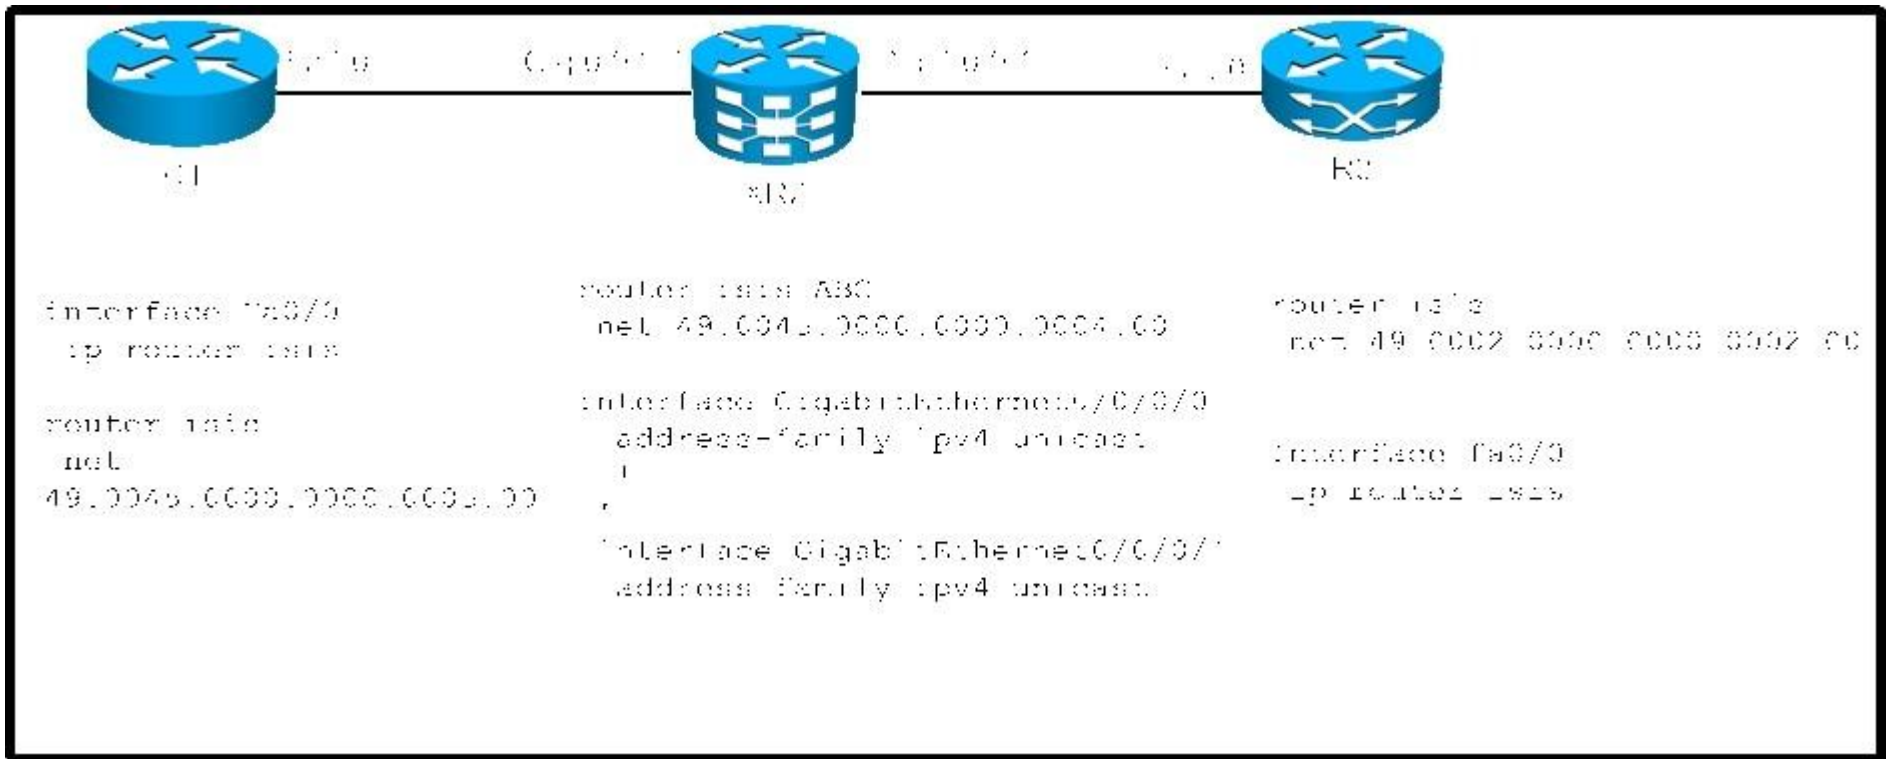
Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 53



Which configuration on XR2 provides R1 with only the minimum routing information to reach the rest of the network?

- A. Router isis ABC
interface GigabitEthernet0/0/0/1
circuit-type level-2-only
- B. Router isis ABC
interface GigabitEthernet0/0/0/0
circuit-type level-1-2
- C. Router isis ABC
interface GigabitEthernet0/0/0/1
circuit-type level-1
- D. Router isis ABC
interface GigabitEthernet0/0/0/0
circuit-type level-1

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 54

Refer to the exhibit.

R2 does not have R1 1.1.1.1/32 subnet in its BGP table. Which issue can be causing this problem?

- A. XR4 does not have next-hop-self for a R2 iBGP session.
- B. R2 does not have no auto-summary configured.
- C. XR4 drops any update coming from R1.
- D. IPv4 address-family is not enabled on R1.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 55

Refer to the exhibit.



Which configuration is required on XR2 to enable MD5 IS-IS Authentication between R1 and XR2 using the password CISCO?

- A. Router isis ABC
interface GigabitEthernet0/0/0/0
hello-password encrypted CISCO

- B. Router isis ABC
interface GigabitEthernet0/0/0/0
hello-password hmac-md5 CISCO
- C. Router isis ABC
lsp-password hmac-md5 CISCO
- D. Router isis ABC
lsp-password encrypted CISCO

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 56

Refer to the exhibit.

IS-IS adjacency is not established between XR2 and R1. Which action fixes this issue?

- A. unify IS-IS process IDs on each router
- B. configure on R1 under IS-IS is-type level-1-2
- C. configure on R1 under IS-IS area-password CISCO
- D. configure on XR2 under IS-IS under Gi0/0/0/0 address-family ipv4 unicast

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 57

An engineer wants to define a route policy that sets MED to a given value in case the extended community attributes of a network prefix match at least one of the values: 65000:10, 65000:20, or 65000:30. Which configuration accomplishes this task?

- A. extcommunity-set rt my-list
65000:10
65000:20
65000:30
end-set

```
route-policy my-list ($med)
if extcommunity rt matches-within my-list then
set met $med
```

- B. extcommunity-set rt my-list
65000:10
65000:20
65000:30
end-set
route-policy my-list (\$med)
if extcommunity rt matches-any my-list then
set met \$med
- C. extcommunity-set rt my-list
65000:10
65000:20
65000:30
end-set
route-policy my-list (\$med)
if extcommunity rt matches-within my-list or med is \$med is then set met \$med
- D. extcommunity-set rt my-list
65000:10
65000:20
65000:30
end-set
route-policy my-list (\$med)
if extcommunity rt matches-any my-list or med is \$med then pass

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 58

Under which two conditions does an enterprise require BGP peering with its upstream Internet provider? (Choose two.)

- A. The enterprise has two exit points that are geographically separated and routing must be controlled so that delay variations are reduced.
- B. The enterprise has one exit point and must announce the IP class it has been assigned from its ISP.
- C. The enterprise has two Layer 2 links to the upstream ISP in the same location and wants to achieve redundancy.
- D. The enterprise needs a simple and efficient way to ensure proper routing for its AP subnets on its single ISP link.
- E. The enterprise is multihomed to two ISPs and has its own IP space AS number.

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 59

An engineer is configuring an eBGP peering session. What is the default TTL value?

- A. 1
- B. 64
- C. 127
- D. 255

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 60

Which two features are specific only to Tier 3 providers? (Choose two.)

- A. interconnects with Tier Level 1 and Tier Level 2 ISPs in public peering points
- B. only purchases transit from other networks to reach the Internet
- C. purchases transit and pays settlements to other service providers
- D. uses strict route filtering to allow only customer and local routes to be sent to upstream providers
- E. receives a full Internet routing table via BGP from its upstream provider, which it can use for load-balancing traffic

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 61

Which option is a tool that is used to ensure that BGP AS does not become a transit AS?

- A. as-path filter-list
- B. local-preference
- C. ttl-security
- D. confederations

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 62

An engineer is working for a service provider who is strategically important to the Internet data traffic. The CTO is promoting a core infrastructure upgrade to connect via optical links network service providers. Which three concerns must be addressed while planning the new BGP features and functionalities? (Choose three.)

- A. BGP suboptimal routing
- B. IPv4 exhaustion
- C. routers and switches hardware capabilities
- D. route dampening of current customers
- E. multihoming policies of new customers
- F. IGP reconvergence time on the inter-AS links
- G. BGP blackholing when acting as transit autonomous system

Correct Answer: ACG

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 63

A network specialist is working for a large enterprise that has a BGP-enabled multihomed core infrastructure. Recently, the ISP that provides Internet connectivity to the company had to temporarily disable one of the BGP sessions as a result of routers blackholing BGP prefixes in transit. Which option can the specialist use to avoid this incident in the future?

- A. filtering local prefixes outbound
- B. using only one router as primary BGP speaker

- C. permitting only local prefixes outbound
- D. filtering external prefixes inbound

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 64

Refer to the exhibit.

Which configuration is correct for XR4 for an e-BGP session with R1?

- A. router bgp 100
neighbor 10.10.14.1
remote-as 100
- B. router bgp 100
neighbor 10.10.14.1
remote-as 10
address-family ipv4 unicast
- C. router bgp 100
address-family ipv4 unicast
neighbor 10.10.14.1
remote-as 100
address-family ipv4 unicast
- D. router bgp 100
address-family ipv4 unicast
neighbor 10.10.14.1
remote-as 100

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 65

An engineer wants to use an address family to configure internal BGP peer 10.1.1.1 as a route reflector client for unicast and multicast prefixes. Which option accomplishes this configuration?

- A. router bgp 140
address-family ipv4 unicast
neighbor 10.1.1.1
remote-as 140
address-family ipv4 unicast
route-reflector-client
exit
address-family ipv4 multicast
route-reflector-client
- B. router bgp 140
neighbor 10.1.1.1
address-family ipv4 unicast
address-family ipv4 multicast
remote-as 150
route-reflector-client
exit
- C. router bgp 140
address-family ipv4 unicast
route-reflector-client
address-family ipv4 multicast
remote-as 140
neighbor 10.1.1.1
- D. router bgp 140
address-family ipv4 unicast
neighbor 10.1.1.1
remote-as 150
address-family ipv4 unicast
route-reflector-client
exit
address-family ipv4 multicast
route-reflector-client

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 66

A service provider is running BGP with clients at the edge of the network. The service provider sees that routing updates from one site are being dropped when the other site receives them.

Which feature fixes this issue?

- A. EBGp multihop
- B. inter-AS peering
- C. AS-override
- D. allow-AS in

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 67

Which statement best describes the function of a BGP af-group?

- A. Use templates in different neighbors so it makes the BGP implementation easier.
- B. Group different address family neighbors inside BGP.
- C. Create a group of neighbors that can inherit address family configuration inside BGP.
- D. Substitute the class maps inside BGP on Cisco IOS XR devices.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 68

For a Cisco IOS XR router, under which BGP configuration group can default-originate be added?

- A. session-group
- B. af-group
- C. neighbor-group
- D. peer-session template

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 69

A network engineer cannot connect different routers by eBGP using peer groups. Which two actions establish an adjacency between both routers? (Choose two.)

- A. Ensure that the peer group statements match across all routers.
- B. Ensure that ebgp-multihop is configured.
- C. Ensure that the remote-as statement is configured with the correct peer AS.
- D. Configure a weight value that is lower than the default value of 32768.
- E. Specify the weight attribute first.

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 70

In Cisco IOS XE, which option is the default threshold value for slow peer detection?

- A. 120 seconds
- B. 180 seconds
- C. 240 seconds
- D. 300 seconds

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 71

Which three valid session commands are supported by peer session templates? (Choose three.)

- A. local-as

- B. as-override
- C. inherit peer-session
- D. inherit peer-policy
- E. disable-connected-check
- F. route-reflector-client

Correct Answer: ACE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 72

Refer to the exhibit.

R1 is advertising subnets 1.1.1.X/32 via eBGP. XR4 must accept only the first seven subnets in its BGP table. Which configuration must occur on XR2 to meet these parameters?

- A. route-policy INBOUND
if destination in SUBNETS then pass endif
router bgp 100 neighbor 10.10.14.1 address-family ipv4 unicast route-policy INBOUND in
- B. route-policy INBOUND
if destination in FILTER then pass endif
router bgp 100 neighbor 10.10.14.1 address-family ipv4 unicast route-policy INBOUND in
- C. router bgp 100
neighbor 10.10.14.1 address-family ipv4 unicast
route-policy FILTER in
- D. router bgp 100
neighbor 10.10.14.1 address-family ipv4 unicast
route-policy SUBNETS in

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 73

Refer to the exhibit

XR4 must be configured to advertise only AS 100 local subnets to AS 10. Which configuration on XR4 achieves this goal?

- A. as-path-set LOCAL
ios-regex ^\$
route-policy ADV
if as-path in LOCAL then pass endif
- B. as-path-set LOCAL
ios-regex '^100\$'
route-policy ADV
if as-path in LOCAL then pass endif
- C. as-path-set LOCAL
ios-regex '\$'
route-policy ADV
if as-path in LOCAL then pass endif
- D. as-path-set LOCAL
ios-regex ^100\$
route-policy ADV
if as-path in LOCAL then pass endif

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 74

Refer to the exhibit.

XR1 must have XR4 as the primary exit point for only the first three subnets received from AS 100 and may not impact the rest of the subnets. Which configuration on XR1 is correct?

- A. prefix-set PREFER
100.100.100.0/30 le 32
route-policy LP
if destination in PREFER then set local-preference 200 endif
- B. route-policy LP
if destination in (100.100.100.0/30 le 32) then set local- preference 200 else pass endif
- C. prefix-set PREFER
100.100.100.0/29 le 32
route-policy LP
if destination in PREFER then set local-preference 200 else pass endif

D. route-policy LP
if destination in (100.100.100.0/29 le 32) then set local preference 200 else set local-preference 100 endif

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 75

A network engineer wants to set the local preference to 100 for the specific BGP community sets of 100:1, 101:1, and 102:1. Which option can the engineer use to implement this through RPL, where the BGP community set list is made by the name of lowpri-communities?

- A. route-policy ch-pref
if community matches lowpri-communities then
set local-preference 100
endif
end-policy
- B. route-policy ch-pref
if community eq lowpri-communities then
set local-preference 100
endif
end-policy
- C. route-policy ch-pref
if community matches-any lowpri-communities then
set local-preference 100
endif
end-policy
- D. route-policy ch-pref
if community neq lowpri-communities then
set local-preference 100
endif
end-policy

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 76

Which RPLs can be used to conditionally advertise a default route based on the prefix 1.1.1.1 being in the routing table?

- A. route-policy RPL_Default_Originate
if rib-has-route in (1.1.1.1/32)
pass
else
drop
endif
end-policy
- B. route-policy RPL_Default_Originate
if rib-has-route in (1.1.1.1/32)
drop
else
pass
endif
end-policy
- C. route-policy RPL_Default_Originate
if route in (1.1.1.1/32)
pass
else
drop
endif
end-policy
- D. route-policy RPL_Default_Originate
if not rib-has-route in (1.1.1.1/32)
pass
else
drop
endif
end-policy

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 77

Which configuration is used when matching a prefix-list <name> in a route map?

- A. match ip address prefix-list mylist
- B. match ip address mylist
- C. match ip next-hop mylist
- D. ip local-policy route-map mylist
- E. ip-policy route-map mylist

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 78

An engineer is using configuration blocks to build policy-based routing. All other traffic must not be affected and the policy must be applied to GigabitEthernet1/0. Which block correctly implements traffic that is destined for 10.10.10.0/24 using a next hop of 192.168.55.1?

- A. access-list 110 permit ip any 10.10.10.0 0.0.0.255
route-map PBR permit 10
match ip address 110
set next-hop 192.168.55.1
interface GigabitEthernet1/0
ip policy route-map PBR
- B. access-list 110 permit ip any 10.10.10.0 0.0.0.255
route-map PBR permit 10
match ip address 110
set next-hop 192.168.55.1
route-map PBR permit 20
interface GigabitEthernet1/0
ip policy route-map PBR
- C. access-list 110 permit ip any 10.10.10.0 255.255.255.0 route-map PBR permit 10
match ip address 110
set next-hop 192.168.55.1
interface GigabitEthernet1/0
ip policy route-map PBR
- D. access-list 110 permit ip any 10.10.10.0 255.255.255.0 route-map PBR permit 10
match ip address 110
set next-hop 192.168.55.1
route-map PBR permit 20
interface GigabitEthernet1/0
ip policy route-map PBR

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 79

An engineer is working on routers within AS 100. Which regular expression can be used in an AS path access list to match locally originated routes from AS 100?

- A. ^100\$
- B. ^\$
- C. _100\$
- D. ^[0-9]*\$

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 80

Which option is a mechanism that can be implemented between two eBGP peers to communicate the routes each peer needs from the other?

- A. ORF
- B. ACL
- C. prefix list
- D. route map

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 81

A Cisco IOS XR router is a member in OSPF 1 and EIGRP 100 domains, and needs to redistribute OSPF learned routes into EIGRP. Which

configuration achieves this goal?

- A. router eigrp 100
address-family ipv4
redistribute ospf 1
- B. router eigrp 100
redistribute ospf 1 route-policy OS_INT0_EIG
route-policy OS_INT0_EIG set eigrp-metric 100 10 255 1 155
- C. router eigrp 100
address-family ipv4
redistribute ospf 1 route-policy OS_INT0_EIG
route-policy OS_INT0_EIG set eigrp-metric 100 10 255 1 155
- D. router eigrp 100
default-metric 100 1 255 1 1500
redistribute ospf 1

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 82

Which value must be configured when redistributing OSPFv2 into RIP?

- A. metric
- B. bandwidth
- C. delay
- D. MTU
- E. reliability

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 83

Which option can a network specialist use to configure connected route redistribution inside VRF "TEST" on Cisco IOS XR and allow only the prefix

10.10.10.0/24?

- A. route-policy ALLOW-CONN
if destination in PERMIT_PREFIX then
pass
else
drop
end-policy
prefix-set PERMIT_PREFIX
10.10.10.0/24
end-set
router bgp 65000
vrf TEST
rd 65000:10000
address-family ipv4 unicast
redistribute connected route-policy ALLOW-CONN
- B. route-policy ALLOW-CONN
if source in PERMIT_PREFIX then
pass
else
drop
end-policy
prefix-set PERMIT_PREFIX
10.10.10.0/24
end-set
router bgp 65000
vrf TEST
rd 65000:10000
address-family ipv4 unicast
redistribute connected route-policy ALLOW-CONN
- C. route-policy ALLOW-CONN
if protocol is connected and source in PERMIT_PREFIX then pass
else
drop
end-policy
prefix-set PERMIT_PREFIX
10.10.10.0/24
end-set
router bgp 65000
vrf TEST
rd 65000:10000
address-family ipv4 unicast
redistribute connected route-policy ALLOW-CONN"

D. route-policy ALLOW-CONN
if route-type is local and destination in PERMIT_PREFIX then pass
else
drop
end-policy
prefix-set PERMIT_PREFIX
10.10.10.0/24
end-set
router bgp 65000
vrf TEST
rd 65000:10000
address-family ipv4 unicast
redistribute connected route-policy ALLOW-CONN

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 84

A static default route has been inserted into the configuration of a Cisco IOS XE router. Which option redistributes the route into the local OSPF process?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 85

Refer to the exhibit.

Routers B and C in transit AS 200 are the exit points toward AS 100 and AS 300. Router B is running Cisco IOS XR and Router C IOS XE. Internally, AS 200 runs OSPF only for internal routing. Which two methods can an engineer use to ensure proper reachability between AS 100 and AS 300 for the needed routes? (Choose two.)

- A. Router B:
router ospf 200
redistribute bgp 200 tag 20
router bgp 200
address-family ipv4 unicast
redistribute ospf 200 metric-type 2
- B. Router B:
router ospf 200
redistribute bgp 200
router bgp 200
address-family ipv4 unicast
redistribute ospf 200
- C. Router C:
router ospf 200
redistribute bgp 200 metric-type 1 subnets tag 200
router bgp 200
address-family ipv4
redistribute ospf 200
- D. Router C:
router ospf 200
redistribute bgp 200 metric-type 1 subnets tag 200
router bgp 200
address-family ipv4
redistribute ospf 200 match external 2
- E. Router C:
router ospf 200
redistribute bgp 200 metric-type 1 subnets tag 200
router bgp 200
address-family ipv4
redistribute ospf 200 route-map match-tag match external route-map match-tag deny 10
match tag 200

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 86

Which two methods implement high availability in OSPF on Cisco IOS XR? (Choose two.)

- A. NSF
- B. NSR
- C. BFD
- D. fast hellos
- E. sham-links

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 87

Nonstop Forwarding works with the Stateful Switchover feature in Cisco IOS and IOS XR software. Which three restrictions for OSPF are true? (Choose three.)

- A. OSPF NSF for virtual links is not supported.
- B. OSPF NSF for virtual links is supported.
- C. OSPF NSF for sham links is not supported.
- D. OSPF NSF for sham links is supported.
- E. OSPF NSF supports NSF/SSO for IPv6 traffic only.
- F. OSPF NSF supports NSF/SSO for IPv4 traffic only.

Correct Answer: ACF

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 88

Which two mandatory tasks must an IS-IS NSF-capable router perform for RP switchover? (Choose two.)

- A. Relearn the available IS-IS neighbors.
- B. Reacquire the contents of the LSD.
- C. Reset peering with the available IS-IS neighbors.
- D. Keep the existing contents of the LSD.
- E. Rediscover DIS for each link segment.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 89

Which option is where BGP graceful restart is negotiated between peers?

- A. open message
- B. keep-alive message
- C. notification message
- D. update message

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 90

A Cisco IOS XR router must be configured with BFD for OSPF. Which configuration is correct?

- A. interface GigabitEthernet0/0/0/0
ip ospf bfd
- B. router ospf 1
area 0
interface GigabitEthernet0/0/0/0
bfd fast-detect
- C. router ospf 1

```
area 0
interface GigabitEthernet0/0/0/0
bfd fast-detect ipv4
```

- D. router ospf 1
bfd all-interfaces

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 91

A network engineer is configuring IS-IS in Cisco IOS XR. Where is BFD configured?

- A. RP/0/RSP0/CPU0:router(config-isis-if)#
- B. RP/0/RSP0/CPU0:router(config)#
- C. RP/0/RSP0/CPU0:router(config-bfd)#
- D. RP/0/RSP0/CPU0:router(config-isis-if-af)#

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference: