

400-101.exam

Number: 400-101
Passing Score: 800
Time Limit: 120 min
File Version: 1.0



VCE to PDF Converter : <https://vceplus.com/vce-to-pdf/>

Facebook: <https://www.facebook.com/VCE.For.All.VN/>

Twitter : https://twitter.com/VCE_Plus

Google+ : <https://plus.google.com/+Vcepluscom>

LinkedIn : <https://www.linkedin.com/company/vceplus>

<https://vceplus.com/>

Cisco

400-101

CCIE Routing and Switching Written Exam

Version 1.0

Exam A

QUESTION 1

You issue the following commands on Router1:

```
Router1(config)#classmap voice
Router1(config-cmap)#match protocol rtp audio
Router1(config-cmap)#exit
Router1(config)#policymap exsim
Router1(config-pmap)#class voice
Router1(config-pmap-c)#priority percent 10
Router1(config-pmap-c)#exit
Router1(config-pmap)#class classdefault
Router1(configpmap-c)#bandwidth percent remaining
Router1(config-pmap-c)#exit
Router1(config-pmap)#exit
Router1(config)#policymap boson
Router1(configpmap)#class class-default
Router1(config-pmap-c)#shape peak 50000000
Router1(config-pmap-c)#service-policy exsim
Router1(config-pmap-c)#exit Router1(config-
pmap)#exit
Router1(config)#interface fa0/1
Router1(config-if)#service-policy output boson
```



Which of the following statements are correct? (Select 2 choices.)

- A. Voice traffic is given priority up to a bandwidth of 5 Mbps.
- B. Voice traffic is given priority up to a bandwidth of 10 Mbps.
- C. Voice traffic is given priority up to a bandwidth of 50 Mbps.
- D. Traffic is configured to strictly conform to the CIR.
- E. Traffic might be dropped when the network becomes congested.

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Voice traffic is given priority up to a bandwidth of 5 Mbps. However, traffic might be dropped when the network is congested. The FastEthernet 0/1 interface is configured with a service policy named boson, which references a child service policy named exsim. All traffic is considered first by the parent policy, then by the child policy.

The parent service policy is configured with the shape peak 50000000 command. The peak keyword indicates that peak shaping should occur. Peak shaping allows higher bursts than average shaping allows, but occasional packet drops can occur when the network is congested. The average keyword can be used with the shape command to ensure that the traffic strictly conforms to the committed information rate (CIR). The variable for the shape command is the CIR, which is specified in bits per second.

Voice traffic is matched in the child policy named exsim. The priority percent 10 command indicates that 10 percent of the bandwidth is guaranteed to the traffic class. In this scenario, the traffic is first shaped by the parent policy to a value of 50 Mbps. Therefore, 10 percent of this value is 5 Mbps.

Reference: <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/command/qos-cr-book/qos-s1.html#wp3296491684>

QUESTION 2

You boot a router that contains an RXBOOT image and no startup configuration file. After a few moments, streamlined Setup mode begins.

Which of the following will you be prompted to configure? (Select the best answer.)



<https://vceplus.com/>

- A. the host name
- B. the VTY password
- C. the enable secret password
- D. interface parameters

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You will be prompted to configure interface parameters, such as IP addresses and subnet masks. Streamlined Setup mode enables a router to load an image from a network server. A router will enter streamlined Setup mode if any of the following conditions are met:

- The startup configuration has been removed by issuing the erase startup-config command. -
The startup configuration file has become corrupted.
- Bit 6 of the configuration register is set, which specifies that the router should ignore the contents of nonvolatile random access memory (NVRAM).
- The last four bits of the configuration register are equal to a value of 0 or 1.

However, if an RXBOOT image is not installed on the router, the router will enter ROM monitor (ROMmon) mode instead.

Streamlined Setup is faster than standard Setup. In addition to interface parameters, standard Setup will prompt you to configure various global router parameters, such as the host name, the virtual terminal (VTY) password, and the enable secret password.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/configfun/configuration/guide/ffun_c/fcf002.html#wp1001707

QUESTION 3

You issue the aaa authentication login default group tacacs+ local command.

Which of the following statements is correct? (Select the best answer.)

- A. If a user's account is not found on the TACACS+ servers, the user will automatically be allowed access.
- B. If the TACACS+ servers are unavailable; the user will automatically be allowed access.
- C. If the TACACS+ servers are unavailable; the user will automatically be denied access.
- D. The default authentication method is applied to all lines for which no other login method has been specified.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When you issue the aaa authentication login default group tacacs+ local command, the default authentication method is applied to all lines for which no other login method has been specified. The syntax of the aaa authentication login command is aaa authentication login {default | listname} method1 [method2...], where listname is an identifier for a list of authentication methods and method1 is at least one of the following authentication methods:

Authentication Method	Description
enable	uses the enable password
krb5	uses Kerberos 5
krb5-telnet	uses Kerberos 5 for Telnet
line	uses the line password
local	uses the local user database
local-case	uses the case-sensitive local user database
none	uses no authentication
group radius	uses a group of Remote Authentication Dial-In User Service (RADIUS) servers
group tacacs+	uses a group of Terminal Access Controller Access Control System Plus (TACACS+) servers
group <i>group-name</i>	uses a subset of RADIUS or TACACS+ servers named <i>group-name</i>

In this scenario, the router will first attempt to authenticate a user by checking a group of TACACS+ servers. If the TACACS+ servers do not respond, the router will use the local user database for authentication. To access the router if the TACACS+ server is unavailable, the user must authenticate to the local database. Configuring a secondary authentication such as the enable password or the local database is useful because administrators can connect to the router even if the authentication server is unavailable.

If a user's account is not found on the TACACS+ servers, the user will be denied access. As long as a TACACS+ server responds, the router will not use the next authentication method on the list.

If the TACACS+ servers are unavailable, the user will not be automatically allowed or denied access. The user can still access the router by using the local database. To ensure that the user will be denied access if the TACACS+ servers are unavailable, you should issue the `aaa authentication login default group tacacs +` command without the `local` keyword.

Reference: <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/a1/sec-a1-cr-book/sec-cr-a1.html#wp4227342077>

QUESTION 4

Which of the following OSPF LSAs are also called ASBR summary LSAs? (Select the best answer.)

- A. Type 1 LSAs
- B. Type 2 LSAs
- C. Type 3 LSAs
- D. Type 4 LSAs
- E. Type 5 LSAs

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Open Shortest Path First (OSPF) Type 4 link-state advertisements (LSAs) are also called autonomous system boundary router (ASBR) summary LSAs. ASBR summary LSAs are used to advertise the location of an ASBR so that routers can determine the best next-hop path. Type 4 LSAs are generated by area border routers (ABRs) and are flooded throughout an area except into stub areas.

Type 1 LSAs, which are also called router LSAs, contain router ID and interface IP address information for a single router. Router LSAs, which are generated by all OSPF routers, are not propagated outside the area? they are flooded only within the local area.

Type 2 LSAs, which are also called network LSAs, contain subnet and neighbor router information. Network LSAs, which are generated by designated routers (DRs), are not propagated outside the area in which they originate? they are flooded only within the local area.

Type 3 LSAs, which are also called network summary LSAs, contain subnet information for an entire area. Network summary LSAs, which are generated by ABRs, are advertised between areas throughout an autonomous system (AS) except into totally stubby areas.

Type 5 LSAs, which are also called AS-external LSAs, contain subnet information for an external AS. AS-external LSAs, which are generated by ASBRs, are advertised throughout an AS except into stub areas, totally stubby areas, and not-so-stubby areas (NSSAs).

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_2/nx-os/unicast/configuration/guide/l3_cli_nxos/l3_ospf.html#pgfId-1243056

https://www.cisco.com/c/en/us/td/docs/security/fwsm/fwsm32/asdm52f/user/guide/asdmug/mon_rtg.html#wp1046958

VCE To PDF - Free Practice Exam

QUESTION 5

Which of the following statements are true when an RSTP switch detects a topology change? (Select 2 choices.)

- A. It starts the TC While timer with a value equal to the hello timer for all its non-edge designated ports and its root port.
- B. It starts the TC While timer with a value equal to twice the hello timer for all its non-edge designated ports and its root port.
- C. It starts the TC While timer with a value equal to three times the hello timer for all its non-edge designated ports and its root port.
- D. It flushes the MAC addresses associated with its edge designated ports and its root port from the CAM table.
- E. It flushes the MAC addresses associated with its non-edge designated ports and its root port from the CAM table.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When a Rapid Spanning Tree Protocol (RSTP) switch detects a topology change, it starts the TC While timer with a value equal to twice the hello timer for all its non-edge designated ports and its root port. Additionally, the switch flushes the Media Access Control (MAC) addresses associated with its non-edge designated

ports and its root port from the Content Addressable Memory (CAM) table. A switch detects a topology change by receiving a bridge protocol data unit (BPDU) with the topology change (TC) bit set? these BPDUs are called topology change notification (TCN) messages.

With RSTP, topology changes are detected for the sole purpose of updating RSTP switching tables. RSTP does not consider loss of connectivity a topology change? consequently, only non-edge ports that transition into the forwarding state are considered topology changes, which results in TCN messages being disseminated throughout the network. When a switch detects a topology change, it starts the TC While timer and generates TCN messages for all its non-edge designated ports and its root port. Additionally, all MAC addresses linked to the non-edge designated ports and the root port are flushed from the CAM table. Flushing the CAM table requires that the MAC addresses be relearned after the topology change in the event that a host now appears on a different link. It is important to note that a flood of TCN messages could cause repeated flushing of the CAM table and a spike in CPU utilization, which could cause performance problems on the switch.

RSTP is used to significantly increase convergence speed after a topology change. RSTP is based on the 802.1w standard developed by the Institute of Electrical and Electronics Engineers (IEEE) to address the slow transition of a Spanning Tree Protocol (STP) port to the forwarding state. Unlike STP, which has five port states, RSTP has only three: discarding, learning, and forwarding. The disabled, blocking, and listening states of STP are combined into the discarding state in RSTP. RSTP uses the STP root port and designated port roles but splits up the STP blocking port role into the alternate port and backup port roles. An alternate port is a blocked port that receives more useful BPDUs from a port on another device, and a backup port is a blocked port that receives more useful BPDUs from a port on the same device. RSTP is backward compatible with switches that can only use STP, but the convergence benefits provided by RSTP are lost when RSTP interacts with STP devices.

RSTP does not flush the MAC addresses associated with its edge ports. Because the edge port is connected to a single host, the port cannot form a loop and is immediately placed into the forwarding state. If an edge port ever receives a BPDU, the port will lose its edge port designation.

When a switch detects a topology change, it does not start the TC While timer with a value equal to the hello timer, nor does it start the TC While timer with a value equal to three times the hello timer. The switch will start the TC While timer with a value equal to twice the hello timer and generate TCN messages for its non-edge designated ports and its root port.

Reference: <https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/24062-146.html#topchnq>

QUESTION 6

Which of the following command sets will cause an EEM applet to finish before the show runningconfig command is executed? (Select the best answer.)

- A. Router(config)#event manager applet boson
Router(configapplet)#event cli pattern "show runningconfig" sync no skip no
Router(configapplet)#action 1.0 syslog msg "Running configuration displayed"
- B. Router(config)#event manager applet boson
Router(configapplet)#event cli pattern "show runningconfig" sync no skip yes
Router(configapplet)#action 1.0 syslog msg "Running configuration displayed"
- C. Router(config)#event manager applet boson
Router(configapplet)#event cli pattern "show runningconfig" sync yes

```
Router(configapplet)#action 1.0 syslog msg "Running configuration displayed"  
Router(configapplet)#set 2.0 _exit_status 0  
D. Router(config)#event manager applet boson  
Router(configapplet)#event cli pattern "show runningconfig" sync yes  
Router(configapplet)#action 1.0 syslog msg "Running configuration displayed"  
Router(configapplet)#set 2.0 _exit_status 1
```

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following command set will cause an Embedded Event Manager (EEM) applet to finish before the show runningconfig command is executed:

```
Router(config)#event manager applet boson  
Router(configapplet)#event cli pattern "show runningconfig" sync yes  
Router(configapplet)#action 1.0 syslog msg "Running configuration displayed" Router(configapplet)#set 2.0 _exit_status 1
```

This command set configures an EEM applet named boson that is triggered when the show runningconfig command is issued. The applet writes a message to syslog and sets the _exit_status to a value of 1. The applet then exits, and the show runningconfig command is executed.

The event cli command configures EEM to monitor commandline interface (CLI) commands and to trigger the event when a specified pattern is matched one or more times. Events can be processed synchronously or asynchronously. The sync yes keywords are used with the event cli command to configure synchronous processing. With synchronous processing, the EEM applet must finish before the CLI command can be executed, and the _exit_status variable determines whether the CLI command is executed or skipped. If the _exit_status variable is set to a value of 0 or is not configured, the CLI command will not execute after the EEM applet is finished; if the _exit_status variable is set to a value of 1, the CLI command will execute after the EEM applet is finished. For example, the following command set will cause the show runningconfig command to be skipped because the _exit_status variable is set to a value of 0:

```
Router(config)#event manager applet boson  
Router(configapplet)#event cli pattern "show runningconfig" sync yes  
Router(configapplet)#action 1.0 syslog msg "Running configuration displayed" Router(configapplet)#set 2.0 _exit_status 0
```

The sync no keywords are used with the event cli command to configure asynchronous processing. With asynchronous processing, the EEM applet is processed at the same time the CLI command is executed. When you issue the event cli command with the sync nokeywords, you must also include the skip no or skip yes keywords to indicate whether the CLI command should be executed or skipped, respectively. For example, the following command set will cause the EEM applet to run at the same time the show runningconfig command is executed:

```
Router(config)#event manager applet boson
```

```
Router(configapplet)#event cli pattern "show runningconfig" sync no skip no  
Router(configapplet)#action 1.0 syslog msg "Running configuration displayed"
```

The following command set will cause the EEM applet to run and to not execute the show runningconfig command:

```
Router(config)#event manager applet boson  
Router(configapplet)#event cli pattern "show runningconfig" sync no skip yes Router(configapplet)#action 1.0 syslog msg "Running configuration displayed"
```

Reference:

Cisco: Cisco IOS Embedded Event Manager Command Reference: event cli
Cisco: Understanding Cisco EEM by examples Part 2

QUESTION 7

Which of the following commands will track whether IP routing is enabled, the interface line protocol is up, and the interface IP address is configured? (Select the best answer.)

- A. track 1 interface FastEthernet1/1 lineprotocol
- B. track 1 interface FastEthernet1/1 ip routing
- C. track 1 ip route 10.10.10.0/24 reachability
- D. track 1 ip route 10.10.10.0/24 metric threshold



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The track 1 interface FastEthernet1/1 ip routing command will track whether IP routing is enabled, the interface line protocol is up, and the interface IP address is configured. If any of the three criteria are not met, the tracked interface is considered to be down. To track whether IPv6 routing is enabled, the line protocol is up, and the interface IPv6 address is configured, you could issue the track 1 interface FastEthernet1/1 ipv6 routing command.

The track 1 interface FastEthernet1/1 lineprotocol command will track whether the interface line protocol is up. However, it will not track whether IP routing is enabled and the interface IP address is configured.

The track 1 ip route 10.10.10.0/24 reachability command will track whether the destination network is reachable. The subnet address must be issued in Classless InterDomain Routing (CIDR) notation.

The track 1 ip route 10.10.10.0/24 metric threshold command will track whether the metric threshold is exceeded. By default, a metric value of 254 or less is considered to be accessible and a metric value of 255 is considered to be inaccessible. The threshold metric up upvalue down downvalue command can be used to change the default metric threshold values.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/command/iap-cr-book/iap-t1.html#wp9828037020> <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/command/iap-cr-book/iap-t1.html#wp4199553004>

QUESTION 8

Which of the following NetFlow features is not unique to version 9 and later? (Select the best answer.)

- A. template-based export protocol
- B. support for IPv6 flows
- C. support for MPLS flows
- D. support for multicast flows
- E. support for router-based aggregation

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation/:

Support for router-based aggregation is not unique to NetFlow version 9 and later. Router-based aggregation is available with NetFlow version 8 and version 9. NetFlow is a reporting tool that can be used to measure network bandwidth, Quality of Service (QoS), and performance, among other things.

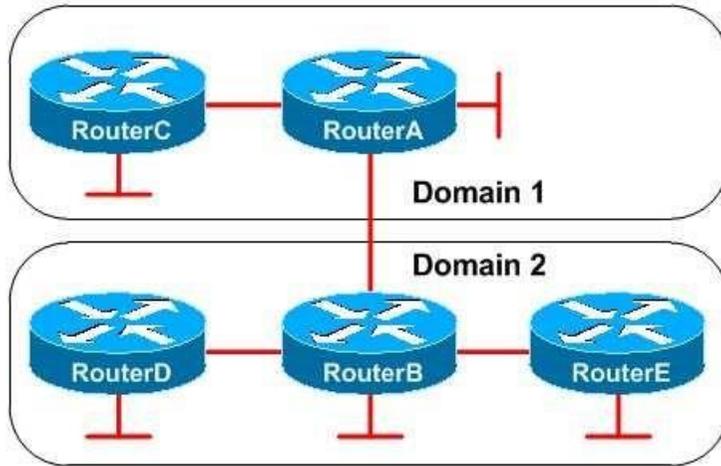
NetFlow version 9 and later use a template-based export protocol. The purpose of the template-based format is to enable future versions of NetFlow to extend services without significantly altering the basic NetFlow export format.

NetFlow version 9 and later contain support for IPv6 flows, Multiprotocol Label Switching (MPLS) flows, and multicast flows. Therefore, it is possible with version 9 and later to measure and report on network characteristics beyond basic IPv4 functionality. All versions of NetFlow support IPv4 flows. Reference:

Cisco: NetFlow Services Solutions Guide: NetFlow Export Version Formats

Cisco: NetFlow Export Datagram Format

QUESTION 9



You administer the network shown above. RouterA and RouterB are Anycast RPs that are configured as MSDP multicast peers. The following partial output is from the show running-config command on RouterA:

```
interface Loopback0 ip address
192.168.1.1 255.255.255.255
!
interface Loopback1 ip address
192.168.1.2 255.255.255.255
!
ip msdp peer 192.168.1.3 connect-source loopback1 remote-as 2
ip msdp originator-id loopback1 ip pim rp-address 192.168.1.1
```

The following partial output is from the show runningconfig command on RouterB:

```
interface Loopback0 ip address
192.168.1.1 255.255.255.255
!
interface Loopback1 ip address
192.168.1.3 255.255.255.255
!
ip msdp peer 192.168.1.2 connect-source loopback1 remote-as 1
ip msdp originator-id loopback1 ip pim rpaddress 192.168.1.1
```

RouterC is configured to use Auto-RP to discover the Anycast RP. RouterD is configured to use BSR to discover the Anycast RP. RouterE is configured with the ip pim rp-address 192.168.1.3 command.

Which of the following statements is correct? (Select the best answer.)

- A. RouterA and RouterB cannot use the same IP address on Loopback 0.
- B. RouterA and RouterB cannot be in different domains.
- C. RouterC cannot use AutoRP to discover the Anycast RP.
- D. RouterD cannot use BSR to discover the Anycast RP.
- E. RouterE is not configured with the correct IP address of the Anycast RP.

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RouterE is not configured with the correct IP address of the Anycast rendezvous point (RP). Anycast RP enables multiple RPs to provide redundancy and loadsharing capabilities. Each downstream router uses the closest RP. If an Anycast RP fails or is added, the Protocol Independent Multicast (PIM) network will converge as quickly as IP routing converges. You should configure each of the Anycast RPs as Multicast Source Discovery Protocol (MSDP) peers of one another by issuing the ip msdp peer command for each Anycast RP peer.

RouterA and RouterB must use the same IP address on a loopback interface; this address is the Anycast RP address. In this scenario, RouterA and RouterB are correctly using the 192.168.1.1 address on Loopback 0. When an Anycast RP fails, the downstream routers will not have to discover a new RP; they will continue to use the shared IP address of the Anycast RPs.

Each downstream router must be configured with the shared IP address of the AnycastRPs, either statically by using the ip pim raddress command or dynamically by using AutoRP or Bootstrap Router (BSR). In this scenario, RouterE is configured with the ip pimraddress 192.168.1.3 command. However, the Loopback 0 interfaces of RouterA and RouterB are configured with the shared IP address 192.168.1.1. Therefore, RouterE should be configured with the ip pim raddress 192.168.1.1 command. RouterC is correctly configured to use AutoRP to discover the RP address, and RouterD is correctly configured to use BSR to discover the RP address.

RouterA and RouterB can be in different domains. In fact, MSDP enables Anycast RPs to share information about multicast sources across domains. Without MSDP, an Anycast RP would be able to know about multicast sources only within its own domain.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/solutions_docs/ip_multicast/White_papers/anycast.html

QUESTION 10

What does an asterisk indicate in the output of the show ip pim tunnel command? (Select the best answer.)



<https://vceplus.com/>

- A. The router is an RP.
- B. Fast switching is enabled.
- C. The entry includes all multicast sources.
- D. The neighbor has been learned through an assert.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

An asterisk in the output of the show ip pim tunnel command indicates that the router is a rendezvous point (RP). An RP is a well-connected, centrally located router that is responsible for keeping track of multicast group membership information. Protocol Independent Multicast sparse mode (PIMSM) requires an RP, whereas PIM dense mode (PIMDM) does not. When PIMSM is used, each multicast receiver must be able to reach the RP through a connected tree of PIMSM routers. If a router along the path is not configured for PIMSM, multicast receivers will not be able to register with the RP and multicast traffic will not flow to those receivers.

When you issue the show ip pim tunnel command on an RP, you will receive output that is similar to the following:

```
Router1#show ip pim tunnel
Tunnel0
  Type   : PIM Encap
  RP     : 10.1.14.1*
  Source: 10.1.14.1
Tunnel1
  Type   : PIM Decap
  RP     : 10.1.14.1*
  Source: -
```

An RP will always have a PIM Encap and a PIM Decap tunnel interface. Additionally, an asterisk will appear next to the RP IP address.

An asterisk in the output of the show ip pim interface count command indicates that fast switching is enabled. The show ip pim interface count command also displays how many multicast packets have been received and sent by each interface. The following output is from the show ip pim interface count command:

```
Router1#show ip pim interface count
Address      Interface      FS  Mpackets In/Out
10.1.14.7    Ethernet0      *   512486521/14522485
10.1.15.7    Serial0        *   7823469/38427838
```

An asterisk in the output of the show ip mroute command indicates that an entry includes all multicast sources or that a neighbor has been learned through an assert. The show ip mroute command displays the multicast routing table. Shared distribution trees are specified by a (*,G) notation? the * indicates all sources, and the G indicates the multicast group address. Source distribution trees, which are also known as shortest path trees (SPTs), are specified by an (S,G) notation? the S indicates the address of the multicast source, and the G indicates the multicast group address. PIMSM supports both shared distribution trees and source distribution trees and can use both (S,G) and (*,G) routes. PIMDM supports only source distribution trees, and PIMDM groups use only (S,G) routes. The following output from the show ip mroute command shows both types of routes:

```
Router1#show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report, Z - Zero-Flap Multicast Tunnel,
       Y - Joined MDT-data group, y - Sending to MDT-data group
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode
(*, 224.0.0.3), uptime 5:29:15, RP is 192.168.99.5, flags: SC
  Incoming interface: Tunnel0, RPF nbr 192.168.50.1, Dvmrp
  Outgoing interface list:
    FastEthernet0, Forward/Sparse, 3:00:10/0:01:50
(192.168.55.0/24, 224.0.0.3), uptime 3:00:10, expires 0:01:50, flags: C
  Incoming interface: Tunnel0, RPF nbr 192.168.50.1
  Outgoing interface list:
    FastEthernet0, Forward/Sparse, 3:00:10/0:01:50
```

If an asterisk appears next to the RPF nbr IP address, the neighbor has been learned through an assert. Asserts are used to elect a designated forwarder (DF). DFs are elected to ensure a loop-free tree with the root at the RP. The router with the lowest cost to the RP will become the DF on a network segment.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/command/imc-cr-book/imc_s1.html#wp3622554730 https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/command/imc-cr-book/imc_s1.html#wp2605288306 https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/command/imc-cr-book/imc_s1.html#wp9533023710

QUESTION 11

Your company has configured STP with the timers set to their default values.

For what duration will the TC bit be set by the root bridge after it receives a TCN BPDU? (Select the best answer.)

- A. two seconds
- B. 15 seconds
- C. 20 seconds
- D. 35 seconds

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The topology change (TC) bit will be set for a duration of 35 seconds on configuration bridge protocol data units (BPDUs) that are sent by the root bridge after it receives a topology change notification (TCN) BPDU. TCN BPDUs are a Spanning Tree Protocol (STP) mechanism sent from a designated bridge back to the root bridge to inform the root bridge of a change in the network topology. TCN BPDUs are sent under the following circumstances: when a switch receives a TCN from another non-root bridge, when a link has failed, or when a port begins to forward packets despite the bridge already having a designated port. Ports that have PortFast enabled will not send TCN BPDUs when entering or leaving the forwarding state.

After a bridge receives a TCN BPDU, it will send a BPDU with the topology change acknowledgment (TCA) bit set back to the bridge that sent the TCN BPDU. When the root bridge receives a TCN BPDU, it begins sending configuration BPDUs to the downstream bridge devices. Configuration BPDUs contain a TC bit indicating that a change has occurred. This bit is set for a period of the max_age timer plus the forward_delay timer. STP defaults the max_age timer to 20 seconds and the forward_delay timer to 15 seconds. Therefore, the TC bit will be set for 35 seconds.

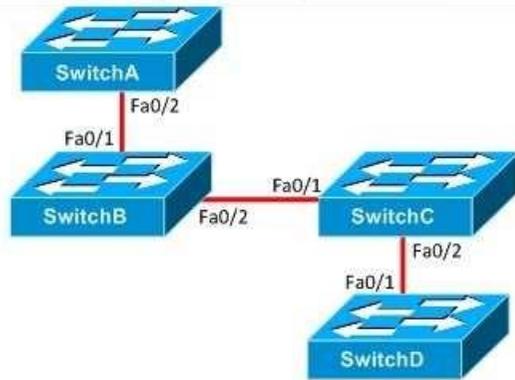
TC BPDUs will cause a switch to shorten the aging time for Media Access Control (MAC) addresses from 300 seconds to the forward_delay value, which is 15 seconds by default. Excessive flooding can occur as the switch receives packets and forwards the traffic out all ports.

The TC bit will not be set for a duration of two seconds in this scenario. The default hello timer is two seconds. The hello_time is the interval at which BPDUs are sent.

Reference:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/12013-17.html#event>

QUESTION 12



You administer the network shown above. You issue the show spanning-tree command on SwitchC and receive the following output:

```
SwitchC#show spanning-tree vlan 42
VLAN0042
```

```
Spanning tree enabled protocol rstp
```

```
Root ID      Priority      31574
             Address      001c.6e5e.7aa3
             Cost        19
             Port        2 (FastEthernet0/1)
             Hello Time  2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID   Priority      32810 (priority 32768 sys-id-ext 42)
             Address      001c.223b.0717
             Hello Time  2 sec Max Age 20 sec Forward Delay 15 sec
             Aging Time  300
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/1	Root	FWD	19	128.2	P2p
Fa0/2	Desg	FWD	19	128.3	P2p Peer (STP)

Which of the following statements are true regarding the switches on the network? (Select 2 choices.)

- A. SwitchA is the root bridge.
- B. SwitchC is using the default STP timer values.
- C. 802.1D STP is running on SwitchB.
- D. 802.1D STP is running on SwitchD.

- E. The BID of SwitchB is 32810.001c.223b.0717.
 F. The BID of SwitchC is 001c.223b.0717.32810.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

SwitchC is using the default Spanning Tree Protocol (STP) timer values, and 802.1D STP is running on SwitchD. STP uses three timer values: the hello timer value, the forward_delay timer value, and the max_age timer value. The hello timer value is the time between the sending of bridge protocol data units (BPDUs)? this value is set to two seconds by default. The forward_delay timer value is the time spent in the listening state and the learning state? this value is set to 15 seconds by default. The max_age timer value is the maximum length of time before BPDU information is aged out? this value is set to 20 seconds by default. The output of the show spanningtree command indicates that SwitchC is using these default STP timer values.

Traditional STP is defined in the Institute of Electrical and Electronics Engineers (IEEE) 802.1D standard. Rapid STP (RSTP), which is defined in the IEEE 802.1w standard, is used to improve the slow transition of an STP port to the forwarding state. The line of text Spanning tree enabled protocol rstp in the output of the show spanningtree command indicates that SwitchC is running RSTP. If SwitchC were running 802.1D STP, the line of text Spanning tree enabled protocol ieee would be displayed in the output of the show spanningtree command.

To determine the type of STP running on neighbor switches, you should analyze the Type field in the show spanningtree command output. If the Type field displays P2p, the neighbor switch is running RSTP. If the Type field displays P2p Peer(STP), the neighbor switch is running traditional 802.1D STP. Therefore, SwitchB is running RSTP and SwitchD is running 802.1D STP.

SwitchA is not the root bridge. The output of the show spanningtree command indicates that the root bridge is reachable through the Fa0/1 interface. Therefore, either SwitchA or SwitchB is the root bridge. You can determine which switch is the root by analyzing the link cost, which is used to determine the best path to the root bridge. The link cost is based on the bandwidth of a link. The higher the bandwidth, the lower the cost. STP uses the following link costs by default:

Bandwidth	Cost
4 Mbps	250
10 Mbps	100
16 Mbps	62
45 Mbps	39
100 Mbps	19
155 Mbps	14
622 Mbps	6
1 Gbps	4
10 Gbps	2

FastEthernet links have a bandwidth of 100 Mbps, so the link cost of traversing a FastEthernet link is 19. If packets must traverse two FastEthernet links, the link cost is 38. In this scenario, the output of the show spanningtree command indicates that the link cost is 19. Because packets must traverse only one FastEthernet link to reach the root bridge, SwitchB must be the root bridge.

The bridge ID (BID) of SwitchB is not 32810.001c.223b.0717. The BID is composed of a 2-byte bridge priority prefix and a 6-byte Media Access Control (MAC) address suffix. The output of the show spanningtree command indicates that the root bridge has a priority value of 31574 and a MAC address of 001c.6e5e.7aa3. Therefore, the BID of the root bridge, SwitchB, is 31574.001c.6e5e.7aa3.

The BID of SwitchC is not 001c.223b.0717.32810. The Bridge ID field in the output of the show spanningtree command indicates the priority value and MAC address of the local switch, SwitchC. In the BID, the bridge priority comes before the MAC address. Therefore, the BID of SwitchC is 32810.001c.223b.0717.

Reference:

<https://www.cisco.com/c/en/us/support/docs/switches/catalyst-6500-series-switches/72836-rapidpvst-mig-config.html>

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/19120-122.html>

QUESTION 13

You issue the show ipv6 ospf command on RouterA and receive the following output:

```
RouterA#show ipv6 ospf
Routing Process "ospfv3 22" with ID 10.20.30.40
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
Number of external LSA 0. Checksum Sum 0x000000
Number of areas in this device is 1. 1 normal 0 stub 0 nssa
  Area 2
    Number of interfaces in this area is 1
    Area has no authentication
    SPF algorithm last executed 00:00:20.532 ago
    SPF algorithm executed 7 times
    Number of LSA 0. Checksum 0x000000
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
```

Which of the following statements are correct? (Select 3 choices.)

- A. RouterA is an ASBR.
- B. RouterA is an ABR.

- C. RouterA is an internal router.
- D. RouterA is a backbone router.
- E. RouterA is part of a stub area.
- F. Router A is part of a totally stubby area.
- G. RouterA is part of a standard area.
- H. Route recalculation has occurred seven times.

Correct Answer: CGH

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The output of the show ipv6 ospf command indicates that RouterA is an internal router, RouterA is part of a standard area, and route recalculation has occurred seven times. The output of the show ipv6 ospf command is similar to the show ip ospf command in that they both display the following information:

- Shortest path first (SPF) timer values and statistics
- Linkstate advertisement (LSA) timer values and statistics
- Number and type of areas in the router
- Whether authentication is enabled for an area

The router output Number of areas in this device is 1 indicates that there is only one Open Shortest Path First (OSPF) area in the router. Internal routers belong to a single OSPF area? therefore, RouterA is an internal router. The router output 1 normal 0 stub 0 nssa indicates that the area is a normal area, which is also called an ordinary or standard area. Therefore, RouterA is part of a standard area. Finally, the router output indicates that route recalculation has occurred seven times because the SPF algorithm has executed seven times on the router.

RouterA is not an autonomous system boundary router (ASBR). ASBRs connect two or more autonomous systems and redistribute routes between them. If RouterA were an ASBR, the output of the show ipv6 ospf command would display the line of text It is an autonomous system boundary router. In addition, the output would display the route sources that RouterA is redistributing into OSPF.

RouterA is not an area border router (ABR). ABRs connect two or more OSPF areas? a separate linkstate database (LSDB) is maintained for each area. If RouterA were an ABR, the output of the show ipv6 ospf command would display the line of text It is an area border device. In addition, the router output would indicate that multiple OSPF areas were configured on the router and statistical information would exist for each of those areas. In the output of the show ipv6 ospf command, statistical information exists only for Area 2.

RouterA is not a backbone router. A backbone router is a router with at least one interface in Area 0, the backbone area. The router output Area 2 indicates that RouterA is within Area 2, not Area 0. If RouterA were a backbone router, the output of the show ipv6 ospfcommand would display the line of text Area BACKBONE (0) instead of Area 2.

RouterA is not part of a stub area. A stub area is an area that does not accept Type 5 summary LSAs. If RouterA had at least one interface in a stub area, the output of the show ipv6 ospf command would indicate at least one stub area within the line of text 1 normal 0 stub 0 nssa. Additionally, if Area 2 were a stub area, the router output It is a stub area would appear within the statistical information for Area 2.

RouterA is not part of a totally stubby area. A totally stubby area is an area that does not accept Type 3, 4, or 5 summary LSAs. If RouterA had at least one interface in a totally stubby area, the output of the show ipv6 ospf command would indicate at least one stub area within the line of text 1 normal 0 stub 0 nssa. Additionally, if Area 2 were a stub area, the router output It is a stub area, no summary LSA in this area would appear within the statistical information for Area 2.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/command/ipv6-cr-book/ipv6-s4.html#wp1570786054>

QUESTION 14

DRAG DROP

Drag the OSPF neighbor relationship states on the left to the corresponding reasons on the right.

Select and Place:

Stuck in Down or Init state	224.0.0.5 is blocked
Stuck in Loading state	LSAs are corrupted
Stuck in Exchange or Exstart state	MTU settings are mismatched
Stuck in 2-Way state	Normal for non-DR neighbors

Correct Answer:



Section: (none)

Explanation

Explanation/Reference:

Explanation:

When an Open Shortest Path First (OSPF) neighbor router is powered on, it transitions through the following neighbor states: -

Down

-Init

-2-Way

-Exstart

-Exchange

-Loading

-Full

An OSPF neighbor router begins in the Down state. A neighbor in the Down state has not yet sent a hello packet. When a hello packet is received from the neighbor router but the hello packet does not contain the receiving router's ID, the neighbor router is in the Init state. The receiving router replies to the neighbor router with a hello packet that contains the neighbor router's ID as an acknowledgment that the receiving router received the neighbor's hello packet. If a router is stuck in the Init state, it has sent hello packets but has not received them from the neighbor router. If a router is stuck in the Down or Init state, you should check to see whether an access list is blocking 224.0.0.5, which is used by OSPF to send hello packets. Additionally, you should ensure that Layer 1 and Layer 2 connectivity exists and that authentication is disabled or enabled on both neighbors.

The neighbor router replies with a hello packet that contains the receiving router's ID. When this occurs, the neighbor router is in the 2Way state. At the end of the 2Way state, the designated router (DR) and backup designated router (BDR) are elected for broadcast and nonbroadcast multiaccess (NBMA) networks. On broadcast and NBMA networks, neighbor routers will proceed to the Full state only with the DR and BDR; routers will remain in the 2Way state with all other

neighbor routers. Routers that remain in the 2Way state will contain 2WAY/DROTHER in the output of the show ip ospf neighbor command. If all routers on a segment remain in the 2Way state, you should verify whether all routers on the segment are set to a priority of 0, which prevents any of them from becoming the DR or BDR.

After the DR and BDR are elected, neighbor routers form master-slave relationships in order to establish the method for exchanging linkstate information. Routers in this state are in the Exstart state. Neighbor routers then exchange database descriptor (DBD) packets. These DBD packets contain linkstate advertisement (LSA) headers that describe the contents of the linkstate database. Routers in this state are in the Exchange state. If a router is stuck in the Exstart or Exchange state, you should determine whether there is a problem with mismatched maximum transmission unit (MTU) settings.

Routers then send linkstate request packets to request the contents of the neighbor router's OSPF database. The neighbor router replies with linkstate update packets that contain the routing database information. Routers in this state are in the Loading state. If a router is stuck in the Loading state, you should determine whether there is a problem with corrupted LSAs.

After the OSPF databases of neighbor routers are fully synchronized, the routers transition to the Full state, which is the normal OSPF router state. A router will periodically send hello packets to its neighbors to indicate that it is still functional. If a router does not receive a hello packet from a neighbor within the dead timer interval, the neighbor router will transition back to the Down state.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13699-29.html> <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13685-13.html>

QUESTION 15

RouterA and RouterB are connected by their Serial 0/0 interfaces. You are able to successfully ping from one router to the other? however, the routers are not establishing an OSPF neighbor relationship.

You issue the debug ip ospf hello command on each router and notice that RouterA is sending hello packets but not receiving them.

Which of the following statements best describes why the routers are unable to form a neighbor relationship? (Select the best answer.)

- A. RouterA should be configured as the DR.
- B. RouterB should be configured as the DR.
- C. The Serial 0/0 interface of RouterA is configured as a passive interface.
- D. The Serial 0/0 interface of RouterB is configured as a passive interface.
- E. RouterA and RouterB are configured with different OSPF process IDs.
- F. RouterA and RouterB are configured with different hello timer or dead timer intervals.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Serial 0/0 interface of RouterB is configured as a passive interface, so RouterA and RouterB are unable to form a neighbor relationship. A passive interface does not send or receive any routing information, regardless of the routing protocol. Additionally, an Open Shortest Path First (OSPF) passive interface does not send hello packets. OSPF uses the periodic exchange of hello packets to maintain neighbor relationships. If an OSPF router does not receive a hello packet from a neighbor after a specified amount of time, the neighbor relationship is terminated and no further routing information is exchanged. In this scenario, RouterA is sending periodic hello packets; however, RouterA is not receiving hello packets, because RouterB is no longer sending them. RouterB also ignores the incoming hello packets from RouterA.

To configure a single interface to be a passive interface, you should issue the `passive-interface interface-type interface-number` command from router configuration mode. Alternatively, you can issue the `passive-interface default` command from router configuration mode to configure all interfaces to be passive? you must then issue the `no passive-interface interface-type interface-number` command from router configuration mode to allow an interface to participate in the routing protocol and to establish neighbor relationships.

The Serial 0/0 interface of RouterA is not configured as a passive interface. The output of the `debug ip ospf hello` command indicates that RouterA is sending periodic hello packets. An OSPF passive interface does not send or receive routing information, including hello packets.

There is no designated router (DR) elected on an OSPF point-to-point network segment. RouterA and RouterB are connected by their Serial 0/0 interfaces and do not attempt to elect a DR. A DR is elected by using hello packets on a multiaccess network, such as a LAN. If RouterA and RouterB were connected by an Ethernet switch, RouterA would become the DR because RouterB is not sending hello packets.

The OSPF process ID is an identifier that is locally significant to the router and is used to distinguish between multiple OSPF processes running on the router.

Because the OSPF process ID is locally significant, two routers could have different process IDs and still establish a neighbor relationship. Therefore, the OSPF process ID has no effect on whether RouterA and RouterB establish a neighbor relationship.

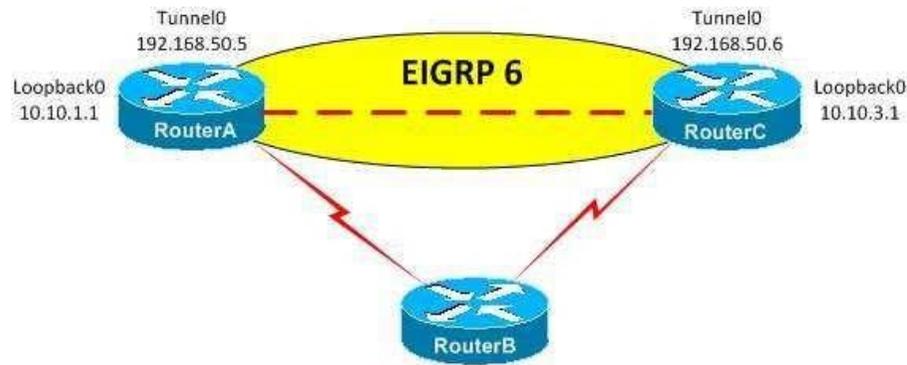
The hello timer and dead timer intervals are used to determine how frequently a router should expect a hello packet from a neighbor. If a router does not receive a hello packet from a neighbor within the dead timer interval, the relationship with that neighbor is terminated. If a router receives hello packets with a different hello timer interval or dead timer interval, the hello packets will be ignored and a neighbor relationship will not be established. However, mismatched timer intervals do not prevent a router from sending or receiving hello packets. Because the output of the `debug ip ospf hello` command reveals that RouterA is not receiving hello packets from RouterB, RouterA cannot compare timer values with RouterB to determine whether a neighbor relationship can be established.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/iproute_pi/command/reference/iri_book/iri_pi1.html#wp1034440

QUESTION 16

You administer the network in the following exhibit:



You issue the show runningconfig command on RouterA and receive the following partial output:

```
interface Loopback0 ip address
10.10.1.1 255.255.255.0
!
interface Tunnel0
ip address 192.168.50.5
255.255.255.0 tunnel source
Loopback0 tunnel destination
10.10.3.1
```



RouterA and RouterC are both configured to use RouterB as a gateway of last resort. Additionally, static routes to the Loopback0 interfaces on RouterA and RouterC have been configured on RouterB.

You configure EIGRP on RouterA and then issue the show ip route command, which produces the following partial output:

Gateway of last resort is 172.15.1.2 to network 0.0.0.0

```
172.15.0.0/24 is subnetted, 1 subnets
C 172.15.1.0 is directly connected, Serial0/0
10.0.0.0/24 is subnetted, 2 subnets
10.10.1.0 is directly connected, Loopback0
10.10.3.0 [90/297372416] via 192.168.50.6, 00:00:01, Tunnel0
C 192.168.50.0/24 is directly connected, Tunnel0
S* 0.0.0.0/0 [1/0] via 172.15.1.2
```

Which of the following is true? (Select the best answer.)

- A. The Tunnel0 interface and EIGRP adjacency on RouterA will flap.
- B. The Tunnel0 interface and EIGRP adjacency on RouterA will function properly.
- C. The Tunnel0 interface on RouterA will function properly, but EIGRP will flap.
- D. The Tunnel0 interface on RouterA will flap, but the EIGRP adjacency will function properly.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Tunnel0 interface and Enhanced Interior Gateway Routing Protocol (EIGRP) adjacency on RouterA will flap because the preferred route to the Tunnel0 destination interface is through the tunnel itself, which results in recursive routing. When recursive routing occurs, the Tunnel0 interfaces on both RouterA and RouterC will be temporarily disabled, which breaks the EIGRP adjacency.

The EIGRP adjacency will reestablish when the tunnel interfaces return to the up state. Therefore, if you were to issue the show ip route command on RouterA while the adjacency is established, you would see that the preferred route to the Loopback0 interface on RouterC from RouterA is through Tunnel0, even though the destination interface for Tunnel0 on RouterA is the Loopback0 interface on RouterC.

If the cause of the recursive routing is not fixed, the Tunnel0 interfaces will flap and errorssimilar to the following will be displayed on RouterA:

```
*Mar 1 00:26:15.379: %TUN5RECURDOWN: Tunnel0 temporarily disabled due to recursive routing
*Mar 1 00:26:16.379: %LINEPROTO5UPDOWN: Line protocol on Interface Tunnel0, changed state to down
*Mar 1 00:26:16.487: %DUAL5NBRCHANGE: IPEIGRP(0) 6: Neighbor 192.168.50.6 (Tunnel0) is down: interface down
```

In this scenario, an EIGRP adjacency has been established between the Tunnel0 interfaces on RouterA and RouterC. When the EIGRP adjacency comes up, the show ip route command displays Tunnel0 as the preferred route to 192.168.50.0 instead of the gateway of last resort. Therefore, the EIGRP 6 domain has been configured to include the 10.10.1.0/24 and 192.168.50.0/24 networks on RouterA and the 10.10.3.0/24 and 192.168.50.0/24 networks on RouterC. As a result, recursive routing to the 10.10.3.0 network through Tunnel0 occurs on RouterA and recursive routing to the 10.10.1.0 network occurs on RouterC.

There are two ways to resolve the recursive routing issue on both RouterA and RouterC in this scenario: remove the 192.168.50.0/24 network from the EIGRP 6 domain, or add a static route to the Tunnel0 destination IP addresses on both RouterA and RouterC. A static route has a lower administrative distance (AD) than EIGRP. Therefore, a static route would fix the recursive routing problem.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/22327-gre-flap.html>

QUESTION 17

Which of the following routes cannot be redistributed into a dynamic routing protocol? (Select 2 choices.)

- A. C 192.168.1.0/30 is directly connected, Ethernet0/0
- B. C 192.168.1.1/32 is directly connected, Loopback0
- C. L 192.168.1.1/32 is directly connected, Ethernet0/0
- D. C 2001:1234::/64 [0/0]
- E. L 2001:1234::1/128 [0/0]
- F. LC 2001:1234::1/128 [0/0]

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following routes cannot be redistributed into a dynamic routing protocol:

- L 192.168.1.1/32 is directly connected, Ethernet0/0
- L 2001:1234::1/128 [0/0]

Routes that are marked with an L in the output of the show ip route command or the show ipv6 route command are local host routes. Local host routes cannot be redistributed into a dynamic routing protocol. IPv4 host routes always have a /32 mask, and IPv6 host routes always have a /128 mask.

Not all IPv4 routes with a /32 mask are considered host routes. IPv4 addresses that are manually configured with a /32 mask are considered to be connected addresses and are marked with a C in the output of the show ip route command. Connected routes can be redistributed into a dynamic routing protocol. The following routes are connected IPv4 routes:

- C 192.168.1.0/30 is directly connected, Ethernet0/0
- C 192.168.1.1/32 is directly connected, Loopback0

The following route is a normal, connected IPv6 route that can be redistributed into a dynamic routing protocol:

C 2001:1234::/64 [0/0]

IPv6 addresses that are manually configured with a /128 mask are marked with an LC in the output of the show ipv6 route command. These LC routes can be redistributed into a dynamic routing protocol. The following route is a local connected route: LC 2001:1234::1/128 [0/0]

Reference:

Cisco: Local Host Routes Installed in the Routing Table on Cisco IOS and Cisco IOS-XR

QUESTION 18

Which of the following can be applied on a switch to filter inbound traffic on nonrouted ports? (Select the best answer.)

- A. VACLs
- B. RACLs
- C. PACLs
- D. both VACLs and RACLs
- E. both VACLs and PACLs

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

VLAN access control lists (VACLs) and port access control lists (PACLs) can be applied on a switch to filter inbound traffic on nonrouted ports. Access control lists (ACLs) are security mechanisms that are used to determine whether inbound and outbound packets should be forwarded or blocked. Unlike standard and extended ACLs, which are typically used to filter Layer 3 traffic, VACLs and PACLs can be used to filter nonrouted Layer 2 traffic. However, PACLs cannot filter outbound traffic; they can filter only inbound traffic.

VACLs are used to filter traffic within a virtual LAN (VLAN). VACLs can be used to prevent malicious users from gaining access to other resources on the same VLAN. Unlike most ACLs, VACLs do not filter packets as they reach an interface. Instead, VACLs filter packets across the entire VLAN, even if it spans multiple interfaces. PACLs are used to filter inbound traffic on Layer 2 switch ports. When PACLs are applied on a switch, all packets are reviewed as they reach a port. PACLs take precedence over VACLs and Layer 3 ACLs. Like VACLs, PACLs can be used to filter VLAN traffic, including voice and data VLAN traffic, if the PACLs are applied to a trunk port.

Router ACLs (RACLs) cannot be applied on a switch to filter inbound traffic on nonrouted ports. RACLs provide similar functionality as VACLs and PACLs, except they cannot be applied to Layer 2 traffic. RACLs are limited to use on Layer 3 interfaces, such as those on routers or multilayer switches configured for Layer 3 routing.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/access-lists/13608-21.html#anc118>

QUESTION 19

Which of the following functions is the data plane responsible for? (Select the best answer.)

- A. forwarding packets
- B. exchanging routing information
- C. exchanging label information
- D. reserving bandwidth

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The data plane is responsible for forwarding packets. Packets are forwarded based on destination address or label information. The Cisco Express Forwarding (CEF) Forwarding Information Base (FIB), which is part of the data plane, is built from information in the routing table. When the routing table is updated, the nexthop information in the FIB is also updated. The Label Forwarding Information Base (LFIB), which is also part of the data plane, contains inbound-to-outbound label mappings. These label mappings are used by Multiprotocol Label Switching (MPLS) to forward packets to the correct destination. When a label switch router (LSR) receives an unlabeled packet destined for an MPLS-enabled interface, it consults the FIB, adds the appropriate label for the destination address, and forwards the packet.

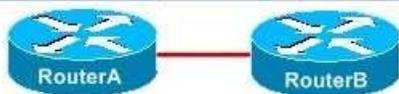
MPLS includes two primary components: the control plane and the data plane. The control plane is responsible for exchanging routing information by using a routing protocol, such as Border Gateway Protocol (BGP), Enhanced Interior Gateway Routing Protocol (EIGRP), Intermediate System to Intermediate System (ISIS), or

Open Shortest Path First (OSPF). Additionally, the control plane is responsible for exchanging label information by using a label exchange protocol, such as Resource Reservation Protocol (RSVP), Tag Distribution Protocol (TDP), or Label Distribution Protocol (LDP). LDP is a newer standard that includes features of the Cisco proprietary TDP. RSVP is used by MPLS Traffic Engineering (MPLS TE) to also reserve network bandwidth. Bandwidth is reserved on demand based on destination address or traffic type so that enough bandwidth is available for the traffic.

Reference:

<https://www.cisco.com/c/en/us/products/collateral/security/ios-network-foundation-protection->

[nfp/prod_white_paper0900aecd805ffde8.html](https://www.cisco.com/c/en/us/products/collateral/security/ios-network-foundation-protection-nfp/prod_white_paper0900aecd805ffde8.html) Cisco: Control Plane Security Overview in Cisco IOS Software **QUESTION 20**



You are creating a point-to-point serial connection between RouterA and RouterB. RouterA has been configured as DTE, and RouterB has been configured as DCE.

Which of the following statements is correct? (Select the best answer.)

- A. You should issue the clock rate command on the serial interface of RouterA.
- B. You should issue the clock rate command on the serial interface of RouterB.
- C. You should issue the clock rate command on the serial interface of both routers.
- D. You should not issue the clock rate command on either router.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Explanation:

You should issue the clock rate command on the serial interface of RouterB. RouterB has been configured as data communications equipment (DCE), and the DCE device must provide clocking to establish the data communication speed for the link. A DCE device is the device that provides a clocking signal. A device that is not capable of providing or not configured to provide a clocking signal is considered the data terminal equipment (DTE) device. Although Cisco routers can be configured as DCE devices, they are typically considered DTE devices when connected to a Channel Service Unit (CSU)/Data Service Unit (DSU). You can verify whether a router has been configured as DTE or DCE by issuing the show controllers serial command.

When issuing the clock rate command, you should specify the parameter in bits per second. Thus, if you were to issue the clock rate 64000 command, you would configure the interface to operate at 64 Kbps.

RouterA has been configured as DTE. The DTE device accepts clocking parameters from the DCE device. Therefore, you should not issue the clock rate command on the serial interface of RouterA.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/command/ir-cr-book/ir-c2.html#wp3930272930>

QUESTION 21

Which of the following statements is correct regarding EIGRPv6? (Select the best answer.)

- A. The ipv6 router eigrp asnumber command enables EIGRPv6 on all router interfaces, including passive interfaces and loopback interfaces.
- B. The network command configures the networks that should be advertised by EIGRPv6.
- C. The auto-summary command enables automatic summarization for EIGRPv6.
- D. The distribute-list route-map command enables route filtering for EIGRPv6.
- E. The no shutdown command enables the EIGRPv6 routing process.
- F. The v6routerid command configures an EIGRPv6 router ID.

Correct Answer: E



<https://vceplus.com/>

Section: (none)
Explanation

Explanation/Reference:

Explanation:

The no shutdown command enables the Enhanced Interior Gateway Routing Protocol version 6 (EIGRPv6) routing process. EIGRPv6 is also referred to as EIGRP for IPv6. To enable EIGRPv6 on a router, you should issue the ipv6 router eigrp as-number command in global configuration mode, where as-number is the autonomous system number (ASN), then issue the no shutdown command in router configuration mode to start the routing process.

The ipv6 router eigrp as-number command does not enable EIGRPv6 on any router interfaces; it only creates the EIGRPv6 routing process. EIGRPv6 must be enabled on each interface that should participate in EIGRP routing. To enable EIGRPv6 on an interface, you should issue the ipv6 eigrp as-number command in interface configuration mode. You need not configure EIGRPv6 on any interfaces that are configured as passive interfaces.

The network command does not configure the networks that should be advertised by EIGRPv6, because EIGRPv6 is configured directly on each participating interface. The network command is used with EIGRP for IPv4 to specify the networks that should be advertised out the router's interfaces.

The autosummary command does not enable automatic summarization for EIGRPv6. IPv6 does not use classful routing like IPv4 does, so automatic summarization is not possible with EIGRPv6.

The distribute-list routemap command cannot be used to filter routes in EIGRPv6. However, you can filter the EIGRPv6 routing updates by prefix list. To implement prefix list route filtering, you should issue the distribute-list prefix-list list-name command in router configuration mode.

The v6routerid command does not configure an EIGRPv6 router ID. The EIGRPv6 router ID is the same as the EIGRPv4 router ID, which is automatically configured unless there are no IPv4 addresses configured on the router. If there are no IPv4 addresses configured on the router, you must issue the routerid id command in router configuration mode to manually configure a router ID, where id is a 32bit value similar to an IPv4 address.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ipv6-15-2mt-book/ip6-eigrp.html>

Cisco: Implementing EIGRP for IPv6

QUESTION 22

What do dual stacks enable a host to do? (Select the best answer.)

- A. Dual stacks enable a host to pass IPv4 traffic over an IPv6only network.
- B. Dual stacks enable a host to pass IPv6 traffic over an IPv4only network.
- C. Dual stacks enable an IPv4only host to communicate with an IPv6only host.
- D. Dual stacks enable an IPv6only host to communicate with an IPv4only host.
- E. Dual stacks enable a host to pass IPv4 and IPv6 traffic.

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation/:

Dual stacks enable a host to send IPv4 and IPv6 traffic. Dualstack devices are configured with an IPv4 address and an IPv6 address; thus a dualstack device can communicate directly with both IPv4 devices and IPv6 devices without requiring protocol translation. However, a network infrastructure capable of routing both IPv4 and IPv6 traffic is required. The following partial output displays a router that is configured with dual stacks:

```
ipv6 unicast routing interface
fastethernet 0/1 ip address
10.1.14.7 255.255.255.0 ipv6
address 2001:0:0:1::2/64
```

Dual stacks alone do not enable a host to pass IPv4 traffic over an IPv6only network. A tunneling method must be implemented for IPv4 traffic to be passed over an IPv6only network. The implemented tunneling method should encapsulate an IPv4 packet inside an IPv6 header, thereby allowing the packet to travel across an IPv6 network. Because routers on the IPv6only network recognize only the IPv6 header information, the IPv4 packet is carried as the data payload of the IPv6 packet.

Similarly, dual stacks alone do not enable a host to pass IPv6 traffic over an IPv4only network. The 6to4 tunneling method is one method that is used to pass IPv6 traffic over an IPv4only network. The 6to4 tunneling method is the reverse of the 4to6 tunneling method; it encapsulates an IPv6 packet inside an IPv4 header. Network Address Translation 64 (NAT64), not dual stacks, enables an IPv4only host to communicate with an IPv6only host and enables an IPv6only host to communicate with an IPv4only host. NAT64 translates IPv4 packets to IPv6 packets and translates IPv6 packets to IPv4 packets. However, a NAT64 router must contain address mappings so that the router can correctly translate IPv4 and IPv6 addresses. NAT64 supports stateless and stateful address translation. When performing stateless translation, NAT64 uses algorithms to create a one-to-one relationship between IPv6 addresses on the inside network and IPv4 addresses on the outside network. When performing stateful translation, NAT64 maps multiple IPv6 addresses to a single IPv4 address and keeps track of the state of each connection. Static mappings can also be applied manually.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ipv6-15-2mt-book/ip6-addrg-bsc-con.html#GUID-1CE3CFDD-8889-4C0C-84BE56505EBC8517>

QUESTION 23

At which security level does SNMPv3 use CBCDES to encrypt authentication? (Select the best answer.)

- A. at the noAuthNoPriv security level
- B. at the authNoPriv security level
- C. at the authPriv security level
- D. at none of the security levels

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Simple Network Management Protocol version 3 (SNMPv3) uses Cipher Block Chaining Data Encryption Standard (CBCDES) to encrypt authentication at the authPriv security level. SNMP is used to monitor and manage network devices by collecting statistical data about those devices. The authPriv security level authenticates by matching a Message Digest 5 (MD5) or Secure Hash Algorithm (SHA) hash of the user name. The authentication process is also encrypted by using either Data Encryption Standard (DES), Triple DES (3DES), or Advanced Encryption Standard (AES).

Three versions of SNMP currently exist. SNMPv1 and SNMPv2C do not provide encryption? password information, known as community strings, is sent as plain text with messages. If an attacker intercepts the message, the attacker can view the password information. SNMPv3 improves upon SNMPv1 and SNMPv2 by providing encryption, authentication, and message integrity to ensure that the messages are not tampered with during transmission.

Two SNMPv3 security levels, authNoPriv and authPriv, authenticate by matching Hashbased Message Authentication CodeSHA (HMACSHA) or HMACMD5 authentication strings. The authPriv security level is the only SNMPv3 security level that can encrypt the authentication process.

The noAuthNoPriv security level in SNMPv3 authenticates by matching a user name sent as clear text. Earlier versions of SNMP, such as SNMPv1 and SNMPv2C, match community strings instead of user names.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/software/15-0_2_EX/network_management/configuration_guide/b_nm_15ex_2960-x_cg/b_nm_15ex_2960-x_cg_chapter_0100.html#reference_160326642C03413B92A68E856426EABA

QUESTION 24

DRAG DROP

Select the metrics from the left, and place them in the correct order that BGP will use them to determine the best path to a destination. Not all metrics will be used.

Select and Place:

longest AS path	highest origin type	First
shortest AS path	lowest origin type	
highest IGP cost	locally originated paths	
lowest IGP cost	externally originated paths	
eBGP paths over iBGP paths	newest eBGP path	
iBGP paths over eBGP paths	oldest eBGP path	
highest local preference	highest BGP RID	
lowest local preference	lowest BGP RID	
highest MED	highest weight	
lowest MED	lowest weight	Last

Help Reset Done

Correct Answer:

The screenshot displays a list of BGP path selection criteria organized into three columns:

- Column 1 (Left):** longest AS path, highest IGP cost, iBGP paths over eBGP paths, lowest local preference, highest MED.
- Column 2 (Middle):** highest origin type, externally originated paths, newest eBGP path, highest BGP RID, lowest weight.
- Column 3 (Right):** highest weight, highest local preference, locally originated paths, shortest AS path, lowest origin type, lowest MED, eBGP paths over iBGP paths, lowest IGP cost, oldest eBGP path, lowest BGP RID.

At the bottom of the interface, there are three buttons: **Help**, **Reset**, and **Done**.

Section: (none)
Explanation

Explanation/Reference:

Explanation:

Border Gateway Protocol (BGP) uses a complex method of selecting the best path to the destination. The following list displays the criteria used by BGP for path selection:

1. Highest weight
2. Highest local preference

3. Locally originated paths over externally originated paths
4. Shortest autonomous system (AS) path
5. Lowest origin type
6. Lowest multiexit discriminator (MED)
7. External BGP (eBGP) paths over internal BGP (iBGP) paths
8. Lowest Interior Gateway Protocol (IGP) cost
9. Oldest eBGP path
10. Lowest BGP router ID (RID)

When determining the best path, a BGP router first chooses the route with the highest weight. Weight is a Cisco proprietary BGP path attribute that is significant only to the local router; it is not advertised to neighbor routers. To configure the weight value, you should issue the `neighbor {ipaddress | peergroupname} weightweightvalue` command, where `ipaddress` is the IP address of a neighbor router, `peergroupname` is the name of a BGP peer group, and `weightvalue` is a locally significant weight value from 0 through 65535. By default, routes generated by the local router are assigned a weight of 32768 and routes learned from another BGP router are assigned a weight of 0.

When weight values are equal, a BGP router chooses the route with the highest local preference. The local preference value is advertised to iBGP neighbor routers to influence routing decisions made by those routers. To configure the local preference, you should issue the `bgp default localpreference number` command, where `number` is a value from 0 through 4294967295.

When local preferences are equal, a BGP router chooses locally originated paths over externally originated paths. Locally originated paths that have been created by issuing the `network` or `redistribute` command are preferred over locally originated paths that have been created by issuing the `aggregate-address` command. If multiple paths to a destination still exist, a BGP router chooses the route with the shortest AS path attribute. The AS path attribute contains a list of the AS numbers (ASNs) that a route passes through.

If multiple paths have the same AS path length, a BGP router chooses the lowest origin type. An origin type of `i`, which is used for IGP, is preferred over an origin type of `e`, which is used for Exterior Gateway Protocols (EGPs). These origin types are preferred over an origin type of `?`, which is used for incomplete routes where the origin is unknown or the route was redistributed into BGP.

If origin types are equal, a BGP router chooses the route with the lowest MED. A MED value is basically the external metric of a route that is advertised to eBGP routers in order to specify a preferred path into an AS with multiple entry points. To configure the MED value, you should issue the `defaultmetric number` command, where `number` is a value from 1 through 4294967295. Routes redistributed into BGP are assigned this MED value; redistributed connected routes are assigned a MED value of 0 regardless of the `defaultmetric` setting.

If MED values are equal, a BGP router chooses eBGP routes over iBGP routes. If there are multiple eBGP paths, or multiple iBGP paths if no eBGP paths are available, a BGP router chooses the route with the lowest IGP metric to the nexthop router. If IGP metrics are equal, a BGP router chooses the oldest eBGP path, which is typically the most stable path.

Finally, if route ages are equal, a BGP router chooses the path that comes from the router with the lowest RID. The RID can be manually configured by issuing the `bgp router-id` command. If the RID is not manually configured, the RID is the highest loopback IP address on the router. If no loopback address is configured, the RID is the highest IP address from among a router's available interfaces.

Reference:

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

QUESTION 25

Which of the following steps in the NAT order of operation typically occur after NAT outside-to-inside translation? (Select 4 choices.)

- A. decryption
- B. encryption
- C. redirect to web cache
- D. check inbound access list
- E. check outbound access list
- F. inspect CBAC
- G. IP routing

Correct Answer: BEFG

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following steps of the Network Address Translation (NAT) order of operation typically occur after NAT outside-to-inside translation:

- Encryption
- Check outbound access list
- IP routing
- Inspect Contextbased Access Control (CBAC)

NAT enables a network to communicate with a separate network, such as the Internet, by translating traffic from IP addresses on the local network to another set of IP addresses that can communicate with the remote network. NAT outside-to-inside translation, which is also known as global-to-local translation, occurs when the NAT router maps an outside destination IP address to an inside destination IP address. When a NAT router performs NAT outside-to-inside translation, the following operations occur in order:

- 1.If IP Security (IPSec) is implemented, check inbound access list
- 2.Decryption
- 3.Check inbound access list
- 4.Check inbound rate limits
- 5.Inbound accounting
- 6.Redirect to web cache
- 7.NAT outside-to-inside translation

8. Policy routing
9. IP routing
10. Check crypto map and mark for encryption
11. Check outbound access list
12. Inspect CBAC
13. Transmission Control Protocol (TCP) intercept
14. Encryption
15. Queueing

Conversely, when a NAT router performs NAT insidetooutside, or localtoexternal, translation, the NAT insidetooutside translation operation immediately follows the IP routing operation. Otherwise, the order of operation is the same:

1. If IPSec is implemented, check inbound access list
2. Decryption
3. Check inbound access list
4. Check inbound rate limits
5. Inbound accounting
6. Redirect to web cache
7. Policy routing
8. IP routing
9. NAT insidetooutside translation
10. Check crypto map and mark for encryption
11. Check outbound access list
12. Inspect CBAC
13. TCP intercept
14. Encryption
15. Queueing

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/network-address-translation-nat/6209-5.html#topic1>

QUESTION 26

Which of the following statements is true regarding VRRPv3? (Select the best answer.)

- A. VRRPv3 supports IPv6, but previous versions of VRRP do not.
- B. VRRPv3 supports authentication, but previous versions of VRRP do not.
- C. VRRPv3 supports multiple master virtual routers, but previous versions of VRRP do not.

D. VRRPv3 is a standards-based protocol, but previous versions of VRRP are not.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Virtual Router Redundancy Protocol version 3 (VRRPv3) supports IPv6, but previous versions of VRRP do not. VRRP enables a group of routers to appear like a single default gateway. VRRPv3 uses the IPv4 or IPv6 address of a physical interface on the master virtual router, which is the router in the group with the highest VRRP priority. The other routers in the group are backup virtual routers. If the master virtual router fails, the backup virtual router with the highest priority will assume the role of the master virtual router, thereby providing uninterrupted service for the network. When the original master virtual router comes back online, it reestablishes its role as the master virtual router.

VRRPv3 is a standardsbased protocol, but so are previous versions of VRRP. VRRPv1 is defined in Request for Comments (RFC) 2338. VRRPv2 is defined in RFC 3768. VRRPv3 is defined in RFC 5798.

Previous versions of VRRP included support for authentication, but VRRPv3 as defined in RFC 5798 does not include support for authentication. Cisco's implementation of VRRP supports plaintext and Message Digest 5 (MD5) authentication. When a router receives a VRRP packet for its VRRP group, it validates the authentication string. If the authentication string does not match the string that is configured on the router, the VRRP packet is discarded. When plaintext authentication is configured, the authentication string is sent unencrypted. When MD5 authentication is configured, each VRRP packet is sent with a keyed MD5 hash of that packet? if the receiving device does not generate the same hash, the packet is ignored.

Neither VRRPv3 nor previous versions of VRRP support multiple master virtual routers. All versions of VRRP allow only a single device to become the master virtual router for a group. Therefore, a VRRP group cannot be configured to use multiple devices in a load-balancing configuration.

Reference:

IETF: RFC 2338: Virtual Router Redundancy Protocol

IETF: RFC 3768: Virtual Router Redundancy Protocol (VRRP)

IETF: RFC 5798: Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6

QUESTION 27

For which of the following routes does PfR use PIRO? (Select 2 choices.)

- A. BGP
- B. EIGRP
- C. IS-IS
- D. OSPF
- E. static

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Performance Routing (PfR) uses Protocol Independent Routing Optimization (PIRO) to control path selection for Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) routes. PfR can control path selection directly for Border Gateway Protocol (BGP), Enhanced Interior Gateway Routing Protocol (EIGRP), and static routes.

PfR enhances traditional routing methods by dynamically selecting the best path for traffic classes based on network performance. The path selection procedure can be influenced by several factors, including delay, packet loss, reachability, throughput, jitter, and mean opinion score (MOS). When PfR wants to modify a path for a traffic class, it will search for a parent route, which is an exactmatching route or a lessspecific route. PfR will search for a parent route in the following locations, in order:

1. BGP routing database
2. EIGRP routing database
3. Static route database

PIRO extends the capabilities of PfR by searching for a parent route within the IP Routing Information Base (RIB) after the other locations have been searched.

Reference:

https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/performance-routing-pfr/product_data_sheet0900aecd806c4ee4.html

http://docwiki.cisco.com/wiki/Performance_Routing_FAQs

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/pfr/configuration/xe-3s/pfr-xe-3s-book/pfr-simple-ph1.html>

https://www.cisco.com/c/dam/global/bg_bg/assets/ciscoexpo2011/pdf/Next_Generation_Routing_Architectures-Gerd_Pflueger.pdf

QUESTION 28

You issue the following commands on RouterA:

```
RouterA(config)#policy-map boson
```

```
RouterA(config-pmap)#class applications
```

```
RouterA(config-pmap-c)#police 100000 5000 8000 conform-action transmit exceed-action set-qos-transmit 4 violate-action drop
```

When will RouterA begin to drop packets? (Select the best answer.)

- A. when the burst rate exceeds 5,000 bits
- B. when the burst rate exceeds 8,000 bits
- C. when the burst rate exceeds 40,000 bits
- D. when the burst rate exceeds 64,000 bits
- E. when the burst rate exceeds 100,000 bits

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RouterA will begin to drop packets when the burst rate exceeds 64,000 bits. You can issue the police command to explicitly configure a maximum bandwidth limit. The syntax of the police command is `policebps [burstnormal] [burstmax] conform action actionexceedaction action [violateaction action]`. The bps parameter is the average rate specified in bits per second, and the optional burstnormal and burstmax parameters are specified in bytes. When traffic exceeds the burstnormal rate, the router will perform the exceedaction action, and when traffic exceeds the burstmax rate, the router will perform the violateaction action.

Traffic policing is used to slow down traffic to a value that the medium can support, to monitor bandwidth utilization, to enforce bandwidth limitations at the service provider edge, and to remark traffic that exceeds the Service Level Agreement (SLA). Excess traffic and outofprofile packets are dropped or remarked and transmitted. By contrast, traffic shaping buffers excess traffic and outofprofile packets in memory and drops traffic only if the queue is full. Because traffic shaping does not remark traffic, it can create queuing delay, particularly when queues are large and traffic flow is heavy.

In this scenario, the burstmax rate is set to a value of 8,000 bytes, which is equal to 64,000 bits. The action that corresponds to the violateaction keyword is drop. The drop keyword configures the router to silently drop packets. Therefore, when burst traffic exceeds 64,000 bits, some packets will be dropped.

RouterA will not begin to drop packets when the burst rate exceeds 5,000 bits or 8,000 bits. The burstnormal and burstmax parameters are specified in bytes, not bits.

RouterA will not begin to drop packets when the burst rate exceeds 40,000 bits. The burstnormal rate is set to a value of 5,000 bytes, which is equal to 40,000 bits. The action that corresponds to the exceedaction keyword is `setqostransmit 4`. Therefore, when burst traffic exceeds 40,000 bits, some packets will begin to be reclassified with a Quality of Service (QoS) value of 4 and will be transmitted.

RouterA will begin to drop packets before the burst rate exceeds 100,000 bits. The bit rate indicates the average rate of burst traffic, not the rate at which packets will begin to be dropped.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/qos/configuration/guide/fqos_c/qcfdpoli.html https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_plcshp/configuration/xs-3s/qos-plcshp-xe-3s-book/qos-plcshp-trfc-plc.html

QUESTION 29

In which of the following situations are redistributed routes not entered into the routing table by default? (Select 2 choices.)

- A. when BGP routes are redistributed into OSPF
- B. when OSPF routes are redistributed into EIGRP
- C. when EIGRP routes are redistributed into BGP
- D. when RIP routes are redistributed into OSPF
- E. when EIGRP routes are redistributed into RIP

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Open Shortest Path First (OSPF) routes redistributed into Enhanced Interior Gateway Routing Protocol (EIGRP) and EIGRP routes redistributed into Routing Information Protocol (RIP) are not entered into the routing table by default. EIGRP uses a complex metric based on bandwidth, delay, reliability, and load. Because of its complex metric, EIGRP requires that redistributed routes be assigned a metric before they are entered into the routing table. To assign a default metric for routes redistributed into EIGRP, you should issue the defaultmetric bandwidth delay reliability loading mtu command. To assign a metric to an individual route redistributed into EIGRP, you should issue the redistribute protocol[processid | autonomoussystemnumber] metric bandwidth delay reliability loading mtucommand. If no metric is assigned during redistribution and no default metric is configured for EIGRP, the routes are assigned an infinite metric and are ignored by EIGRP. RIP uses hop count as a metric. Valid hop count values are from 1 through 15; a value of 16 is considered to be infinite. The hop count metric increases by 1 for each router along the path. Cisco recommends that you set a low value for the hop count metric for redistributed routes. To assign a default metric for routes redistributed into RIP, you should issue the defaultmetric hopcount command. To assign a metric to an individual route redistributed into RIP, you should issue the redistribute protocolhopcount command. If no metric is assigned during redistribution and no default metric is configured for RIP, the routes are assigned an infinite metric and are ignored by RIP.

Border Gateway Protocol (BGP) routes and RIP routes redistributed into OSPF are entered into the routing table as external routes. The default metric that OSPF assigns to redistributed routes is 20; however, BGP is an exception and is assigned a default metric of 1. OSPF uses a cost metric based on the bandwidth of each participating interface. OSPF prefers internal routes with the lowest cost. By default, all routes redistributed into OSPF are designated as Type 2 external (E2) routes. E2 routes have a metric that remains constant throughout the routing domain. Alternatively, routes redistributed into OSPF can be designated as Type 1 external (E1) routes. With E1 routes, the internal cost of the route is added to the initial metric assigned during redistribution.

EIGRP routes redistributed into BGP are entered into the routing table without the metric being changed. BGP uses the EIGRP metric as a multi-exit discriminator (MED). The MED is one of several variables BGP considers before making a path selection. BGP considers weight, local preference, origin, and autonomous system (AS) path length before using the MED for path selection.

Reference:

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

<https://www.cisco.com/networkers/nw04/presos/docs/CERT-2100.pdf#page=6>

QUESTION 30

Which of the following does a BGP cluster ID identify? (Select the best answer.)

- A. the originator of a route
- B. a group of route reflectors
- C. a route reflector in a cluster
- D. the clusters that a route has passed through

E. a group of peers with the same update policies

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A Border Gateway Protocol (BGP) cluster ID identifies a group of route reflectors. Internal BGP (iBGP) routes are not advertised to iBGP peers. In order to avoid having to create a fullmesh configuration, you can configure one or more route reflectors to pass iBGP routes between iBGP routers. A route reflector and its peers form a cluster, and the route reflector is configured with a 4byte cluster ID. To increase redundancy, a cluster can have multiple route reflectors. Each route reflector in the cluster should be fully meshed and configured with the same cluster ID so that the route reflector can recognize routing updates from other route reflectors in the cluster. To configure a route reflector with a cluster ID, you should issue the `bgp cluster-id cluster-id` command from BGP router configuration mode. A BGP cluster ID does not identify a single route reflector in a cluster. Each route reflector is identified by its router ID. When a cluster has only a single route reflector, the cluster ID is often configured with the route reflector's router ID. When a cluster has multiple route reflectors, the cluster ID must be the same on all of the route reflectors.

A BGP cluster ID does not identify the clusters that a route has passed through; this is the function of a cluster list. When a route reflector sends a route to or receives a route from a nonclient peer router, the route reflector appends its cluster ID to the cluster list. If no cluster list exists, a cluster list is created with the cluster ID of the route reflector. If a route reflector receives a routing update with its cluster ID in the cluster list, the routing update is ignored.

A BGP cluster ID does not identify a group of peers with the same update policies; this is the function of a peer group. Peer groups can simplify administration by enabling an administrator to simultaneously configure a group of peers with the same update policies, such as route maps, filter lists, and distribute lists. Any configuration options that are configured with the specified peer group name will be applied to members of the peer group. To define a peer group, you should issue the `neighbor peer-group-name-peer-group` command.

A BGP cluster ID does not identify the originator of a route; this is the function of the originator ID. A route reflector that originates a route in a local autonomous system (AS) will insert its router ID as the originator ID. If a route reflector receives a routing update with its router ID as the originator ID, the routing update is ignored.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfbgp.html#wp1001965 <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/26634-bgp-toc.html#routereflectors>

QUESTION 31

When optimizing throughput, which of the following formulas would you use to calculate BDP? (Select the best answer.)

- A. $BW \times rwin = BDP$
- B. $BW \div rwin = BDP$
- C. $BW \times RTT = BDP$

D. $BW \div RTT = BDP$

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When optimizing throughput, you would use the formula $BW \times RTT = BDP$ to calculate bandwidth delay product (BDP), where BW is the bandwidth and RTT is the round-trip time, also known as latency. BDP is the maximum amount of data that can exist on a network path at any given time. This value is often used to establish a maximum goal when optimizing the bandwidth used by Transmission Control Protocol (TCP) traffic flows.

When calculating BDP, you should ensure that the unit measurements match for each variable. For example, BW is typically measured in bits per second, whereas RTT is typically measured in milliseconds. Therefore, you should convert the RTT value to seconds before multiplying the values. For example, a Fast Ethernet link with a latency of 40 milliseconds (ms) would have a BDP of 4 Mb: $100 \text{ Mbps} \times 0.040 \text{ seconds} = 4 \text{ Mb}$

The BDP measurement will match the measurement that is used for the BW variable. In this example, the BDP is expressed in megabits. If you need to convert from bits (lowercase b) to bytes (uppercase B), you should divide by 8. Conversely, if the bandwidth is expressed in bytes and you need to convert to bits, you should multiply by 8.

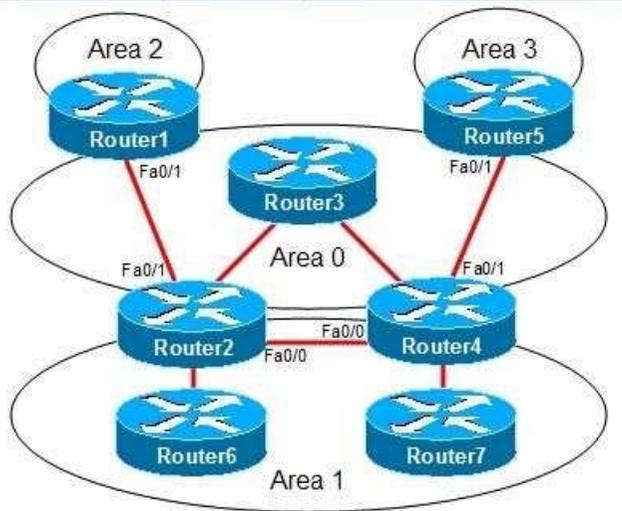
The rwin variable is not used when calculating BDP. However, rwin is related to BDP. The TCP receive window size, sometimes expressed as rwin, is often a limiting factor in optimizing throughput. In order to achieve maximum throughput, you should set the TCP receive window size to a value equal to or greater than BDP, thereby ensuring that either bandwidth or latency is the limiting factor.

Reference:

https://www.cisco.com/application/pdf/en/us/guest/tech/tk277/c1482/ccmigration_09186a00801b1259.pdf#page=26

https://www.cisco.com/c/en/us/td/docs/nsite/enterprise/wan/wan_optimization/wan_opt_sg/chap06.html#wp1053392

QUESTION 32



You administer the OSPF network shown in the diagram. Area 1 is configured as a standard area. Area 2 and Area 3 are configured as stub areas. Router3 fails. Several routes are lost throughout the network. Which of the following actions can you take to restore the lost routes? (Select 2 choices.)

- A. Configure Area 1 as a stub area.
- B. Configure Area 2 and Area 3 as standard areas.
- C. Create a virtual link between Router1 and Router5.
- D. Create a virtual link between Router2 and Router4.
- E. Configure the Fa0/0 interfaces on Router2 and Router4 to be part of Area 0.
- F. Configure the Fa0/1 interfaces on Router2 and Router4 with IP addresses that were configured on Router3.

Correct Answer: DE

Section: (none)

Explanation

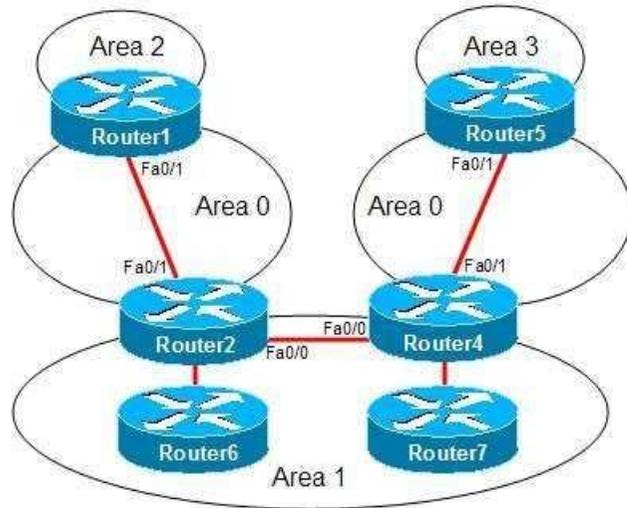
Explanation/Reference:

Explanation:

You can take either of the following actions to restore the lost routes:

- Create a virtual link between Router2 and Router4.
- Configure the Fa0/0 interfaces on Router2 and Router4 to be part of Area 0.

In this scenario, the backbone area, Area 0, has become discontinuous, or partitioned, as shown in the following network diagram:



To connect a backbone area that has become discontinuous because of the loss of a router or the loss of a link between two routers, you can create a virtual link. The routers at each end of the virtual link must adhere to the following restrictions:

- Both routers must connect to the backbone area.
- Both routers must share another common area, which is used as a transit area.
- The transit area cannot be a stub area.
- The transit area cannot be the backbone area.

To create a virtual link, you should issue the `area area-id virtual-link router-id` command in router configuration mode on the routers at each end of the virtual link, where `area-id` is the transit area ID and `routerid` is the router ID of the router at the other end of the virtual link. For example, if the router ID of Router4 were 1.2.3.4, you would issue the `area 1 virtual-link 1.2.3.4` command on Router2. You would also issue a similar command on Router4 by using the router ID of Router2 as the `router-id` parameter.

Alternatively, you can configure the Fa0/0 interfaces on Router2 and Router4 to be part of Area 0. Doing so would make Area 1 discontinuous. This is acceptable because interarea traffic must pass through the backbone or a transit area; therefore, nonbackbone areas can be discontinuous. The discontinuous Area 1 partitions would be advertised to one another through inter-area routes instead of intra-area routes.

Configuring Area 1 as a stub area will not restore the lost routes. Additionally, configuring Area 1 as a stub area eliminates the possibility of using a virtual link to connect the discontinuous backbone areas.

Configuring Area 2 and Area 3 as standard areas will not restore the lost routes. Changing a stub area to a standard area will only allow Type 5 external summary routes to be advertised throughout the area.

You cannot create a virtual link between Router1 and Router5. For a virtual link to be created, both routers must share a common area. If Router1 and Router5 shared a nonstub area, you could create a virtual link between them and the lost routes would be restored. Configuring the Fa0/1 interfaces on Router2 and Router4 with IP addresses that were configured on Router3 will not restore the lost routes. The routes were not lost because of the unavailability of the IP addresses on Router3; the routes were lost because of the discontinuous backbone area.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html#t14> <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html#t17>

QUESTION 33

Which of the following ICMPv6 message types is sent by an IPv6capable host at startup? (Select the best answer.)

- A. router solicitation
- B. router advertisement
- C. neighbor solicitation
- D. neighbor advertisement

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

An Internet Control Message Protocol version 6 (ICMPv6) router solicitation message is sent by an IPv6capable host at startup. When IPv6 is enabled on a router interface, a linklocal address is created. Before the address is assigned to the interface, duplicate address detection (DAD) is performed to determine whether the IPv6 address is unique on the link. If DAD determines that the address is unique, the linklocal address is assigned to the interface and the router solicitation message is sent to the allrouters multicast address FF02::2. Hosts use router solicitation messages to request an immediate router advertisement.

A router advertisement that is sent in response to a router solicitation message is sent directly to the host that sent the router solicitation. Routers also send unsolicited router advertisements periodically to the allnodes multicast address FF02::1. Router advertisements contain the following information:

- The IPv6 address of the router interface attached to the link
- One or more IPv6 prefixes for the local link
- The lifetime for each prefix
- Flags that specify whether stateless or stateful autoconfiguration can be used
- The hop limit and maximum transmission unit (MTU) that the host should use
- Whether the router is a default router
- The amount of time that the router can be used as a default router

When a host receives a router advertisement, the IPv6 link-local prefix is added to the host's interface identifier to create the host's full IPv6 address. The first three octets of the interface identifier are set to the Organizationally Unique Identifier (OUI) of the Media Access Control (MAC) address of the interface. The fourth and fifth octets are set to FFFE. The sixth, seventh, and eighth octets are equal to the last three octets of the MAC address.

A host will send a neighbor solicitation message to determine the link-layer address of another host on the local link. Neighbor solicitation messages are sent with the sender's own link-layer address to the solicited-node multicast address. The solicited-node multicast address is created by adding the FF02::1:FF00/104 prefix to the last 24 bits of the destination host's IPv6 address. After a destination host's link-layer address is discovered, neighbor solicitations can be used to verify the reachability of a destination host.

When a host receives a neighbor solicitation message, it will reply with a neighbor advertisement message that contains the link-layer address of the host. The neighbor advertisement is sent directly to the host that sent the neighbor solicitation. A host will send an unsolicited neighbor advertisement whenever its address changes. Unsolicited neighbor advertisements are sent to the allnodes link-local multicast address FF02::1.

Reference:

Cisco: Implementing IPv6 Addressing and Basic Connectivity: IPv6 Router Advertisement Message IETF: RFC 4861: Neighbor Discovery for IP version 6 (IPv6)

QUESTION 34

Which of the following is used to encrypt data between GET VPN group members? (Select the best answer.)

- A. KEK
- B. SAR
- C. TEK
- D. TSK



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A traffic encryption key (TEK) is used to encrypt data between Group Encrypted Transport (GET) virtual private network (VPN) group members. GET VPN is a connectionless, nontunneling VPN technology based on the Group Domain of Interpretation (GDOI) standard proposed in Request for Comments (RFC) 3547. Nontunneling VPNs such as GET VPN can be used on a variety of networks, including IP, Frame Relay, Multiprotocol Label Switching (MPLS), and Asynchronous Transfer Mode (ATM) networks. Although GET VPN does not use tunneling, it does rely upon Internet Key Exchange (IKE) and IP Security (IPSec) security associations (SAs).

GET VPN requires a key server. The key server maintains the policy, creates and maintains group keys, and services registration requests. When a group member registers with the key server, the group member downloads the IPSec policy and encryption keys from the key server. If a group member fails to register with a key server, all traffic is sent unencrypted through the group member unless the FailClose feature is activated.

A key encryption key (KEK) is used to encrypt data between the key server and group members. Periodically, the key server will send rekey messages to group members in order to refresh the IPsec SA before it expires. The KEK protects the rekey message, which contains new encryption keys that the group members should use, thereby securing the control plane.

Synchronous Antireplay (SAR) provides antireplay protection for GET VPN group members. The key server keeps track of time by maintaining a pseudotime clock. Group members regularly synchronize to the pseudotime on the key server. If an intercepted message is replayed, the replayed message will likely fall outside the pseudotime window. A group member will detect the pseudotime discrepancy and will therefore reject the replayed message.

A transmission security key (TSK) is used by directsequence spread spectrum (DSSS) or frequencyhopping radios. TSKs are not used by GET VPN group members.

Reference:

https://www.cisco.com/en/US/docs/ios-xml/ios/sec_conn_getvpn/configuration/15-2mt/sec-get-vpn.html#GUID-E692E019-DFFD-4763-AD2F-2C3080844581

QUESTION 35

Which of the following commands will cause a router to be prone to CEF polarization? (Select the best answer.)

- A. no ip cef loadsharing algorithm
- B. ip cef loadsharing algorithm original
- C. ip cef loadsharing algorithm tunnel
- D. ip cef loadsharing algorithm universal
- E. ip cef loadsharing algorithm includeports source destination



Correct Answer: B

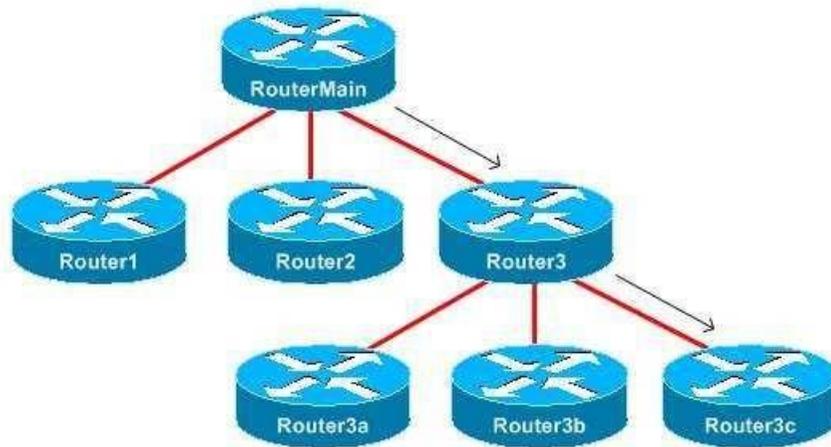
Section: (none)

Explanation

Explanation/Reference:

Explanation:

The ip cef loadsharing algorithm original command will cause a router to be prone to Cisco Express Forwarding (CEF) polarization, which occurs when multiple routers in sequence use the same loadbalancing mechanism. To understand CEF polarization, consider the following topology:



RouterMain will run the load-balancing algorithm on a flow and, based on the hash result, will send the flow to Router1, Router2, or Router3. If Router1, Router2, and Router3 run the same loadbalancing algorithm as RouterMain uses, those routers will get the same hash result and will therefore no longer load balance. For example, flows that are sent from RouterMain to Router3 will always be forwarded to Router3c because Router3 generates the same hash for each flow that RouterMain does.

The ip cef loadsharing algorithm original command configures CEF to load balance based only on the source and destination. Universal mode improves on the original CEF loadbalancing algorithm by using a source, a destination, and a 32bit Universal ID as a hashing seed. Because each router uses a different Universal ID, each router will produce different hashing values, thereby avoiding CEF polarization by enabling each router to load balance differently. Universal mode is enabled by default or by issuing the ip cef loadsharing algorithm universal command. Because universal mode is enabled by default, the no ip cef loadsharing algorithm command enables universal mode, thereby avoiding CEF polarization.

The ip cef loadsharing algorithm tunnel command avoids CEF polarization. The tunnelmode algorithm uses an improved universalmode algorithm that works well in environments with a small number of source and destination pairs, which often occurs with tunnels.

The ip cef loadsharing algorithm includeports source destination command avoids CEF polarization. This command configures CEF to not only use the universal loadbalancing algorithm but also to consider Layer 4 source and destination port information.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipswitch/command/isw-cr-book/isw-i1.html#wp5609710740>

QUESTION 36

Which of the following commands configures primary and fallback link groups for Cisco PfR? (Select the best answer.)



<https://vceplus.com/>

- A. border
- B. link-group
- C. pfr-map
- D. set link-group

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The set link-group command configures primary and fallback link groups for Cisco Performance Routing (PfR). The syntax of the set linkgroup command is set linkgroup primarylinkgroupname [fallbackfallbacklinkgroupname]. Link groups enable you to configure a set of exit interfaces as preferred or standby links to optimize traffic classes in a PfR policy. Primary and fallback link groups are configured on the master controller.

The border command does not configure primary and fallback link groups. Instead, the border command is issued on the master controller in order to configure a border router. A master controller can control up to 10 border routers. In small environments, the master controller and the border router will be the same device.

The syntax of the bordercommand is borderipaddress [keychain keychainname], where ipaddress is the IP address of the border router.

The linkgroup command does not configure primary and fallback link groups. Instead, the linkgroup command configures a border router exit interface as a member of a link group. Up to three link groups can be specified on an interface. The syntax of the linkgroup command is linkgroup linkgroupname [linkgroupname [linkgroupname]].

The pfrmap command does not configure primary and fallback link groups. Instead, the pfrmap command configures a PfR map. The syntax of the pfrmap command is pfrmap mapname sequencenumber. Issuing the pfrmap command places the router into pfrmap configuration mode. In this mode, you will configure a match clause with the match command and specify the primary and fallback link groups with the set linkgroupcommand. Only one match command can be configured for a PfR map sequence.

Reference:

<https://search.cisco.com/search?query=Cisco%20IOS%20Performance%20Routing%20Configuration%20Guide&locale=enUS&tab=Cisco>

QUESTION 37

DRAG DROP

Select the terms on the left, and drag them to the corresponding features on the right.

Select and Place:

BGP Enhanced Route Refresh	makes the ABR an RR and sets the next hop to self
BGP PIC	stores an alternate path in the RIB, FIB, and CEF
RTC	synchronizes peers without a hard reset
Unified MPLS	uses the rfilter address family

Correct Answer:

	Unified MPLS
	BGP PIC
	BGP Enhanced Route Refresh
	RTC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Route Target Constraint (RTC) uses the rfilter address family. In a normal Multiprotocol Label Switching (MPLS) virtual private network (VPN), the route reflector (RR) sends all of its VPN version 4 (VPNv4) and VPNv6 prefixes to the Provider Edge (PE) router. The PE router then drops the prefixes for which it does not have a matching VPN routing and forwarding (VRF). RTC sends only prefixes that the PE router wants. When RTC is enabled, the PE router sends its route target (RT) membership data to the RR within an address family named rfilter. The RR then uses rfilter to determine which prefixes to send to the PE. In order for RTC to work, both the RR and the PE need to support RTC.

Border Gateway Protocol (BGP) Enhanced Route Refresh finds route inconsistencies, and if inconsistencies exist, peers are synchronized without a hard reset. If two BGP peers support Enhanced Route Refresh, each peer will send a RouteRefresh StartofRIB (SOR) message and a RouteRefresh EndofRIB (EOR) message before and after an AdjRIBOut message, respectively. After a peer receives an EOR message, or after the EOR timer expires, the peer will check to see whether it has any routes that were not readvertised. If any stale routes remain, they are deleted and the route inconsistency is logged.

Unified Multiprotocol Label Switching (MPLS) makes the area border router (ABR) an RR and sets the next hop to self. Unified MPLS increases scalability for an MPLS network by extending the label switched path (LSP) from end to end, not by redistributing interior gateway protocols (IGPs) into one another, but by distributing some of the IGP prefixes into BGP. BGP then distributes those prefixes throughout the network.

BGP PrefixIndependent Convergence (PIC) improves convergence by creating and storing an alternate path in the Routing Information Base (RIB), Forwarding Information Base (FIB), and Cisco Express Forwarding (CEF). As soon as a failure is detected, BGP uses the alternate path. BGP PIC works on IPv4, IPv6, and MPLS networks.

Reference: <https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/116062-technologies-technote-restraint-00.html>

<https://search.cisco.com/search?query=Cisco%20IOS%20BGP%20Configuration%20Guide&locale=enUS&tab=Cisco>

<https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/116127-configure-technology-00.html>

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

QUESTION 38

Which of the following is true regarding the structure of a multiprotocol BGP VPNIPv4 address? (Select the best answer.)

- A. It begins with a 2byte Type field and ends with a 6byte Value field.
- B. It begins with an 8byte RD and ends with a 4byte IPv4 address.
- C. It begins with a 4byte VPN ID and ends with an 8byte RD.
- D. It begins with a 4byte ASN and ends with a 2byte Assigned Number.
- E. It begins with a 6byte MAC address and ends with a 4byte IPv4 address.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A multiprotocol Border Gateway Protocol (BGP) virtual private network (VPN) IPv4 address begins with an 8byte route distinguisher (RD) and ends with a 4byte IPv4 address. The RD consists of a 2byte Type field and a 6byte Value field. The value of the Type field determines what the structure of the Value field is. The following table lists the Type values along with their corresponding Value field structures:

Type Field	Value Field
0	2-byte Administrator subfield and 4-byte Assigned Number subfield
1	4-byte Administrator subfield and 2-byte Assigned Number subfield
2	4-byte Administrator subfield and 2-byte Assigned Number subfield

If the Type field is 0, the Administrator subfield is a 2byte autonomous system number (ASN). If the Type field is 1, the Administrator subfield is an IP address. If the Type field is 2, the Administrator subfield is a 4byte ASN. In all cases, the Assigned Number subfield contains a number assigned by the administrator.

The BGP VPNIPv4 address does not contain a VPN ID or a Media Access Control (MAC) address.

Reference:

<https://tools.ietf.org/html/rfc4364>

QUESTION 39

Which of the following is required by MPLS? (Select the best answer.)

- A. BGP
- B. CDP
- C. CEF
- D. IS-IS
- E. TDP

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Cisco Express Forwarding (CEF) is required by Multiprotocol Label Switching (MPLS). MPLS relies on CEF to forward labeled packets through the network. If CEF is not enabled, an MPLS router cannot switch labeled packets and MPLS functionality is lost. By default, CEF is enabled on Cisco routers; however, if CEF is disabled, you should enable it by issuing the ip cef command in global configuration mode. CEF depends on the IP routing functionality of the router and cannot be enabled unless IP routing is enabled. By default, IP routing is enabled on Cisco routers; however, if IP routing is disabled, you should enable it by issuing the ip routing command in global configuration mode.

The CEF Forwarding Information Base (FIB) is built from information contained in the IP routing table. When the routing table is updated, the nexthop information in the FIB is also updated. A routing protocol, such as Border Gateway Protocol (BGP) or Intermediate System to Intermediate System (ISIS), can be used to populate the routing table, and, therefore, the FIB. However, neither BGP nor ISIS is required by MPLS; another internal gateway protocol (IGP), such as Open Shortest Path First (OSPF) or Enhanced Interior Gateway Routing Protocol (EIGRP), can be used instead.

A label exchange protocol, such as Tag Distribution Protocol (TDP), is used by MPLS to exchange label information. However, TDP itself is not required by MPLS; another label exchange protocol, such as Label Distribution Protocol (LDP) or Resource Reservation Protocol (RSVP), can be used instead.

Cisco Discovery Protocol (CDP) is used to collect information about neighboring Cisco devices, such as the host name, network address, port information, device type, and IOS version. CDP is enabled by default on Cisco devices. However, CDP is not required by MPLS.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/12_4t2/ftldp41.html#wp1632254

QUESTION 40

Which of the following actions cannot be performed in VRF configuration mode? (Select the best answer.)

- A. associating an SNMP context with the VRF
- B. assigning an RD
- C. configuring shared route targets between IPv4 and IPv6
- D. defining the VRF instance name
- E. updating a VPN ID



Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You cannot define the VPN routing and forwarding (VRF) instance name in VRF configuration mode. A VRF instance name is defined by issuing the vrf definition command in global configuration mode. The syntax of the vrf definition command is vrf definition vrfname, where vrfname is any name that you want to assign to the VRF, except default. After you issue the vrf definition vrfname command, the router will be placed into VRF configuration mode, where you can issue commands to configure other features of the VRF.

If you issue the vrf definition default command on a Cisco router, a VRF instance will be created with the name configured to a NULL value. The NULL value acts as a VRF name placeholder until a default VRF name can be defined.

You can assign a route distinguisher (RD) in VRF configuration mode. To assign an RD to a VRF, issue the rd routedistinguisher command. Configuring an RD creates the routing tables and the forwarding tables for the VRF instance.

You can configure shared route targets between IPv4 and IPv6 in VRF configuration mode. To configure a shared route target, issue the routetarget [import | export | both] routetargettextcommunity command. Specifying the import keyword imports routing information from the community specified by the

routetargettextcommunity parameter. Conversely, the export keyword sends routing information to the specified community. The both keyword imports routing information to and exports routing information from the specified community. To configure separate route target policies for IPv4 and IPv6, you should first issue the addressfamily command, which places the router into address family configuration mode.

You can associate a Simple Network Management Protocol (SNMP) context in VRF configuration mode. To associate an SNMP context, issue the contextcontextnamecommand, where contextname is the name of the SNMP context you want to assign.

You can update a virtual private network (VPN) ID in VRF configuration mode. The VPN ID consists of a 3byte Organizationally Unique Identifier (OUI) and a 4byte VPN index. To update a VPN ID for a VRF, issue the vpn id oui: vpn-index command.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/ipv6/command/reference/ipv6_book/ipv6_17.html#wp2449861

QUESTION 41

DRAG DROP

Select the features on the left, and place them in the appropriate category on the right.

Select and Place:



- CDP
- dot1ad
- EoMPLS
- LACP
- Layer 2 protocol tunneling
- LLDP
- MAC address security
- MST
- PAgP
- Q-in-Q tagging
- split horizon
- UDLD

Required for EVCs	Supported on EVCs	Unsupported on EVCs

Help

Reset

Done

Correct Answer:

The following features are supported on EVCs:

- EtherChannel, including Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol(LACP)
- UniDirectional Link Detection (UDLD)
- Link Layer Discovery Protocol (LLDP)
- Cisco Discovery Protocol (CDP)
- Media Access Control (MAC) address security

The following features are not supported on EVCs:

- Layer 2 multicast frame flooding
- Layer 2 protocol tunneling
- QinQ tagging
- Virtual LAN (VLAN) translation
- Ethernet over Multiprotocol Label Switching (EoMPLS)
- Split horizon
- Bridge domain routing

Reference: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SY/configuration/guide/sy_swcg/ethernet_virtual_connection.html#54887

QUESTION 42

You want to create a named ACL to use in a route map that allows redistribution of the following subnets:

- 192.168.3.0/24
- 192.168.4.0/24
- 192.168.5.0/24
- 192.168.6.0/24
- 192.168.7.0/24
- 192.168.8.0/24
- 192.168.9.0/24

Which of the following commands should you issue in order to fulfill your objective? (Select the best answer.)

- A. permit 192.168.3.0 0.0.7.255
- B. permit 192.168.0.0 0.0.15.255
- C. permit 192.168.0.0 255.255.240.0
- D. permit 192.168.3.0 255.255.248.0

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the permit 192.168.0.0 0.0.15.255 command. The basic syntax of the permit command is permit source wildcard-mask. A 20bit subnet mask, which corresponds to the wildcard mask 0.0.15.255, will aggregate 16 contiguous 24bit subnets. Therefore, a 20bit mask can be used to allow redistribution of addresses from 192.168.0.0 through 192.168.15.255, which includes all of the subnets in this scenario.

You should not issue the permit 192.168.3.0 0.0.7.255 command. A 21bit subnet mask, which corresponds to the wildcard mask 0.0.7.255, would aggregate eight contiguous 24bit subnets. Although there are only seven subnets in this scenario, the 21bit mask boundary falls between the 192.168.7.0/24 subnet and the 192.168.8.0/24 subnet. Therefore, the permit 192.168.3.0 0.0.7.255 command would allow redistribution of only the addresses from 192.168.0.0/24 through 192.168.7.255/24.

You should not issue the permit 192.168.0.0 255.255.240.0 command or the permit 192.168.3.0 255.255.248.0 command. The permit command accepts wildcard masks, not subnet masks.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/m1/sec-m1-cr-book/sec-cr-p1.html#wp1501248131>

QUESTION 43

You issue the show running-config command on a router and receive the following partial output:

```
router eigrp boson
```

```
addressfamily ipv4 autonomoussystem 1 afinterface default bandwidthpercent 75 afinterface Ethernet 0/0 hellointerval  
15 afinterface Ethernet 0/1 holdtime 45
```

What is the hold time for the Ethernet 0/0 interface? (Select the best answer.)

- A. five seconds
- B. 15 seconds
- C. 45 seconds
- D. 60 seconds
- E. 180 seconds

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The hold time for the Ethernet 0/0 interface is 15 seconds. Settings for Enhanced Interior Gateway Routing Protocol (EIGRP) named mode are configured under each address family. The afinterface default command enables you to configure settings that apply to all EIGRP interfaces unless a conflicting setting is explicitly

configured on the interface. For example, the bandwidthpercent 75 command has been issued under the afinterface default command; therefore, the bandwidthpercent 75 command would apply to all interfaces that are not explicitly configured with the bandwidthpercent command.

In this scenario, the holdtime command has not been configured under the afinterface default command or under the afinterface Ethernet 0/0 command? therefore, the default EIGRP hold time for Ethernet interfaces will be used, which is 15 seconds.

If the holdtime 5 command had been issued under the afinterface default command, the Ethernet 0/0 interface would have had a hold time of five seconds.

However, that setting would not have applied to the Ethernet 0/1 interface, because the holdtime 45 command has been issued under the afinterface Ethernet 0/1 command.

The hello interval for the Ethernet 0/0 interface is 15 seconds. Even if a different hellointerval command had been issued under the afinterface default command, the

Ethernet 0/0 interface would still be configured with a hello interval of 15 seconds because the hellointerval 15 command has been explicitly issued on the Ethernet 0/0 interface. By default, the hello interval is five seconds for Ethernet interfaces; therefore, the Ethernet 0/1 interface would have a hello interval of five seconds because the hellointervalcommand has not been issued on that interface.

The hold time for the Ethernet 0/0 interface is not 60 seconds. The default hello interval for lowspeed nonbroadcast multiaccess (NBMA) interfaces is 60 seconds. The hold time for the Ethernet 0/0 interface is not 180 seconds. The default hold time for lowspeed NBMA interfaces is 180 seconds. Reference:

Cisco: Cisco IOS IP Routing: EIGRP Command Reference: afinterface

Cisco: Cisco IOS IP Routing: EIGRP Command Reference: hellointerval

Cisco: Cisco IOS IP Routing: EIGRP Command Reference: holdtime

CCIE Routing and Switching v5.0 Certification Guide, Volume 1, Chapter 8, EIGRP Named Mode, pp. 410417

QUESTION 44

When you are enabling AutoQoS for the Enterprise on a router, which of the following commands should you ensure has been issued before you issue the auto discovery qoscommand? (Select the best answer.)

- A. mls qos
- B. auto qos voip
- C. ip cef
- D. auto qos

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When you are enabling AutoQoS for the Enterprise on a router, you should ensure the ip cef command has been issued before you issue the auto discovery qos command. Cisco Express Forwarding (CEF) must be enabled before you can implement AutoQoS for the Enterprise. AutoQoS uses Network Based Application Recognition (NBAR) for packet classification. If you find that CEF has not been enabled, you should enable it by issuing the ip cef command. After ensuring that

CEF is enabled, you can issue the auto discovery qos interface configuration command, which will initiate the autodiscovery phase of the AutoQoS for the Enterprise implementation process. The autodiscovery phase profiles the traffic on the network to determine the volume and type of traffic being sent on the network. By default, autodiscovery will run for three days to determine as accurately as possible the volume and type of traffic sent on the network. However, you can configure autodiscovery to take more or less time, depending on the needs of the network.

You should issue the auto qos command after you issue the auto discovery qoscommand on an interface. After the auto discovery qos command has profiled the traffic, the auto qos command uses that data to generate Quality of Service (QoS) templates, which are then used to create class maps and policy maps. You can use the auto qos voip command if you are configuring AutoQoS Voice over IP (VoIP) rather than AutoQoS for the Enterprise.

The mls qos command enables QoS. This command is automatically enabled when you issue the auto qos command or the auto qos voip command, so it is not necessary to issue the mls qos command when you are configuring AutoQoS.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_auto/configuration/15-mt/qos-auto-15-mt-book/qos-auto-ent.html#GUID-72FB0F57-E978-4DD2-A314-9E0C677768FA

https://www.cisco.com/en/US/technologies/tk543/tk879/technologies_qas0900aecd8020a589.html

QUESTION 45

Which of the following TLVs are specific to LLDPMED? (Select 2 choices.)

- A. location
- B. management address
- C. port description
- D. power management
- E. system capabilities
- F. system description
- G. system name



Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The location and power management type, length, and value (TLV) descriptions are specific to Link Layer Discovery Protocol for Media Endpoint Devices (LLDPMED). LLDP-MED is an extension of Link Layer Discovery Protocol (LLDP). LLDP is a Layer 2 openstandard discovery protocol that is used to facilitate interoperability between Cisco devices and nonCisco devices. LLDP-MED operates between endpoint devices, such as a PC or a Voice over IP (VoIP) phone, and vendorneutral network devices. By contrast, LLDP does not operate between endpoint devices and network devices; LLDP operates only between network devices, such as routers, switches, and access servers.

Attributes that can be learned from neighboring devices are contained within TLVs. The following TLVs are supported by LLDP:

- Port description
- System name
- System description
- System capabilities
- Management address

In addition, the following LLDP TLVs are advertised to support LLDP-MED:

- Port VLAN ID
- MAC/PHY configuration status

The following TLVs are supported by LLDP-MED:

- LLDP-MED capabilities
- Network policy
- Power management
- Inventory management
- Location

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3560/software/release/12-2_55_se/configuration/guide/3560_scg/swlldp.html



QUESTION 46

DRAG DROP

Select the default timer settings from the left, and drag them to the corresponding RIP timers on the right. Some values might be used more than once.

Select and Place:

Holddown timers are used by RIP to specify the amount of time to suppress information regarding a better path to a route. When a router receives a routing update stating that a route is unreachable, the router waits a specified amount of time before accepting routes advertised by other sources. By default, the holddown timer is set to 180 seconds.

The flush timer is used to specify the amount of time to wait before deleting a route from the routing table. By default, the flush timer is set to 240 seconds, and it should always be set to a value greater than the invalid timer.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/iproute/command/reference/fiprrp_r/1rfrip.html#wp1018019

QUESTION 47

Which of the following is an IETF-standard FHRP that can use object tracking and preemption to provide Layer 3 failover? (Select the best answer.)

- A. GLBP
- B. HSRP
- C. LACP
- D. VRRP

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Virtual Router Redundancy Protocol (VRRP) is an Internet Engineering Task Force (IETF) standard FirstHop Redundancy Protocol (FHRP) that can use object tracking and preemption to provide Layer 3 failover. FHRPs are protocols that are used to provide Layer 3 gateway redundancy, such as failover and load balancing. Providing Layer 3 redundancy ensures that hosts on a LAN will have a backup path to external networks should a primary path fail or become too congested to forward traffic. Layer 3 devices in an FHRP configuration typically share a virtual IP address that is then configured as the default gateway on each host for which the device is to forward traffic. The FHRP devices might also share a virtual Media Access Control (MAC) address or multiple virtual MAC addresses, depending on the protocol. FHRPs typically use a priority system to elect a primary Layer 3 forwarding device, which is known as an active virtual gateway (AVG), an active router, or a master router, depending on the protocol. The same priority system elects either a single or multiple backup-forwarding devices.

VRRP can be configured to use object tracking to influence the priority of a router in a group and therefore force the election of a different master router when certain conditions are met. When combined with VRRP preemption, which enables a VRRP router to automatically assume the master router role when priority values change, object tracking enables VRRP to adjust the priority of a router based on the line protocol status of a specific interface or the availability of a given route to a destination. For example, if RouterA and RouterB in a VRRP configuration had different paths to the Internet, VRRP could be configured to monitor RouterA's outbound interface and to automatically set RouterA's VRRP priority to a value lower than RouterB's if RouterA's outbound interface were to go down. RouterB would then become the master router, and Layer 3 traffic would be forwarded through its outbound interface instead of through RouterA's.

Gateway Load Balancing Protocol (GLBP) is a Cisco proprietary FHRP, not an IETF standard. However, GLBP does support object tracking and preemption. GLBP is different from both VRRP and Hot Standby

Router Protocol (HSRP) in that it is, by default, capable of load balancing traffic between all routers in a GLBP group. VRRP and HSRP are primarily failover protocols. GLBP elects an AVG and up to four primary active virtual forwarders (AVFs). The routers in a GLBP group receive traffic sent to a virtual IP address that is configured for the group. Each GLBP group contains an AVG that is elected based on which router is configured with the highest priority value or the highest IP address value if multiple routers are configured with the highest priority value. The other routers in the GLBP group are configured as primary or secondary AVFs. The AVG in a GLBP group assigns a virtual MAC address to up to four primary AVFs? all other routers in the group are considered secondary AVFs and are placed in the listen state. When the AVG receives ARP requests that are sent to the virtual IP address for the GLBP group, the AVG responds with different virtual MAC addresses. This provides load balancing, because each of the primary AVFs will participate by forwarding a portion of the traffic sent to the virtual IP address. HSRP is a Cisco proprietary FHRP, not an IETF standard. However, like VRRP, HSRP is capable of using object tracking and preemption to modify the priority of the HSRP active router and force the standby router to take over if a specific interface goes down or a path to the destination becomes unavailable.

Link Aggregation Control Protocol (LACP) is not an FHRP. LACP is an Institute of Electrical and Electronics Engineers (IEEE) protocol that is used to enable link aggregation on EtherChannel links.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp_fhrp/configuration/12-4/fhp-12-4-book/fhp-vrrp.html#GUID-ECF180A0-1633-4DB2-AD31-9D807B5833AD

QUESTION 48

You manage a StackWise stack of nine switches.

Which of the following switches will be elected as the stack master? (Select the best answer.)

- A. the switch with stack ID 1
- B. the switch with stack ID 9
- C. the switch with the lowest MAC address
- D. the switch with the highest MAC address
- E. the switch with the lowest configured IP address
- F. the switch with the highest configured IP address
- G. the switch with the lowest priority value
- H. the switch with the highest priority value

Correct Answer: H

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The switch with the highest priority value will be elected as the stack master. StackWise is a Cisco proprietary technology that is used to provide Layer 2 or Layer 3 connectivity between switches so that the stack of switches acts as a single device. When a StackWise configuration is used, the failure of a single switch will not result in an outage. Instead, the other switches in the stack will compensate for the failed switch. The switches are connected sequentially by stack cables: the first switch is connected to the second, the second switch is connected to the third, and so on until the last switch is connected to the first. If a stack cable is broken, the bandwidth of the stack will be reduced by 50 percent until the cable is fixed.

The stack master controls the operation of the stack. From the stack master, you can configure global features that apply to all switches in the stack as well as interface-level features for individual stack members. The stack priority is a value from 1 through 15; by default, the priority value is set to 1. To change the stack priority, you should issue the `switch stackid priority value` command from global configuration mode. The following checklist is used to elect a stack master:

- The current stack master is elected as the stack master. Otherwise, the switch with the highest priority is elected stack master.
- If multiple switches have the same priority, the switch with a nondefault saved interface-level configuration is elected stack master.
- If multiple switches have a nondefault saved interface-level configuration, the switch with the highest feature set priority is elected stack master, based on the following hierarchy:
 - IP services with cryptographic image
 - IP services with noncryptographic image
 - IP base with cryptographic image
 - IP base with noncryptographic image
- If multiple switches have the same feature set, the switch with the shortest startup time is elected stack master.
- If multiple switches are still eligible, the stack member with the lowest Media Access Control (MAC) address is elected stack master.

The stack ID is not used to elect a stack master. Each stack member has a unique stack ID. By default, all switches use stack ID 1. However, if two switches attempt to take the same stack ID, the switch with the higher priority will retain the stack ID number and the other switch will automatically be assigned a new stack ID. Therefore, you need not make any configuration changes before adding a switch to a StackWise stack. However, you can manually configure a stack ID by issuing the `switch current-stack-id renumber new-stack-id` command from global configuration mode. The change will not take effect until the switch is reloaded.

IP addresses are not used to elect a stack master. In fact, a switch need not be configured with any IP addresses to become a stack member or stack master.

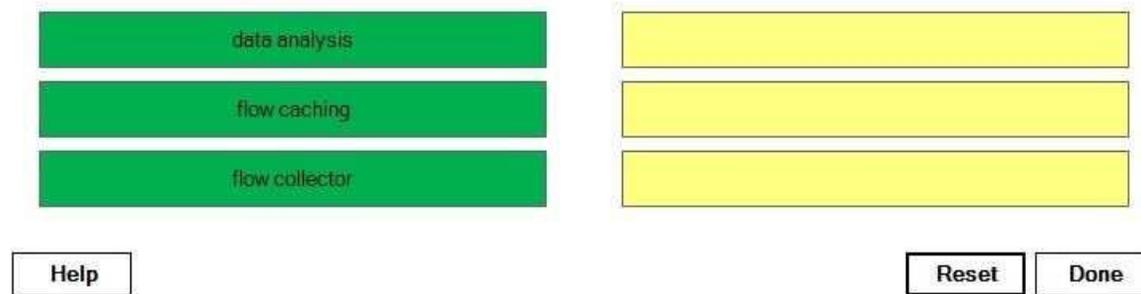
Reference: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750x_3560x/software/release/12-2_55_se/configuration/guide/3750xscg/swstack.html#pgfId-1228109

QUESTION 49

DRAG DROP

Select the NetFlow components from the left, and place them on the right in the order in which they are used by NetFlow.

Select and Place:



Correct Answer:



Section: (none)
Explanation

Explanation/Reference:

Explanation:

NetFlow is a Cisco IOS feature that can be used to monitor traffic flows. A traffic flow is defined as a series of packets with the same source IP address, destination IP address, protocol, and Layer 4 information. NetFlow gathers flowbased statistics such as packet counts, byte counts, and protocol distribution. The data gathered by NetFlow is typically exported to management software. You can then analyze the data to facilitate network planning, customer billing, and traffic engineering. Flow caching collects IP data flow information and prepares the information for export. A traffic flow can be identified based on the combination of the following attributes:

- Source IP address
- Destination IP address
- Source port number
- Destination port number
- Protocol value

- Type of Service (ToS) value
- Input interface

A flow collector collects the exported data from multiple devices so that it can be aggregated and stored for analysis by a data analyzer. NetFlow can be used to perform all three functions, or it can export the data to a thirdparty product that can read the data that is stored within FlowCollector files.

Reference:

https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/ios-netflow/product_data_sheet0900aecd80173f71.html

QUESTION 50

You issue the following commands on the FastEthernet 0/1 interface of SwitchA:

```
SwitchA(config)#switchport portsecurity
SwitchA(config)#switchport portsecurity maximum 3
SwitchA(config)#switchport portsecurity macaddress sticky
SwitchA(config)#switchport portsecurity violation restrict
```

Which of the following are true? (Select 2 choices.)



<https://vceplus.com/>

- A. Up to three MAC addresses will be stored in the running configuration.
- B. Up to three MAC addresses will be stored in the address table but not in the running configuration.
- C. The switch will silently discard the traffic when a security violation occurs.
- D. The switch will discard the traffic, log the unauthorized entry attempt, increment the SecurityViolationcounter, and send an SNMP trap message when a security violation occurs.
- E. The switch will discard the traffic, log the unauthorized entry attempt, increment the SecurityViolation counter, and place the port into the error-disabled state when a security violation occurs.

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Up to three Media Access Control (MAC) addresses will be stored in the running configuration. In this scenario, the switchport portsecurity maximum 3 command specifies that three MAC addresses are authorized to send traffic on port FastEthernet 0/1. MAC addresses can be configured statically or learned dynamically by port security. Dynamically learned MAC addresses are converted to sticky addresses and stored in the running configuration when the switchport portsecurity macaddress stickycommand is issued on a port. Any MAC addresses that are not configured statically will be learned dynamically from incoming traffic, up to the maximum number of MAC addresses allowed to communicate on the port.

Because no MAC addresses have been statically configured in this scenario, all three MAC addresses will be learned dynamically. If the switchport portsecurity macaddress sticky command had not been issued, the switch would retain dynamically learned MAC addresses in the MAC address table but not in the running configuration.

Additionally, the switch will discard the traffic, log the unauthorized entry attempt, increment the SecurityViolation counter, and send a Simple Network Management Protocol (SNMP) trap message when a security violation occurs in this scenario. You can configure a switch to perform the following actions when a switch port with port security enabled receives traffic from a host with an unauthorized MAC address:

- Protect: The switch will discard the traffic.
- Restrict: The switch will discard the traffic, log the unauthorized entry attempt, increment the SecurityViolation counter, and send an SNMP trap message.
- Shutdown: The switch will discard the traffic, log the unauthorized entry attempt, increment the SecurityViolation counter, and place the port into the error-disabled state.

To configure the action that a switch will perform when unauthorized traffic is received on a switch port, you should issue the switchport portsecurity violation {protect | restrict | shutdown} command in interface configuration mode. By default, a switch port with port security enabled will be configured for shutdown mode. For example, the following commands would configure port security on SwitchA to use the default violation behavior:

```
SwitchA(config-if)#switchport port-security
SwitchA(config-if)#switchport port-security maximum 3
SwitchA(config-if)#switchport port-security mac-address sticky
```

Because no switchport portsecurity violation command is issued in the output above, the switch will discard the traffic, log the unauthorized entry attempt, increment the SecurityViolation counter, and place the port into the error-disabled state when an unauthorized MAC address attempts to use the port.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960/software/release/12-2_25_see/configuration/guide/scg_1/swtrafc.html#wp1038501

QUESTION 51

You issue the show framerelay map command on Router2 and receive the following output:

```
Serial2/0 (up): ip 10.11.12.13 dlcI 20(0x14,0x440), dynamic,
                CISCO, status defined, active
```

What protocol was used to dynamically create this PVC? (Select the best answer.)

- A. DHCP
- B. ARP
- C. Inverse ARP
- D. RARP

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Inverse Address Resolution Protocol (ARP) is used to dynamically create permanent virtual circuits (PVCs). A PVC is a virtual connection between a source and a destination through which data is transmitted as if over a physical connection. The destination becomes the next hop from the source over the PVC. PVCs that are created dynamically by Inverse ARP are marked as dynamic in the output of the show frame-relay map command. The show frame-relay map command can be used to verify the local datalink connection identifier (DLCI) numbers that have been assigned to remote IP addresses, the status of the PVC, the encapsulation format that is used by the PVC, and whether the PVC was manually configured or created dynamically.

Inverse ARP is used to find a Layer 3 address when the Layer 2 address is known. In the case of Frame Relay, Inverse ARP maps Layer 2 DLCIs to Layer 3 IP addresses. DLCIs uniquely identify a PVC connection in a Frame Relay circuit.

ARP is not used to dynamically create PVCs. ARP is used to find a Layer 2 address when the Layer 3 address is known. Because Inverse ARP is an extension of ARP, ARP packets are structured the same as Inverse ARP packets.

Reverse ARP (RARP) is not used to dynamically create PVCs. A device uses RARP to obtain an IP address for itself based on its Layer 2 Media Access Control (MAC) address. RARP has the same packet structure as ARP and Inverse ARP. However, RARP has been largely replaced with Dynamic Host Configuration Protocol (DHCP).

DHCP is not used to dynamically create PVCs. DHCP dynamically assigns network configuration information to client computers. This network configuration information can include the IP address, subnet mask, default gateway, and Domain Name System (DNS) servers that the client computer will use.

Reference:

<https://www.cisco.com/en/US/docs/internetworking/troubleshooting/guide/tr1918.html#wp1020791>

https://www.cisco.com/c/en/us/td/docs/ios/12_2/wan/configuration/guide/fwan_c/wcffrely.html#wp1001012

QUESTION 52

Which of the following best defines a FEC? (Select the best answer.)

- A. a group of packets that are forwarded similarly
- B. a table that is built from information in the routing table
- C. a table that contains inboundtooutbound label mappings
- D. a value that enables MPLS VPN customers to use overlapping IP address ranges

Correct Answer: A

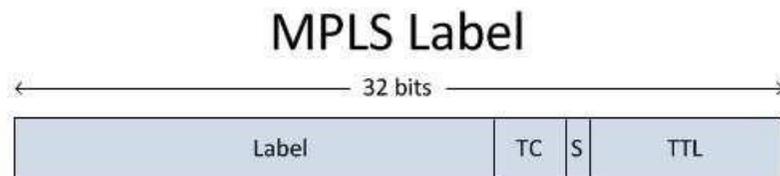
Section: (none)

Explanation

Explanation/Reference:

Explanation:

A Forwarding Equivalence Class (FEC) is a group of packets that are forwarded similarly. A FEC is generally associated with a destination IP network, although it can also be associated with a Layer 2 circuit or an IP precedence value. The Label field of a Multiprotocol Label Switching (MPLS) label is a 20bit field that is used to represent the FEC. The structure of an MPLS label is shown below:



The Forwarding Information Base (FIB) is a table that is built from information in the routing table. When a label switch router (LSR) receives an unlabeled packet destined for an MPLS-enabled interface, it consults the FIB, adds the appropriate label for the destination address, and forwards the packet. However, if the packet's destination address is not contained in the FIB, the packet is dropped.

The Label Forwarding Information Base (LFIB) is a table that contains inbound-to-outbound label mappings. If a route becomes unavailable, the LFIB information will be modified based on information in the Label Information Base (LIB) and FIB tables; the LIB contains all of the labels received from neighboring LSRs. When an LSR receives a labeled packet, it consults the LFIB, swaps or removes the label, and forwards the packet. However, if the label mapping is not contained in the LFIB, the packet is dropped.

A route distinguisher (RD) is a value that enables MPLS virtual private network (VPN) customers to use overlapping IP address ranges; MPLS VPNs are described in Request for Comments (RFC) 4364. An ingress LSR creates a globally unique VPN version 4 (VPNv4) address by adding the RD to the beginning of an IP address. The LSR then assigns a label to the VPNv4 address prefix and stores the inbound-to-outbound label mapping in the LFIB. Authentication to the MPLS VPN is provided based on logical port and RD information.

Reference:

<https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/4649-mpls-faq-4649.html#anc4>

QUESTION 53

You issue the show ip nhrp detail command on RouterA and receive the following output:

```
RouterA#show ip nhrp detail
10.10.10.1/32 via 10.10.10.1, Tunnel0 created 00:44:05, never expire
  Type: static, Flags: used
  NBMA address: 192.168.51.50
10.10.10.5/32 via 10.10.10.5, Tunnel0 created 00:05:40, expire 00:00:41
  Type: dynamic, Flags: authoritative unique nat registered used
  NBMA address: 172.16.0.44
```

Which of the following flags indicates that the entry cannot be overwritten by a different NBMA entry with the same IP address? (Select the best answer.)

- A. authoritative
- B. nat
- C. registered
- D. unique
- E. used

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Of the available choices, the unique flag in the output of the show ip nhrp detail command indicates that the entry cannot be overwritten by another nonbroadcast multiaccess (NBMA) entry with the same IP address. Similar to Address Resolution Protocol (ARP) on broadcast networks, Next Hop Resolution Protocol (NHRP) enables devices on an NBMA network to dynamically discover the physical addresses of other devices on the network. The show ip nhrp command displays information about NHRP mappings. When issued with the detail keyword, this command displays more detailed information about those mappings, including a list of flags.

The show ip nhrp detail command might display any of the following flags for a given entry:

- authoritative -The mapping was obtained directly from the nexthop router or server.
- implicit -The mapping was obtained from an NHRP resolution request or packet.
- local -The mapping is for networks that are local to the router.
- nat -The remote device supports NHRP Network Address Translation (NAT) extensions.
- negative -A mapping could not be obtained for negative caching.
- (no socket) -IP Security (IPSec) will not set up encryption, because data traffic does not require this tunnel.
- registered -The mapping was created in response to an NHRP registration.
- router -The mappings for a remote router are marked with the router flag.
- unique -The mapping cannot be overwritten by a different NBMA entry with the same IP address.

-used-Data packets are being processswitched for the given mapping.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/command/ipaddr-cr-book/ipaddr-s1.html#wp2302625547>

QUESTION 54

DRAG DROP

Select the capabilities on the left, and place them underneath the corresponding protocols on the right. Fill all boxes. Some capabilities will be used more than once.

Select and Place:

VTP management	LLDP-MED	CDP
MTU size		
native VLAN		
port speed		
port duplex		
third-party devices		
topology change notifications		

Correct Answer:

	LLDP-MED	CDP
VTP management	port speed	VTP management
MTU size	port duplex	MTU size
native VLAN	third-party devices	native VLAN
port speed	topology change notifications	port duplex
port duplex		
third-party devices		
topology change notifications		

 **VCEplus**
VCE To PDF - Free Practice Exam

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Link Layer Discovery Protocol (LLDP) is a Layer 2 openstandard discovery protocol that is used to facilitate interoperability between Cisco devices and thirdparty devices. LLDP for Media Endpoint Devices (LLDPMED) is an extension of LLDP that operates between endpoint devices, such as a PC or a Voice over IP (VoIP) phone, and vendorneutral network devices. By contrast, LLDP does not operate between endpoint devices and network devices; LLDP operates only between network devices, such as routers, switches, and access servers.

Cisco Discovery Protocol (CDP) is a Layer 2 Cisco proprietary discovery protocol that is used to collect information about neighboring Cisco devices. However, CDP cannot be used to collect information about thirdparty devices.

Both LLDPMED and CDP can determine whether a neighboring port is running full or half duplex, but only LLDPMED can also determine the speed capabilities of a port. In addition, LLDPMED supports topology change notifications, whereas CDP does not. CDP supports VLAN Trunking Protocol (VTP) management, native virtual LAN (VLAN) detection, and maximum transmission unit (MTU) detection, whereas LLDPMED does not.

Reference:

https://www.cisco.com/en/US/technologies/tk652/tk701/technologies_white_paper0900aecd804cd46d.html

QUESTION 55

Which of the following commands should you issue to increase the number of protocols that NBAR can classify and inspect? (Select the best answer.)

- A. ip nbar pdlm
- B. ip nbar port-map
- C. ip nbar protocol-discovery
- D. ip nbar resources

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the ip nbar pdlm command to increase the number of protocols that Network Based Application Recognition (NBAR) can classify and inspect. NBAR enables a router to perform deep packet inspection for all packets that pass through an NBAR-enabled interface. Although NBAR supports several common applications and protocols, you can update or expand the base protocol support by installing Packet Description Language Modules (PDLMs). Cisco provides many PDLMs for download on its support website. PDLMs are stored in Flash memory.

Issuing the ip nbar portmap command does not increase the number of protocols that NBAR can classify and inspect; it modifies the mapping between NBAR-recognized applications and their associated ports. NBAR supports a limited number of protocols and applications based on their well-known port numbers. However, if an application or protocol has been configured to use nonstandard port numbers, you can issue the ip nbar portmap command to modify the NBAR configuration accordingly. For example, if Secure Shell (SSH) servers on the network are configured to listen on ports 22 and 2222, you should issue the ip nbar portmap ssh tcp 22 2222 command to configure NBAR to search for SSH on those ports.

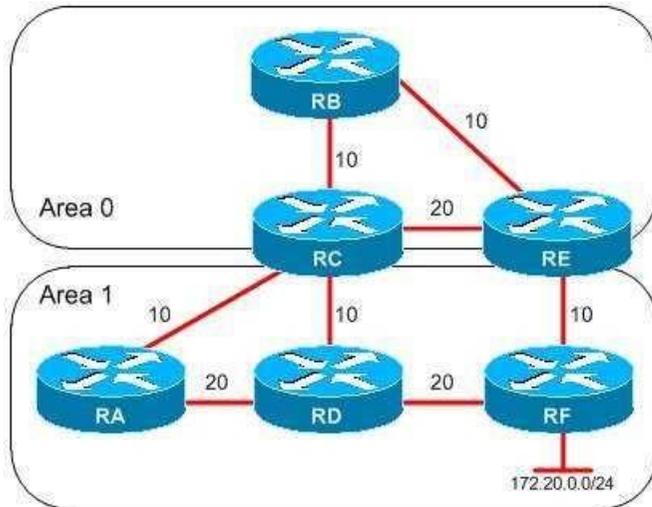
Issuing the ip nbar protocol-discovery command does not increase the number of protocols that NBAR can classify and inspect; it records traffic statistics on an interface based on packet content. After NBAR has been enabled on an interface, you can issue the service-policy input command to configure NBAR to classify inbound traffic or you can issue the service-policy output command to configure NBAR to classify outbound traffic.

Issuing the ip nbar resources command does not increase the number of protocols that NBAR can classify and inspect; it tunes NBAR's memory usage. You can issue the ip nbar resources command to adjust how quickly state information expires and how much system memory is available to NBAR-recognized applications.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/qos/command/reference/qos_book/qos_i1.html#wp1022981

QUESTION 56



You administer the OSPF network shown above. The cost values are displayed next to each link.

RA receives packets destined for the 172.20.0.0/24 network.

How many paths will RA use to send the packets? (Select the best answer.)

- A. one
- B. two
- C. three
- D. four

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RA will use two paths to send the packets: the path from RA through RD to RF and the path from RA through RC and RD to RF. The total cost from RA through RD to RF is $20 + 20 = 40$, and the total cost from RA through RC and RD to RF is $10 + 10 + 20 = 40$. Open Shortest Path First (OSPF) can load balance traffic across equal-cost paths; since both paths have a total cost of 40, RA can use both paths to send the packets. RA will not always prefer the route through the least number of routers. Instead, RA prefers the intra-area route with the lowest total cost, regardless of the number of routers the packets must pass through.

RA will not use the two paths through Area 0 to send the packets. Although the total cost from RA through both paths is 40, OSPF prefers intra-area routes over inter-area routes, regardless of the total path cost. OSPF uses the following preference order when selecting the best route to a destination:

1. Intra-area routes
2. Inter-area routes
3. External Type 1 routes
4. External Type 2 routes

Therefore, RA prefers an intra-area route with a cost of 40 over an inter-area route with a cost of 40. If all of the routers were within the same area, RA would use all four paths to send the packets.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html#7>

QUESTION 57

Which of the following IPv6 addresses is used as the source address in OSPFv3 packets? (Select the best answer.)

- A. A link-local unicast address is used on all OSPF interfaces.
- B. A link-local unicast address is used on all OSPF interfaces except virtual links.
- C. A global scope address is used on all OSPF interfaces.
- D. A global scope address is used on all OSPF interfaces except virtual links.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A link-local unicast address is used as the source address in Open Shortest Path First version 3 (OSPFv3) packets on all OSPF interfaces except virtual links. Linklocal unicast addresses use the FE80::/10 address range; these addresses begin with the hexadecimal characters FE80 through FEBF. Because link-local addresses are unique only on the local segment, link-local addresses are not routable. An IPv6capable host typically creates a link-local unicast address automatically at startup.

OSPFv3 virtual link interfaces must use a global scope IPv6 address as the source address for OSPFv3 packets. If a router has one or more virtual links configured, it includes the global scope IPv6 address in the LSA, sets the LA-bit in the PrefixOptions field, configures the PrefixLength field to a value of 128, and sets the Metric field to a value of 0. Global aggregatable unicast addresses use the 2000::/3 address range; these addresses begin with the hexadecimal characters 2000 through 3FFF. Global aggregatable unicast address prefixes are distributed by the Internet Assigned Numbers Authority (IANA) and are globally routable over the Internet.

Reference:

<https://tools.ietf.org/html/rfc5340#section-2.5>

QUESTION 58

Which of the following commands will cause iBGP routes to have the same AD as internal EIGRP routes? (Select 2 choices.)

- A. distance bgp 90 110 120
- B. distance bgp 90 120 110
- C. distance bgp 110 90 120
- D. distance bgp 110 120 90
- E. distance bgp 120 90 110
- F. distance bgp 120 110 90

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The distance bgp 110 90 120 command and the distance bgp 120 90 110 command will cause internal Border Gateway Protocol (iBGP) routes to have the same administrative distance (AD) as internal Enhanced Interior Gateway Routing Protocol (EIGRP) routes. When multiple routes to the same destination network exist and each route uses a different routing protocol, a router prefers the routing protocol with the lowest AD. The following list contains the most commonly used ADs:

Route Source	Distance
Connected route	0
Static route	1
EIGRP summary route	5
eBGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
iBGP	200
Unknown	255

ADs for a routing protocol can be manually configured by issuing the distance command in router configuration mode. The syntax of the distance bgp command is distance bgp external-distance internaldistance localdistance. The externaldistance value configures the AD for external BGP (eBGP) routes. The internaldistance value configures the AD for iBGP routes. The localdistance value configures the AD for local BGP routes. For example, the distance bgp 110 90 120 command configures an AD of 110 for eBGP routes, an AD of 90 for iBGP routes, and an AD of 120 for local BGP routes. Internal EIGRP routes have an AD of 90; therefore, the distance bgp 110 90 120 command configures iBGP routes to have the same AD as internal EIGRP routes.

The distance bgp 90 110 120 command and the distance bgp 120 110 90 command will cause iBGP routes to have the same AD as Open Shortest Path First (OSPF), not internal EIGRP. The distance 90 120 110 command and the distance bgp 110 120 90 command will cause iBGP routes to have the same AD as Routing Information Protocol (RIP), not internal EIGRP.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-c1.html#wp1296277485

<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/15986-admin-distance.html>

QUESTION 59

You issue the show adjacency command on RouterA and receive the following output:

```
RouterA#show adjacency
Protocol Interface Address
IP FastEthernet0/0 192.168.51.49 (7) (incomplete)
IP FastEthernet0/1 192.168.51.54 (7)
```

Which of the following is not a step in troubleshooting the incomplete marker in the output? (Select the best answer.)

- A. waiting 60 seconds and issuing the show adjacency command again
- B. issuing the debug arp command to verify that ARP requests are sent
- C. issuing the show ip arp command to verify the ARP table's contents
- D. issuing the show ip cef command to verify the contents of the FIB

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Issuing the show ip cef command to verify the contents of the Cisco Express Forwarding (CEF) Forwarding Information Base (FIB) is not a step in troubleshooting the incomplete marker when it appears in the output of the show adjacency command. The show ip cef command is used to display the contents of the FIB. The FIB, which is similar in concept to a routing table, is one of two databases used by CEF. The other database is the adjacency table, which is used to store nexthop information that is discovered by using Address Resolution Protocol (ARP) requests. The incomplete marker in this scenario appears in the contents of the adjacency table.

When an incomplete marker appears in the adjacency table, the cause is typically either a failed ARP request or a failure to clear a transient incomplete marker that resulted from issuing the clear ip arp command or the clear adjacency command. Therefore, you could wait 60 seconds and issue the show adjacency command again if you are troubleshooting the appearance of the incomplete marker. When the clear ip arp command or the clear adjacency command is issued, the nexthop entry will initially be marked as incomplete in the adjacency table. However, under normal circumstances, that marker should be cleared from the table after 60

seconds. In addition, you could choose to verify that ARP is sending requests by issuing the debug arp command or verify that the ARP table is populated with correct information by issuing the show ip arpcommand.

Reference:

Cisco: Cisco IOS IP Switching Command Reference: show ip cef

Cisco: Troubleshooting Incomplete Adjacencies with CEF: Reasons for Incomplete Adjacencies

QUESTION 60

Which of the following entries in the output of the show ip eigrp topology command indicates that a route is undergoing recomputation? (Select the best answer.)

- A. A 192.168.13.0/24, 1 successor, FD is Inaccessible, Q
- B. P 192.168.13.0/24, 1 successor, FD is Inaccessible, Q
- C. R 192.168.13.0/24, 1 successor, FD is Inaccessible, Q
- D. r 192.168.13.0/24, 1 successor, FD is Inaccessible, Q

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following show ip eigrp topology output indicates that a route is undergoing recomputation:

A 192.168.13.0/24, 1 successor, FD is Inaccessible, Q

The letter A indicates a route that is in the active state. A route is in the active state when a successor becomes unavailable and no feasible successor exists. When a route transitions from the passive state to the active state, the router will send multicast query packets to its neighbors to find an alternate route to the destination network. The route will remain in the active state until replies are received for each of the neighbor queries. When all queries have been received, the router will calculate the best route to the destination network. If a neighbor router does not respond before the active timer expires, the querying router will become stuck in active (SIA) and the neighbor router will be removed from the querying router's neighbor table. A router that is SIA because of missing replies will generate %DUAL3-SIA debug error messages.

The letter P indicates a route that is in the passive state. A route is in the passive state if it has connectivity to the successor, which is the best nexthop router to a destination network. If all of the routes in the topology table display a P, the network is stable and is not undergoing recomputation.

The letters R and r are not displayed at the beginning of routes in the topology table. The letter R is displayed after the IP address of a neighbor router that has responded to a query. If a neighbor router has not yet responded to the query, the letter r is displayed after the neighbor's IP address.

Reference:

Cisco: EIGRP Commands: show ip eigrp topology

Cisco: What Does the EIGRP DUAL3SIA Error Message Mean?

QUESTION 61

The MPLS TTL field of a packet is set to 0.

Which of the following statements is accurate? (Select the best answer.)

- A. The packet is discarded.
- B. MPLS TTL propagation is disabled.
- C. The MPLS label is the last label in the stack.
- D. The packet has a low priority.

Correct Answer: A

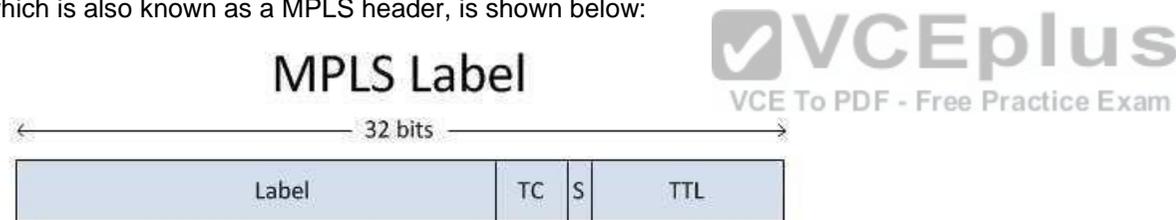
Section: (none)

Explanation

Explanation/Reference:

Explanation:

When the Multiprotocol Label Switching (MPLS) TimeToLive (TTL) field of a packet is set to 0, the packet is discarded. The structure of a typical 4byte MPLS label, which is also known as a MPLS header, is shown below:



The TTL field is an 8bit field in the MPLS label that is used to control the propagation of packets through an MPLS network. Thus the MPLS TTL field is similar to the TTL field in an IP header. When an IP packet enters an MPLS network, the ingress router decrements the IP TTL value by 1 and copies that value to the MPLS TTL field. Each MPLS router along the path decrements the MPLS TTL field by 1. When the packet reaches the egress router, the MPLS TTL value is decremented by 1 and copied to the IP TTL field.

If the MPLS TTL field of a packet is set to 0, MPLS TTL propagation is more likely to be enabled than disabled. When MPLS TTL propagation is disabled, the MPLS TTL field is set to 255 and decrements as the packet passes through the MPLS network. When the packet reaches the egress router, the MPLS TTL value is not copied to the IP TTL field. By default, MPLS TTL propagation is enabled, but you can disable it by issuing the no mpls ip propagate-ttl command.

The MPLS TTL field does not indicate whether an MPLS label is the last label in the stack. The BottomofStack field, sometimes called the S field or Stack bit, is a 1bit field that indicates whether the label is the last MPLS label in a packet. A BottomofStack field set to 0 indicates that one or more MPLS labels follow this label. A BottomofStack field set to 1 indicates that this label is the last label in the stack.

The MPLS TTL field does not indicate whether a packet has a low priority. Cisco routers use the 3bit Traffic Class (TC) field in the MPLS label to carry the IP precedence value, which is used to classify and prioritize network traffic. The TC field was formerly designated as the Experimental (EXP) field in Request for Comments (RFC) 3032. However, RFC 3032 did not officially designate the use of the EXP field, so some nonCisco routers use this field for other purposes. RFC 5462 officially renames the EXP field as the TC field and designates it to carry traffic class information, such as IP precedence values. Lowpriority traffic might be assigned an IP precedence value of 0, and highpriority traffic might be assigned an IP precedence value of 7.

Reference:

https://www.cisco.com/c/en/us/support/docs/multi-protocol-label-switching-mpls/mpls/26585-mpls-traceroute.html#no_mpls

QUESTION 62

Which of the following commands should you issue to manually attach a traffic policy to an interface in an NBAR configuration? (Select the best answer.)

- A. class-map
- B. policy-map
- C. service-policy
- D. ip nbar protocol-discovery
- E. auto qos

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the service-policy command to manually attach a traffic policy to an interface. Network Based Application Recognition (NBAR) is a Quality of Service (QoS) feature that classifies application traffic that flows through a router interface. You can use the Cisco IOS modular QoS commandline interface (MQC) to manually configure NBAR on a router or a switch. Before NBAR can classify any traffic, Cisco Express Forwarding (CEF) must be enabled on the router. CEF is enabled by default on Cisco routers. If CEF has been disabled by the no ip cef command, you can reenale CEF by issuing the ip cefcommand.

There are three mandatory steps in a typical NBAR configuration:

1. Define a class map.
2. Configure a policy map.
3. Attach the policy map to an interface.

The first step in an NBAR configuration is to define a class map, also known as a traffic class. A class map is used to identify packets based on the parameters that you specify. Packets that match the parameters are considered to be part of a particular traffic class. You should issue the classmap command to create a class map and to place the router in classmap configuration mode. From classmap configuration mode, you can use match protocol statements to identify the traffic that

should be discovered and classified by NBAR. For example, the command set below creates the class map named `seureshell`, which identifies incoming Secure Shell (SSH) packets:

```
Router(config)#classmap seureshell
Router(configmap)#match protocol ssh
Router(configmap)#exit
```

However, if an application or protocol has been configured to use nonstandard port numbers, you can issue the `ip nbar portmap` command to modify the NBAR configuration accordingly. For example, if SSH servers on the network are configured to listen on ports 22 and 2222, you should issue the `ip nbar portmap ssh tcp 22 2222` command to modify the default NBAR port mapping for SSH.

Next, you should issue the `policy-map` command to configure a policy map and to enter `policy-map` configuration mode. A policy map ties a traffic class to a QoS policy and is used to define actions that are performed on packets identified in a particular class map. For example, the command set below creates a policy map named `NBARpolicy` and then specifies that any packets identified by the class map named `secure-shell` should be rate-limited to 128 Kbps:

```
Router(config)#policy-map NBAR-policy
Router(config-pmap)#class secure-shell
Router(config-pmap-c)#bandwidth 128
Router(config-pmap-c)#exit
Router(config-pmap)#exit
```

Then you should issue the `service-policy` command from interface configuration mode to apply the QoS policy to a particular interface. A service policy can be applied in either the inbound or the outbound direction. For example, the command set below applies the service policy named `NBARpolicy` to the `Serial1/0` interface in the inbound direction:

```
Router(config)#interface serial 1/0
Router(config-if)#servicepolicy input NBARpolicy
Router(config-if)#exit
```

The `ip nbar protocol-discovery` command can be issued from interface configuration mode to record traffic statistics based on packet content. Either or both inbound and outbound traffic can be monitored. To monitor only IPv4 traffic, you should issue the `ip nbar protocol-discovery ipv4` command; to monitor only IPv6 traffic, you should issue the `ip nbar protocol-discovery ipv6` command.

The `auto qos` command enables AutoQoS, which automatically configures QoS settings on an interface. However, if you have manually configured and attached a service policy to an interface by issuing the `service-policy` command, you cannot use AutoQoS to automatically configure QoS.

Reference:

Cisco: Configuring NBAR Using the MQC: Attaching a Traffic Policy to an Interface or Subinterface

Cisco: Cisco AutoQoS White Paper: Considerations, Caveats, and Restrictions for AutoQoS VoIP

QUESTION 63

Which of the following statements are true regarding the IGMPv3 source filtering feature? (Select 2 choices.)

- A. It enables hosts to specify the systems to which they will send multicast traffic.
- B. It enables hosts to specify the systems from which they want to receive multicast traffic.
- C. It enables hosts to specify the systems to which they do not want to send multicast traffic.
- D. It enables hosts to specify the systems from which they do not want to receive multicast traffic.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Internet Group Management Protocol version 3 (IGMPv3) source filtering feature enables hosts to specify the systems from which they want to receive multicast traffic? it also enables hosts to specify the systems from which they do not want to receive multicast traffic. IGMPv3 hosts operate in either INCLUDE mode or EXCLUDE mode. In INCLUDE mode, an IGMPv3 host specifies that it wants to join only the multicast groups listed in the INCLUDE list. In EXCLUDE mode, an IGMP host specifies that it wants to join all multicast groups except those listed in the EXCLUDE list.

IGMPv3 source filtering does not enable hosts to specify the systems to which they will send multicast traffic. Similarly, IGMPv3 source filtering does not enable hosts to specify the systems to which they will not send multicast traffic. IGMPv3 hosts are typically multicast receivers, not multicast sources? multicast sources send traffic, and multicast receivers receive traffic.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfmulti.html#wp1046118

QUESTION 64

What extra information does the loginput keyword provide in an ACL log that the logkeyword does not? (Select 2 choices.)

- A. destination IP address
- B. source IP address
- C. destination MAC address
- D. source MAC address



<https://vceplus.com/>

- E. ingress interface
- F. egress interface

Correct Answer: DE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The log-input keyword provides source Media Access Control (MAC) address and ingress interface information in an access control list (ACL) log? the log keyword does not provide that information. Apart from this information, the log-input keyword logs everything that the log keyword logs, including the source and destination IP address and port numbers.

Neither the log keyword nor the log-input keyword provides the destination MAC address or egress interface information. Both the log keyword and the log-input keyword provide the message identifier, the ACL name or number, whether the packet was permitted or denied, the protocol, the source IP address and port, the destination IP address and port, and the number of similar packets logged during the log update threshold. By default, the log update threshold is five minutes. If multiple matching packets are received during the log update threshold, only one instance is reported every five minutes? additional instances will increment a packet counter and will be reported when the log update threshold expires. The following sample output is generated by an ACL with the log keyword:

*Mar 16 17:02:24.519: %SEC6IPACCESSLOGP: list 101 permitted tcp
10.1.14.3(1234) > 192.168.17.6(6543), 1 packet

The following sample output is generated by an ACL with the log-input keyword? note the addition of the source MAC address and ingress interface:

*Mar 16 17:02:24.519: %SEC6IPACCESSLOGP: list 101 permitted tcp
10.1.14.3(1234) (FastEthernet0/1 0000.0c12.3456) > 192.168.17.6 (6543), 1 packet

You can uniquely identify a particular ACL log message by enabling ACL hash generation. When you enable hash generation by issuing the ip accesslist logging hashgeneration command, an MD5 hash is appended to each ACL log entry.

The following sample output is generated by an ACL when the ip access-list logging hash-generation command has been issued on the router:

*Mar 16 17:02:24.519: %SEC6IPACCESSLOGP: list 101 permitted tcp
10.1.14.3(1234) (FastEthernet0/1 0000.0c12.3456) > 192.168.17.6 (6543), 1 packet Hash code is 0xCE87F535

Reference:

Cisco: Understanding Access Control List Logging

QUESTION 65

Which of the following is best suited for many-to-many applications? (Select the best answer.)

- A. SSM
- B. PIM-SM
- C. PIM-DM
- D. Bidirectional PIM

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Bidirectional Protocol Independent Multicast (bidir-PIM) is best suited for many-to-many applications, such as conferencing and multiplayer gaming. Bidir-PIM enables designated forwarder (DF) routers to forward multicast traffic up the shared tree directly to multicast receivers; the router with the lowest cost to the rendezvous point (RP) is elected as the DF for that network segment. By contrast, unidirectional PIM implementations, such as PIM sparse mode (PIMSM), use a designated router (DR), which forwards multicast traffic from the multicast sources directly to the RP. The RP then sends the multicast traffic down the shared tree. The router with the highest IP address is elected as the DR for that network segment.

Source Specific Multicast (SSM) is best suited for one-to-many applications, which are also called broadcast applications. One-to-many applications include streaming multimedia and other push-based applications. Each application must use a separate multicast group. The Internet Assigned Numbers Authority (IANA) has reserved the IPv4 multicast address range 232.0.0.0/8 and the IPv6 multicast address range of FF3x::/32 for use with SSM. When SSM is used, a multicast host can specify the source addresses from which they will accept multicast traffic.

Cisco provides no specific recommendations for applications to be used with PIM dense mode (PIMDM) or PIMSM. PIMDM routers initially add all the dense mode interfaces to the multicast routing table, flood multicast traffic out all available interfaces, and then prune back those interfaces that have no multicast receivers. By contrast, PIMSM routers add an interface to the multicast routing table only when a device connected through that interface joins the multicast group.

Reference:

Cisco: Source Specific Multicast: SSM Components

Cisco: Bidirectional PIM Deployment Guide (PDF)

IETF: RFC 4607: SourceSpecific Multicast for IP: 4.3. Allocation of SourceSpecific Multicast Addresses

QUESTION 66

Which of the following statements is true regarding EEM? (Select the best answer.)

- A. The Watchdog System Monitor can monitor interface errors.
- B. EEM cannot be configured to restart a router.
- C. EEM cannot be configured to send an email message.
- D. EEM cannot be configured to generate an SNMP trap.
- E. EEM must publish events to subsystem number 798.

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Embedded Event Manager (EEM) must publish events to subsystem number 798. EEM enables routers to monitor events and perform actions if those events are triggered. To configure EEM to publish an application specific event when the EEM event is triggered, you should issue the action publish event command. The Watchdog System Monitor (IOSWDSysMon) cannot monitor interface errors; it is used to monitor memory and processor usage. To configure the Watchdog System Monitor, you should issue the event ioswdsysmon command. To monitor interface errors, you should issue the event interface command. To trigger a command when you manually run a policy event, you can issue the event none command.

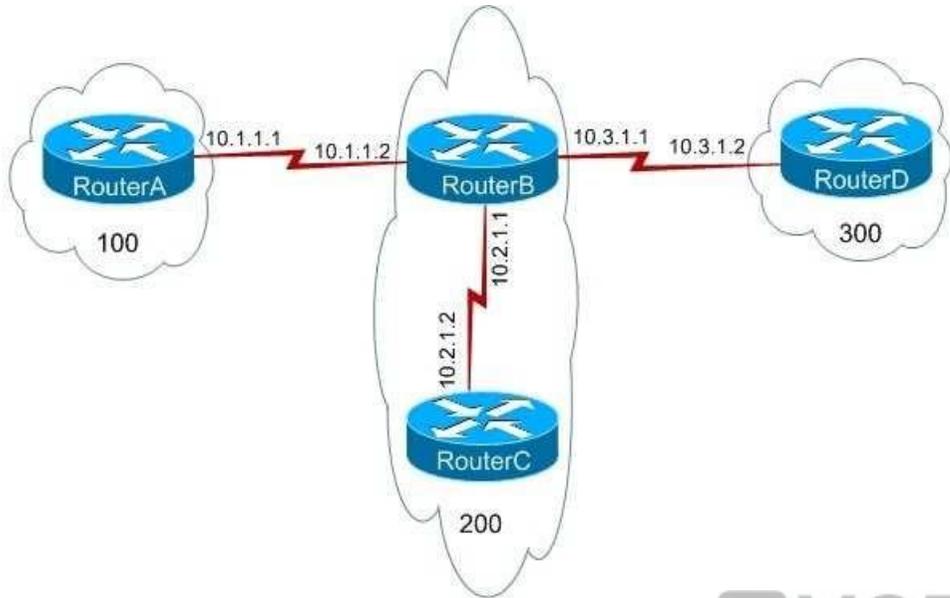
EEM can be configured to restart a router with the action reload command, to send an email message with the action mail command, or to generate a Simple Network Management Protocol (SNMP) trap with the action snmptrap command. The following keywords can be used with the action command:

- cli
- cnsevent
- counter
- force-switchover
- info
- mail
- policy
- publish-event
- reload
- snmp-trap
- syslog

Reference:

<https://search.cisco.com/search?query=Cisco%20IOS%20Network%20Management%20Configuration%20Guide&locale=enUS&tab=Cisco>

QUESTION 67



You administer the network shown above. You issue the show running-config command on RouterA and receive the following partial output:

```
router bgp 100
  no synchronization
  bgp log-neighbor-changes
  network 192.168.0.0
  network 192.168.1.0
  neighbor 10.1.1.2 remote-as 200
  neighbor 10.1.1.2 send-community
  neighbor 10.1.1.2 route-map map1 out
  no auto-summary
!
access-list 1 permit 192.168.0.0 0.0.0.255
!
route-map map1 permit 10
  match ip address 1
  set community no-advertise
!
route-map map1 permit 20
```

Which of the following routers will receive the route to 192.168.0.0/24? (Select the best answer.)

- A. only RouterB
- B. only RouterB and RouterC
- C. only RouterB and RouterD
- D. RouterB, RouterC, and RouterD
- E. None of the routers will receive the route.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Only RouterB will receive the route to 192.168.0.0/24. The neighbor 10.1.1.2 remote-as 200 command specifies that RouterB, which is in autonomous system (AS) 200, is an external Border Gateway Protocol (eBGP) neighbor of RouterA. The neighbor 10.1.1.2 sendcommunity command configures RouterA to send community attribute settings to RouterB. The neighbor 10.1.1.2 routemap map1 out command applies route map map1 to modify outgoing routes from RouterA. Route map map1 will not affect which routes are advertised from RouterA to RouterB; it will only apply the noadvertise community attribute to routes that match access list 1. Routes that do not match access list 1 are advertised without the attribute. Because route map map1 is configured to apply to only the routes that pass access list 1, the noadvertise community attribute will affect only the route to 192.168.0.0/24.

The neighbor 10.1.1.2 send-community command configures RouterA to send community attribute settings to RouterB. The community attribute is an optional, transitive Border Gateway Protocol (BGP) attribute that is not required to be supported by all BGP implementations. Additionally, BGP implementations that do not support the community attribute are not required to pass the attribute to other routers. By default, Cisco routers do not pass community attributes to BGP neighbors. The community attribute can be modified in a route map by issuing the set community command with one of the following four keywords:

-no-advertise -prevents advertisements to any BGP peer
-no-export -prevents advertisements to eBGP peers

-local-as -prevents advertising outside the AS, or in confederation scenarios, outside the subAS
-internet -advertises the route to any router

The set community no-advertise command configures the BGP community attribute to inform neighbor routers to not advertise routes to any BGP peer. Because the community attribute in this scenario applies only to the 192.168.0.0/24 route, RouterB will advertise the route to 192.168.1.0/24 but not the route to 192.168.0.0/24.

The community attribute does not modify how RouterA advertises the routes? it modifies how neighbor routers advertise the routes received from RouterA.

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/26634-bgp-toc.html#communityattribute> <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/26634-bgp-toc.html#sec3> https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-n1.html#wp2607806244

QUESTION 68

Which of the following statements best describes loop guard? (Select the best answer.)

- A. Loop guard prevents a switch port from transitioning to the forwarding state.
- B. Loop guard prevents a switch port from becoming the root port.
- C. Loop guard prevents a switch port from receiving BPDUs.
- D. Loop guard prevents a switch port from becoming a trunk port.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Loop guard prevents a switch port from transitioning to the forwarding state when it stops receiving bridge protocol data units (BPDUs)? this prevents Layer 2 switching loops from occurring. A port that is configured with loop guard that stops receiving BPDUs will be put into the loopinconsistent state, as shown in the following output:

```
%SPANTREE-4-LOOPGUARDBLOCK: No BPDUs were received on port 0/1 in vlan 4. Moved to loop inconsistent state
```

After the port starts receiving BPDUs again, loop guard enables the port to transition through the normal Spanning Tree Protocol (STP) states. Loop guard is only used on interfaces that STP considers to be pointtopoint links.

Root guard, not loop guard, prevents a switch port from becoming the root port, thereby influencing where the root bridge is located on the network. When a port receives a superior BPDU, it will normally attempt to become a root port. However, a root guard port that receives a superior BPDU will be put into the rootinconsistent state, as shown in the following output:

```
%SPANTREE-2-ROOTGUARDBLOCK: Port 0/1 tried to become non-designated in VLAN 4. Moved to root-inconsistent state
```

When the port stops receiving superior BPDUs, the port will be enabled, as shown in the following output:

```
%SPANTREE-2-ROOTGUARDUNBLOCK: Port 0/1 restored in VLAN 4
```

BPDU guard disables a switch port that receives BPDUs. Access ports should never receive BPDUs? to prevent access ports from receiving BPDUs, you can enable BPDU guard on the access ports, thereby defining the edge of the STP domain. When a port that is configured with BPDU guard receives a BPDU, BPDU guard immediately puts the port into the errdisable state and shuts down the port, as shown in the following output:

```
%SPANTREE-2-RX_PORTFAST: Received BPDU on PortFast enabled port.
```

```
Disabling FastEthernet0/1.
```

```
%PM-4-ERR_DISABLE: bpduguard error detected on Fa0/1, putting Fa0/1 in err-disable state
```

%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to down %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down

The port must be manually re-enabled, or it can be recovered automatically through the errdisable timeout function.

Although loop guard can be enabled on trunk ports, it does not prevent a switch port from becoming a trunk port. Loop guard should not be enabled on access ports. By contrast, BPDU guard and root guard can be enabled on access ports.

Reference:

CCIE Routing and Switching v5.0 Certification Guide, Volume 1, Chapter 3, Protecting and Optimizing STP, pp. 148-154

QUESTION 69

Which of the following port states exist in both 802.1D and 802.1w? (Select 2 choices.)

- A. blocking
- B. disabled
- C. discarding
- D. forwarding
- E. learning
- F. listening

Correct Answer: DE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The forwarding and learning port states exist in both traditional 802.1D Spanning Tree Protocol (STP) and 802.1w Rapid STP (RSTP). An STP switch port will pass through the following port states after a switch is turned on:

- Blocking
- Listening
- Learning
- Forwarding

When STP is enabled and a switch is turned on, each port first enters the blocking state. The switch port then transitions to the listening state, in which it begins processing bridge protocol data units (BPDUs) as it listens for information to determine whether it should transition to the learning state. After entering the learning state, a switch port begins to transmit BPDUs and learn addressing information with which to build the switching database. Finally, a switch port transitions to the forwarding state, in which the switch port forwards frames. If a switch port determines at any time during the STP state process that a switching loop would be

caused by entering the forwarding state, the switch port enters the disabled state, in which the switch receives BPDUs but does not direct them to the system module.

RSTP combines the STP disabled, blocking, and listening states into a single port state called the discarding state. An RSTP switch port will pass through the following port states after a switch is turned on: -Discarding

-Learning

-Forwarding

When RSTP is enabled and a switch is turned on, each port first enters the discarding state, in which a port receives BPDUs and directs them to the system module; however, the port neither sends BPDUs nor forwards any frames. The switch port then transitions to the learning state, in which it begins to transmit BPDUs and learn addressing information. Finally, a switch port transitions to the forwarding state, in which the switch port forwards frames. If a switch port determines at any time during the RSTP state process that a switching loop would be caused by entering the forwarding state, the switch port again enters the discarding state, in which the switch receives BPDUs and directs them to the system module but does not send BPDUs or forward frames.

Reference:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/24062-146.html#states>

QUESTION 70

Which of the following commands creates a capture point named boson? (Select the best answer.)

- A. monitor capture point boson size 256 circular
- B. monitor capture point start boson
- C. monitor capture point associate boson exsim
- D. monitor capture point ip cef boson fastethernet 0/1 both

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The monitor capture point ip cef boson fastethernet 0/1 both command creates a capture point named boson on the FastEthernet 0/1 interface. The syntax for the monitor capture point command is monitor capture point {ip | ipv6} {cefcapturepointname interfacename interfacetype {both | in | out} | processswitched capturepointname {both | fromus | in | out}}.

Cisco IOS Embedded Packet Capture (EPC) is a feature that you can implement to assist with tracing packets and troubleshooting issues with packet flow in and out of Cisco devices. You can create multiple capture points with unique names and parameters on a single interface? however, you can associate each capture point with only one capture buffer. To implement Cisco IOS EPC, you must perform the following steps:

1. Create a capture buffer.

2. Create a capture point.
3. Associate the capture point with the capture buffer.
4. Enable the capture point.

The monitor capture point boson size 256 circular command will not create a capture point named boson? however, the monitor capture buffer boson size 256 circularcommand would create a capture buffer named boson. The syntax for the monitor capture buffer command is monitor capture bufferbuffername [clear | exportexportlocation | filteraccesslist {ipaccesslist | ipexpandedlist | accesslistname} | limit {allownthpknthpacket | duration seconds | packetcounttotalpackets | packetspersecpackets} | [maxsizeelements] [size buffersize] [circular | linear]]. When creating capture buffers, you can adjust several items, including buffer type, sampling interval, buffer size, and packet capture rate. Specifying the sampling interval and the buffer type will allow for the maximum number of pertinent packets to be stored in the buffer. The capture buffer contains packet data and metadata? the metadata contains a timestamp indicating when the packet was added to the buffer, the direction of transmission of the packet, the switch path, and the encapsulation type.

The monitor capture point associate boson exsim command will not create a capture point. The monitor capture point associate boson exsim command will associate a capture point named boson with a capture buffer named exsim. The syntax of the command to associate a capture point with a capture buffer is monitor capture point associate capture-point-name capture-buffer-name.

The monitor capture point start boson command will not create a capture point. The monitor capture point start boson command will enable a capture point named boson and begin the process of capturing packet data. The capture point must first be created before it can be enabled.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/epc/configuration/15-mt/epc-15-mt-book/nm-packet-capture.html#GUID-B343AF14-7CA2-45EC-BF9C5DA65AEAD7A3>

QUESTION 71

IS-IS encapsulates its data at which layer of the OSI model? (Select the best answer.)

- A. the Data Link layer
- B. the Network layer
- C. the Transport layer
- D. the Application layer

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Intermediate System-to-Intermediate System (ISIS) encapsulates its data at the Data Link layer of the Open Systems Interconnection (OSI) model and can therefore be used with both IP and Connectionless Network Protocol (CLNP). When ISIS encapsulates packets, it uses 0xFE and 0xFEFE in the Layer 2 header to identify the Layer 3 protocol as OSI. ISIS is specified in International Organization for Standardization (ISO) 10589.

ISIS is a linkstate routing protocol. Routers that use a linkstate routing protocol maintain a complete topology of the network by flooding the state of each router's links across the entire network until each of the routers has information about all of the other routers in the autonomous system (AS). ISIS uses the following Data Link layer multicast addresses to send hello packets and linkstate information:

0180.C200.0014 -All Level 1 (L1) Intermediate Systems

0180.C200.0015 -All Level 2 (L2) Intermediate Systems

0900.2B00.0005 -All Intermediate Systems

0900.2B00.0004 -All End Systems

IS-IS uses a designated intermediate system (DIS) in a broadcast multiaccess network. All ISIS routers on the network segment establish adjacencies with the DIS. The DIS serves as a focal point for the distribution of ISIS routing information. Once elected, the DIS must relinquish its duties if another router with a higher priority joins the network. If the DIS is no longer detected on the network, a new DIS is elected based on the priority of the remaining routers on the network segment. If a new DIS cannot be elected based solely on router priority, the highest Media Access Control (MAC) address is used. If there is still a tie, the highest system ID is the deciding factor. Every ISIS router is required to have a unique system ID.

Reference:

<https://www.cisco.com/c/en/us/products/index.html#wp38435>

QUESTION 72

Which of the following events would cause the overload bit to be set on a router running IS-IS? (Select the best answer.)

- A. Multiple routes to the same destination exist on the router.
- B. The router comes online for the first time.
- C. The router has recently been reloaded.
- D. The router goes offline.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Of the available choices, the overload bit would be set on a router running Intermediate System to Intermediate System (ISIS) that has recently been reloaded. Routers that use a linkstate routing protocol maintain a complete topology of the network by flooding the state of each router's links across the entire network until each of the routers has information about all of the other routers in the autonomous system (AS). ISIS uses Data Link layer multicast addresses to send hello packets and linkstate information. Once a recently reloaded router comes back online, if the overload bit is not set, the adjacent routers will begin forwarding packets to the router before the router can completely populate its routing table. Because the routing table is not complete, the router will drop packets to destinations that have not been written in the table yet? this describes a blackhole issue.

Multiple routes to the same destination existing on a router running IS-IS would not cause the overload bit to be set. Multiple routes typically do not exist in a routing table? however, when a network is using shortest path first (SPF), multiple routes can exist. Linkstate routing protocols consider the links that represent the shortest path to a destination as the best paths. After a router has collected linkstate information for every destination in a topology, the router uses an SPF algorithm to construct an SPF tree. The best paths from the SPF tree are then inserted into the router's routing table.

A router going offline or coming online for the first time would not cause the overload bit to be set. When a router is offline, the adjacent routers will not attempt to send packets to that router. An alternate path, if one exists, will be used to deliver packets to a destination? otherwise, the packet will be returned to the originator as undeliverable. A router coming online for the first time does not have the configuration required for the overload bit to be set.

Reference:

Cisco: Intermediate System to Intermediate System Protocol: Fast Convergence at Adjacency Setup

QUESTION 73

Which of the following are E-line services? (Select 2 choices.)

- A. E-LAN
- B. EPL
- C. E-Tree
- D. EVPL, VPLS

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Both Ethernet private line (EPL) and Ethernet virtual private line (EVPL) are E-line services. E-line services are Ethernet point-to-point Ethernet virtual connection (EVC) services that can be used to connect two User Network Interfaces (UNIs). A UNI is the physical demarcation between a service provider and a subscriber. The difference between an EPL and an EVPL is that an EVPL is capable of service multiplexing. In addition, an EPL requires full service frame transparency. An EVPL does not.

An ELAN service is a multipoint-to-multipoint EVC. Therefore, an ELAN service is not an E-line service. ELAN services fully mesh two or more UNIs and follow a specific set of rules for delivering service frames to a UNI. Each UNI in an ELAN can communicate with any other UNI in the ELAN. ELANs typically have a distance limitation of 50 miles (80 kilometers). Layer 2 Virtual Private Networks (L2VPNs) and multipoint L2VPNs are examples of ELANs.

A Virtual Private LAN Service (VPLS) enables multipoint ELAN services on a Multiprotocol Label Switching (MPLS) network. In a VPLS configuration, MPLS pseudowires are used to link virtual switch instances (VSIs), which emulates an Ethernet switch. A VPLS can then be used to provide EVC services and Transparent LAN Service (TLS).

An E-Tree is a point-to-multipoint EVC that resembles a hubandspoke configuration. Therefore, an E-Tree is not an E-line service. An E-Tree service connects more than one UNI to a single root UNI or leaf UNI. Root UNIs can send data to any leaf UNI. However, a leaf UNI can send traffic only to a root UNI. E-Trees are typically used to provide Internet access to multiple sites.

Reference:

https://www.cisco.com/c/en/us/td/docs/net_mgmt/active_network_abstraction/3-7/service_activation/user/guide/anansaug/tech_overview.html#wp1106296

QUESTION 74

Which of the following best describes a feasible successor? (Select the best answer.)

- A. the best metric along a path to a destination
- B. the total metric along a path to a destination
- C. the highest metric along a path to a destination
- D. a reported distance lower than the feasible distance of the current best path

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A feasible successor is the Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor that has a reported distance that is lower than the feasible distance of the current best path. A feasible successor must have a valid, loop-free path to the destination. If these conditions are met, EIGRP immediately installs the feasible successor in the Routing Information Base (RIB) in order to speed up convergence should the best path become unavailable.

Feasible distance is the EIGRP term for the best metric along a path to a destination. The feasible distance includes the metric to the EIGRP neighbor that is advertising the path.

Reported distance, not feasible distance, is the total metric along a path to a destination. The reported distance is determined by using the metric of the path as advertised by an upstream EIGRP neighbor. The reported distance might also be the highest metric or the lowest metric along a path to a destination.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html>

QUESTION 75

You want to move several company functions to the cloud, including software development and CRM. You decide to use an IaaS vendor.

Which of the following will you most likely have to provide and manage? (Select 3 choices.)

- A. the CRM application
- B. the operating system
- C. the networking infrastructure

- D. the software development platform
- E. the computing and storage resources

Correct Answer: ABD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You will most likely have to provide and manage the customer relationship management (CRM) application, the operating system, and the software development platform. An Infrastructure as a Service (IaaS) vendor provides computing and storage resources as well as the network infrastructure. The customer is responsible for everything else, including operating systems, software development platforms, database platforms, and software applications. With IaaS, the customer has a great deal of control and flexibility. However, IaaS places a larger management burden on the customer than the other cloud-based services do.

A Platform as a Service (PaaS) vendor provides the same services as an IaaS vendor does. In addition, a PaaS vendor also provides operating systems, software development platforms, and database platforms. PaaS is often used by companies that want to migrate their application development to a cloud-based solution. However, a PaaS customer must use whatever software development platform is supported by the PaaS vendor, so a degree of control and flexibility is lost. The PaaS vendor is responsible for maintaining the operating systems, software development platforms, and database platforms, as well as any underlying hardware infrastructure. If you were to use a PaaS vendor, you would have to provide and manage only the CRM application.

A Software as a Service (SaaS) vendor typically provides a complete software application package to customers. For example, a company might contract with an SaaS vendor to provide hosted email services. The software application, the operating system on which the application runs, the hardware on which the operating system runs, and the network infrastructure on which the hardware communicates are maintained by the SaaS vendor, thereby lowering the management burden for the customer. Access to the software application is often provided through a web browser interface. If you were to use an SaaS vendor, you would not have to provide or manage anything; however, you would have to use whatever platforms and CRM applications that the SaaS vendor has available.

Reference:

<https://www.cisco.com/c/en/us/about/press/internet-protocol-journal/back-issues/table-contents-45/123-cloud1.html>

https://www.cisco.com/en/US/services/ps2961/ps10364/ps10370/ps11104/Migration_of_Enterprise_Apps_to_Cloud_White_Paper.pdf

https://www.cisco.com/c/en/us/td/docs/solutions/Hybrid_Cloud/PaaS/1-0/PaaS/PaaS1.pdf

QUESTION 76

In Cisco ACI, what is a collection of VRF instances or IP address spaces? (Select the best answer.)

- A. an ANP
- B. an EPG
- C. a context
- D. a contract

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In Cisco Application Centric Infrastructure (ACI), a context is a collection of VPN routing and forwarding (VRF) instances or IP address spaces. Each customer, or tenant, can have one or more contexts. Endpoints and endpoint groups (EPGs) define the application within each context.

An EPG is a collection of endpoints that provide a similar function, such as an application tier or a set of services. The endpoints within an EPG are defined by network interface card (NIC), virtual NIC (vNIC), port group, IP address, or Domain Name System (DNS) name.

A contract is a collection of rules and policies that define how endpoints and EPGs can communicate. For example, a contract can be created so that a web server can be accessed only by Hypertext Transfer Protocol (HTTP) or HTTP Secure (HTTPS).

An Application Network Profile (ANP) is a collection of EPGs, their connections, and related policies. To create an ANP, you should perform the following steps:

1. Create EPGs.
2. Create policies that define connectivity rules.
3. Create contracts between EPGs by applying policies.

Reference: <https://www.cisco.com/c/en/us/products/cloud-systems-management/index.html>

QUESTION 77

In a three-node OpenStack architecture, the network node consists of services from which of the following OpenStack components? (Select the best answer.)



<https://vceplus.com/>

- A. Glance
- B. Horizon
- C. Keystone
- D. Neutron
- E. Nova

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In a three-node OpenStack architecture, the network node consists of services from the Neutron component. OpenStack is an open-source cloud-computing platform. Each OpenStack modular component is responsible for a particular function, and each component has a code name. The following list contains several of the most popular OpenStack components:

- Nova -OpenStack Compute: manages pools of computer resources
- Neutron -OpenStack Networking: manages networking and addressing
- Cinder -OpenStack Block Storage: manages blocklevel storage devices
- Glance -OpenStack Image: manages disk and server images
- Swift -OpenStack Object Storage: manages redundant storage systems
- Keystone -OpenStack Identity: is responsible for authentication
- Horizon -OpenStack Dashboard: provides a graphical user interface (GUI)
- Ceilometer -OpenStackTelemetry: provides counterbased tracking that can be used for customer usage billing

A three-node OpenStack architecture consists of the network node, the controller node, and the compute node. The network node consists of the following Neutron services:

- Neutron Modular Layer 2 (ML2) PlugIn
- Neutron Layer 2 Agent
- Neutron Layer 3 Agent
- Neutron Dynamic Host Configuration Protocol (DHCP) Agent

The controller node consists of the following services:

- Keystone
- Glance
- Nova Management
- Neutron Server
- Neutron ML2 Plug-In
- Horizon
- Cinder
- Swift
- Ceilometer Core

The compute node consists of the following services:

- Nova Hypervisor
- Kernel-based Virtual Machine (KVM) or Quick Emulator (QEMU)
- Neutron ML2 Plug-In
- Neutron Layer 2 Agent -
- Ceilometer Agent

Reference: <https://www.redhat.com/archives/rdo-list/2014-November/pdfzGvyHATdWc.pdf#page=12>

QUESTION 78

Which of the following statements is true regarding hypervisors? (Select the best answer.)

- A. Both KVM and Xen are Type1 hypervisors.
- B. Both KVM and Xen are Type2 hypervisors.
- C. Type1 hypervisors are generally slower than Type2 hypervisors.
- D. Type2 hypervisors are also called native hypervisors.

Correct Answer: A

Section: (none)

Explanation



Explanation/Reference:

Explanation:

Both Kernel based Virtual Machine (KVM) and Xen are Type1 hypervisors. A hypervisor is used to create and run virtual machines (VMs). A Type1 hypervisor runs directly on the host computer's hardware. Other Type1 hypervisors include HyperV and VMware ESX/ESXi.

KVM and Xen are not Type2 hypervisors. A Type2 hypervisor runs within an operating system on the host computer. VMware Workstation, Parallels Desktop for Mac, and Quick Emulator (QEMU) are Type2 hypervisors.

Type-1 hypervisors are generally faster than Type-2 hypervisors because Type-1 hypervisors run directly on the host computer's hardware and because Type-2 hypervisors have a host operating system that consumes system resources.

Type-1 hypervisors are also called native hypervisors or baremetal hypervisors. Type-2 hypervisors are also called hosted hypervisors.

Reference:

<https://www.ibm.com/developerworks/library/l-hypervisor/>

<https://www.xenproject.org/users/virtualization.html>

QUESTION 79

Which of the following routing protocols can be used for routing on IoT networks? (Select the best answer.)

- A. EIGRP
- B. IS-IS
- C. OSPF
- D. RPL

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Routing Protocol for Lowpower and Lossy Networks (RPL) can be used for routing on Internet of Things (IoT) networks. RPL is an IP version 6 (IPv6) routing protocol that is defined in Request for Comments (RFC) 6550. An IoT network is considered to be a Low-power and Lossy Network (LLN).

IoT networks connect embedded devices. Embedded devices, or smart objects, are typically lowpower, lowmemory devices with limited processing capabilities. These devices are used in a variety of applications, such as environmental monitoring, healthcare monitoring, process automation, and location tracking. Many embedded devices can transmit data wirelessly, and some are capable of transmitting over a wired connection. However, connectivity is generally unreliable and bandwidth is often constrained.

IoT networks require a routing protocol that can handle the limitations of embedded devices. Neither Enhanced Interior Gateway Routing Protocol (EIGRP), Intermediate System to Intermediate System (IS-IS), nor Open Shortest Path First (OSPF) meets the requirements for routing an IoT network, as specified by the Internet Engineering Task Force (IETF) Routing over LLNs (ROLL) working group. In addition to RPL, IPv6 over Low Power Wireless Personal Area Networks (6LoWPAN) and Constrained Application Protocol (CoAP) have been created to address the challenges of routing an IoT network.

Reference:

<https://tools.ietf.org/html/rfc6550>

<https://datatracker.ietf.org/wg/roll/charter/>

QUESTION 80

Which of the following statements are true regarding Ansible, Salt, Chef, and Puppet? (Select 2 choices.)

- A. All have a web UI.
- B. All are written in Ruby.
- C. All are written in Python.
- D. All require client installation.
- E. All are configuration management tools.

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Ansible, Salt, Chef, and Puppet are all configuration management tools and all have a web user interface (UI). Configuration management tools are used to automate the installation, configuration, and maintenance of multiple computer systems, including the software that runs on those systems.

Not all of these configuration management tools are written in the same language. Puppet and Chef are written in Ruby, whereas Salt and Ansible are written in Python.

Not all of these configuration management tools require client installation. Ansible does not use client agent software on managed nodes. By contrast, Puppet and Chef require client agent software on managed nodes. Salt nodes can use client agent software but do not require it.

Of the four major configuration management tools, Puppet is the most mature and the most widely used. Puppet operates on Linux distributions, UNIXlike systems, and Microsoft Windows. Puppet uses a client/ server architecture; managed nodes running the Puppet Agent application can receive configurations from a master server running Puppet Server. Modules are written in Ruby or by using a Rubylike Puppet language.

Like Puppet, Chef operates on Linux distributions, UNIXlike systems, and Microsoft Windows. Chef can use a client/server architecture or a standalone client configuration. Configuration information is contained within cookbooks that are written in Ruby and are stored on a Chef Server. Managed nodes running the Chef Client can pull cookbooks from the server. Standalone clients that do not have access to a server can run chefsolo and pull cookbooks from a local directory or from a tar.gz archive on the Internet.

Salt also operates on Linux distributions, UNIXlike systems, and Microsoft Windows. Salt can use a client/ server architecture by installing Salt master software on the server and Salt minion software on managed nodes. Masters and minions communicate by using ZeroMQ. Salt can also be used without installing Salt minion software by using Salt Secure Shell (SSH). However, Salt SSH is much slower than ZeroMQ.

Configuration information is stored primarily in state modules that are typically written in YAML; however, Python or Python Domain Specific Language (PyDSL) can also be used for complex configuration scripts.

Like the other configuration management software packages, Ansible also operates on Linux distributions, UNIXlike systems, and Microsoft Windows. However, unlike the other configuration management software packages, Ansible does not use agent software on managed nodes. Configurations are stored on the Ansible server in playbooks that are written in YAML. Managed nodes can download scripted modules from an Ansible server by using SSH.

Reference:

<https://www.infoworld.com/article/2609482/data-center/data-center-review-puppet-vs-chef-vs-ansible-vs-salt.html?page=4>

QUESTION 81

Which of the following features influences root bridge selection by putting a designated port into an inconsistent state when it receives a superior BPDU? (Select the best answer.)

- A. BPDU guard
- B. PortFast
- C. loop guard

D. root guard

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Root guard influences root bridge selection by putting a designated port into an inconsistent state when it receives a superior bridge protocol data unit (BPDU). Normally, a port that receives a superior BPDU will become the root port. However, if a designated port configured with root guard receives a superior BPDU, the port transitions to the root-inconsistent state and no data will flow through that port until it stops receiving superior BPDUs. As a result, root guard can be used to influence the placement of the root bridge on a network by preventing other switches from propagating superior BPDUs throughout the network and becoming the root bridge.

Loop guard does not put a port into an inconsistent state when it receives a superior BPDU; loop guard puts a port into an inconsistent state when it stops receiving BPDUs. This prevents the trunk port from transitioning to the forwarding state, thereby preventing a Layer 2 switching loop.

BPDU guard does not put a port into an inconsistent state when it receives a superior BPDU; BPDU guard puts a port into the errdisable state when it receives any BPDUs. This is useful for host ports, which should never receive BPDUs. BPDU guard defines the edge of the Spanning Tree Protocol (STP) domain by limiting the advertisement of BPDUs to a port.

PortFast does not put a port into an inconsistent state when it receives a superior BPDU; PortFast enables a port to immediately access the network by transitioning the port into the STP forwarding state without passing through the listening and learning states.

Reference:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10588-74.html>

QUESTION 82

Which of the following commands could you issue to configure a router interface to use unicast packets to establish OSPF neighbor relationships? (Select 2 choices.)

- A. ip ospf network broadcast
- B. ip ospf network non-broadcast
- C. ip ospf network point-to-point
- D. ip ospf network point-to-multipoint
- E. ip ospf network point-to-multipoint non-broadcast

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You could issue either the `ip ospf network non-broadcast` command or the `ip ospf network point-to-multipoint non-broadcast` command to configure a router interface to use unicast packets to establish Open Shortest Path First (OSPF) neighbor relationships. These commands are useful for creating an OSPF network over a non-broadcast multiaccess (NBMA) network, such as a Frame Relay network, where broadcast and multicast traffic is not allowed. However, OSPF neighbor discovery relies on multicast traffic, so routers on an NBMA network cannot automatically discover neighbor routers. Therefore, manual configuration of neighbor routers with the `neighbor` command is required with NBMA networks. All OSPF traffic is then sent to the unicast IP addresses configured in the neighbor commands. The `ip ospf network broadcast` command, the `ip ospf network point-to-point` command, and the `ip ospf network point-to-multipoint` command all use multicast packets to establish OSPF neighbor relationships. As a result, these OSPF network types do not require manual configuration of neighbor routers. OSPF uses multicast addresses 224.0.0.5 and 224.0.0.6 to send linkstate advertisements (LSAs) and hello packets.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html#t24> https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/ospf-i1.html#wp3564440872

QUESTION 83

The OSPF process on RouterA, which is not a PE router, is associated with a VRF instance. Which of the following commands will disable PEspecific checks? (Select the best answer.)

- A. `address-family ipv4 vrf`
- B. `capability vrf-lite`
- C. `ip vrf forwarding`
- D. `ip vrf`

**Correct Answer: B****Section: (none)****Explanation****Explanation/Reference:****Explanation:**

The `capability vrf-lite` command will disable provider edge (PE)specific checks on RouterA. PE checks are used to prevent loops when the Open Shortest Path First (OSPF) process on a PE is associated with the VPN routing and forwarding (VRF) instance and the router is mutually redistributing OSPF and Border Gateway Protocol (BGP). VRF enables multiple instances of a routing table to exist on a router. When OSPF is associated with a VRF instance and is being mutually redistributed with BGP, it is possible for routing loops to occur. PE checks examine linkstate advertisements (LSAs) to determine whether a specific path should be considered for insertion into the routing table. However, PE checks are not necessary on a router that is not running BGP and is therefore not a PE router. The `addressfamily ipv4 vrf` command does not disable PE checks. Instead, the `address-family ipv4 vrf vrf-name` command is used to configure VRF contexts under a routing process. For example, to configure the routing context for VRF boson under Enhanced Interior Gateway Routing Protocol (EIGRP) process 65000, you would issue thefollowing commands:

```
RouterA(config)#router eigrp 65000
RouterA(config-router)#address-family ipv4 vrf boson
```

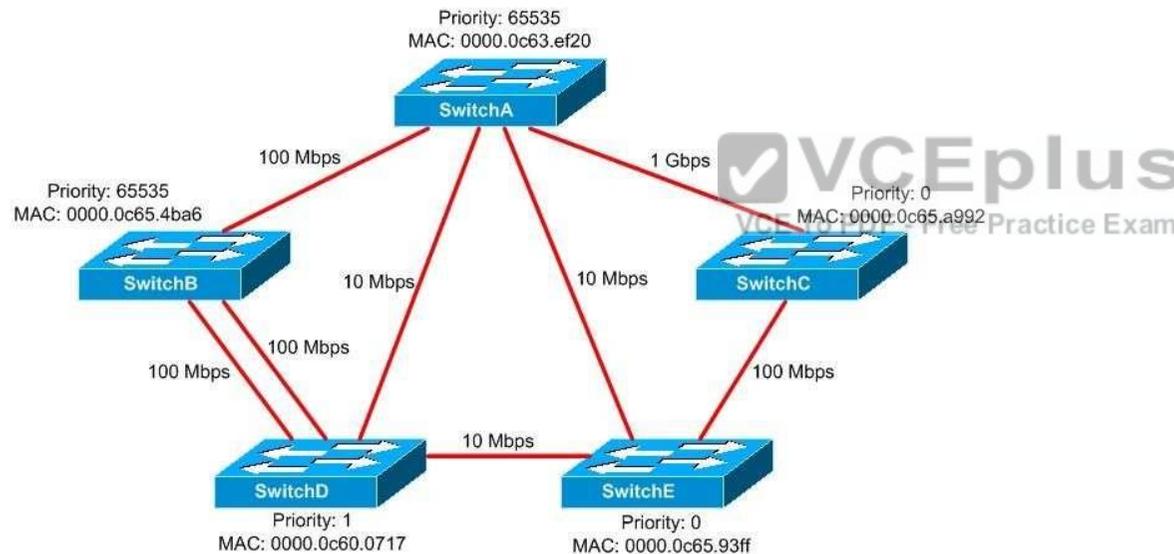
Neither the ip vrf command nor the ip vrf forwarding command disables PE checks. Instead, the ip vrf forwarding vrfname command adds a VRF instance to an interface. For example, to add the VRF boson to the FastEthernet0/1 interface on RouterA, you would issue the following commands:

```
RouterA(config)#interface FastEthernet0/1
RouterA(config-if)#ip vrf forwarding boson
```

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/ospf-a1.html#wp2582896905

QUESTION 84



You administer the network shown above. No VLANs are configured on any of the switches. Which of the following switches is the root bridge for the network? (Select the best answer.)

- A. SwitchA
- B. SwitchB
- C. SwitchC

- D. SwitchD
- E. SwitchE

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

SwitchE is the root bridge for the network. The root bridge is the switch with the lowest bridge ID. The bridge ID is composed of a 2byte bridge priority and a 6byte Media Access

Control (MAC) address. The bridge priority is considered first in the determination of the lowest bridge ID. The bridge priority can be set by issuing the spanning-tree priorityvalue command, where value is a number from 0 through 65535? the default priority is 32768.

SwitchC and SwitchE both have a priority of 0. When two or more switches have the lowest priority, the switch with the lowest MAC address becomes the root bridge. MAC addresses are written in hexadecimal format. With MAC addresses, numbers are lower than letters and the hexadecimal value A is lower than the hexadecimal value F. Because SwitchE has a lower MAC address than SwitchC, SwitchE is the root bridge.

SwitchA is not the root bridge for the network, because it has the highest priority value, not the lowest priority value. Although link speed is somewhat relevant in determining the root port for a switch, link speed is irrelevant in determining the root bridge.

SwitchB is not the root bridge for the network; like SwitchA, SwitchB also has the highest priority value, not the lowest priority value. SwitchB contains redundant links to SwitchD, but redundant links are irrelevant in determining the root bridge. To avoid a switching loop, at least one of the redundant links between SwitchB and SwitchD will be blocked.

SwitchC is not the root bridge for the network. If the bridge priority of SwitchE were higher than 0, SwitchC would be the root bridge because a priority of 0 is the lowest configurable priority value.

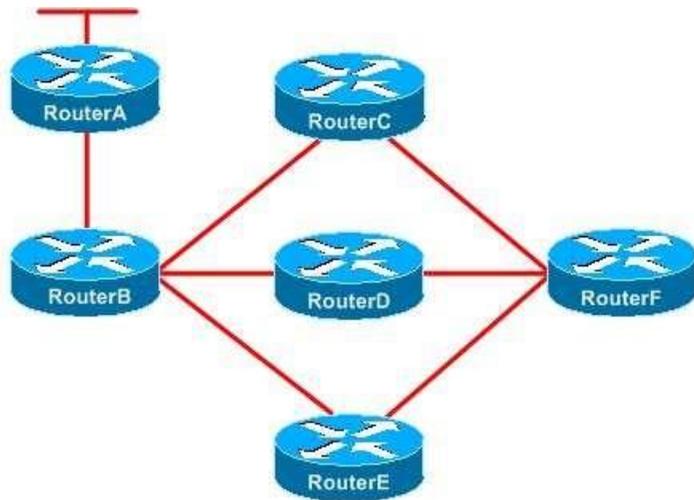
SwitchD is not the root bridge for the network. Although SwitchD has the lowest MAC address on the network, the bridge priority is considered first in the determination of the root bridge. If all of the switches on the network had the same bridge priority values, SwitchD would be the root bridge because it has the lowest MAC address.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960/software/release/12-2_53_se/configuration/guide/2960scg/swstp.html#wp1157719

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10588-74.html>

QUESTION 85



You administer the EIGRP network shown above. RouterB is configured to send only a summary route to RouterE. RouterC is configured as a stub router. The link between RouterA and RouterB fails. Which of the following routers will send a query to RouterF? (Select the best answer.)

- A. only RouterC
- B. only RouterD
- C. only RouterE
- D. only RouterC and RouterD
- E. only RouterD and RouterE
- F. RouterC, RouterD, and RouterE

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Only RouterD will send a query to RouterF. Query packets are sent to find routes to a destination network. When a router loses the best route to a destination and does not have a feasible successor, it floods query packets to its neighbors. If a neighbor has a route to the destination network, it replies with the route. However, if a neighbor does not have a route to the destination network, it queries its neighbors, those neighbors query their neighbors, and so on. This process continues

until either a router replies with the route or there are no routers left to query. The network cannot converge until all the replies have been received, which can cause a router to become stuck in active (SIA).

Limiting Enhanced Interior Gateway Routing Protocol (EIGRP) queries prevents queries from consuming bandwidth and processor resources and prevents routers from becoming SIA. You can display which routers have not yet replied to a query by issuing the `show ip eigrp topology active` command, as shown in the following output:

```
RouterA#show ip eigrp topology active
IP-EIGRP Topology Table for AS(1)/ID(192.168.99.2)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

A: 192.168.99.3/32, 1 successors, FD is Inaccessible, tag is 1
   1 replies, active 00:05:15, query-origin: Local origin
     via Redistributed (2297856/0)
   Remaining replies:
     via 192.168.99.3, r, Serial10/0.223
```

The `eigrp stub` command limits EIGRP queries by creating a stub router. Stub routers advertise only a specified set of routes and therefore typically need only a default route from a hub router. A hub router detects that a router is a stub router by examining the Type-Length-Value (TLV) field within EIGRP hello packets sent by the router. The hub router will specify in its neighbor table that the router is a stub router and will no longer send query packets to that stub router, thereby limiting how far EIGRP queries spread throughout a network. Because RouterC is configured as a stub router, RouterB will not send queries to RouterC, and RouterC will therefore not propagate those queries to RouterF. Although hub routers will not send queries to stub routers, stub routers can initiate queries of their own. The `ip summary address eigrp as number address mask` command limits EIGRP queries by configuring route summarization. If a neighbor router has a summarized route but does not have the specific route to the destination network in the query, the neighbor router will reply that it does not have a route to the destination network and will not query its neighbors. Thus route summarization creates a query boundary that prevents queries from propagating throughout the network. In this scenario, RouterB is configured to send only a summary route to RouterE; therefore, RouterE will not send queries to RouterF. RouterD is not configured as a stub router, and RouterB is not sending RouterD a summarized route. Therefore, when RouterB sends a query to RouterD, RouterD will send a query to RouterF.

Reference:

https://www.cisco.com/en/US/technologies/tk648/tk365/technologies_white_paper0900aecd8023df6f.html

QUESTION 86

You are configuring a serial link on one of your company's routers. You want to enable encapsulation for the link and ensure that your configuration will support either asynchronous or synchronous communications. You also want to configure authentication for the link, and you want to use the most secure authentication mechanism available for the link.

Which of the following commands should you use to accomplish your goal? (Select the best answer.)

- A. Router1(config)#interface serial 1 Router1(config)#encapsulation ppp Router1(config)#ppp authentication pap
- B. Router1(config)#interface serial 1
Router1(config)#encapsulation hdlc
Router1(config)#ppp authentication pap
- C. Router1(config)#interface serial 1 Router1(config)#encapsulation ppp Router1(config)#ppp authentication chap
- D. Router1(config)#interface serial 1
Router1(config)#encapsulation hdlc
Router1(config)#ppp authentication chap

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should enable PointtoPoint Protocol (PPP) as the encapsulation protocol for the link and Challenge Handshake Authentication Protocol (CHAP) as the authentication protocol for the link by issuing the following command sequence:

```
Router1(config)#interface serial 1
Router1(config)#encapsulation ppp
Router1(config)#ppp authentication chap
```



PPP is a wide area network (WAN) protocol used on pointtopoint serial links. PPP supports both synchronous and asynchronous communications. HighLevel Data Link Control (HDLC) is another WAN protocol that can be used on pointtopoint serial links, but HDLC supports only synchronous communications. Unlike PPP, HDLC does not support authentication. On Cisco routers, HDLC is the default serial interface encapsulation protocol.

PPP supports two types of authentication mechanisms: Password Authentication Protocol (PAP) and CHAP. When PAP is used, the user name and password of the originating router are sent over the link in plain text. By contrast, when CHAP is used, a hash of the user name and password combination, as well as a random number, is sent to the destination router? the user name and password are not sent across the link. Thus CHAP is more secure than PAP for authentication.

Reference:

http://docwiki.cisco.com/wiki/Point-to-Point_Protocol <https://www.cisco.com/c/en/us/support/docs/wan/point-to-point-protocol-ppp/25647-understanding-ppp-chap.html>

QUESTION 87

You use PPP on your Cisco router to allow users to access the network remotely. Which of the following protocols can you use for authentication? (Select 2 choices.)

- A. PPP

- B. CHAP
- C. SLIP
- D. PAP
- E. HDLC

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You can use Password Authentication Protocol (PAP) or Challenge Handshake Authentication Protocol (CHAP) for authentication with PointtoPoint Protocol (PPP). PPP relies on an authentication protocol to provide authentication. PPP is a wide area network (WAN) protocol used to provide remote connectivity. PPP is easy to configure and can transport multiple Layer 3 protocols, such as IP, Internetwork Packet Exchange (IPX), and AppleTalk. To configure an interface for PPP, you should issue the encapsulation ppp command in interface configuration mode.

When PAP is used, the user name and the password of the originating router are sent over the link in plain text. By contrast, when CHAP is used, a hash of the user name and password combination, as well as a random number, is sent to the destination router; the user name and password are not sent across the link. Thus CHAP is more secure than PAP for authentication.

To configure a router interface to use CHAP authentication, you should issue the ppp authentication chap command from interface configuration mode. To configure a router interface to use PAP authentication, you should issue the ppp authentication pap command. To configure a router interface to use both CHAP and PAP, you should issue the ppp authentication chap pap command or the ppp authentication pap chap command? the authentication methods will be used in the order they are listed.

You can also use Microsoft CHAP (MSCHAP), MSCHAP version 2 (MSCHAP v2), and Extensible Authentication Protocol (EAP) with PPP. To configure a router interface to use MSCHAP, you should issue the ppp authentication mschap command. To configure a router interface to use MSCHAP v2, you should issue the ppp authentication mschapv2 command. To configure a router interface to use EAP, you should issue the ppp authentication eap command.

You cannot use HighLevel Data Link Control (HDLC) for authentication. Similar to PPP, HDLC is a WAN encapsulation protocol, and it can be used with multiple Layer 3 protocols. However, HDLC does not support authentication.

Like PPP, the older Serial Line Internet Protocol (SLIP) can also be used for remote connectivity and relies on an authentication mechanism to provide authentication. Unlike PPP and HDLC, SLIP can transport only IP traffic. To configure an interface for SLIP, you should issue the encapsulation slip command in interface configuration mode.

Reference:

<https://www.cisco.com/c/en/us/support/docs/wan/point-to-point-protocol-ppp/25647-understanding-ppp-chap.html> http://docwiki.cisco.com/wiki/Point-to-Point_Protocol

QUESTION 88

Which device or devices receive packets destined for FF02::5? (Select the best answer.)

- A. all IPv6capable nodes on the segment
- B. all IPv6capable routers on the segment
- C. the single device that is configured with the address FF02::5
- D. all RIPv6 routers
- E. all OSPFv3 routers
- F. all EIGRPv6 routers
- G. all DRs and BDRs

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

All Open Shortest Path First version 3 (OSPFv3) routers receive packets destined for FF02::5, which is similar to the OSPFv2 allrouters multicast address 224.0.0.5. OSPFv3designated routers (DRs) and backup designated routers (BDRs) receive packets destined for FF02::6, which is similar to the OSPFv2 allDR/BDR multicast address 224.0.0.6. These multicast addresses are used to exchange hello messages and linkstate advertisements (LSAs) among OSPF routers.

The IPv6 address FF02::5 is a multicast address, which is used for onetomany communication. IPv6 multicast addresses begin with the hexadecimal characters FF. Individual devices cannot be configured with a particular multicast address. Therefore, packets destined for FF02::5 will not be received by a single device configured with that address.

All Routing Information Protocol version 6 (RIPv6) routers receive packets destined for FF02::9, which is similar to the RIPv2 allrouters multicast address 224.0.0.9. This address is used to exchange hello packets and routing updates among RIP routers.

All Enhanced Interior Gateway Routing Protocol version 6 (EIGRPv6) routers receive packets destined for FF02::A, which is similar to the EIGRP IPv4 allrouters address 224.0.0.10. This address is used to exchange hello packets and routing updates among EIGRP routers.

All IPv6 nodes on a segment receive packets destined for FF02::1, which is similar to the IPv4 allhosts multicast address 224.0.0.1. Traffic sent to FF02::1 is received by all hosts on the local segment. IPv6 nodes automatically join the FF02::1 multicast group at startup.

All IPv6 routers on a segment receive packets destined for FF02::2, which is similar to the IPv4 allrouters multicast address 224.0.0.2. Traffic sent to FF02::1 is received by all routers on the local segment. IPv6 routers automatically join the FF02::2 multicast group at startup.

Reference:

<https://tools.ietf.org/html/rfc5340#page-57>

QUESTION 89

Which of the following must you do to enable OSPFv3 to function on an interface? (Select 3 choices.)

- A. You must issue the ipv6 router ospf process id command in global configuration mode.

- B. You must issue the router ospfv3 [processid] command in global configuration mode.
- C. You must issue the ipv6 unicastrouting command in global configuration mode.
- D. You must issue the ipv6 address command or the ipv6 enable command in interface configuration mode.
- E. You must issue the ipv6 ospf processidarea areaid [instanceinstanceid] command in interface configuration mode.
- F. You must issue the networknetworkidwildcardmaskarea areaid command in router configuration mode.

Correct Answer: CDE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You must perform the following steps to enable Open Shortest Path First version 3 (OSPFv3) to function on an interface:

- Issue the ipv6 unicastrouting command in global configuration mode.
- Issue the ipv6 address command or the ipv6 enable command in interface configuration mode.
- Issue the ipv6 ospf processidarea areaid [instanceinstanceid] command in interface configuration mode.

The ipv6 unicast-routing command enables IPv6 unicast routing on the router. IPv6 unicast routing must be enabled globally on the router so that the router can forward IPv6 packets.

The ipv6 address command enables IPv6 on the interface and assigns an IPv6 address to the interface.

The ipv6 enable command enables IPv6 on an interface but does not assign an address to the interface.

The ipv6 ospf processidarea areaid [instance instanceid] command enables OSPFv3 on the interface if the previous two steps have been completed. You can also issue the ospfv3 processidarea areaid {ipv4 | ipv6} [instanceinstanceid] command to enable OSPFv3 on an interface for a particular address family. Address families allow OSPFv3 to support both IPv4 and IPv6, but only one address family can be enabled on an OSPFv3 instance. Unlike OSPF version 2 (OSPFv2), OSPFv3 allows multiple instances to be enabled on a link.

Issuing the ipv6 router ospfprocessid command or the router ospfv3 [processid] command creates an OSPFv3 routing process on the router, but neither command is required to enable OSPFv3 to function on an interface. Enabling OSPFv3 on an interface automatically creates an OSPFv3 routing process on the router, so neither the ipv6 router ospfprocessid command nor the router ospfv3 command is required unless you need to configure global OSPFv3 parameters.

You are not required to issue the networknetworkidwildcardmaskarea areaid command in router configuration mode. The network command in OSPFv2 is used to specify which networks should participate in OSPF. Because OSPFv3 is configured directly on the interface, the network command is no longer necessary and is therefore unavailable in OSPFv3.

Reference:

<https://search.cisco.com/search?query=Cisco%20IOS%20IPv6%20Configuration%20Guide&locale=enUS&tab=Cisco> <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2s/ipv6-15-2s-book.html#GUID-F7619122-1D7B-4433-8F1B-012F5978AFFE>

QUESTION 90

Which of the following must you do before IP source guard can be used on a switch port? (Select 2 choices.)

- A. Configure static IP bindings on the switch.
- B. Enable DHCP snooping on the switch.
- C. Enable uRPF on the switch port.
- D. Enable IP routing on the switch port.
- E. Enable CEF on the switch.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You must configure static IP bindings or enable Dynamic Host Configuration Protocol (DHCP) snooping on the switch before IP source guard can be used on a switch port. To configure a static IP binding, you should issue the `ip source binding mac addressvlan vlanid ipaddressinterface interfaceid` command. To enable DHCP snooping, you should issue the `ip dhcp snooping` command.

IP source guard prevents all IP traffic except for the following packets:

-DHCP packets allowed by DHCP snooping

-Traffic that matches entries in the IP source binding table

The IP source binding table is populated by static bindings or by DHCP snooping. If you enable IP source guard on a switch port but do not configure static IP bindings or DHCP snooping, all IP traffic will be dropped by the switch.

IP source guard mitigates DHCP spoofing attacks. In a DHCP spoofing attack, an attacker installs a rogue DHCP server on the network in an attempt to intercept DHCP requests. The rogue DHCP server can then respond to the DHCP requests with its own IP address as the default gateway address? hence all traffic is routed through the rogue DHCP server. As a result, a host that has obtained an IP address from a rogue DHCP server could become the victim of a maninthemiddle attack in which a malicious individual eavesdrops on a network conversation between two hosts. Enabling DHCP snooping with IP source guard helps to mitigate DHCP spoofing attacks.

You do not need to enable unicast Reverse Path Forwarding (uRPF) on the switch port. Like

IP source guard, uRPF can mitigate spoofing attacks. uRPF checks the source IP address of a packet to determine whether the packet arrived on the best path back to the source based on routing table information. If the IP address information is spoofed, the uRPF check will fail and the packet will be dropped.

You do not need to enable Cisco Express Forwarding (CEF) on the switch. Unlike uRPF, IP source guard does not rely on CEF to function? CEF must be enabled for uRPF to function.

You should not enable IP routing on the switch port. In fact, enabling routing on a switch port by issuing the `no switchport` command prevents you from enabling IP source guard on the switch port.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/blades/3120/software/release/12-2_40_ex/configuration/guide/3120scg/swdhcp82.pdf

QUESTION 91

You have enabled CEF and have issued the ip verify unicast source reachable via rx command to enable uRPF in strict mode on a router. A TCP packet with a source address of 10.11.12.1 arrives on the router's FastEthernet0/1 interface. A route to 10.11.12.1 exists in the FIB, but the path through the FastEthernet0/1 interface is not the best path to the source.

Which of the following will occur? (Select the best answer.)

- A. The packet will be dropped.
- B. The packet will be forwarded through a valid path.
- C. The packet will be forwarded through the best path.
- D. The packet will be logged as suspicious.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The packet will be dropped because unicast Reverse Path Forwarding (uRPF) is operating in strict mode. When you enable uRPF in strict mode, the router checks packets upon arrival at an interface to determine whether those packets arrived through the best path to the source. If a packet did not arrive from the best path, the packet is dropped. Implementing uRPF in strict mode can cause legitimate traffic to be dropped in asymmetric routing configurations.

For uRPF to be used in either strict or loose mode, Cisco Express Forwarding (CEF) must be enabled. The router uses the information in the Forwarding Information Base (FIB) to perform the reverse lookup. The FIB is generated by CEF. In strict mode, the router checks to see whether a path to the source exists in the FIB and whether the packet arrived on the interface with the best path to the source. In loose mode, the router checks to see whether the source exists in the FIB and is a valid forwarding entry, not just the best path.

There are two network addresses that uRPF always allows to pass even though they might not be present in the FIB: 0.0.0.0 and 255.255.255.255. Not allowing those addresses to pass would cause problems with both Bootstrap Protocol (BOOTP) and Dynamic Host Configuration Protocol (DHCP).

The packet would not be forwarded through any path. Because you have enabled strict mode and the packet did not arrive on the best path back to the source, the packet is dropped. If the packet had arrived on the best path to the source, the best path criteria would have been met and the packet would have been forwarded.

If you had issued the ip verify unicast source reachable via any command, which enables uRPF in loose mode, the packet would have been forwarded. In loose mode, the router checks the FIB to determine whether the packet arrived on a valid path back to the source. uRPF in loose mode forwards the packet as long as the reverse path is a valid path, even if it is not the best path back to the source.

The packet will not be logged as suspicious. uRPF is a reverse path checking tool and not a logging tool for suspicious activity. However, uRPF can mitigate spoofing attacks.

Reference:

<https://www.cisco.com/c/en/us/about/security-center/unicast-reverse-path-forwarding.html>

https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/srpf_gsr.html#wp1053391

https://www.cisco.com/c/en/us/td/docs/ios/sec_data_plane/configuration/guide/convert/sec_data_urpf_15_1_book/sec_cfg_unicast_rpf.html#wp1000928

QUESTION 92

Which of the following types of authentication is configured with the isis password command? (Select the best answer.)



<https://vceplus.com/>

- A. interface authentication
- B. area authentication
- C. domain authentication
- D. router authentication

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Interface authentication is configured with the isis password command. Intermediate System-to-Intermediate System (ISIS) has three methods of authentication:

- Interface authentication
- Area authentication
- Domain authentication

Interface authentication configures ISIS to add authentication for hello messages sent from that interface. Hello messages are used to establish and maintain adjacencies. To configure ISIS interface authentication, you should issue the isis password command from interface configuration mode. The syntax of the isis password command is isis password password [level1 | level2]. If the level is not specified, the password is used for both L1 and L2 routing. The isis password command is not used to configure router authentication.

Area authentication configures ISIS to add authentication for L1 Link State Packets (LSPs),

Complete Sequence Number PDUs (CSNPs), and Partial Sequence Number PDUs (PSNPs). LSPs distribute routing information between nodes. The designated intermediate system (DIS), which is similar to the designated router (DS) in Open Shortest Path First (OSPF), sends CSNPs and PSNPs that describe the LSPs in

the linkstate database. CSNPs, which describe all of the LSPs in the database, are multicast periodically by the DIS. PSNPs, which describe a subset of the LSPs, are used to acknowledge received LSPs and to request missing LSPs. To configure area authentication, you should issue the area password command from ISIS router configuration mode. The syntax of the area password command is area password password.

Domain authentication configures ISIS to add authentication for L2 LSPs, CSNPs, and PSNPs. To configure domain authentication, you should issue the domain password command from ISIS router configuration mode. The syntax of the domain password command is domain password password.

It is possible to configure interface authentication, area authentication, and domain authentication on the same router. The following partial configuration is from a router that has been configured with all three ISIS authentication methods:

```
interface fastethernet0/1
 ip address 10.1.1.1 255.255.255.0
 ip router isis
 isis password b0$on
!
router isis
 net 49.2222.7777.7777.00
 domain-password clsc0
 area-password cCnpR&S
```

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

QUESTION 93

RouterA and RouterB are connected routers.

You issue the show runningconfig command on RouterA and receive the following partial output:

```
router isis net
49.1741.0000.0000.000a.00
```

You issue the show runningconfig command on RouterB and receive the following partial output:

```
router isis net
48.1741.0000.0000.000b.00
```

Neither the istype command nor the isis circuit-type command has been issued on either router.

Which of the following output would you expect to see after issuing the show clns neighbors command on RouterA? (Select the best answer.)

A.	<input type="radio"/>	System Id	Interface	SNPA	State	Holdtime	Type	Protocol
		RouterB	Et0/0	0000.0000.000b	Up	23	L1/L2	IS-IS
B.	<input type="radio"/>	System Id	Interface	SNPA	State	Holdtime	Type	Protocol
		RouterB	Et0/0	0000.0000.000b	Up	23	IS	ES-IS
C.	<input checked="" type="radio"/>	System Id	Interface	SNPA	State	Holdtime	Type	Protocol
		RouterB	Et0/0	0000.0000.000b	Up	23	L2	IS-IS
D.	<input type="radio"/>	System Id	Interface	SNPA	State	Holdtime	Type	Protocol
		RouterB	Et0/0	0000.0000.000b	Up	23	L1	IS-IS

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You would expect to see the following output after issuing the show clns neighbors command on RouterA:

```
System Id  Interface  SNPA          State  Holdtime  Type  Protocol
RouterB   Et0/0     0000.0000.000b  Up    23        L2   IS-IS
```

Routers running the Intermediate System to Intermediate System (ISIS) routing protocol are placed into administrative domains called areas. Each ISIS router resides in only one area. The collection of all areas managed by a single organization is called a routing domain. RouterA and RouterB are in separate areas, as indicated by the different net commands in their running configurations. RouterA is part of area 49.1741, and RouterB is part of area 48.1741.

Each ISIS router is configured with a routing level. Level 1 (L1) routers are capable of intraarea routing, which delivers data within a single area. Level 2 (L2) routers are capable of interarea routing, which delivers data between areas. Level 1/Level 2 (L1/L2) routers are capable of both intraarea and interarea routing and maintain a separate linkstate database for each. You can configure the routing level for an ISIS process by issuing the `istype {level1 | level12 | level2only}` command, and you can configure the routing level for an ISIS interface by issuing the `isis circuittype {level1 | level12 | level2only}` command. By default, all ISIS routing processes and interfaces are configured for L1/L2 routing. Therefore, both RouterA and RouterB are configured for L1/L2 routing. Additionally, RouterA and RouterB will establish an L2 adjacency because the routers are in different areas.

If RouterA and RouterB were in the same area, you would expect to see the following output after issuing the show clns neighbors command on RouterA:

```
System Id  Interface  SNPA          State  Holdtime  Type  Protocol
RouterB   Et0/0      0000.0000.000b Up     23        L1L2  IS-IS
```

If RouterA and RouterB were in the same area and if either router were configured for L1 routing only, you would expect to see the following output after issuing the show clns neighbors command on RouterA:

```
System Id  Interface  SNPA          State  Holdtime  Type  Protocol
RouterB   Et0/0      0000.0000.000b Up     23        L1    IS-IS
```

If RouterA and RouterB were in different areas and if either router were configured for L1 routing only, you might see the following output after issuing the show clns neighbors command on RouterA:

```
System Id  Interface  SNPA          State  Holdtime  Type  Protocol
RouterB   Et0/0      0000.0000.000b Up     23        IS    ES-IS
```

L1 routers cannot form adjacencies between areas. When an ISIS routing level mismatch, authentication mismatch, or maximum transmission unit (MTU) mismatch occurs, an ISIS adjacency will not form, but the output of the show clns neighbors command might instead show an End System-to-Intermediate System (ESIS) adjacency. ESIS is used to discover end systems.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/isoclns/command/reference/iso_book/iso_m1.html#wp1023033

QUESTION 94

You issue the following commands on RouterA:

```
router bgp 400
neighbor 192.168.1.1 remote-as 400
neighbor 192.168.1.1 route-map map-in in
neighbor 192.168.1.1 route-map mapout out
neighbor 192.168.1.1 filter-list 1 in neighbor
192.168.1.1 password jdsfr39oo26 neighbor
192.168.1.2 remoteas 400 neighbor
192.168.1.2 route-map map-in in neighbor
192.168.1.2 route-map map-out out neighbor
192.168.1.2 filter-list 1 in neighbor
192.168.1.2 password jdsfr39oo26 neighbor
192.168.1.3 remote-as 400 neighbor
192.168.1.3 route-map map-in in neighbor
192.168.1.3 route-map map-out out neighbor
```

```
192.168.1.3 filter-list 1 in neighbor
192.168.1.3 password jdsfr39oo26 neighbor
192.168.1.4 remote-as 400 neighbor
192.168.1.4 route-map map-in in neighbor
192.168.1.4 route-map map-out out neighbor
192.168.1.4 filter-list 1 in neighbor
192.168.1.4 password jdsfr39oo26 neighbor
192.168.1.5 remote-as 400 neighbor
192.168.1.5 route-map map-extra in
neighbor 192.168.1.5 route-map map-out out
neighbor 192.168.1.5 route-map map-in in
neighbor 192.168.1.5 filter-list 1 in neighbor
192.168.1.5 password jdsfr39oo26
```

Which of the following statements are correct? (Select 2 choices.)

- A. RouterA uses eBGP to communicate with its neighbors.
- B. You can implement a peer group to simplify the configuration.
- C. Each neighbor is configured with a different number of route maps.
- D. Each neighbor is configured with the same number of AS path filters.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You can implement a peer group to simplify and shorten the configuration. Additionally, each neighbor is configured with the same number of autonomous system (AS) path filters.

You can use a peer group to easily configure multiple neighbor statements that contain the same policy information for each neighbor. Using a peer group simplifies the configuration, reduces the chance for typographical errors, reduces the CPU and memory load on a router, and enables updates to be replicated more efficiently. To configure a peer group, you should issue the `neighbor peer-group-name-peer-group` command in router configuration mode. To configure a peer group member, you should issue the `neighboripaddresspeergroup peergroupname` command in router configuration mode.

The following commands configure a peer group named `pg1` and effectively replace the commands previously issued in this scenario:

```
router bgp 400 neighbor pg1
peergroup neighbor pg1
```

```
remoteas 400 neighbor pg1
routemap mapin in
neighbor pg1 routemap mapout out
neighbor pg1 filterlist 1 in neighbor pg1
password jdsfr39oo26 neighbor
192.168.1.1 peergroup pg1 neighbor
192.168.1.2 peergroup pg1 neighbor
192.168.1.3 peergroup pg1 neighbor
192.168.1.4 peergroup pg1 neighbor
192.168.1.5 peergroup pg1 neighbor
192.168.1.5 routemap mapextra in
```

Each neighbor is configured with one AS path filter. To create an AS path filter, you should issue the neighbor {ipaddress | peergroupname} filterlist accesslist {in | out} command. The in keyword specifies an inbound path filter, and the out keyword specifies an outbound path filter.

Each neighbor is not configured with a different number of route maps. Each neighbor router is configured with an inbound route map named mapin and an outbound route map named mapout. Although the neighbor router at 192.168.1.5 was configured with another inbound route map named mapextra, as indicated by the neighbor 192.168.1.5 routemap mapextra in command, it was replaced by the neighbor 192.168.1.5 routemap mapin in command. A router can be configured with only one inbound and one outbound route map per neighbor.

When converting the individual neighbor ipaddress routemap commands as a peer group configuration, you can specify each route map that is shared by all of the members of the peer group in a single neighbor peergroupname routemap command. Any extra route maps that are used by only a few members of the peer group can be specified by using individual neighbor ipaddress routemap commands.

RouterA does not use external Border Gateway Protocol (eBGP) to communicate with its neighbors, because all of the routers share the same AS number, as indicated by the neighbor ipaddress remoteas 400 command. Routers that share the same AS number are internal BGP (iBGP) routers.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13755-29.html> https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-m1.html#wp4505123410

QUESTION 95

You have issued the ip multicasting command and the ip pim ssm command from global configuration mode.

Which of the following commands can you issue on each multicast interface to complete the SSM configuration? (Select 3 choices.)

- A. ip pim densemode
- B. ip pim passive
- C. ip pim sparsemode
- D. ip pim sparsedensemode
- E. ip igmp version 2
- F. ip igmp version 3

Correct Answer: CDF

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You can issue the ip igmp version 3 command and either the ip pim sparsemodecommand or the ip igmp sparsedensemode command on each interface. Source Specific Multicast (SSM) enables a device to specify the source addresses from which it will accept multicast traffic. Internet Group Management Protocol version 3 (IGMPv3) improves upon IGMPv2 by adding support for SSM; therefore, IGMPv3 is required for SSM to function. You can enable IGMPv3 on an interface by issuing the ip igmp version 3command. IGMPv2 does not support SSM; therefore, you cannot issue the ip igmp version 2 command on an SSM interface. SSM is derived from Protocol Independent Multicast sparse mode (PIMSM). Therefore, PIMSM and PIM sparse dense mode (PIMSDM) are the only modes that can be used with SSM. You can enable PIMSM on an interface by issuing the ip pim sparse mode command, and you can enable PIMSDM on an interface by issuing the ip pim sparse dense mode command. You cannot use PIM dense mode (PIMDM) or PIM passive mode with SSM; therefore, you should not issue the ip pim dense mode command or the ip pim passive command on an SSM interface.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750x_3560x/software/release/12-2_53_se/configuration/guide/3750xscg/swmcast.html#wp1308052

QUESTION 96

You want to protect the control plane of RouterA from SNMPbased DoS attacks.

Additionally, you want to ensure that SNMP traffic from your management station,192.168.1.111, to the control plane of RouterA is not restricted.

You have created a class map, created a policy map, and applied a service policy on the control plane in the inbound direction, as indicated by the partial running configuration shown below:

```
<output omitted>
!
class-map match-all limit-snmpp
  match access-group name boson
!
policy-map snmp-policy
  class limit-snmpp
    drop
!
<output omitted>
!
control-plane
  service-policy input snmp-policy
!
<output omitted>
```

Which of the following command sets should you issue to complete the configuration? (Select the best answer.)

- A. RouterA(config)#ip access-list extended boson
RouterA(config-ext-nacl)#deny udp host 192.168.1.111 any eq snmp
RouterA(config-ext-nacl)#permit udp any any eq snmp
- B. RouterA(config)#ip access-list extended boson
RouterA(config-ext-nacl)#permit udp host 192.168.1.111 any eq snmp
RouterA(config-ext-nacl)#deny udp any any eq snmp
- C. RouterA(config)#ip accesslist extended boson
RouterA(config-ext-nacl)#deny udp any host 192.168.1.111 eq snmp
RouterA(config-ext-nacl)#permit udp any any eq snmp
- D. RouterA(config)#ip access-list extended boson
RouterA(config-ext-nacl)#permit udp any host 192.168.1.111 eq snmp
RouterA(config-ex-tnacl)#deny udp any any eq snmp

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the command set below to create a named extended access control list (ACL) named boson that completes the configuration:

```
RouterA(config)#ip access-list extended boson
```

```
RouterA(config-ext-nacl)#deny udp host 192.168.1.111 any eq snmp
```

```
RouterA(config-ext-nacl)#permit udp any any eq snmp
```

To create a named extended ACL, you should issue the ip access-list extended acl-name command. Issuing this command will place the router in extendednamedACL mode. Extended ACL entries can be created by using the following basic syntax:

```
[sequence-number] {deny | permit} protocol source source-wildcard [operator [port]] destination destination-wildcard [operator [port]]
```

In this scenario, the deny udp host 192.168.1.111 any eq snmp command creates an extended ACL entry that denies User Datagram Protocol (UDP) packets that have a source IP address of 192.168.1.111 and a destination port number of 161. Simple Network Management Protocol (SNMP) traffic uses UDP port 161 for control traffic and UDP port 162 for SNMP trap traffic. The permit udp any any eq snmp command adds a second entry to the ACL that permits all other SNMP traffic.

AC boson is used in a match statement in the class map named limit-snm, as shown by the following partial command output:

<output omitted>

```
!  
Class-map match-all limit-snm  
  match access-group name boson  
!  
<output omitted>
```

A class map defines a traffic class and specifies the criteria used to identify packets belonging to that class. In this scenario, the match access-group name boson command specifies that all packets permitted by the ACL named boson will belong to the traffic class named limitsnmp. The traffic class named limitsnmp is then used to identify traffic in a policy map named snmp-policy, as shown by the following partial command output:

```
<output omitted>  
!  
Policy-map snmp-policy  
class limit-snm  drop  
!  
<output omitted>
```

A policy map specifies the actions that are taken on packets that match a particular traffic-class. In this scenario, the drop keyword specifies that packets identified as members of the traffic class named limits-nmp are discarded by the router. Finally, the service policy named snmp-policy is applied to the control plane in the inbound direction, as shown by the partial command output below:

```
<output omitted>  
!  
Controlplane servicepolicy  
input snmp-policy  
!  
<output omitted>
```

When this service policy is applied to the control plane in the inbound direction, only SNMP packets sourced from the management station, 192.168.1.111, are permitted to pass to the control plane. Because the service policy instructs the router to discard the remainder of the SNMP packets that are destined to the router's control plane, the router is protected from SNMP based Denial of Service (DoS) attacks.

The command set below does not complete the configuration, because it incorrectly permits only SNMP packets sourced from the management station:

```
RouterA(config)#ip accesslist extended boson  
RouterA(configextnacl)#permit udp host 192.168.1.111 any eq snmp  
RouterA(configextnacl)#deny udp any any eq snmp
```

When the above ACL is used with the service policy in this scenario, only SNMP packets sourced from the management station, 192.168.1.111, are denied access to pass to the control plane. Because the service policy instructs the router to permit the remainder of the SNMP packets that are destined to the router's control plane, the router is not protected from SNMP-based DoS attacks.

The command set below does not complete the configuration, because it incorrectly permits all SNMP packets, regardless of their source IP address:

```
RouterA(config)#ip accesslist extended boson
RouterA(configextnacl)#deny udp any host 192.168.1.111 eq snmp
RouterA(configextnacl)#permit udp any any eq snmp
```

Although the above ACL can be used with the service policy in this scenario to protect the router's control plane from SNMP-based DoS attacks, the ACL does not enable the management station to access the router's control plane.

Conversely, the command set below does not complete the configuration, because it incorrectly denies all SNMP packets, regardless of their source IP address:

```
RouterA(config)#ip accesslist extended boson
RouterA(configextnacl)#permit udp any host 192.168.1.111 eq snmp
RouterA(configextnacl)#deny udp any any eq snmp
```

When the above ACL is used with the service policy in this scenario, all SNMP packets can access the router's control plane.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/qos/configuration/guide/fgos_c/qcfccli2.html

QUESTION 97

You have implemented OSPF on your network. However, your supervisor asks you to try to reduce the time to detect a failed link to less than one second. Which of the following should you do? (Select the best answer.)

- A. Enable STP UplinkFast.
- B. Implement OSPF fast hellos.
- C. Reduce OSPF hello timers.
- D. Enable BFD.
- E. Switch to EIGRP, and reduce hello and hold timers.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should enable Bidirectional Forwarding Detection (BFD) to try to reduce the time to detect a failed link to less than one second. BFD is a detection protocol that is designed to detect forwarding path failures at a consistent rate, thereby providing network administrators with predictable reconvergence times.

Additionally, BFD is designed to work regardless of media type, encapsulation, or routing protocol, providing network administrators with a uniform forwarding failure detection method across a network. BFD supports Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), Border Gateway Protocol (BGP), and Intermediate System to Intermediate System (IS-IS).

The detection of a forwarding path failure causes BFD to notify the routing protocol that a link has failed, which causes the routing protocol to recalculate the routing table. BFD works by sending control packets between two adjacent routers to create a BFD neighbor session. Once the neighbor relationship is established, the two adjacent routers send control packets to each other to maintain the neighbor relationship, similarly to how routing protocols maintain neighbor relationships. However, BFD sends packets at a much faster rate than routing protocols do. In addition, BFD can distribute some functionality from the control plane to the data plane, thereby requiring fewer CPU resources than routing protocol timers do. Only a single BFD session is established per interface regardless of how many routing protocols are running on that interface.

You should not enable Spanning Tree Protocol (STP) UplinkFast to try to reduce the time to detect a failed link to less than one second. STP UplinkFast is a Cisco proprietary STP convergence enhancement. You can enable STP UplinkFast to reduce STP convergence time from the standard 14 to 30 seconds down to one second, but not to less than one second.

You should not implement OSPF fast hellos to try to reduce the time to detect a failed link to less than one second. OSPF fast hellos can send multiple hello packets per second, which results in a faster convergence time. However, the detection and notification of link failures also depend on the OSPF dead interval, which can be set to a minimum of one second.

You should not reduce OSPF hello timers to try to reduce the time to detect a failed link to less than one second. Although reducing the OSPF dead interval and increasing the number of hellos sent during the dead interval can result in hellos being sent at a subsecond rate, the OSPF dead interval can be set to a minimum of one second; therefore, reducing OSPF hello timers can reduce the time to detect a failed link to a minimum of one second.

You should not switch to EIGRP and reduce hello and hold timers to try to reduce the time to detect a failed link to less than one second. Reducing EIGRP timers can reduce the time to detect a failed link to less than two seconds.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html#wp1053332

https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fasthelo.html#wp1036997

QUESTION 98

You issue the ip as-path access-list 1 permit ^[09]*_222\$ command on a BGP router.

Which of the following paths are allowed by the AS path filter? (Select the best answer.)

- A. paths that are learned from AS 222 and originate from any directly attached AS
- B. paths that originate from AS 222 and are learned from AS 222 or any AS directly attached to AS 222
- C. paths that are learned from AS 222 and originate from a directly attached AS numbered from 0 through 9
- D. paths that originate from AS 222 and are learned from a directly attached AS numbered from 0 through 9

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:

Explanation:

Paths that originate from Border Gateway Protocol (BGP) autonomous system (AS) 222 and are learned from AS 222 or any AS directly attached to AS 222 are allowed by the AS path filter. Regular expressions are used to locate character strings that match a particular pattern.

The caret (^) character indicates that the subsequent characters should match the start of the string. Each router in the path prepends its AS number to the beginning of the AS path? therefore, the first AS number in the AS path is the AS from which the path is learned. The expression [09] indicates a single digit from 0 through 9. The asterisk (*) character indicates zero or more sequences of the previous expression. When combined, the expression [09]* indicates any number of digits, including a set of no digits. Therefore, the ip aspath accesslist 1 permit ^[09]*_222\$ command allows paths that are learned from any AS number, including AS 222.

The dollar sign (\$) character indicates that the preceding characters should match the end of the string. The originating router will insert its AS number into the AS path, and subsequent routers will prepend their AS numbers to the beginning of the AS path string. The last AS number in the AS path is the originating AS; therefore, the ip aspath accesslist 1 permit ^[09]*_222\$ command allows paths that originate from AS 222.

The underscore (_) character is used to indicate a comma, a brace, the start or end of an input string, or a space. When used between two AS path numbers, the _ character indicates that the ASes are directly connected. Therefore, the ip aspath accesslist 1 permit ^[09]*_222\$ command indicates that if the path is not learned from AS 222 directly, the AS from which the path is learned must be directly connected to AS 222.

The ip aspath accesslist 1 permit ^[09]*_222\$ command does not permit paths that are learned from AS 222 and originate from any directly attached AS. To permit paths that are learned from AS 222 and originate from AS 222 or any directly attached AS, you should issue the ip aspath accesslist 1 permit ^222_[09]*\$ command.

The ip aspath accesslist 1 permit ^[09]*_222\$ command does not permit paths that are learned from AS 222 and originate from a directly attached AS numbered from 0 through 9. To permit paths that are learned from AS 222 and originate from a directly attached AS numbered from 0 through 9, you should issue the ip aspath accesslist 1 permit ^222_[09]\$ command.

The ip aspath accesslist 1 permit ^[09]*_222\$ command does not permit paths that originate from AS 222 and are learned from a directly attached AS numbered from 0 through 9. To permit paths that originate from AS 222 and are learned from a directly attached AS numbered from 0 through 9, you should issue the ip aspath accesslist 1 permit ^[09]_222\$ command.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13754-26.html>

https://www.cisco.com/c/en/us/td/docs/ios/12_2/termserv/configuration/guide/ftersv_c/tcfaapre.html <https://supportforums.cisco.com/t5/other-service-provider-subjects/bgp-regular-expression-as-path-filter/td-p/1821020>

QUESTION 99

Which of the following commands will cause EEM to check the value of the _exit_status variable after an applet is finished? (Select the best answer.)

A. event cli pattern "show ip interface brief" sync yes

- B. event cli pattern "show ip interface brief" sync no skip no
- C. event cli pattern "show ip interface brief" sync no skip yes
- D. set 1 _exit_status 0
- E. set 1 _exit_status 1

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The event cli pattern "show ip interface brief" sync yes command will cause

Embedded Event Manager (EEM) to check the value of the _exit_status variable after an applet is finished. The event cli command configures EEM to monitor commandline interface (CLI) commands and to trigger the event when a specified pattern is matched one or more times. Events can be processed synchronously or asynchronously. The sync yes keywords are used with the event cli command to configure synchronous processing. With synchronous processing, the EEM applet must finish before the CLI command can be executed, and the _exit_status variable determines whether the CLI command is executed or skipped. If the _exit_status variable is set to a value of 0 or is not configured, the CLI command will not execute after the EEM applet is finished? if the _exit_status variable is set to a value of 1, the CLI command will execute after the EEM applet is finished.

The set 1 _exit_status 0 command sets the _exit_status variable to a value of 0, which will cause EEM to not execute the CLI command after the applet is finished. The set 1 _exit_status 1 command sets the _exit_status variable to a value of 1, which will cause EE to execute the CLI command after the applet is finished. The sync no keywords are used with the event cli command to configure asynchronous processing. With asynchronous processing, the EEM applet is processed at the same time the CLI command is executed. Asynchronous processing does not check the value of the _exit_status variable. Instead, asynchronous processing uses the skip no or skip yes keywords to indicate whether the CLI command should be executed or skipped, respectively.

The event cli pattern "show ip interface brief" sync no skip no command will cause EE to execute the CLI command when the applet runs. The event cli pattern "show ipinterface brief" sync no skip yes command will not execute the CLI command when the applet runs.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/eem/command/eem-cr-book/eem-cr-e1.html#wp1886141985> <https://learningnetwork.cisco.com/docs/DOC-19468>

QUESTION 100

Which of the following values is the default TCP MSS setting for a Cisco router that is originating data destined for a remote IP network? (Select the best answer.)

- A. 68 bytes
- B. 536 bytes
- C. 1460 bytes
- D. 10000 bytes

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The default Transmission Control Protocol (TCP) maximum segment size (MSS) setting for a Cisco router that is originating data destined for a remote IP network is 536 bytes. The MSS is specified in the TCP SYN packet during the TCP handshake. MSS values can be used to restrict outgoing TCP segments to a segment size that is small enough to avoid fragmenting the IP datagram, thereby avoiding the performance problems that can occur as a result of IP fragmentation. The smallest maximum transmission unit (MTU) that can be used on an IPv4 network is 576 bytes. The 536byte default MSS value is therefore derived by subtracting the 20byte TCP header and the 20byte IP header from that MTU value.

It is important to note that some firewall rules are capable of stripping TCP options from a segment. If a firewall is configured to strip TCP options from a segment, the MSS value that is applied to a TCP segment by the router will not be used. If you have NetFlow enabled, you can issue the show ip cache flow command to view statistics that include IP packet size distribution.

The default TCP MSS setting for a Cisco router that is sending data destined for a local LAN is 1460 bytes. The typical default MSS value for PCs communicating on a LAN is 1500 bytes.

The lowest value you can use to enable an MSS for TCP connections that originate from a router is 68 bytes. To configure an MSS value for TCP segments that originate from a router, you should issue the ip tcp mssmssvalue command in global configuration mode.

By issuing the ip tcp adjustmssmssvalue command (where mssvalue is a value in the range from 500 through 1460) in interface configuration mode, you can configure an MSS for TCP segments that do not originate from the router but that are being forwarded by the router.

The highest value you can use to enable an MSS for TCP connections that originate from a router is 10000 bytes. To configure a 10000byte MSS value, you should issue the ip tcp mss 10000 command in global configuration mode. However, you cannot configure a 10000byte MSS for TCP segments that are simply being forwarded on a router interface, because the maximum MSS you can configure for TCP segments that are being forwarded on a router interface is 1460.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/ipaddr/command/reference/fipras_r/1rfip2.html#wp1103772

QUESTION 101

DRAG DROP

Drag the functions on the left underneath the corresponding VSLP protocols on the right.

Select and Place:

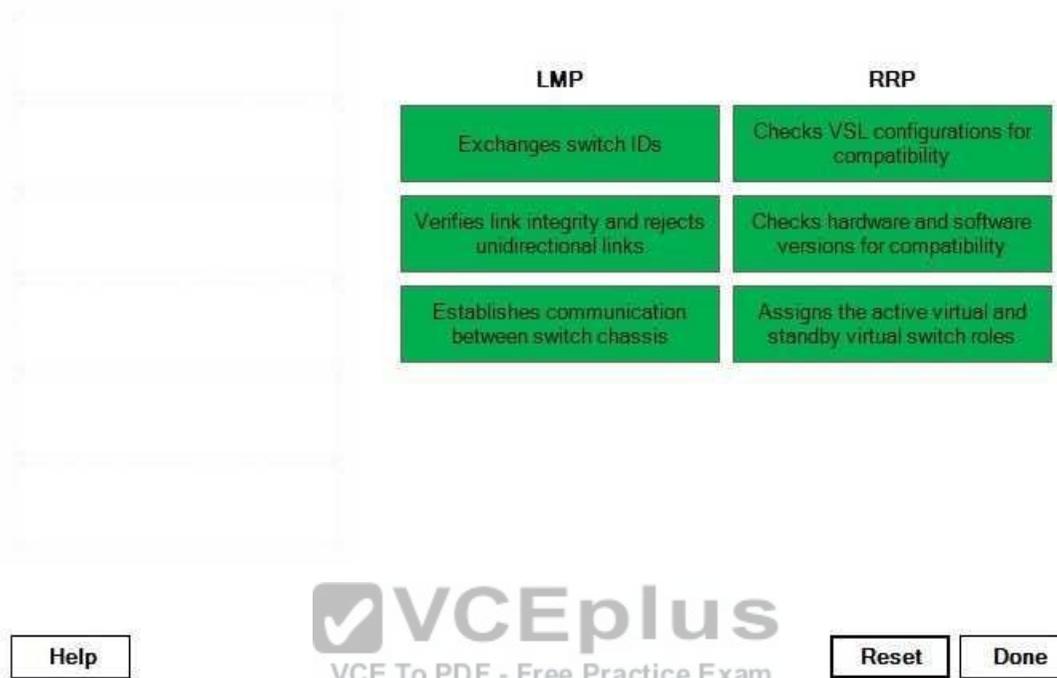
	LMP	RRP
Exchanges switch IDs		
Checks VSL configurations for compatibility		
Checks hardware and software versions for compatibility		
Verifies link integrity and rejects unidirectional links		
Establishes communication between switch chassis		
Assigns the active virtual and standby virtual switch roles		

Help

Reset

Done

Correct Answer:



Section: (none)

Explanation

Explanation/Reference:

Explanation:

Virtual Switching System (VSS) combines two physical Cisco Catalyst switches into a single virtual switch, which can result in greater network efficiency and bandwidth capacity. One switch chassis becomes the active virtual switch, and the other switch becomes the standby virtual switch. The switch chassis are connected together by a virtual switch link (VSL), which is implemented as an EtherChannel of up to eight physical interfaces.

Configuration, monitoring, and troubleshooting must be performed on the active virtual switch; console access is disabled on the standby virtual switch. The active virtual switch is responsible for all control plane functions, such as Simple Network Management Protocol (SNMP), Telnet, Secure Shell (SSH), Spanning Tree Protocol (STP), Link Aggregation Control Protocol (LACP), and Layer 3 routing. The data plane is active on both switches.

Virtual Switch Link Protocol (VSLP) is responsible for establishing the VSS. VSLP has two component protocols: Link Management Protocol (LMP) and Role Resolution Protocol (RRP). The VSS initialization process consists of the following steps:

- 1.The configuration file is pre-parsed for VSL configuration commands.
- 2.The VSL member interfaces are brought online.

3. LMP verifies link integrity, rejects unidirectional links, and establishes bidirectional communication between switch chassis.
4. LMP exchanges switch IDs in order to detect duplicate IDs.
5. RRP checks hardware versions, software versions, and VSL configurations for compatibility.
6. RRP assigns the active virtual and standby virtual switch roles.
7. Switches come up in Nonstop Forwarding/Stateful Switchover (NSF/SSO) mode or route-processor redundancy (RPR) mode.
8. Switches continue the normal boot process.

If RRP determines that both switches are compatible, both chassis will come up in NSF/SSO mode, in which all modules are powered up and can forward traffic. If RRP determines that an incompatibility exists, the standby virtual switch will come up in RPR mode, in which all modules are powered down. The switch chassis that is started first will always become the active virtual switch unless preemption is configured. If both chassis are started simultaneously, the switch with the highest priority will become the active virtual switch. By default, the priority is set to a value of 100. If priorities are equal, the switch with the lower switch ID will become the active virtual switch.

Reference:

https://www.cisco.com/c/dam/en/us/products/collateral/interfaces-modules/network-modules/white_paper_c11_429338.pdf

QUESTION 102

Which of the following queuing methods is the most appropriate for handling voice, video, mission-critical, and lower-priority traffic? (Select the best answer.)

- A. FIFO
- B. LLQ
- C. WFQ
- D. CBWFQ

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Of the choices provided, low latency queuing (LLQ) is the most appropriate for handling voice, video, missioncritical, and lowerpriority traffic. LLQ supports the creation of up to 64 user defined traffic classes as well as one or more strict priority queues that can be used specifically for delay sensitive traffic, such as voice and video traffic. Each strictpriority queue can use as much bandwidth as possible but can only use its guaranteed minimum bandwidth when other queues have traffic to send, thereby avoiding bandwidth starvation. You can also implement weighted random early detection (WRED) on each of the user defined traffic classes to mitigate packet loss? WRED is particularly useful for networks with a large amount of Transmission Control Protocol (TCP) traffic.

On voice networks, you can implement LLQ to help reduce jitter. Additionally, you can configure the voice class with a smaller queue size. Although a smaller queue size could result in dropped packets, voice traffic is more tolerant of dropped packets than of delayed packets. A small amount of packet loss is not

noticeable to the human ear. Additionally, some codecs can correct small amounts of packet loss. Therefore, a smaller queue size combined with the use of LLQ could reduce delay and jitter.

Class based weighted fair queuing (CBWFQ) provides bandwidth guarantees, so it can be used for voice, video, missioncritical, and lowerpriority traffic. However, CBWFQ does not provide the delay guarantees provided by LLQ, because CBWFQ does not provide support for strict priority queues. CBWFQ improves upon weighted fair queuing (WFQ) by enabling the creation of up to 64 custom traffic classes, each with a guaranteed minimum bandwidth.

Although WFQ can be used for voice, video, and missioncritical traffic, it does not provide the bandwidth guarantees or the strictpriority queues that are provided by LLQ. WFQ is used by default on Cisco routers for serial interfaces at 2.048 Mbps or lower. Traffic flows are identified by WFQ based on source and destination IP address, port number, protocol number, and Type of Service (ToS). Although WFQ is easy to configure, it is not supported on high speed links.

First-in-first-out (FIFO) queuing is the least appropriate for voice, video, and missioncriticaltraffic. By default, Cisco uses FIFO queuing for interfaces faster than 2.048 Mbps. FIFO queuing requires no configuration, because all packets are arranged into a single queue. As the name implies, the first packet received is the first packet transmitted, without regard for packet type, protocol, or priority.

Reference:

<https://www.cisco.com/c/en/us/support/docs/voice/voice-quality/7111-voip-mlppp.html>

QUESTION 103

You administer Cisco routers in a PIM-SSM environment. You issue the ip multicast multipath s-g-hash basic command on a router that has multiple equalcost paths to a multicast source.

Which of the following will occur? (Select the best answer.)



- A. The router will loadsplit based on the source address only.
- B. The router will loadsplit based on the source and group addresses.
- C. The router will loadsplit based on the source, group, and nexthop addresses.
- D. The router will send traffic to the PIM neighbor with the highest IP address.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The router will loadsplit based on the source and group addresses if you issue the ip multicast multipath sghash basic command on a router that has multiple equalcost paths to a multicast source. The ip multicast multipath sghash basic command uses the SGhash algorithm so that Protocol Independent Multicast SourceSpecific Multicast (PIMSSM) can reply to traffic by using either the Reverse Path Forwarding (RPF) interface or the source group address when equalcost paths exist. The basic SGhash configuration is subject to polarization because the hash is the same no matter which router calculates it.

The router will not send traffic to the Protocol Independent Multicast (PIM) neighbor with the highest IP address. By default, when equalcost paths to a multicast source exist, PIMSSM will send traffic to the neighbor with the highest IP address. However, issuing the ip multicast multipath command with or without keywords enables loadsplitting and disables the default behavior.

The router will not loadsplit based on the source address only. You should issue the ip multicast multipath command to enable Equal Cost Multipath (ECMP) loadsplitting based on only the source address. By issuing the ip multicast multipath command, you can configure PIMSSM to loadsplit traffic between equalcost paths by using the Shash algorithm, which selects the interface on which the traffic arrives as the interface on which to send a response.

The router will not loadsplit based on the source, group, and nexthop addresses. The ip multicast multipath sghash nexthopbased command uses the SGhash algorithm and configures PIMSSM to loadsplit traffic based on the source address, the group address, and the nexthop address. Unlike the ip multicast multipath sghash basiccommand, the nexthopbased SGhash configuration is not subject to polarization.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/ip_mcast/configuration/guide/mctlsplt.html#wp1061381

QUESTION 104

Which of the following statements best describes the purpose of the wellknown BGP AS 23456? (Select the best answer.)

- A. It facilitates the gradual transition from 4-byte ASes to 2-byte ASes.
- B. It facilitates the gradual transition from 2byte ASes to 4byte ASes.
- C. It is used by default on BGP routers that support 4byte ASes.
- D. It will be used by default on new BGP routers when 2byte allocations are exhausted.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The well-known Border Gateway Protocol (BGP) autonomous system (AS) 23456 facilitates the gradual transition from 2byte ASes to 4byte ASes. Similar to the creation of IPv6, which is intended to combat the threat of exhaustion of IPv4 addresses, 4byte BGP ASes were introduced to combat the eventual exhaustion of 2byte ASes. The well-known 2byte AS 23456, which is also known as AS_TRANS, can be used by a 4byte BGP router to peer with a BGP router that supports only 2byte ASes.

When a 4byte BGP router must advertise an AS value larger than 2 bytes to a 2byte BGP router, the 4byte router will advertise the AS number 23456. Therefore, if the AS number 23456 appears in the output of the show ip bgp command, the router is not compatible with 4byte ASes.

When a 4byte BGP router peers with another 4byte BGP router, the AS is displayed in the output of the show ip bgp command in as plain or as dot format. As plain format displays the 4byte AS number as a decimal value from 65536 through 4294967295; this format is used by default. As dot format displays the 4byte AS number as a dotted decimal value from 1.0 through 65535.65535. To change the output to display in as dot format, you should issue the bgp as notation dot command from BGP router configuration mode.

AS 23456 does not facilitate the gradual transition from 4byte ASes to 2byte ASes, because of the impending exhaustion of the pool of 2byte ASes. A 2byte BGP router cannot use AS 23456 as its AS number when transitioning from 2byte ASes to 4byte ASes, because AS 23456 is reserved exclusively for use by 4byte routers peering with 2byte routers. Additionally, AS 23456 is not used by default on any BGP router. It is possible to configure a 4byte BGP router to use a 2byte AS other than 23456 in order to peer with 2byte BGP routers.

Reference: https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/border-gateway-protocol-bgp/white_paper_c11-516826.html <https://www.ietf.org/rfc/rfc4893.txt>

QUESTION 105

DRAG DROP

Drag the metrics on the left to the corresponding PfR monitoring modes on the right. Factors can be used multiple times.

Select and Place:

delay	Active Monitoring Mode	Passive Monitoring Mode
jitter		
MOS		
packet loss		
reachability		
throughput		

Help

Reset

Done

Correct Answer:

	Active Monitoring Mode	Passive Monitoring Mode
delay	delay	delay
jitter	jitter	packet loss
MOS	MOS	reachability
packet loss	reachability	throughput
reachability		
throughput		

 **VCEplus**
 VCE To PDF - Free Practice Exam

Help
Reset
Done

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Cisco Performance Routing (PfR) enhances traditional routing methods by dynamically selecting the best path for applications based on network performance. The path selection procedure can be influenced by several factors, including delay, packet loss, reachability, throughput, jitter, and mean opinion score (MOS). PfR passive monitoring mode relies on NetFlow to capture performance metrics. Metrics used by passive mode include delay, packet loss, reachability, and throughput. Throughput can be measured for all traffic flows. Delay, packet loss, and reachability can be measured only for Transmission Control Protocol (TCP) flows.

PfR active monitoring mode relies on IP Service Level Agreement (SLA) probes that generate traffic to capture performance metrics. Metrics used by active mode include delay, jitter, MOS, and reachability. Short term monitoring uses the last five probe results; long term monitoring uses the last 60 probe results.

A third PfR monitoring mode, fast mode, is similar to active mode. Active mode generates probes only for the active exit path. By contrast, fast mode continuously generates probes for all possible exit paths, not just the active exit path. Fast mode allows route changes to be made within three seconds. However, the

performance benefits of fast mode require significant processor overhead? therefore, Cisco recommends that you use fast mode only for performance sensitive traffic, such as Voice over IP (VoIP) or video traffic.

Reference:

https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/performance-routing-pfr/product_data_sheet0900aecd806c4ee4.html

QUESTION 106

You issue the show pfr master border detail command on RouterA and receive the following output:

```
RouterA#show pfr master border detail
```

```
Border      Status    UP/DOWN   AuthFail  Version
10.20.30.1  ACTIVE   UP        11:22:19   0   3.0
  Et2/0      EXTERNAL UP
  Et0/0      INTERNAL UP
  Et1/0      EXTERNAL UP
```

```
External    Capacity  Max BW    BW Used    Load Status  Exit Id
Interface   (kbps)    (kbps)    (kbps)     (%)
-----
Et2/0      Tx        1200      900        578      48 UP        2
           Rx                1200      0         0
Et1/0      Tx        1200      900        386      32 UP        1
           Rx                1200      231      19
```

What is the link utilization OOP threshold for traffic exiting RouterA? (Select the best answer.)

- A. 19%
- B. 32%
- C. 48%
- D. 75%
- E. 90%
- F. 100%

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The link utilization outofpolicy (OOP) threshold for traffic exiting RouterA is 75%. Cisco Performance Routing (PfR) enhances traditional routing methods by dynamically selecting the best path for applications based on network performance. The show pfr master border detail command displays whether the PfR state is up or down. Additionally, the show pfr master border detail command displays statistics related to inbound and outbound traffic. The information next to Tx is related to outbound traffic, which is traffic that is transmitted by the router. The information next to Rx is related to inbound traffic, which is traffic that is received by the router.

If outbound traffic exceeds the link utilization OOP threshold, PfR will attempt to conform to policy levels by shifting traffic to other exit links. The link utilization OOP threshold for outbound traffic can be determined by issuing the show pfr master border detail command and dividing the maximum bandwidth by the capacity. In this scenario, the maximum bandwidth is 900 Kbps and the capacity is 1,200 Kbps? therefore, the link utilization OOP threshold is 75%. By default, the link utilization OOP threshold is set to 90%.

The link utilization OOP threshold is not 19%. The Ethernet1/0 interface is receiving traffic at 19% of the maximum bandwidth. The load percentage is calculated by dividing the bandwidth used by the maximum bandwidth.

The link utilization OOP threshold is not 32%. The Ethernet1/0 interface is sending traffic at 32% of the maximum bandwidth. Therefore, the traffic exiting the Ethernet1/0 interface is not exceeding the link utilization OOP threshold.

The link utilization OOP threshold is not 48%. The Ethernet2/0 interface is sending traffic at 48% of the maximum bandwidth. Therefore, the traffic exiting the Ethernet2/0 interface is not exceeding the link utilization OOP threshold.

Reference:

http://docwiki.cisco.com/wiki/PfR:Technology_Overview#Link_Utilization <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/pfr/command/pfr-cr-book/pfr-s1.html#wp3007750440>

QUESTION 107

Which of the following statements regarding traffic flooding in a VPLS network is true? (Select the best answer.)

- A. Only broadcast traffic is flooded through the network.
- B. Only multicast traffic is flooded through the network.
- C. Both broadcast traffic and multicast traffic are flooded through the network.
- D. Neither broadcast nor multicast traffic is flooded through the network.

Correct Answer: C



<https://vceplus.com/>

Section: (none)
Explanation

Explanation/Reference:

Explanation:

Broadcast traffic and multicast traffic are flooded through a Virtual Private LAN Service(VPLS) network. VPLS is a Metro Ethernet (ME) technology that is used to implement

Ethernet Multipoint Service (EMS) and Ethernet Relay Multipoint Service (ERMS) over a Multiprotocol Label Switching (MPLS) network. With VPLS, the service provider (SP) network emulates a single Layer 2 switch, or Ethernet bridge. Customer edge (CE) devices in a VPLS network connect directly to an SPprovided user provider edge (UPE) device, and the UPEs use virtual connections known as pseudowires (PWs) to interconnect through the SP network. The PWs form a fullmesh topology that creates a virtual switch instance (VSI), which emulates an 802.1 bridge.

The VPLS architecture creates a topology wherein each CE device can function as though it were a member of a virtual LAN (VLAN) on a physical switch. Because the SP network functions as a Layer 2 switch, broadcast and multicast packets received by the SP are always flooded through the network. Additionally, packets destined to unknown Media

Access Control (MAC) addresses are initially flooded through the network until their MAC addresses and associated ports are correlated.

Reference:

https://www.cisco.com/en/US/tech/tk436/tk891/technologies_g_and_a_item09186a00801ed3bf.shtml

QUESTION 108

Which of the following scenarios could cause the error message %STP-2-DISPUTE_DETECTED to appear? (Select the best answer.)

- A. An interface has been administratively shut down.
- B. A unidirectional link failure exists between two switches.
- C. A BPDU has been received on an interface with BPDU guard enabled.
- D. A BPDU has not been received on an interface with Bridge Assurance enabled.
- E. An interface has received a BPDU that is tagged with the same VLAN ID as the interface's native VLAN.

Correct Answer: B

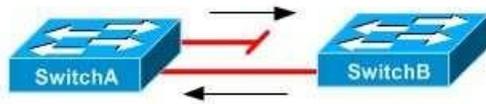
Section: (none)

Explanation

Explanation/Reference:

Explanation:

The error message %STP-2-DISPUTE_DETECTED could appear if a unidirectional link failure exists between two switches. A unidirectional link failure exists when a defective cable causes one device to not receive what the other device sends. Consider the following topology in which SwitchB does not receive packets sent by the root bridge, SwitchA:



SwitchB cannot receive the superior bridge protocol data units (BPDUs) sent by SwitchA. Therefore, SwitchB will continue to send inferior BPDUs that are marked as designated and learning or forwarding. SwitchA will record the %STP-2-DISPUTE_DETECTED error and shut down the interface to prevent a bridging loop. The error message %STP-2-DISPUTE_DETECTED would not appear if an interface has been administratively shut down. An interface that has been administratively shut down does not participate in Spanning Tree Protocol (STP).

The error message %STP-2-DISPUTE_DETECTED would not appear if a BPDU has been received on an interface with BPDU guard enabled; instead, the error message %STP-2-BLOCK_BPDU GUARD would appear. When an interface that is configured with BPDU guard receives a BPDU, BPDU guard immediately puts the interface into the errdisable state and shuts down the interface. Afterward, the interface must be manually reenabled, or it can be recovered automatically through the errdisable timeout function.

The error message %STP2DISPUTE_DETECTED would not appear if a BPDU has not been received on an interface with Bridge Assurance enabled? instead, the error message %STP2-

BRIDGE_ASSURANCE_BLOCK would appear. Bridge Assurance ensures that BPDUs are sent bidirectionally on all network interfaces. If an interface with Bridge Assurance does not receive a BPDU, or if the connected interface does not have Bridge Assurance enabled, the interface is put into an inconsistent state and is blocked. Bridge Assurance is supported only with Rapid PerVLAN Spanning Tree Plus (RPVST +) and Multiple Spanning Tree (MST) and only on point-to-point links.

The error message %STP-2-DISPUTE_DETECTED would not appear if an interface has received a BPDU that is tagged with the same virtual LAN (VLAN) ID as the interface's native VLAN? instead, the error message %STP-2-BLOCK_PVID_LOCAL would appear on the local switch and the error message %STP-2-BLOCK_PVID_PEER would appear on the remote switch. Native VLAN BPDUs are sent untagged, so if a switch receives BPDUs that are tagged with the native VLAN for that interface, the switch will block the interface.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/system_messages/reference/sys_Book.html#wp1400041

[https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_2/nx-](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_2/nx-os/layer2/configuration/guide/Cisco_Nexus_7000_Series_NXOS_Layer_2_Switching_Configuration_Guide_Release_4-2_chapter6.html#con_1490082)

[os/layer2/configuration/guide/Cisco Nexus 7000 Series NXOS Layer 2 Switching Configuration Guide Release 4-2 chapter6.html#con_1490082](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_2/nx-os/layer2/configuration/guide/Cisco_Nexus_7000_Series_NXOS_Layer_2_Switching_Configuration_Guide_Release_4-2_chapter6.html#con_1490082)

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/24063-pvid-inconsistency-24063.html#topic1>

QUESTION 109

Which of the following is not a feature of VTP version 3? (Select the best answer.)

- A. It can send PVLAN information in addition to normal VLANs.
- B. It can use encrypted passwords for authentication.
- C. It uses primary servers.
- D. VLANs must be in the range from 1 through 1,000.

E. VTP version 3 is backward compatible with VTP version 2.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

VLAN Trunking Protocol (VTP) version 3 does not require virtual LANs (VLANs) to be in the range from 1 through 1,000. VTP version 3 improves on VTP version 2 by increasing the number of supported VLANs from 1,000 to 4,095, which is the same range specified in the Institute of Electrical and Electronics Engineers (IEEE) 802.1Q standard.

VTP version 3 uses primary servers. VTP version 2 relies on a configuration revision number to determine whether the VLAN configuration should be modified on a switch. By contrast, VTP version 3 uses configuration revision numbers and a primary server system to determine which configurations should be changed and which devices are allowed to implement changes. The intended purpose of the primary server is to mitigate accidental overwrites of the VLAN database. However, because VTP version 2 does not support primary servers, Cisco recommends that VTP version 2 devices that are to be connected to a VTP version 3 network be placed into VTP client mode.

VTP version 3 can use encrypted passwords for authentication. VTP version 2 supports cleartext passwords only.

VTP version 3 is backward compatible with VTP version 2 and can therefore support normal VLANs. When a switch configured with VTP version 3 receives a VTP version 2 advertisement on a port, it sends VTP version 2 messages on that port and VTP version 3 messages on the other ports.

VTP version 3 improves on VTP version 2 by adding support for private VLANs (PVLANS). In addition, VTP version 3 adds support for databases other than VLAN databases, such as Multiple Spanning Tree (MST) databases.

Reference:

https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-6500-series-switches/solution_guide_c78_508010.html

QUESTION 110

Which of the following does not typically occur during the PPP Session stage of a PPPoE session? (Select the best answer.)

- A. The MAC address of the peer is obtained.
- B. LCP negotiates configuration options.
- C. NCP configures Network layer protocols.
- D. PPP authenticates by using CHAP or PAP.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Media Access Control (MAC) address of the peer is obtained during the Discovery stage of a Point-to-Point Protocol over Ethernet (PPPoE) session, not the Session stage. The Discovery stage is also sometimes called the Active Discovery stage. PPPoE sessions are divided into two distinct stages: the Discovery stage and the Session stage. Because an Ethernet host must first establish a connection to the remote peer before it can send data, the PPPoE Discovery stage must retrieve the MAC address of the remote peer and establish a Point-to-Point Protocol (PPP) session ID before establishing a PPP session.

Following the Discovery stage, the Session stage behaves mostly the same as a normal PPP session over a WAN link or dialup connection behaves. Therefore, the

Session stage is also sometimes called the PPP Session stage. During the Session stage, PPP negotiates configuration options by sending Link Control Protocol (LCP) frames. Next, PPP sends out Network Control Protocol (NCP) frames to configure Network layer protocol information on the link and enable the link for packet traversal. The Session stage is also the stage in which PPP authentication occurs by using either Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).

Reference:

<https://tools.ietf.org/html/rfc2516#section-5>

QUESTION 111

You issue the show ip ospf interface fastethernet 0/1 command on RouterE and receive the following output:

```
FastEthernet0/1 is up, line protocol is up
Internet Address 10.2.16.43/24, Area 0
Process ID 1, Router ID 10.0.0.4, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State DROTHER, Priority 50
Designated Router (ID) 10.0.0.7, Interface Address 10.2.16.1
Backup Designated router (ID) 10.0.0.11, Interface Address 10.2.16.17
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:04
Neighbor Count is 5, Adjacent neighbor count is 2
  Adjacent with neighbor 10.0.0.7 (Designated Router)
  Adjacent with neighbor 10.0.0.11 (Backup Designated Router)
Suppress hello for 0 neighbor(s)
```

Which of the following statements is correct? (Select the best answer.)

- A. RouterE is connected to a point-to-multipoint network.
- B. RouterE is the DR for the segment.
- C. The BDR has a priority higher than 50.
- D. RouterE is configured with incorrect timer settings.
- E. RouterE can establish adjacencies with only two routers on this interface.

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RouterE can establish adjacencies with only two routers on this interface. The output of the show ip ospf interface fastethernet 0/1 command shows that RouterE is in the DROTHER state. A router in the DROTHER state can only establish adjacencies with the designated router (DR) and the backup designated router (BDR). Therefore, RouterE is neither the DR nor the BDR. The DR has a router ID of 10.0.0.7 and an IP address of 10.2.16.1, and the BDR has a router ID of 10.0.0.11 and an IP address of 10.2.16.17.

RouterE is not connected to a point-to-multipoint network, because the network segment contains a DR and a BDR. A DR and a BDR are not elected on point-to-multipoint or point-to-point networks; they are elected only on multi-access networks.

The BDR might or might not have a priority higher than 50. If RouterE were started before the DR and BDR were elected, RouterE would not be eligible to become the DR or the BDR, regardless of the priority value of RouterE, until the existing DR and BDR failed or were powered off. If RouterE were started at the same time as the existing DR and BDR, the BDR would have a priority of at least 50 because RouterE has a priority of 50. If the BDR and RouterE have the same priority, the BDR will be elected before RouterE because it has a higher router ID than RouterE. RouterE is not configured with incorrect timer settings. The hello timers and dead timers between two routers must match for the routers to establish a neighbor adjacency. Therefore, if RouterE were configured with incorrect timer settings, RouterE would not be able to establish adjacencies with the DR and the BDR. By default, the hello timer is set to 10 seconds and the dead timer is set to 40 seconds on point-to-point and broadcast links.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13689-17.html>

QUESTION 112

Which of the following LSA types contains subnet and router information for all the routers on a segment? (Select the best answer.)

- A. Type 1
- B. Type 2
- C. Type 3
- D. Type 5

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A Type 2 link-state advertisement (LSA) contains subnet and router information for all the routers on a segment. Type 2 LSAs, which are also called network LSAs, are generated by only the designated router (DR) to each of the segments connected to the DR. These LSAs are not propagated outside the area in which they originate; they are flooded only within the local area.

A Type 1 LSA contains router ID and interface IP address information for a single router. Type 1 LSAs, which are also called router LSAs, are generated by all Open Shortest Path First (OSPF) routers on a segment. Like Type 2 LSAs, Type 1 LSAs are not propagated outside the area in which they originate; they are flooded only within the local area.

A Type 3 LSA contains subnet information for an entire area. Type 3 LSAs, which are also called network summary LSAs, are generated by area border routers (ABRs). Unlike Type 1 and Type 2 LSAs, Type 3 LSAs are advertised between areas throughout an autonomous system (AS) except into totally stubby areas. A Type 5 LSA contains subnet information for an external AS. Type 5 LSAs, which are also called AS external LSAs, are generated by autonomous system boundary routers (ASBRs). Therefore, Type 5 LSAs are advertised throughout an AS except into stub areas, totally stubby areas, and not so stubby areas (NSSAs).

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_2/nx-os/unicast/configuration/guide/l3_cli_nxos/l3_ospf.html#pgfId-1243056

https://www.cisco.com/c/en/us/td/docs/security/fwsm/fwsm32/asdm52f/user/guide/asdmug/mon_rtg.html#wp1046958

QUESTION 113

Which of the following is true of NAT64? (Select the best answer.)

- A. It uses DNSALG for name resolution.
- B. It can be deployed in a stateful configuration.
- C. It translates unique IPv6 prefixes to other unique IPv6 prefixes.
- D. It cannot be deployed in a stateless configuration.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Network Address Translation 64 (NAT64), which is typically used to enable communication between IPv6-only hosts and IPv4-only hosts, can be deployed in a stateful configuration. When configured as stateful, NAT64 maps multiple IPv6 addresses to a single IPv4 address and keeps track of the state of each connection. Static mappings can also be applied manually.

NAT64 can also be deployed in a stateless configuration. When operating in a stateless configuration, NAT64 uses algorithms to create a one-to-one relationship between IPv6 addresses on the inside network and IPv4 addresses on the outside network. Although this technique preserves end-to-end connectivity at the Network layer, it does not conserve IPv4 addresses the way a stateful many-to-one configuration does.

NAT-Protocol Translation (NATPT), not NAT64, uses Domain Name System Application Level Gateway (DNSALG) for name resolution. NATPT is a predecessor of NAT64; it supports bidirectional translation of addresses between IPv6 and IPv4 networks. In NATPT, the DNSALG function is included, along with the address

family translation (AFT) function, as part of the protocol. NAT64, on the other hand, is typically deployed alongside an independent Domain Name System (DNS) solution, such as DNS64, to facilitate name resolution.

Network Prefix Translation version 6 (NPTv6), not NAT64, translates unique IPv6 prefixes to other unique IPv6 prefixes. As the name implies, NPTv6 enables the stateless translation of inside IPv6 prefixes to outside IPv6 prefixes at the Internet edge. NPTv6 creates a one-to-one relationship between addresses on each side of the translating device in order to maintain end-to-end reachability at the Network layer. NPTv6 does not modify the interface identifier portion of an IPv6 address.

Reference:

https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/enterprise-ipv6-solution/white_paper_c11-676277.html

QUESTION 114

You issue the ip rip triggered command on the Serial 0 interface of RouterA.

When will RouterA send a partial routing database? (Select 2 choices.)

- A. when RouterA is first powered on
- B. when the Serial 0 interface comes up or goes down
- C. when RouterA receives a specific request for a routing table update
- D. when information from another interface modifies the database

Correct Answer: BD

Section: (none)

Explanation



Explanation/Reference:

Explanation:

RouterA will send a partial routing database when information from another interface modifies the database and when the Serial 0 interface comes up or goes down. By default, routes are advertised by Routing Information Protocol (RIP) every 30 seconds, not just when the routing database is updated. However, you can configure RIP to send triggered advertisements by issuing the ip rip triggered command from interface configuration mode. The following events will trigger a partial database update:

- When the configured interface comes up or goes down
- When information from another interface modifies the routing table

The following events will trigger a full database update:

- When the router is first powered on
- When the router receives a specific request for a routing table update

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_rip/command/irr-cr-book/irr-cr-rip.html#wp3830876380

QUESTION 115

Which of the following statements describe how LACP determines whether an interface is an active interface or a standby interface? (Select 2 choices.)

- A. Ports with higher LACP port priorities are preferred over ports with lower LACP port priorities.
- B. Ports with lower LACP port priorities are preferred over ports with higher LACP port priorities.
- C. When LACP port priorities are equal, ports with higher port numbers are preferred over ports with lower port numbers.
- D. When LACP port priorities are equal, ports with lower port numbers are preferred over ports with higher port numbers.
- E. When LACP port priorities are equal, all ports with equal priorities are configured as active interfaces.

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following statements describe how Link Aggregation Control Protocol (LACP) determines whether an interface is an active interface or a standby interface:

- Ports with higher LACP port priorities are preferred over ports with lower LACP port priorities.
- When LACP port priorities are equal, ports with lower port numbers are preferred over ports with higher port numbers.

The `lacp port-priority` value command configures an LACP interface with a port priority, which is used to determine which interfaces are active interfaces and which interfaces are standby interfaces. The value parameter is a value from 1 through 65535; if no priority value is defined, the default port priority value of 32768 is used.

Setting the LACP port priority to a value of 1 will ensure that the port becomes an active port unless a port with a lower port number is also set to a priority of 1.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2sb/feature/guide/gigeth.html

QUESTION 116

Which of the following steps in the NAT order of operation typically occur before inside-to-outside translation but after outside-to-inside address translation? (Select the best answer.)

- A. checking inbound access lists, rate limits, and accounting
- B. policy routing and IP routing
- C. checking outbound access lists and inspecting CBAC
- D. encryption and queuing

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following steps of the Network Address Translation (NAT) order of operation typically occur before inside-to-outside translation but after outside-to-inside translation: - Policy routing

- IP routing

NAT enables a network to communicate with a separate network, such as the Internet, by translating traffic from IP addresses on the local network to another set of IP addresses that can communicate with the remote network. NAT inside-to-outside translation, which is also known as local-to-global translation, occurs when the NAT router maps an inside network source IP address to an outside network source IP address before forwarding the packet to the next hop. When a NAT router performs NAT inside-to-outside translation, the following operations occur in order:

1. If IP Security (IPSec) is implemented, check inbound access list
2. Decryption
3. Check inbound access list
4. Check inbound rate limits
5. Inbound accounting
6. Redirect to web cache
7. Policy routing
8. IP routing
9. NAT inside-to-outside translation
10. Check crypto map and mark for encryption
11. Check outbound access list
12. Inspect Context-based Access Control (CBAC)
13. Transmission Control Protocol (TCP) intercept
14. Encryption
15. Queueing



NAT outside-to-inside translation, which is also known as global-to-local translation, occurs when the NAT router maps an outside destination IP address to an inside destination IP address. When a NAT router performs NAT outside-to-inside translation, the following operations occur in order:

1. If IPSec is implemented, check inbound access list
2. Decryption
3. Check inbound access list
4. Check inbound rate limits
5. Inbound accounting
6. Redirect to web cache
7. NAT outside-to-inside translation
8. Policy routing

9. IP routing
10. Check crypto map and mark for encryption
11. Check outbound access list
12. Inspect CBAC
13. TCP intercept
14. Encryption
15. Queueing

Other than the policy routing and IP routing steps, the other steps in the NAT order of operation are the same for insidetooutside NAT and outsidetoinside NAT. Checking inbound access lists, rate limits, and accounting are performed before insidetooutside address translation and before outsidetoinside address translation. Checking outbound access lists, inspecting CBAC, encryption, and queueing are performed after insidetooutside address translation and after outsidetoinside address translation.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/network-address-translation-nat/6209-5.html#topic1>

QUESTION 117

Which of the following first-hop routing protocols can have up to four primary AVFs provide load balancing across multiple WAN links? (Select the best answer.)

- A. GLBP
- B. HSRP
- C. VRRP
- D. GLBP and HSRP
- E. HSRP and VRRP
- F. GLBP, HSRP, and VRRP

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Gateway Load Balancing Protocol (GLBP) can have up to four primary active virtual forwarders (AVFs) provide load balancing across multiple WAN links. GLBP is a Cisco proprietary First Hop Redundancy Protocol (FHRP) that enables up to four routers to act as a single virtual router. The virtual router has its own virtual IP address and up to four virtual Media Access Control (MAC) addresses, one for each of the four primary AVFs in the group. One of the routers in the GLBP group is elected the active virtual gateway (AVG) and performs the administrative tasks for the standby group, such as responding to Address Resolution Protocol (ARP)

requests. When a client sends an ARP request for the IP address of the default gateway, the AVG responds with one of the virtual MAC addresses in the group. Because multiple routers in the GLBP group can actively forward traffic, GLBP provides load balancing as well as local redundancy.

Additionally, you can control the percentage of traffic that is sent to a specific gateway by configuring weighted load balancing. By default, GLBP uses a round-robin technique to load balance between routers. If you configure weighted load balancing, GLBP can send a higher percentage of traffic to a single GLBP group member based on the weight values assigned to the interfaces of that member.

Hot Standby Router Protocol (HSRP) is a Cisco proprietary protocol that enables two or more routers to act as a single virtual router. However, only one router in the HSRP standby group forwards traffic for the group. Because there is only one AVF in an HSRP standby group, HSRP cannot provide load balancing across multiple WAN links.

Virtual Router Redundancy Protocol (VRRP) is an open-standard protocol that is similar to HSRP. VRRP enables two or more routers to act as a single virtual router, but it does not enable more than a single router to act as an AVF. Therefore, VRRP cannot provide load balancing across multiple WAN links.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp_fhpr/configuration/15-mt/fhp-15-mt-book/fhp-qlbp.html#GUID-26C72408-6183-415A-9949-8B97542246A9

QUESTION 118

How often is an IP SLA operation repeated if a frequency is not configured? (Select the best answer.)

- A. every five seconds
- B. every 60 seconds
- C. every 300 seconds
- D. IP SLA operations are not repeated if a frequency is not configured.



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

An IP Service Level Agreement (SLA) operation is repeated every 60 seconds if a frequency is not configured. IP SLA operations are a suite of tools on Cisco devices that enable an administrator to analyze and troubleshoot IP networks. For example, the following command set configures IP SLA to regularly test and verify the reachability of IP address 10.10.10.2:

```
ip sla 1
  type echo protocol ipIcmpEcho 10.10.10.2
  timeout 1000
  threshold 2
  frequency 3
ip sla schedule 1 life forever start-time now
```

To change how often an IP SLA operation is repeated, you should issue the frequency command from an IP SLA configuration submode. The variable for the frequency command is specified in seconds? therefore, the frequency 3 command specifies that the operation should repeat every three seconds. The frequency 60 command has the same effect as the default frequency of 60 seconds.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/command/sla-cr-book/sla_a1.html#wp4022386755

QUESTION 119

Which of the following statements are correct regarding 802.1X portbased authentication? (Select 3 choices.)

- A. Before authentication occurs, only DHCP traffic is allowed through a port that is configured for 802.1X authentication.
- B. Before authentication occurs, only EAPOL, STP, and CDP traffic is allowed on a port that is configured for 802.1X authentication.
- C. If a host is configured to use 802.1X but a switch is not, the host will be unable to communicate on the network.
- D. If a switch is configured to use 802.1X but a host is not, the host will be unable to communicate on the network.
- E. Multiple hosts can be connected to a port that is configured for 802.1X authentication.
- F. Only one host can be connected to a port that is configured for 802.1X authentication.

Correct Answer: BDE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Of the available choices, the following statements are correct regarding 802.1X portbased authentication:

-Before authentication occurs, only Extensible Authentication Protocol over LANs (EAPOL), SpanningTree Protocol (STP), and Cisco Discovery Protocol (CDP) traffic is allowed on a port that is configured for 802.1X authentication.

-If a switch is configured to use 802.1X but a host is not, no communication will take place. -

Multiple hosts can be connected to a port that is configured for 802.1X authentication.

Port-based authentication that uses the Institute of Electrical and Electronics Engineers (IEEE) 802.1X standard can be used on Cisco switches to ensure that only authenticated users are able to send traffic through the switch. Before authentication occurs, the only traffic that the port allows is EAPOL traffic, STP traffic, and CDP traffic. This ensures that a host connected to the port is authenticated before any other traffic is allowed through the port. The use of 802.1X authentication requires that both the host and the switch be configured for 802.1X. If the host is configured for 802.1X but the switch is not, the host can communicate with the switch but 802.1X authentication will not be used. However, if the switch is configured for 802.1X but the host is not, the host will be unable to send traffic through the switch? the port will remain in the unauthorized state.

Dynamic Host Configuration Protocol (DHCP) traffic is not allowed through a port that is configured for 802.1X authentication before authentication occurs. A host connected to a switch port that is configured for 802.1X authentication can only communicate with the switch in order to authenticate with the switch. After authentication occurs, the host can request an IP address from a DHCP server.

You can connect more than one host to a port that is configured for 802.1X authentication. For example, if multiple hosts are connected to a hub or a switch, you can connect the hub or switch to a port that is configured for 802.1X authentication. To configure the port to accept connections from multiple hosts, you should issue the dot1x host-mode multi-host command on the interface to which the hub or switch will be connected.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/12_4t11/ht_8021x.html#wp1025060

QUESTION 120

You issue the show runningconfig command on a Cisco 3700 series router and receive the following partial output:

```
class-map match-any boson
  match ip precedence 3
  match ip precedence 4
class-map match-any exsim
  match ip precedence 2
```

```
policy-map applications
  class boson
    priority 32
  class exsim
    bandwidth 16
  class class-default
    fair-queue
    random-detect
```



Which of the following classes use FIFO queuing within the class? (Select the best answer.)

- A. only the boson class
- B. only the boson and exsim classes
- C. only the exsim and classdefault classes
- D. the boson, exsim, and classdefault classes

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The boson and exsim classes use firstinfirstout (FIFO) queuing within the class. In this scenario, classbased weighted fair queuing (CBWFQ) and low latency queuing (LLQ) are used, as indicated by the priority and bandwidth commands. However, only the classdefault class can use weighted fair queuing (WFQ); the other classes can use only FIFO queuing within the class. A traffic class may be prioritized over other traffic classes, but traffic within that class is processed in the order the traffic is received, without regard for packet type, protocol, or IP precedence. For example, traffic with an IP precedence value of 3 might be transmitted ahead of traffic with an IP precedence value of 4 even though its precedence value is lower.

Although the classdefault class can use FIFO queuing, it is currently configured to useWFQ. The fairqueue command configures the classdefault class to use WFQ queuing. To configure the classdefault class to use FIFO queuing, you should issue the no fair-queue command.

Reference:

<https://www.cisco.com/c/en/us/support/docs/quality-of-service-qos/qos-policing/22833-qos-faq.html#qc>

Cisco: QoS Frequently Asked Questions: Queueing and Congestion Management

QUESTION 121

You issue the activation character 8 command from line configuration mode.

Which of the following statements is true? (Select the best answer.)



<https://vceplus.com/>

- A. Terminal access will not begin until the 8 key is pressed.
- B. Terminal access will not begin until the Backspace key is pressed.
- C. Terminal access will not begin until the Enter key is pressed.
- D. Terminal access will not begin until eight keys have been pressed.
- E. Terminal access must begin within eight seconds, or the session will be disconnected.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Terminal access will not begin until the Backspace key is pressed. The activation character command can be issued from line configuration mode to help secure access to the console (CON) port, auxiliary (AUX) port, and virtual terminal (VTY) ports. This method should not be used in lieu of implementing strong password security, but it can help to dissuade casual intruders looking for easy device access.

The variable for the activation character command is the decimal value for the ASCII character. By default, the value 13 is used for the activation character, which corresponds to the Enter or Return key. The decimal value 8 corresponds to the Backspace key, not the 8 key? the decimal value 56 corresponds to the 8 key. The `exec-timeout` command is used to configure a terminal line with an idle timeout. If no input is detected on the line within the idle timeout period, the session is disconnected. The syntax of the `exec-timeout` command is `exec-timeout minutes [seconds]`.

Therefore, the `exec-timeout 1 30` command configures the router to automatically disconnect idle sessions after one minute and 30 seconds, which is equal to 90 seconds. If the `exec-timeout` command has not been issued, an idle session will remain established for 10 minutes without input. Issuing the `no exec-timeout` command or the `exec-timeout 0 0` command causes a session to never time out due to inactivity.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fundamentals/command/cf_command_ref/A_through_B.html#wp1723943209

QUESTION 122

Which of the following commands should you issue to ensure that the enable password will be used if a RADIUS server is unavailable? (Select the best answer.)

- A. `aaa accounting exec enable start-stop group radius`
- B. `aaa accounting connection default start-stop group radius`
- C. `aaa authorization exec default group radius local`
- D. `aaa authorization exec default group radius if-authenticated`
- E. `aaa authentication enable default group radius enable`
- F. `aaa authentication login default local`



Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the `aaa authentication enable default group radius enable` command to ensure that the enable password will be used if a Remote Authentication Dial In User Service (RADIUS) server is unavailable. Authentication, Authorization, and Accounting (AAA) is used to control access to a router or switch. When implementing AAA, you can configure users to be authenticated against a local database, against a RADIUS server, or against a Terminal Access Controller Access Control System Plus (TACACS+) server. For AAA authentication to be used with a RADIUS server, a RADIUS server must exist on the network. However, you can configure a router so that if a RADIUS server becomes unavailable, the enable password can be used for authentication. This is accomplished by issuing the `aaa authentication enable default group radius enable` command. The `aaa authentication` command can be used to configure AAA authentication on a router or a switch. The first `enable` parameter specifies that the command applies to the enable mode. The default keyword specifies that the default authentication list

should be used. The group radius keywords specify that the RADIUS server should be used. The final enable keyword specifies that if the RADIUS server is unavailable, the enable password should be used.

The aaa authentication login default local command is used to configure AAA authentication to use the local database for authentication purposes. This command does not ensure that the enable password will be used if a RADIUS server is unavailable.

The aaa accounting command is used to enable AAA accounting on a router. The syntax of the aaa accounting command is aaa accounting {authproxy | system | network | exec | connection | commandslevel} {default | listname} [vrfvrfname] {startstop | stoponly | none} [broadcast] group groupname. Although the aaa accounting exec enable startstop group radius command and the aaa accounting connection default startstop group radius command are valid IOS commands, they do not ensure that the enable password will be used if a RADIUS server is unavailable.

Instead, these commands configure AAA accounting with the defined parameters.

The aaa authorization command is used to configure AAA authorization on a router. The syntax of the aaa authorization command is aaa authorization {network | exec | commandslevel | reverseaccess | configuration} {default | listname} method1[method2...]. Although the aaa authorization exec default group radius local command and the aaa authorization exec default group radius if-authenticated command are valid IOS commands, they do not ensure that the enable password will be used if a RADIUS server is unavailable. Instead, these commands configure AAA authorization with the parameters defined.

Reference:

<https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/10384-security.html#ex3>

QUESTION 123

When creating a VPN tunnel, on which of the following devices should you issue the tunnel mode auto command? (Select the best answer.)

- A. on the responder only
- B. on the initiator only
- C. on both the responder and the initiator
- D. on neither the responder nor the initiator

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the tunnel mode auto command on the responder only. The tunnel mode auto command enables the Tunnel Mode Auto Selection feature, which simplifies the configuration of a virtual private network (VPN) tunnel. When Tunnel Mode Auto Selection is configured, the responder will apply the tunneling protocol and transport protocol that is established by the initiator. Tunneling protocols include Generic Routing Encapsulation (GRE) and IP Security (IPSec). Transport protocols include IPv4 and IPv6.

You should not issue the tunnel mode auto command on the initiator. The tunnel configuration parameters must be statically configured on an initiator.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_vpnips/configuration/xs-3s/sec-sec-for-vpns-w-ipsec-xe-3s-book/sec-ipsec-virtunnl.html#concept_D55B0B7783A441BBB576E9F85693DF39 <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/s1/sec-s1-cr-book/sec-cr-t2.html#wp3291311677>

QUESTION 124

Which of the following seed metrics is assigned by default when OSPF routes are redistributed into EIGRP? (Select the best answer.)

- A. 0
- B. 1
- C. 20
- D. infinity
- E. the metric used by the OSPF route

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A default seed metric with the value of infinity is assigned to Open Shortest Path First (OSPF) routes that are redistributed into Enhanced Interior Gateway Routing Protocol (EIGRP). Routes with an infinite metric are ignored by EIGRP and are not entered into the routing table. There is no direct translation of the OSPF cost-based metric into an EIGRP-equivalent metric; the EIGRP metric is based on bandwidth, delay, reliability, and load. Because the OSPF metric cannot be automatically converted into a metric that EIGRP understands, EIGRP requires that the metric be defined for all redistributed routes before those routes are entered into the routing table. To assign a default metric for routes redistributed into EIGRP, you should issue the `default-metric bandwidth delay reliability loading mtu` command. To assign a metric to an individual route redistributed from OSPF into EIGRP, you should issue the `redistribute ospf process-id metric bandwidth delay reliability loading mtu` command.

A default seed metric of infinity is also assigned to routes that are redistributed into Routing Information Protocol (RIP). Like EIGRP, RIP requires that the metric be defined for all redistributed routes before those routes are entered into the routing table. RIP uses hop count as a metric. Valid hopcount values are from 1 through 15; a value of 16 is considered to be infinite. The hopcount metric increases by 1 for each router along the path. Cisco recommends that you set a low value for the hopcount metric for redistributed routes. To assign a default metric for routes redistributed into RIP, you should issue the `default-metric hopcount` command. To assign a metric to an individual route redistributed into RIP, you should issue the `redistribute protocol hopcount` command. If no metric is assigned during redistribution and no default metric is configured for RIP, the routes are assigned an infinite metric and are ignored by RIP.

A default seed metric of 0 is assigned to routes that are redistributed into Intermediate System to Intermediate System (ISIS). ISIS uses a cost metric assigned to each participating interface. ISIS prefers routes with the lowest cost. Routes redistributed into IS-IS are designated as Level 2 routes unless otherwise specified.

A default seed metric of 1 is assigned to Border Gateway Protocol (BGP) routes that are redistributed into OSPF. OSPF uses a cost metric based on the bandwidth of each participating interface and prefers internal routes with the lowest cost. By default, all routes redistributed into OSPF are designated as Type 2 external (E2)

routes. E2 routes have a metric that remains constant throughout the routing domain. Alternatively, routes redistributed into OSPF can be designated as Type 1 external (E1) routes. With E1 routes, the internal cost of the route is added to the initial metric assigned during redistribution.

A default seed metric of 20 is assigned to routes that are redistributed into OSPF from an internal gateway protocol other than OSPF. When OSPF routes are redistributed from one OSPF routing process to another OSPF routing process, the metrics are preserved and no default seed metric is assigned. Metrics are also preserved when routes are redistributed from one Interior Gateway Routing Protocol (IGRP) or EIGRP routing process into another IGRP or EIGRP routing process.

Reference:

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

<https://www.cisco.com/networkers/nw04/presos/docs/CERT-2100.pdf#page=6>

QUESTION 125

What is the size of the IPv6 fragment header? (Select the best answer.)

- A. 32 bits
- B. 64 bits
- C. 20 bytes
- D. 40 bytes
- E. 1,280 bytes

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The IPv6 fragment header is 64 bits long. The fragment header is used by an IPv6 source to indicate a packet that exceeds the path maximum transmission unit (MTU) size. Unlike IPv4, which enables intervening devices such as routers to fragment packets that exceed the permitted size for a local link, IPv6 requires the traffic originator to ensure that each packet sent is small enough to traverse the entire link without fragmentation. The packet can then be reassembled at the destination.

The IPv6 fragment header is not 32 bits long. However, the IPv6 fragment header contains a 32bit field called the identification field. The identification field is used to uniquely identify each fragmented packet.

The IPv6 fragment header is neither 20 bytes nor 40 bytes long. A basic IPv4 header without options is 20 bytes long, and a basic IPv6 header without extension headers is 40 bytes long. Although an IPv4 header is shorter than an IPv6 header, it is more complex and contains more fields than an IPv6 header. Several fields that exist in an IPv4 header, such as the Header Checksum field and the Fragment Offset field, do not exist in an IPv6 header. Because many protocols at the Data Link and Transport layers contain mechanisms to verify the integrity of the packet, IPv6 does not contain a redundant method to calculate checksum values.

The IPv6 fragment header is not 1,280 bytes long. The default IPv6 MTU size is 1,280 bytes. IPv6 requires that each device have an MTU of 1,280 bytes or greater.

Reference:

<https://www.ietf.org/rfc/rfc2460.txt>

QUESTION 126

Which of the following mutual redistribution scenarios does not require you to configure manual redistribution? (Select the best answer.)

- A. static routes and RIPv2
- B. static routes and EIGRP
- C. OSPF processes with different process IDs
- D. IS-IS and OSPF processes with the same area number
- E. IGRP and EIGRP processes with the same ASN
- F. EIGRP processes with different ASNs

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Interior Gateway Routing Protocol (IGRP) processes and Enhanced IGRP (EIGRP) processes with the same autonomous system number (ASN) do not require manual redistribution. Mutual redistribution of IGRP and EIGRP routing processes occurs automatically if the processes share the same ASN; there is no additional configuration required to enable route redistribution between the IGRP and EIGRP processes. However, you must manually configure route redistribution between IGRP and EIGRP processes with different ASNs.

Routing Information Protocol version 2 (RIPv2) automatically redistributes static routes that point to an interface on the router. However, RIP does not redistribute static routes that point to a nexthop IP address unless you issue the redistribute static command from RIP router configuration mode. RIPv2 assigns static routes a metric of 1 and redistributes them as though they were directly connected. Because there is only one routing protocol involved when static routes are redistributed into a RIPv2 routing domain, this is a one way redistribution of routing information. EIGRP automatically redistributes static routes that point to an interface on the router.

However, EIGRP does not redistribute static routes that point to a nexthop IP address unless you issue the redistribute static command from EIGRP router configuration mode. The static route is redistributed as an external route. Because there is only one routing protocol involved when static routes are redistributed into an EIGRP routing domain, this is a oneway redistribution of routing information.

Open Shortest Path First (OSPF) processes with different process IDs do not redistribute routes without manual configuration. Although it is possible to run multiple OSPF processes on a single router, it is not recommended, because suboptimal routing and routing loops may occur.

Intermediate System-to-Intermediate System (IS-IS) and OSPF processes with the same area number do not redistribute routes without manual configuration. ISIS and OSPF both assign a default metric to redistributed routes unless otherwise specified.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html#sameauto>
<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html>

QUESTION 127

Which of the following Cisco IOS XE subpackages is always different among consolidated packages? (Select the best answer.)

- A. RPAccess
- B. RPBase
- C. RPControl
- D. RPIOS
- E. ESPBase
- F. SIPBase
- G. SIPSPA

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The RPIOS Cisco IOS XE subpackage is always different among consolidated packages. A consolidated package is an image that contains multiple subpackage files. Every consolidated package will contain the following subpackages:

- RPAccess -provides router access software, either with or without cryptologic support
- RPBase -provides the operating system software for the route processor
- RPControl -provides the control plane interface between the IOS software and the platform
- RPIOS -provides the IOS kernel, which stores and runs IOS software features
- ESPBase -provides the Embedded Service Processor (ESP) operating system and control processes
- SIPBase -controls the Session Initiation Protocol (SIP) operating system and control processes
- SIPSPA -provides the shared port adaptor (SPA) driver and field-programmable device (FPD) images

Of these subpackages, only the RPIOS subpackage is always different among consolidated packages. The RPBase, RPControl, ESPBase, SIPBase, and SIPSPA subpackages are always the same regardless of the consolidated package. There are two different versions of RPAccess: a K9 version, which includes cryptographic support, and a nonK9 version, which does not include cryptographic support.

Optional subpackages are also available. However, optional subpackages are not contained within consolidated packages; they must be downloaded directly from Cisco.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/ios_xe/2/release/notes/rnasr21/rnasr21_gen.html#wp2996921

QUESTION 128

Which of the following Cisco Performance Monitor commands is not issued from global configuration mode? (Select the best answer.)

- A. class-map
- B. flow monitor type performance-monitor
- C. flow record type performance-monitor
- D. policy-map type performance-monitor
- E. service-policy type performance-monitor

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The servicepolicy type performancemonitor command is not issued from global configuration mode; it is issued from interface configuration mode. Cisco Performance Monitor enables you to monitor traffic flow information, such as packet count, byte count, drops, jitter, and roundtrip time (RTT). To configure Cisco Performance Monitor, you must perform the following tasks:

1. Create a flow record.
2. Configure a flow monitor.
3. Create one or more classes.
4. Create a policy.
5. Associate the policy with an interface.

First, create a Performance Monitor flow record by issuing the flow record typeperformancemonitor command from global configuration mode. The flow record is used to specify the data that will be collected. To configure the flow record, issue match and collect commands.

Next, configure a Performance Monitor flow monitor by issuing the flow monitor type performancemonitor command from global configuration mode. The flow monitor allows you to associate a flow record with a flow exporter. A flow exporter is used to send Performance Monitor data to a remote system.

Third, create one or more classes by issuing the classmap command from global configuration mode. A Performance Monitor class map is configured like any other class map by issuing match statements to specify the classification criteria.

Fourth, create a Performance Monitor policy by issuing the policymap type performancemonitor command from global configuration mode. A Performance Monitor policy associates a class with a flow monitor.

Finally, associate the Performance Monitor policy with an interface by issuing the servicepolicy type performancemonitor command from interface configuration mode. Issuing this command activates the Performance Monitor policy.

Reference:

<https://search.cisco.com/search?query=Cisco%20IOS%20Media%20Monitoring%20Configuration%20Guide&locale=enUS&tab=Cisco>

QUESTION 129

You issue the ping mpls ipv4 command.

Which of the following return codes indicates a successful ping? (Select the best answer.)

- A. 0
- B. 1
- C. 2
- D. 3
- E. 4

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A return code of 3 indicates a successful ping. The ping mpls ipv4 command can be used to verify Multiprotocol Label Switching (MPLS) label switched path (LSP) connectivity. The output will display symbolic and numeric return codes. The following list contains all of the numeric return codes along with their corresponding symbols and definitions:

Return Code	Symbol	Definition
0	x	No return code
1	M	Malformed echo request
2	m	Unsupported Type-Length-Values (TLVs)
3	!	Proper egress for the Forwarding Equivalence Class (FEC)
4	F	No FEC mapping
5	D	Downstream mapping mismatch
6	I	Unknown upstream interface index
7	U	Reserved
8	L	Labeled output interface
9	B	Unlabeled output interface
10	f	FEC mismatch
11	N	No label entry
12	P	Protocol not associated with interface
13	p	Premature termination of LSP
unknown	X	Undefined return code

A successful MPLS ping will look similar to the following output:

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'X' - unknown return code, 'x' - return code 0
Type escape sequence to abort.
! size 100, reply addr 10.1.2.3, return code 3
! size 100, reply addr 10.1.2.3, return code 3
! size 100, reply addr 10.1.2.3, return code 3
! size 100, reply addr 10.1.2.3, return code 3
! size 100, reply addr 10.1.2.3, return code 3
Success rate is 100 percent (5/5), round-trip min/avg/max = 91/94/96 ms
```

Reference:

<https://www.ietf.org/rfc/rfc4379.txt>

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/12_4t11/ht_lspng.html#wp1054221 https://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-2/mpls/command/reference/b_mpls_cr42crs/b_mpls_cr42crs_chapter_0110.html#wp2685217708

QUESTION 130

DRAG DROP Select the components on the left that create a complete multiprotocol BGP VPN-IPv4 address. Place the components in the order they appear in the address.

Select and Place:

The interface shows a list of components on the left: Administrator field, Assigned Number field, ASN, IPv4 address, RD, Type field, and Value field. On the right, there is a yellow address field divided into two segments. Below the list are 'Help', 'Reset', and 'Done' buttons.

Correct Answer:

The correct answer shows the 'RD' component placed in the first segment of the address field and the 'IPv4 address' component placed in the second segment. The other components remain in the list on the left. Below the list are 'Help', 'Reset', and 'Done' buttons.

Section: (none)
Explanation

Explanation/Reference:

Explanation:

A multiprotocol Border Gateway Protocol (BGP) virtual private network (VPN) IPv4 address begins with an 8byte route distinguisher (RD) and ends with a 4byte IPv4 address. The RD consists of a 2byte Type field and a 6byte Value field. The value of the Type field determines what the structure of the Value field is. The following table lists the Type values along with their corresponding Value field structures:

Type Field	Value Field
0	2-byte Administrator subfield and 4-byte Assigned Number subfield
1	4-byte Administrator subfield and 2-byte Assigned Number subfield
2	4-byte Administrator subfield and 2-byte Assigned Number subfield

If the Type field is 0, the Administrator subfield is a 2byte autonomous system number (ASN). If the Type field is 1, the Administrator subfield is an IP address. If the Type field is 2, the Administrator subfield is a 4byte ASN. In all cases, the Assigned Number subfield contains a number assigned by the administrator. Although the Type field and the Value field are found within a BGP VPN IPv4 address, these fields compose only the RD, not the entire multiprotocol BGP VPN IPv4 address. Although the Administrator subfield and Assigned Number subfield are found within a multiprotocol BG VPN IPv4 address, these subfields compose the Value field of the RD. The ASN is a part of the Administrator subfield.

Reference:

<https://tools.ietf.org/html/rfc4364>

QUESTION 131

DRAG DROP

Select the attributes from the left, and place them on the right in the order they are prioritized by the OSPFv2 Loop-Free Alternate Fast Route feature by default.

Select and Place:

broadcast-interface-disjoint	
interface-disjoint	
linecard-disjoint	
lowest-metric	
node-protecting	
primary-path	
srig	

Help

Reset

Done

Correct Answer:

	srng
	primary-path
	interface-disjoint
	lowest-metric
	linecard-disjoint
	node-protecting
	broadcast-interface-disjoint

Section: (none)
Explanation

Explanation/Reference:

Explanation:

The Open Shortest Path First version 2 (OSPFv2) Loop-Free Alternate Fast Reroute feature is used to reroute traffic if a link fails. Repair paths are calculated and stored in the Routing Information Base (RIB). When a primary path fails, the repair path is used without requiring route recomputation. OSPFv2 Loop-Free Alternate Fast Reroute is not supported on virtual links, but it is supported on VPN routing and forwarding (VRF) OSPF instances. You can configure a traffic engineering (TE) tunnel interface as a repair path but not as a protected interface.

The srng attribute is considered first in the calculation of a repair path. A shared risk link group (SRLG) is a group of next-hop interfaces that are likely to fail simultaneously. You can issue the srng command to assign an interface to an SRLG.

The primary path attribute is considered second. You can configure the primary path attribute so that a particular repair path is used.

The interface-disjoint attribute is considered third. You can set the interface-disjoint attribute to prevent selection of pointtopoint interfaces, which have no alternate next hop for rerouting.

The lowest metric attribute is considered fourth. The lowest cost route might not be the most stable route. However, you can configure the metric attribute to ensure that routes with lower metrics are selected as repair paths.

The linecard-disjoint attribute is considered fifth. Interfaces on the same line card are likely to fail at the same time if there is a problem with the card.

The node-protecting attribute is considered sixth. You can configure the node-protecting attribute so that the primary path gateway router is not selected for the repair path.

The broadcast-interface-disjoint attribute is considered last. You can configure the broadcast-interface-disjoint attribute so that the repair path does not use the broadcast network to which the primary path is connected.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xr-3s/iro-xe-3s-book/iro-lfa-frr.html

QUESTION 132

You issue the show running-config command on RouterA and receive the following partial output:

```
Access-list 101 permit ip host 172.16.223.82 10.17.88.0
0.0.0.255 route-map map1 permit 10 match ip address 101 set
next-hop 192.168.1.1
```

Which of the following packets will RouterA redirect to the nexthop router at 192.168.1.1? (Select the best answer.)

- A. packets sent from the 10.17.88.0/24 network or destined to 172.16.223.82
- B. packets sent from the 10.17.88.0/24 network and destined to 172.16.223.82
- C. packets sent from 172.16.223.82 or destined to the 10.17.88.0/24 network
- D. packets sent from 172.16.223.82 and destined to the 10.17.88.0/24 network

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RouterA will detect packets sent from 172.16.223.82 and destined to the 10.17.88.0/24 network and then redirect them to the nexthop router at 192.168.1.1. Route maps are conditional statements that determine whether a packet is processed normally or modified. A route map can be divided into a series of sequences that are processed in sequential order. If a route matches all the match criteria in a sequence, the route is permitted or denied based on the permit or deny keywords in the routemap command and any set conditions are applied. If a route does not match all the match criteria in any sequence, the route is discarded.

In this scenario, the routemap map1 permit 10 command creates a route map named map1. The permit 10 keywords indicate that any route satisfying all the match statements in route map sequence number 10 will be redistributed. In this sequence, there is only one match statement, match ip address 101, which indicates that packets that match the IP addresses in access list 101 will be processed by the route map.

The accesslist 101 permit ip host 172.16.223.82 10.17.88.0 0.0.0.255 command creates access list 101, which specifies that IP packets sent from 172.16.223.82 and destined to the 10.17.88.0/24 network are processed by the route map. Packets have to match only one accesslist statement in order to be processed by the route map.

RouterA will not redirect packets sent from the 10.17.88.0/24 network and destined to 172.16.223.82. To configure RouterA to match this traffic, you would need to reverse the keywords in the accesslist statement so that the source is the 10.17.88.0/24 network and the destination is the host at 172.16.223.82. The following command set would configure RouterA to detect packets sent from the 10.17.88.0/24 network and destined to 172.16.223.82 and then redirect those packets to the nexthop router at 192.168.1.1:

```
RouterA(config)#accesslist 101 permit ip 10.17.88.0 0.0.0.255 host 172.16.223.82
RouterA(config)#routemap map1 permit 10
RouterA(configroutemap)#match ip address 101
RouterA(configroutemap)#set nexthop 192.168.1.1
```

RouterA will not redirect packets sent from the 10.17.88.0/24 network or destined to 172.16.223.82. To configure RouterA to match either of two access list criteria, you would need to create two separate accesslist statements: one that matches traffic sent from the 10.17.88.0/24 network destined to anywhere, and one that matches traffic sent from anywhere destined to 172.16.223.82. The following command set would configure RouterA to detect packets sent from the 10.17.88.0/24 network or destined to 172.16.223.82 and redirect those packets to the nexthop router at 192.168.1.1:

```
RouterA(config)#accesslist 101 permit ip 10.17.88.0 0.0.0.255 any
RouterA(config)#accesslist 101 permit ip any host 172.16.223.82
RouterA(config)#routemap map1 permit 10
RouterA(configroutemap)#match ip address 101
RouterA(configroutemap)#set nexthop 192.168.1.1
```



RouterA will not redirect packets sent from 172.16.223.82 or destined to the 10.17.88.0/24 network. To configure RouterA to match either of two access list criteria, you would need to create two separate accesslist statements: one that matches traffic sent from 172.16.223.82 destined to anywhere, and one that matches traffic sent from anywhere destined to the 10.17.88.0/24 network. The following command set would configure RouterA to detect packets sent from 172.16.223.82 or destined to the 10.17.88.0/24 network and then redirect those packets to the nexthop router at 192.168.1.1:

```
RouterA(config)#accesslist 101 permit ip any 10.17.88.0 0.0.0.255
RouterA(config)#accesslist 101 permit ip host 172.16.223.82 any
RouterA(config)#routemap map1 permit 10
RouterA(configroutemap)#match ip address 101 RouterA(configroutemap)#set nexthop 192.168.1.1
```

Reference:

CCIE Routing and Switching v5.0 Certification Guide, Volume 1, Chapter 11, Configuring Route Maps with the routemap Command, pp. 638-640

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/a1/sec-a1-cr-book/sec-cr-a2.html#wp4698537840>

QUESTION 133

You administer an MPLS domain. You issue the ip vrf Customer129 command to create a VRF table for a customer. You now want to create an RD and configure RTs.

Which of the following formats can you use? (Select 2 choices.)

- A. nn:AS, where nn is a 16-bit decimal number and AS is a 16-bit ASN
- B. nn:A.B.C.D, where nn is a 16-bit decimal number and A.B.C.D is a 32-bit IP address
- C. nn:MAC, where nn is a 16-bit decimal number and MAC is a 48-bit MAC address
- D. AS:nn, where AS is a 16-bit ASN and nn is a 32-bit decimal number
- E. A.B.C.D:nn, where A.B.C.D is a 32-bit IP address and nn is a 16-bit decimal number
- F. MAC:nn, where MAC is a 48-bit MAC address and nn is a 16-bit decimal number

Correct Answer: DE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You can use the following formats to create a route distinguisher (RD) and configure route targets (RTs):

- AS:nn, where AS is a 16bit autonomous system number (ASN) and nn is a 32bit decimal number
- A.B.C.D:nn, where A.B.C.D is a 32bit IP address and nn is a 16bit decimal number

When you issue the ip vrf name command in global configuration mode, you are placed in VPN routing and forwarding (VRF) configuration mode, where you can configure the VRF.

First, you should create an RD by issuing the rd value command. An RD is a value that is used to create a virtual private network (VPN) prefix to identify the VPN. You can specify the RD by combining an ASN or IP address with any decimal number.

There are three types of RDs: Type 0, Type 1, and Type 2. The type of RD configuration you create depends on how you issue the value parameter of the rd command and whether you are configuring a multicast VPN environment. Type 0 and Type 1 RDs are used in unicast configurations. A Type 0 RD is configured by issuing the value parameter of the rd command with the 16bit ASN in front of the 32bit decimal number. A Type 1 RD is configured by issuing the value parameter of the rd command with the 32bit decimal number in front of the 16bit ASN. A Type 2 RD is configured similarly to a Type 1 RD but applies to only multicast VPN configurations.

To configure RT extended community attributes for the VRF, you should issue the route-target {import | export | both} value command. Like RDs, RTs are specified by combining an ASN or IP address with any decimal number. The import, export, and both keywords specify whether extended community attributes should be imported, exported, or both.

You should also associate an interface with the VRF by issuing the ip vrf forwarding name command, where name is the name of the VRF as specified in the ip vrf name command. When the ip vrf forwarding command is issued, the IP address is removed from the interface. Therefore, you should reconfigure the IP address on the interface after issuing the ip vrf forwarding command.

The following commands can be used to create a VRF table for a customer and apply it to an interface:

```
RouterA(config) #ip vrf Customer129
```

```
RouterA(config-vrf) #rd 123:6
```

```
RouterA(config-vrf) #route-target both 123:6
```

```
RouterA(config-vrf) #route-target export 192.168.14.1:77
RouterA(config-vrf) #exit
RouterA(config) #interface fa0/1
RouterA(config-if) #ip vrf forwarding Customer129
```

You would not use the following formats to create RDs and RTs, because the parameters are in the wrong order:

- nn:AS, where nn is a 16bit decimal number and AS is a 16-bit ASN
- nn:A.B.C.D, where nn is a 16bit decimal number and A.B.C.D is a 32-bit IP address

You would not use the following formats to create RDs and RTs, because Media Access Control (MAC) addresses cannot be used to create them: -nn:MAC, where nn is a 16-bit decimal number and MAC is a 48-bit MAC address - MAC:nn, where MAC is a 48-bit MAC address and nn is a 16-bit decimal number

Reference:

- <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mppls/command/mp-cr-book/mp-m4.html#wp3212018555>
- https://www.cisco.com/c/en/us/td/docs/ios/ipv6/command/reference/ipv6_book/ipv6_11.html#wp2419815

QUESTION 134

DRAG DROP

Select the BFD mode or function from the left, and place it on the corresponding description on the right.

Select and Place:



Asynchronous mode	can reduce round-trip jitter
Demand mode	requires half as many packets for failure detection
Echo function	supports a large number of BFD sessions

Correct Answer:

	Echo function
	Asynchronous mode
	Demand mode

Section: (none)
Explanation

Explanation/Reference:

Explanation:

Bidirectional Forwarding Detection (BFD) is a detection protocol that is designed to detect forwarding path failures in less than one second. Additionally, BFD is designed to work regardless of media type, encapsulation, or routing protocol, providing network administrators with a uniform forwarding failure detection method across a network. BFD supports Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), Border Gateway Protocol (BGP), and Intermediate System to Intermediate System (ISIS).

BF has two operating modes: Asynchronous mode and Demand mode. Asynchronous mode systems periodically send BFD Control packets. If Control packets are not received from a neighbor in a timely fashion, the neighbor is assumed to be down.

Demand mode systems assume that there is an external method of verifying connectivity. When Demand mode is enabled on a system, the system can ask neighbors to stop sending BF Control packets except when absolutely necessary. This enables Demand mode systems to support a large number of BFD sessions, but the relative lack of Control packets can cause failure detection to be slower.

The Echo function can be enabled with Asynchronous mode or Demand mode. When the Echo function is enabled on a system, the system will send a stream of BFD Echo packets to a neighbor. If enough Echo packets are not returned, the neighbor is assumed to be down. The Echo function can reduce roundtrip jitter and can increase the speed of failure detection. Although the Echo function does not send packets as often as Asynchronous mode does, Asynchronous mode requires half as many packets as the Echo function does in order to detect a failure.

Reference:

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

QUESTION 135

Which of the following statements is true regarding BGP soft reconfiguration? (Select the best answer.)

- A. It requires very little memory.
- B. It can be performed by issuing the clear ip bgp command.
- C. It tears down BGP sessions before rebuilding the routing tables.

D. It requires no configuration to support both inbound and outbound updates.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

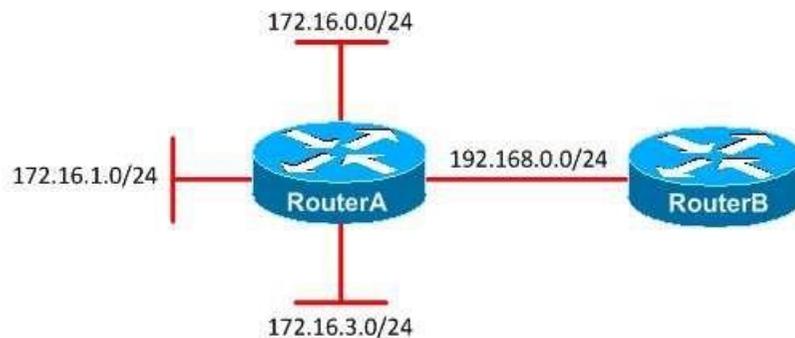
Explanation:

Border Gateway Protocol (BGP) soft reconfiguration can be performed by issuing the clear ip bgp command. The clear ip bgp command rebuilds the BGP routing table. This command can be used to begin a soft or hard reset. A soft reset uses stored prefix information in order to rebuild the BGP routing tables without breaking down any active peering sessions, whereas a hard reset breaks down the active peering sessions and then rebuilds the BGP routing tables. Typically, the clear ip bgp * soft command initiates soft reconfiguration; however, if all BGP routers support soft reconfiguration, the soft keyword is assumed by default. BG soft reconfiguration does not require any configuration to support outbound updates, because outbound updates are stored automatically. However, soft reconfiguration requires that you issue the neighbor rneighbor-id soft-reconfiguration inbound command before it stores inbound updates from a neighbor. These updates are stored in memory, so soft reconfiguration is very memory intensive. When soft reconfiguration inbound is configured, the route will display (received only) in the output of the show ip bgp command.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp2.html#wp1107408

QUESTION 136



You administer the network shown above. RouterA and RouterB are configured to use EIGRP on all interfaces. Automatic summarization is enabled. What network or networks will RouterA advertise to RouterB? (Select the best answer.)



<https://vceplus.com/>

- A. 172.16.0.0/16
- B. 172.16.0.0/22
- C. 172.16.0.0/23 and 172.16.3.0/24
- D. 172.16.0.0/24, 172.16.1.0/24, and 172.16.3.0/24

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RouterA will advertise the 172.16.0.0/16 network to RouterB. When the auto-summary command has been used to enable automatic summarization on a router, Enhanced Interior Gateway Routing Protocol (EIGRP) automatically summarizes networks on classful boundaries. Summarization, which is also referred to as aggregation, minimizes and optimizes the size of routing tables and advertisements and reduces a router's processor and memory requirements. Summarization is also useful in limiting the scope of EIGRP queries.

The 172.16.0.0/24, 172.16.1.0/24, and 172.16.3.0/24 networks in this scenario use Class B addresses. Therefore, these network ranges are summarized to the Class B boundary, which is /16.

To disable automatic summarization, you should issue the no autosummary command. The no autosummary command enables EIGRP to advertise the actual networks, not the classful summary. You should use the no autosummary command when a classful network is divided and portions of the same classful network exist in different parts of the network topology. If you were to issue the no autosummary command on RouterA, RouterA would advertise the individual network ranges and subnet mask information to RouterB.

You can issue the ip summary-address eigrp command to enable manual summarization. Manual summarization is configured on a per-interface basis. The syntax of the ip summary-address eigrp command is ip summary-address eigrp as-number address mask, where as-number is the EIGRP autonomous system number (ASN), address is the summary address, and mask is the subnet mask in dotted decimal notation.

You can also summarize external routes. However, EIGRP will not automatically summarize external routes unless there is an internal route that uses the same classful network. If the 172.16.0.0/24, 172.16.1.0/24, and 172.16.3.0/24 networks in this scenario were external routes redistributed into EIGRP, those networks would not be automatically summarized by RouterA, because there is no internal route that uses the 172.16.0.0/16 network range.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/iproute/command/reference/fiprrp_r/1rfeigrp.html#wp1017389

<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html#auto>

QUESTION 137

Which of the following TCP features can cause TCP starvation in a network with a large amount of UDP traffic and no QoS mechanism? (Select the best answer.)

- A. window scaling
- B. sliding window
- C. MSS adjustment
- D. selective acknowledgment

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Transmission Control Protocol (TCP) sliding window feature can cause TCP starvation in a network with a large amount of User Datagram Protocol (UDP) traffic and no Quality of Service (QoS) mechanism. TCP starvation, which is also known as UDP dominance, occurs when congestion and packet loss cause TCP data streams to scale back their transmission window sizes, thereby enabling UDP data streams to dominate the available network bandwidth. TCP starvation can introduce additional latency and reduce the overall throughput of a network link.

TCP has flow control mechanisms to prevent a sending device from transmitting data faster than the receiver can process it. When TCP detects dropped packets, it reduces the TCP transmission window size and retransmits the dropped packets. Reducing the window size slows the rate at which TCP sends traffic. If multiple TCP data streams exist and no QoS mechanism is in place, the streams typically reduce their window sizes in unison because they experience dropped packets in an equal distribution. If there are a large number of UDP data streams on the same network link as the TCP data streams, they will quickly consume the network bandwidth that was made available by the reduction of TCP traffic.

However, because UDP does not have an inherent flow control mechanism like TCP does, the UDP data streams are not directly affected by dropped packets and the network congestion will likely continue or possibly get worse.

Window scaling is not a TCP feature that can cause TCP starvation in a network with a large amount of UDP traffic and no QoS mechanism. Window scaling enables a router to store the equivalent of a 32bit value in the 16bit TCP header field that specifies the window size. This enables the router to process a significantly larger number of bytes before it is required to send an acknowledgment. Larger window sizes are of particular use on networks with high bandwidth and high delay, which are known as Long Fat Networks (LFNs).

Selective acknowledgment is not a TCP feature that can cause TCP starvation in a network with a large amount of UDP traffic and no QoS mechanism. Selective acknowledgment enables TCP to acknowledge packets that were received out of order. Without selective acknowledgment, the receiving router would only be able to acknowledge packets in order.

For example, if 10 packets were sent and only packets 1, 2, 3, 5, 7, 8, 9, and 10 were received, a router without selective acknowledgment would acknowledge the receipt of only packets 1, 2, and 3. This would likely cause packets 5, 7, 8, 9, and 10 to be retransmitted. However, with selective acknowledgment, the router could

acknowledge the receipt of all of the packets and only the missing packets would be retransmitted. Selective acknowledgment reduces wasted transmissions and increases the overall efficiency of TCP on a particular link.

Maximum segment size (MSS) adjustment is not a TCP feature that can cause TCP starvation in a network with a large amount of UDP traffic and no QoS mechanism. MSS adjustment enables a router to override the MSS value of TCP SYN packets. Hosts use the MSS option in the TCP header to negotiate a maximum size of an IP segment. However, if an intervening device cannot support this size, the packets might get dropped and the TCP session might terminate. With the TCP MSS adjustment feature, you can modify the TCP MSS value for transient packets, which are packets that neither originate from nor terminate on the router. This can ensure that the router will not drop the packets because they exceed the maximum transmission unit (MTU) of one of its interfaces.

Reference:

https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book/VPNQoS.html#pgfId-113408

QUESTION 138

You want to connect four ports on two switches in an EtherChannel configuration. SwitchEast is a Cisco switch, and SwitchWest is a nonCisco switch.

Which of the following command sets should you issue to configure interfaces Gi 2/1 through Gi 2/4 on SwitchEast? (Select the best answer.)

- A. SwitchEast(config)#interface port-channel 1
SwitchEast(config-if)#ip address 10.20.30.40 255.255.255.0
SwitchEast(config-if)#interface range gi 2/1 4
SwitchEast(config-if-range)#no ip address
SwitchEast(config-if-range)#channel-protocol pagp
SwitchEast(config-if-range)#channel-group 1 mode active
- B. SwitchEast(config)#interface port-channel 1
SwitchEast(config-if)#ip address 10.20.30.40 255.255.255.0
SwitchEast(config-if)#interface range gi 2/1 4
SwitchEast(config-if-range)#no ip address
SwitchEast(config-if-range)#channel-protocol pagp
SwitchEast(config-if-range)#channel-group 1 mode desirable non-silent
- C. SwitchEast(config)#interface port-channel 1
SwitchEast(config-if)#ip address 10.20.30.40 255.255.255.0
SwitchEast(config-if)#interface range gi 2/1 4
SwitchEast(config-if-range)#no ip address
SwitchEast(config-if-range)#channel-protocol lacp
SwitchEast(config-if-range)#channel-group 1 mode active
- D. SwitchEast(config)#interface port-channel 1
SwitchEast(config-if)#ip address 10.20.30.40 255.255.255.0

```
SwitchEast(config-if)#interface range gi 2/1 - 4 SwitchEast(config-if-  
range)#no ip address  
SwitchEast(config-if-range)#channel-protocol lacp  
SwitchEast(config-if-range)#channel-group 1 mode desirable non-silent
```

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the following command set to configure interfaces Gi 2/1 through Gi 2/4 on SwitchEast:

```
SwitchEast(config)#interface portchannel 1  
SwitchEast(config-if)#ip address 10.20.30.40 255.255.255.0  
SwitchEast(config-if)#interface range gi 2/1 - 4  
SwitchEast(config-if-range)#no ip address  
SwitchEast(config-if-range)#channel-protocol lacp  
SwitchEast(config-if-range)#channel-group 1 mode active
```

These commands configure a Link Aggregation Control Protocol (LACP) EtherChannel on SwitchEast.

LACP is defined in the Institute of Electrical and Electronics Engineers (IEEE) 802.3ad standard. Because LACP is a standardsbased protocol, it can be used between Cisco and nonCisco switches.

The interface portchannel 1 command creates port channel interface 1; the port channel interface number can be any number from 1 through 64. The ip address 10.20.30.40 255.255.255.0 command assigns the IP address 10.20.30.40/24 to the port channel interface. The interface range gi 2/1 4 command enters interface configuration mode for interfaces Gi 2/1 through Gi 2/4, and the no ip address command ensures that no IP address is configured for any of the interfaces that will belong to the channel group.

The channel-protocol lacp command configures the interfaces for LACP operation. The channelgroup 1 mode active command configures the interfaces for channel group 1 in active mode. The syntax of the channelgroup command is channel group numbermode {on | active | passive | {auto | desirable} [nonsilent]}, where number is the port channel interface number. The auto, desirable, and nonsilentkeywords can be used only with Port Aggregation Protocol (PAgP). The on keyword configures the channel group to unconditionally create the channel with no LACP negotiation. The active keyword configures the channel group to actively negotiate LACP, and the passive keyword configures the channel group to listen for LACP negotiation to be offered.

The following command set configures a PAgP EtherChannel on interfaces Gi 2/1 through 2/4 on SwitchEast:

```
SwitchEast(config)#interface port-channel 1  
SwitchEast(config-if)#ip address 10.20.30.40 255.255.255.0  
SwitchEast(config-if)#interface range gi 2/1 - 4 SwitchEast(config-if-  
range)#no ip address
```

```
SwitchEast(config-if-range)#channel-protocol pagp  
SwitchEast(config-if-range)#channel-group 1 mode desirable non-silent
```

PAgP is a Cisco proprietary protocol. Therefore, PAgP cannot be used to create an EtherChannel between SwitchEast and SwitchWest; it can only be used to create an EtherChannel between two Cisco switches.

The first part of PAgP configuration is identical to LACP configuration. The channel-protocol pagp command configures the interfaces for PAgP operation. The channel-group 1 mode desirable non-silent command configures the interfaces for channel group 1 in desirable mode. The syntax of the channelgroup command is channel-group numbermode {on | active | passive | {auto | desirable} [nonsilent]}, where number is the port channel interface number. The active and passive keywords can be used only with LACP. The on keyword configures the channel group to unconditionally create the channel with no PAgP negotiation. The desirable keyword configures the channel group to actively negotiate PAgP, and the auto keyword configures the channel group to listen for PAgP negotiation to be offered. The optional nonsilent keyword requires that a port receive PAgP packets before the port is added to the channel. The following command set is invalid because the desirable and nonsilent keywords cannot be used with LACP:

```
SwitchEast(config)#interface port-channel 1  
SwitchEast(config-if)#ip address 10.20.30.40 255.255.255.0  
SwitchEast(config-if)#interface range gi 2/1 - 4  
SwitchEast(config-if-range)#no ip address  
SwitchEast(config-if-range)#channel-protocol lacp  
SwitchEast(config-if-range)#channel-group 1 mode desirable non silent
```

The following command set should not be issued on SwitchEast, because PAgP cannot be used on nonCisco switches. Additionally, the command set is invalid because the active keyword cannot be used with PAgP:

```
SwitchEast(config)#interface portchannel 1  
SwitchEast(config-if)#ip address 10.20.30.40 255.255.255.0  
SwitchEast(config-if)#interface range gi 2/1 - 4  
SwitchEast(config-if-range)#no ip address  
SwitchEast(config-if-range)#channel-protocol pagp  
SwitchEast(config-if-range)#channel-group 1 mode active
```

The following table displays the channelgroup configurations that will establish an EtherChannel:

SwitchA	SwitchB	off	auto	desirable	passive	active	on
off		NO	NO	NO	NO	NO	NO
auto		NO	NO	PAgP	NO	NO	NO
desirable		NO	PAgP	PAgP	NO	NO	NO
passive		NO	NO	NO	NO	LACP	NO
active		NO	NO	NO	LACP	LACP	NO
on		NO	NO	NO	NO	NO	ON

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750x_3560x/software/release/12-2_55_se/configuration/guide/3750xscg/swethchl.html

https://www.cisco.com/c/en/us/td/docs/ios/12_2sb/feature/guide/gigeth.html

QUESTION 139

Which of the following commands should you issue on a switch port so that the port will trust the CoS value of the incoming voice and data traffic? (Select the best answer.)

- A. mls qos trust
- B. mls qos trust cos
- C. mls qos trust ip-precedence
- D. switchport priority extend cos
- E. switchport priority extend trust
- F. switchport trunk native vlan tag



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the mls qos trust cos interface configuration command on the switch port connected to an IP phone to configure that port to trust the class of service (CoS) values of incoming voice and data traffic. CoS values are used to prioritize voice and data traffic so the delaysensitive voice traffic receives preferential treatment on the network. By default, the switch port does not trust the CoS values of incoming traffic and reclassifies the traffic with the port's default CoS value of 0. The mls qos trust cos command configures the switch to trust the CoS value of both voice traffic and data traffic that is sent by an IP phone. The multilayer switching (MLS) port trust feature can be used to examine the CoS or differentiated services code point (DSCP) value to classify incoming traffic. To configure the MLS port trust state, you should issue the mls qos trust command. The syntax for configuring the MLS port trust state is mls qos trust [cos | dscp | ipprecedence]. If you do not specify one of the keywords or if Quality of Service (QoS) has been disabled globally, the mls qos trust command defaults to dscp,

which classifies incoming traffic according to the DSCP values in the packet header. If you use the ip-precedence keyword, the incoming packets are classified according to the type of service (ToS) bits in the packet header.

Issuing the switchport priority extend cos interface configuration command on the switch port to which the IP phone is connected configures the IP phone to override the priority of the data packets it receives from the host and assigns new CoS values to the host generated packets. Using the switchport priority extend cos command reclassifies incoming data packets with the default CoS value of 0. Thus an IP phone can prevent the computer from exploiting a high-priority data queue. The switchport priority extend trust interface configuration command configures an IP phone to trust the CoS value of incoming data packets it receives from the attached computer. The switchport priority extend trust command does not configure the switch port to trust the traffic it receives from an IP phone.

The switchport trunk native vlan tag interface configuration command configures native virtual LAN (VLAN) traffic to be tagged. By default, traffic from the native VLAN is sent untagged. Tagging native VLAN traffic is necessary to enable Layer 2 QoS support on the native VLAN.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/command/qos-cr-book/qos-m2.html#wp3022589880>

QUESTION 140

Which of the following algorithms will an OSPF router use to determine the best route for packets? (Select the best answer.)

- A. the Dijkstra algorithm
- B. the Diffie-Hellman algorithm
- C. the Bellman-Ford algorithm
- D. the DUAL algorithm
- E. the path-vector algorithm



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

An Open Shortest Path First (OSPF) router will use the Dijkstra algorithm to determine the best route for packets. The OSPF routing process determines the best route for packets by analyzing its link state database and building a shortest path first (SPF) tree. The SPF tree is a simplified view of the entire network topology and contains the shortest route to any destination on the network. Although every router in an area has the same link state database, each router can calculate its own SPF tree to determine the best route to any location on the network. The SPF tree is typically recalculated only if a router receives a link state packet that indicates a change on the network. Because link state updates are only transmitted within an area, OSPF routers in other areas do not recalculate their SPF trees when a change occurs in an area to which they are not connected.

An OSPF router will not use the Bellman-Ford algorithm to determine the best route for packets. The Bellman-Ford algorithm is used by distance vector routing protocols, such as Routing Information Protocol (RIP), to determine the best routes to locations on the network.

An OSPF router will not use the path-vector algorithm to determine the best route for packets. The path-vector algorithm, which is used by Border Gateway Protocol (BGP) to determine the best routes to locations on the network, is based on the Bellman-Ford algorithm.

An OSPF router will not use the Diffusing Update Algorithm (DUAL) to determine the best route for packets. The DUAL algorithm is a hybrid of distance-vector routing protocols and is used by Enhanced Interior Gateway Routing Protocol (EIGRP) to determine the best routes to locations on the network.

An OSPF router will not use the Diffie-Hellman algorithm to determine the best route for packets. The Diffie-Hellman algorithm is not used by dynamic routing processes. The Diffie-Hellman algorithm is commonly used by IP Security (IPSec) to generate shared keying material and to securely transfer the information necessary to establish a security association (SA) with an IPSec peer.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html#t5>

QUESTION 141

Which of the following OSPF network types have a default hello timer of 10 seconds and dead timer of 40 seconds? (Select 2 choices.)

- A. broadcast
- B. nonbroadcast
- C. point-to-point
- D. point-to-multipoint
- E. point-to-multipoint nonbroadcast

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The broadcast and point-to-point Open Shortest Path First (OSPF) network types have a default hello timer of 10 seconds and dead timer of 40 seconds. The nonbroadcast, point-to-multipoint, and point-to-multipoint nonbroadcast OSPF network types have a default hello timer of 30 seconds and a dead timer of 120 seconds. The hello timer is used to specify the amount of time between sending hello packets, and the dead timer is used to specify the amount of time to wait for hello packets before declaring a neighbor to be down. In order for OSPF routers to establish an adjacency, the hello timer on one router should match the hello timer on the other router, and the dead timer on one router should match the dead timer on the other router. The dead timer is set to four times the hello timer value by default.

To manually configure the hello timer interval, you should issue the `ip ospf hello-interval seconds` command in interface configuration mode. To manually configure the dead timer interval, you should issue the `ip ospf dead-interval seconds` command in interface configuration mode.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/ospf-i1.html#wp4134450560 https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/ospf-a1.html#wp2917383021

QUESTION 142

Which of the following statements is accurate regarding alternate ports? (Select the best answer.)

- A. An alternate port is always in the forwarding state.
- B. An alternate port is the port on a switch that has the best path to the root bridge.
- C. An alternate port is the port on a segment that has the best path to the root bridge.
- D. An alternate port is a blocked port that receives more useful BPDUs from a port on another switch.
- E. An alternate port is a blocked port that receives more useful BPDUs from another port on the local switch.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

An alternate port is a blocked port that receives more useful bridge protocol data units(BPDUs) from a port on another switch. Rapid Spanning Tree Protocol (RSTP) defines the following four port roles:

- Root port
- Designated port
- Alternate port
- Backup port



The root port on a switch is the port that receives the best BPDU, which indicates the best path to the root bridge based on the best root port cost. All switches except the root bridge contain exactly one root port. Because there is only one best path to the root bridge, a switch can have only one root port; only the root bridge does not have a root port. Root ports are always in the forwarding state.

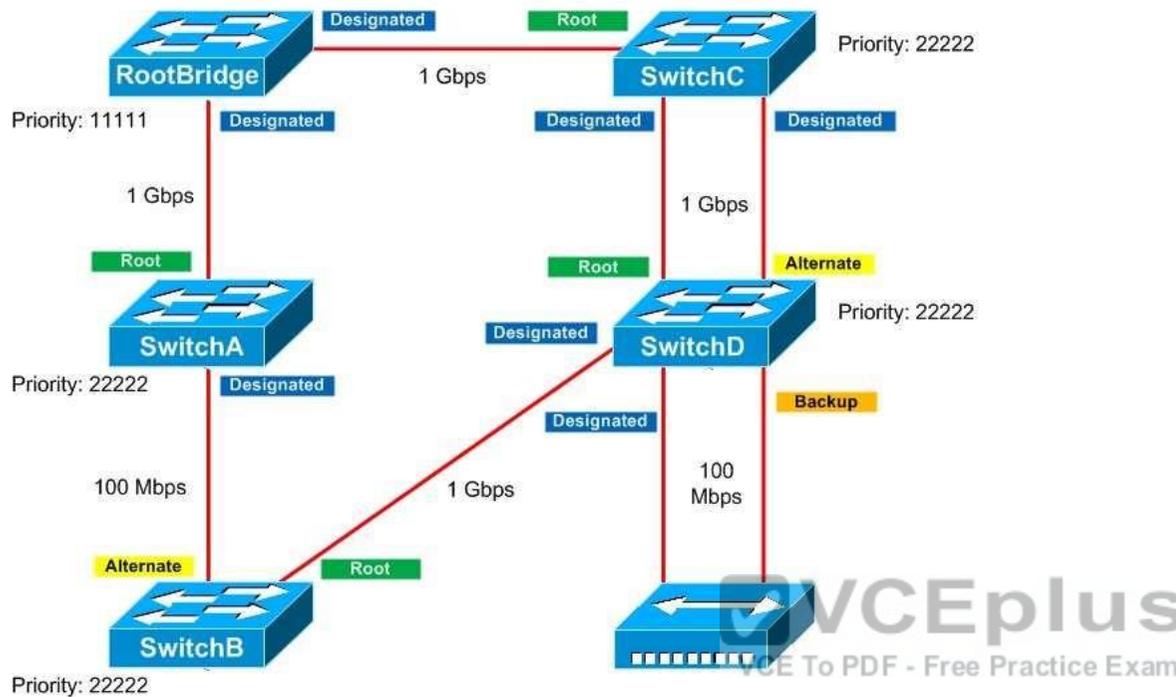
The best path to the root bridge is the one that has the lowest cost to the root bridge. The cost is based on the bandwidth of a link. The higher the bandwidth, the lower the cost. RSTP uses the following link costs by default:

Bandwidth	Cost
4 Mbps	250
10 Mbps	100
16 Mbps	62
45 Mbps	39
100 Mbps	19
155 Mbps	14
622 Mbps	6
1 Gbps	4
10 Gbps	2

In the event of a tie, the port connected to the device with the lowest bridge ID (BID) becomes the root port. If both ports are connected to the same switch, the port that receives the BPDU with the lowest sending port ID becomes the root port. The port ID consists of the port priority and the port number. The port configured with the lowest port priority will also have the lowest port ID? port numbers are considered only when port priorities are equal.

A designated port is the port on a segment that has the best path to the root bridge. One designated port is selected for each segment. If multiple ports on a segment have the same root port cost, the port on the switch with the lowest BID becomes the designated port. Switches can have one or more designated ports, and some switches might not have any designated ports. All the ports on a root bridge are designated ports. Designated ports are normally in the forwarding state. The blocking port role in Spanning Tree Protocol (STP) is split into two RSTP port roles: the alternate port role and the backup port role. An alternate port receives more useful BPDUs from a designated port located on another switch, and a backup port receives more useful BPDUs from a designated port on the switch itself. For a port to receive more useful BPDUs from the same switch, it must be connected to the same collision domain as another port on the same switch or it must be connected to itself by a loopback device, such as a loopback adapter. An alternate port guarantees a path to the root bridge should the current root port become unavailable; however, a backup port only guarantees redundant access to a particular network segment and not necessarily an alternate path to the root bridge. Alternate ports and backup ports are always in the blocking state.

The following graphic displays the correlation among the various RSTP port roles:



Reference:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/24062-146.html>

QUESTION 143

You are configuring EEM on a router. You have issued the event ioswdsysmon sub1 cpu-proc taskname Task1 op ge val 90 period 20 command to trigger an action when the processor usage exceeds 90 percent over a 20second period. You want to configure the router to change the IP address of the s0/0 interface when this event is triggered.

Which of the following commands should you issue? (Select the best answer.)

- A. action cli
- B. action force-switchover
- C. action info
- D. action ipaddress
- E. action publish-event

- F. action reload
- G. action snmp-trap

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the action cli command to configure the router to change the IP address of the s0/0 interface when the Embedded Event Manager (EEM) event is triggered. EEM enables event detectors to monitor events and perform actions if those events are triggered. The event ioswdsysmon command configures the Watchdog System Monitor (IOSWDSysMon) to monitor memory and processor usage.

The syntax of the action cli command is action labelcli command commandstring. The label variable, which is used with all the action commands, is an alphanumeric value that determines the order in which the actions are processed when the event is triggered. The commandstring variable is the IOS command that should be issued; if the command contains spaces, the command must be enclosed in double quotation marks (").

You should not issue the action force-switchover command to change the IP address of the s0/0 interface when the EEM event is triggered. The action forceswitchovercommand configures the router to switch to a secondary processor when the event is triggered.

You should not issue the action info command to change the IP address of the s0/0 interface when the EEM event is triggered. The action info command enables the router to retrieve command history and logging information.

You should not issue the action publish-event command to change the IP address of the s0/0 interface when the EEM event is triggered. The action publish-event command configures the router to publish an application-specific event when the EEM event is triggered. EEM must publish events to subsystem number 798. You should not issue the action reload command to change the IP address of the s0/0 interface when the EEM event is triggered. The action reload command reloads the Cisco IO software when the EEM event is triggered.

You should not issue the action snmp-trap command to change the IP address of the s0/0 interface when the EEM event is triggered. The action snmp-trap command generates a Simple Network Management Protocol (SNMP) trap when the EEM event is triggered.

You should not issue the action ip-address command to change the IP address of the s0/0 interface when the EEM event is triggered, because the ip-address keyword is not supported with the action command. The following keywords can be used with the action command:

- cli
- cns-event
- counter
- force-switchover
- info
- mail
- policy
- publish-event
- reload

-snmp-trap -
syslog

Reference: <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/eem/command/eem-cr-book/eem-cr-a1.html#wp3174969440> **QUESTION 144**

Which of the following commands would you issue to create a unique identifier for a DHCP client that uses Ethernet and that has a MAC address of aaaa.bbbb.cccc? (Select the best answer.)

- A. client-identifier aaaa.bbbb.cccc
- B. client-identifier 01aa.aabb.bbcc.cc
- C. client-identifier aaaa.bbbb.cccc.01
- D. hardware-address aaaa.bbbb.cccc
- E. hardware-address 01aa.aabb.bbcc.ccF. hardware-address aaaa.bbbb.cccc.01

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

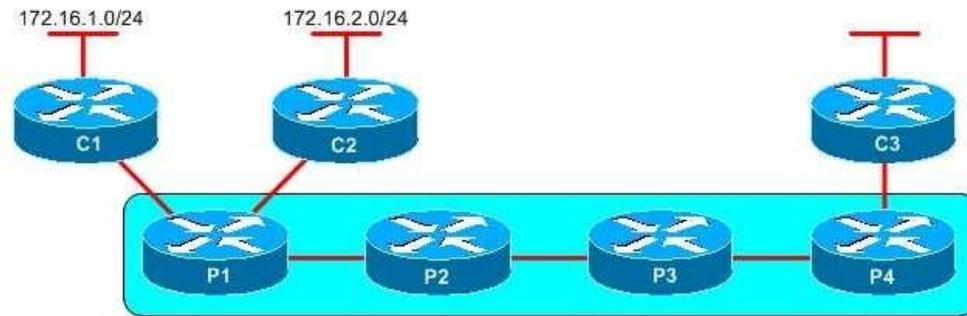
Explanation:

You would issue the clientidentifier 01aa.aabb.bbcc.cc command to create a unique identifier for a Dynamic Host Configuration Protocol (DHCP) client that uses Ethernet and that has a Media Access Control (MAC) address of aaaa.bbbb.cccc. You can specify the client identifier as a sevenbyte hexadecimal notation or as a 27byte dotted hexadecimal notation. To use the sevenbyte version, add the twocharacter media type to the beginning of the MAC address; for Ethernet, the media type is 01. The 27byte version takes an ASCII string that contains the vendor, the MAC address, and the source interface and converts it into dotted hexadecimal. You would not issue the hardwareaddress command to create a unique identifier for a DHCP client. The hardwareaddress command is used to specify the hardware address of a Bootstrap Protocol (BOOTP) client.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/command/ipaddr-cr-book/ipaddr-c1.html#wp2742622403>

QUESTION 145



You administer the MPLS network displayed above. Routing information regarding the networks connected to C1 and C2 is passed along the MPLS core. C3 forwards a packet that is destined for 172.16.1.8 to P4.

Which router will perform PHP? (Select the best answer.)

- A. P1
- B. P2
- C. P3
- D. P4

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Explanation:

P2 will perform penultimate hop popping (PHP). PHP is used for directly connected and summarized routes to optimize Multiprotocol Label Switching (MPLS) networks. When PHP is used, the next-to-last router will remove, or pop off, the MPLS label so that the next-hop egress router does not have to perform two table lookups. The removed label is actually an implicit null label. When the egress router sees the implicit null label, it will forward the packet based on information in the Forwarding Information Base (FIB).

P1 receives routes to 172.16.1.0/24 and 172.16.2.0/24. P1 tells P2 that P2 should remove the label on packets destined for these two networks. P2 will generate a label for the route and advertise it to P3. Similarly, P3 will generate a label for the route and advertise it to P4.

P4 will not perform PHP. When P4 receives a packet destined for 172.16.1.8, it determines that the packet is destined for the 172.16.0.0/16 network. P4 consults the FIB, labels the packet with the label advertised by P3, and forwards the packet to P3.

P3 will not perform PHP. When P3 receives the packet, it consults the Label Forwarding Information Base (LFIB), swaps the label with the label advertised by P2, and forwards the packet to P2. When P2 receives the packet, it will consult the LFIB, pop the label, and forward the packet to P1. P1 receives the packet and forwards the IP packet to C1.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/12_4t2/ftldp41.html#wp1656659

QUESTION 146

DRAG DROP

Select the 802.1D STP states from the left, and drag them to corresponding definitions on the right.

Select and Place:

The interface shows a drag-and-drop activity. On the left, there are five green boxes representing STP states: blocking, disabled, forwarding, learning, and listening. On the right, there are five yellow boxes representing definitions. A 'VCEplus' watermark is visible in the center. At the bottom, there are three buttons: 'Help', 'Reset', and 'Done'.

blocking	forwards frames, learns addresses, receives BPDUs
disabled	discards frames, learns addresses, receives BPDUs
forwarding	discards frames, does not learn addresses, receives BPDUs
learning	discards frames, does not learn addresses, does not receive BPDUs
listening	discards frames, does not learn addresses, receives BPDUs, performs elections

Buttons: Help, Reset, Done

Correct Answer:

	forwarding
	learning
	blocking
	disabled
	listening

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A switch port that is in the Spanning Tree Protocol (STP) listening state can send and receive bridge protocol data units (BPDUs), but the port will not update its Media Access Control (MAC) address table.

There are five possible port states when 802.1D Spanning Tree Protocol (STP) is used: blocking, listening, learning, forwarding, and disabled. After a port is initialized, it enters the blocking state. From the blocking state, the port transitions to either the listening state or the disabled state. From the listening state, the port transitions to either the learning state or the disabled state. From the learning state, the port transitions to either the forwarding state or the disabled state. A port will enter the listening state from the blocking state if STP determines that the port can enter the forwarding state. When a port is in the listening state, the port discards any data frames that it receives; however, it receives and forwards BPDUs. Additionally, the port is able to receive and respond to network management messages. The port does not learn any MAC addresses and does not populate its MAC address table. The root bridge, root port, and designated ports are elected during the listening state.

When a port is in the learning state, it does not forward data frames but does populate the MAC address table based on the frames that it receives. The port responds to network management messages, receives and directs BPDUs to the system module, and processes BPDUs received from the system module. Similar to a port in the learning state, a port in the forwarding state can populate the MAC address table based on the frames that it receives. However, unlike a port in the learning state, a port in the forwarding state can forward data frames as well as receive and process data frames, BPDUs, and network management messages. A port in the blocking state is similar to a port in the listening state in that the port cannot forward data frames or populate the MAC address table. Additionally, a port in the blocking state can receive BPDUs as well as receive and respond to network management messages.

A port in the disabled state does not process or forward data frames, nor does it forward BPDUs or update the MAC address table. A port in the disabled state does not participate in STP.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/cisco_ie2000/software/release/15-0_1_ey/configuration/guide/scg-ie2000/swstp.html#pgfId-1020084

QUESTION 147

DRAG DROP

Select the CHAP packet types from the left, and drag them to the corresponding packet formats on the right.

Select and Place:

The interface shows four green boxes on the left representing CHAP packet types: Challenge, Failure, Response, and Success. On the right, there are four yellow boxes representing packet formats: 01, ID, Length, Value-Size, Value, Name; 02, ID, Length, Value-Size, Value, Name; 03, ID, Length, Message; and 04, ID, Length, Message. A 'Help' button is on the bottom left, and 'Reset' and 'Done' buttons are on the bottom right.

Correct Answer:

The correct answer interface shows four empty white boxes on the left. On the right, the green boxes are ordered: Challenge, Response, Success, and Failure. A 'Help' button is on the bottom left, and 'Reset' and 'Done' buttons are on the bottom right.

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A Challenge Handshake Authentication Protocol (CHAP) packet consists of the following fields:

- A one-octet Code field
- A one octet Identifier field, which helps to match challenges to responses
- A two-octet Length field, which indicates the length of the packet -
- One or more fields that are determined by the Code field

A Challenge packet has a Code field that is set to a value of 1. It also has the following additional fields:

- A one octet Value-Size field, which indicates the length of the Value field
- A variable-length Challenge Value field, which contains a variable, unique stream of octets -
- A variable-length Name field, which identifies the name of the transmitting device

A Response packet has a Code field that is set to a value of 2. It also has the following additional fields:

- A one-octet Value-Size field, which indicates the length of the Response Value field
- A variable-length Response Value field, which contains a concatenated one-way hash of the ID, the secret key, and the Challenge Value -
- A variable-length Name field, which identifies the name of the transmitting device

A Success packet has a Code field that is set to a value of 3, and a Failure packet has a Code field that is set to a value of 4. In addition to the standard fields, the Success packet and the Failure packet have a variable-length Message field, which displays a success or failure message, typically in human-readable ASCII characters.

Reference: <https://www.cisco.com/c/en/us/support/docs/wan/point-to-point-protocol-ppp/25647-understanding-ppp-chap.html> <https://www.ietf.org/rfc/rfc1994.txt>

QUESTION 148

You issue the following commands on RouterA:

```
ARouterA(config) #interface fa0/1
RouterA(config-if) #ip address 10.1.1.1 255.255.255.0
RouterA(config-if) #ip address 10.1.1.2 255.255.255.0
RouterA(config-if) #ip address 10.1.1.3 255.255.255.0 secondary
RouterA(config-if) #ip address 10.1.1.4 255.255.255.0 secondary.
```

Which of the following statements is correct? (Select the best answer.)

- A. Two IP addresses are active on the interface: 10.1.1.1 and 10.1.1.3.
- B. Two IP addresses are active on the interface: 10.1.1.1 and 10.1.1.4.
- C. Two IP addresses are active on the interface: 10.1.1.2 and 10.1.1.3.
- D. Two IP addresses are active on the interface: 10.1.1.2 and 10.1.1.4.
- E. Three IP addresses are active on the interface: 10.1.1.1, 10.1.1.3, and 10.1.1.4.
- F. Three IP addresses are active on the interface: 10.1.1.2, 10.1.1.3, and 10.1.1.4.
- G. All four IP addresses are active on the interface.

Correct Answer: F

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Three IP addresses are active on the interface: 10.1.1.2, 10.1.1.3, and 10.1.1.4. An interface can be configured with one primary IP address and any number of secondary IP addresses. Primary IP addresses are configured by issuing the ip address address mask command. If a primary IP address is configured on an interface, it replaces any previously configured primary IP address.

A secondary IP address is configured by issuing the ip address address mask secondary command. Unlike primary IP addresses, secondary IP addresses do not replace previously configured secondary IP addresses. Although secondary IP addresses work like primary IP addresses, a Cisco device will not use a secondary IP address to generate datagrams other than routing updates.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/command/ipaddr-cr-book/ipaddr-c1.html#wp1930336422>

QUESTION 149

You issue the following commands from OSPF router configuration mode on RouterA:

```
timers throttle lsa all 500 10000 40000
```

```
timers lsa arrival 2000
```

For how long will RouterA ignore identical LSAs that are received from a neighbor router? (Select the best answer.)

- A. One-half second
- B. two seconds
- C. 10 seconds
- D. 40 seconds

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:

Explanation:

RouterA will ignore identical linkstate advertisements (LSAs) that are received from a neighbor router for two seconds. Identical Open Shortest Path First (OSPF) LSAs are LSAs with the same LSA ID number, LSA type, and advertising router ID. To configure the interval at which a router will ignore identical LSAs, you should issue the `timers lsa arrival milliseconds` command? by default, this interval is set to a value of 1000 milliseconds, or one second.

The `timers throttle lsa all` command configures the rate at which LSAs are generated by a router, not received from a neighbor router. The syntax of the `timers throttle lsa all` command is `timers throttle lsa all start interval hold interval max interval`.

The `startinterval` timer is a value expressed in milliseconds that indicates how long the router will wait before generating an LSA. An LSA is generated immediately upon a local topology change, and the second LSA is generated sometime after the `startinterval` timer expires. In this scenario, the `startinterval` timer is set to a value of 500 milliseconds, or onehalf second. By default, the `startinterval` timer is set to a value of 0.

The `holdinterval` timer is a value expressed in milliseconds that indicates how long the router will wait before generating an LSA. In this scenario, the `holdinterval` timer is set to a value of 10000 milliseconds, or 10 seconds. By default, the `holdinterval` timer is set to a value of 5000 milliseconds, or five seconds.

The `maxinterval` timer is a value expressed in milliseconds that indicates how long the router will wait before generating an identical LSA, not receiving an identical LSA. In this scenario, the `maxinterval` timer is set to a value of 40000 milliseconds, or 40 seconds. By default, the `maxinterval` timer is set to a value of 5000 milliseconds, or five seconds.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fsolsath.html https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/ospf-s1.html#wp3611083381 https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/ospf-s1.html#wp3761694958

QUESTION 150

Which of the following statements is correct regarding an MPLS VPN super backbone that connects a customer's OSPF network? (Select the best answer.)

- A. The Area 0 backbone cannot be used on a customer's OSPF network because the backbone area is replaced by the MPLS VPN super backbone.
- B. Only one Area 0 backbone can be used on a customer's OSPF network and must be connected to the MPLS VPN super backbone.
- C. Only one Area 0 backbone can be used on a customer's OSPF network and must not be connected to the MPLS VPN super backbone.
- D. Multiple Area 0 backbones can be used on a customer's OSPF network, but all must be connected to the MPLS VPN super backbone.
- E. Multiple Area 0 backbones can be used on a customer's OSPF network and do not need to be connected to the MPLS VPN super backbone.

Correct Answer: D

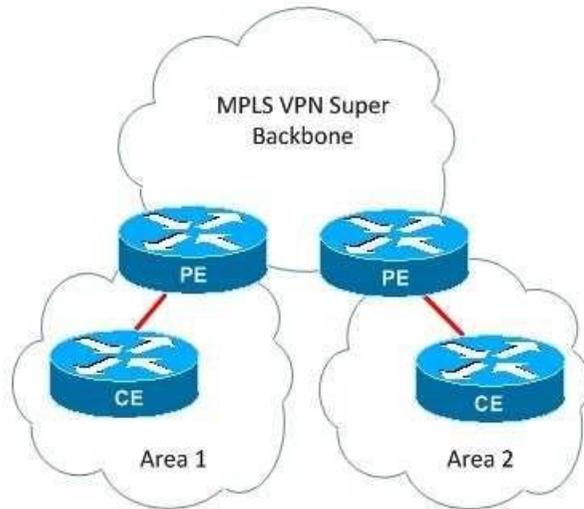
Section: (none)
Explanation

Explanation/Reference:

Explanation:

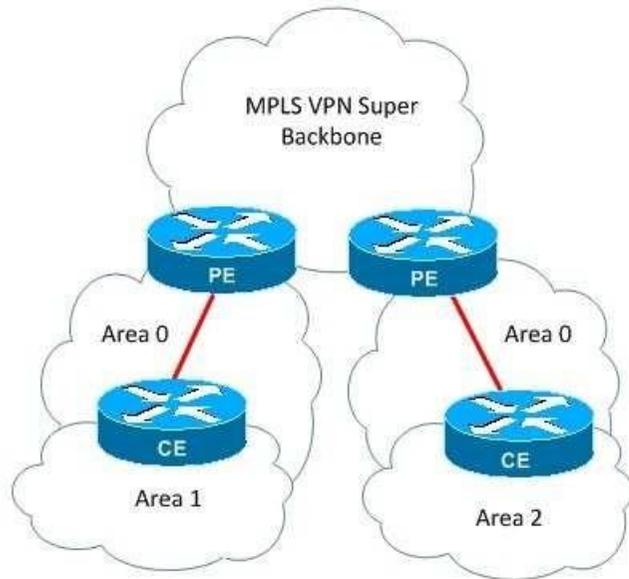
Multiple Area 0 backbones can be used on a customer's Open Shortest Path First (OSPF) network, but all must be connected to the Multiprotocol Label Switching (MPLS) virtual private network (VPN) super backbone. An MPLS VPN super backbone connects multiple sites over a service provider network, becoming either an extension of or a replacement for a customer's OSPF backbone.

The following graphic shows an MPLS VPN super backbone connecting two customer sites:



In this topology, the MPLS VPN super backbone has replaced the customer's OSPF backbone. The provider edge (PE) routers are area border routers (ABRs) and autonomous system boundary routers (ASBRs), and the customer edge (CE) routers are normal intra area routers.

Although an MPLS VPN super backbone can replace a customer's OSPF backbone, it does not have to do so. Instead, the super backbone can become an extension of a customer's OSPF backbone. However, each of the customer's backbone areas must connect to the super backbone, as shown in the following graphic:



In this topology, the customer has multiple Area 0 backbone areas connected to the super backbone. The PE routers are ABRs and ASBRs, and the CE routers are ABRs.

Reference: <https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/multiprotocol-label-switching-over-atm-mpls-over-atm/10472-mplsospf.html#backinfo>

QUESTION 151

DRAG DROP

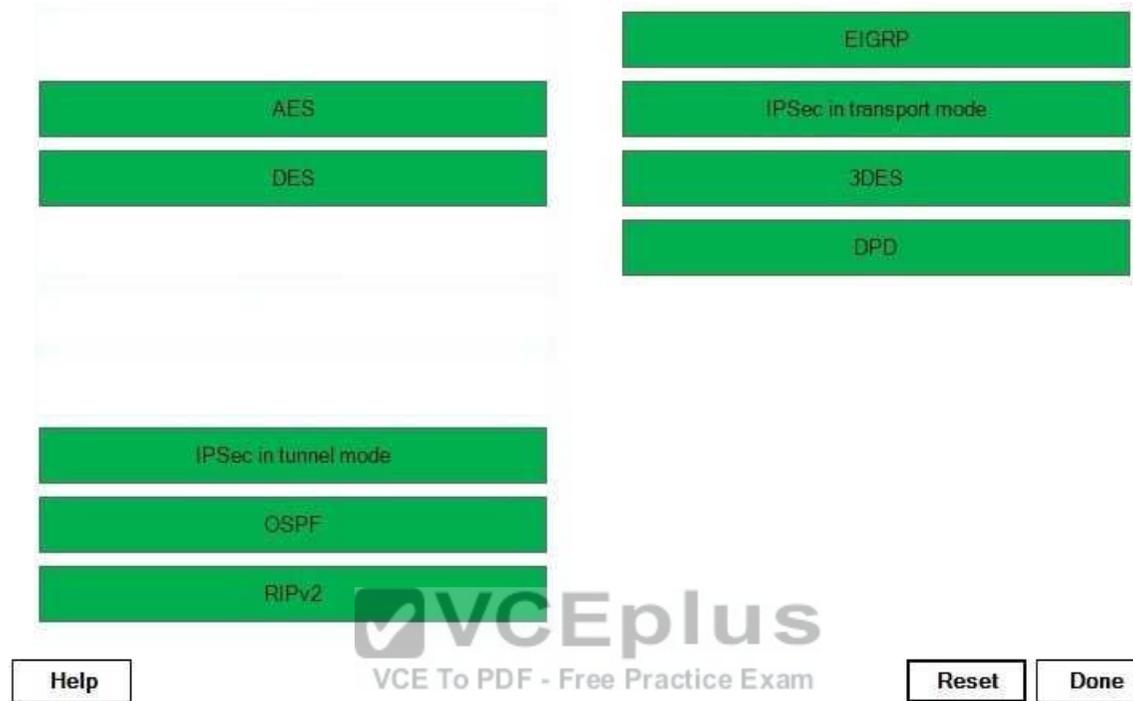
Select features from the left that are recommended by Cisco for implementing a scalable DMVPN, and place them on the corresponding boxes on the right. Multiple correct answers are possible.

Select and Place:

3DES	routing protocol
AES	IPSec mode
DES	encryption
DPD	failure detection
EIGRP	
IPSec in transport mode	
IPSec in tunnel mode	
OSPF	
RIPv2	

 **VCEplus**
VCE To PDF - Free Practice Exam

Correct Answer:



Section: (none)

Explanation

Explanation/Reference:

Explanation:

Dynamic Multipoint virtual private network (DMVPN) enables an administrator to easily configure scalable IP Security (IPSec) virtual private networks (VPNs) using a hub-and-spoke design. The hub router or routers are typically assigned a static IP address? the spoke routers can be dynamically addressed.

DMVPN requires Generic Routing Encapsulation (GRE), Next Hop Resolution Protocol(NHRP), and a dynamic routing protocol. NHRP is used to create a database of tunnel address to real address mappings. Although several routing protocols can be used to create a DMVPN, Cisco recommends that Enhanced Interior Gateway Routing Protocol (EIGRP) be used to enhance scalability.

A multipoint GRE (mGRE) tunnel is used to carry multiple IPSec or GRE tunnels. Although you can use either tunnel mode or transport mode, Cisco recommends that transport mode be used. In addition, strong encryption should be used, such as Triple Data Encryption Standard (3DES) or Advanced Encryption Standard (AES). Data Encryption Standard (DES) is not as strong as 3DES or AES.

You should enable Dead Peer Detection (DPD) to provide failure detection. By default, DPD messages are sent only if there is a 10second lull in traffic from a tunnel peer and only if there is outbound traffic destined for that tunnel peer. For example, if 10 seconds pass and RouterA has not received traffic from RouterB, RouterA prepares a DPD message for transmission. However, the DPD message is not sent to RouterB until RouterA has traffic to send to RouterB.

Reference:

<https://sso.cisco.com/autho/forms/CDClogin.html>

QUESTION 152

Which of the following events will trigger a log entry to be created with a severity level of 5 if syslog has been enabled on a router? (Select the best answer.)

- A. the Cisco IOS software failing to load
- B. a packet being denied by an ACL
- C. a router interface transitioning to the down state
- D. an invalid packet type being received on an interface

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A router interface transitioning to the down state will generate a log entry with a severity level of 5 if syslog has been enabled on a router. Syslog is a management protocol that can be used to transmit logging information from a device to a syslog server. When syslog is enabled on a router, logging messages are generated by the router and sent to the console or to a syslog server that is used to view and process the syslog messages. Log files that are generated by routers are categorized into one of the following severity levels:

- Level 0 -Emergency
- Level 1 -Alert
- Level 2 -Critical
- Level 3 -Errors
- Level 4 -Warnings
- Level 5 -Notifications
- Level 6 -Informational
- Level 7 -Debugging

A router interface transitioning to the down state would cause a log entry to be generated with a severity level of 5. A severity level of 5 indicates that a normal but significant event has occurred. System restart messages are also displayed at this level.

A log entry with a severity level of 6 indicates an informational message. An informational message is generated when an event such as a packet being denied as a result of matching an access control list (ACL) entry occurs. Reload request messages are also displayed at this level.

The Cisco IOS software failing to load would generate a log entry with a severity level of 0. A severity level of 0 indicates that the system is unusable; an emergency condition has occurred that has prevented the router from functioning.

If a router interface receives an invalid packet type, a log entry will be created with a severity level of 7. A severity level of 7 indicates a debugging message.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750/software/release/12-2_55_se/configuration/guide/scg3750/swlog.html#pgfId-1031557

QUESTION 153

On which interface type is turbo flooding not supported? (Select the best answer.)



<https://vceplus.com/>

- A. ARPAencapsulated Ethernet
- B. FDDI
- C. HDLCencapsulated Serial
- D. Token Ring



Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Turbo flooding is not supported on Token Ring interfaces. Turbo flooding is a Cisco feature that speeds up flooding of User Datagram Protocol (UDP) datagrams using the spanningtree algorithm. To enable turbo flooding, you should issue the following commands: ip forward-protocol turbo-flood ip forward-protocol spanning-tree

Turbo flooding is supported on the following interface types:

- Advanced Research Projects Agency (ARPA)encapsulated Ethernet
- Fiber Distributed Data Interface (FDDI)

-HighLevel Data Link Control (HDLC)encapsulated Serial

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/command/iap-cr-book/iap-i1.html#wp4079755394>

QUESTION 154

What is the default WTD maximum threshold for queue 1? (Select the best answer.)

- A. 50 percent
- B. 100 percent
- C. 200 percent
- D. 400 percent

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The default weighted taildrop (WTD) threshold for queue 1 is 400 percent. To configure the WTD thresholds, you should issue the `mls qos queues-et output threshold` command. The syntax of the `mls qos queue-set output threshold` command is `mls qos queue-set output qset-idthreshold [queue-id] drop-threshold1 droptreshold2 reserved-threshold maximum-threshold`.

When Quality of Service (QoS) is enabled, WTD is enabled and uses the default threshold values. The following table displays the default threshold values for WTD:

	Queue 1	Queue 2	Queue 3	Queue 4
Drop Threshold 1	100 percent	200 percent	100 percent	100 percent
Drop Threshold 2	100 percent	200 percent	100 percent	100 percent
Reserved Threshold	50 percent	100 percent	50 percent	50 percent
Maximum Threshold	400 percent	400 percent	400 percent	400 percent

The two drop thresholds are expressed as a percentage of the allocated memory of the queue. The reserved threshold is the percentage of allocated memory that is guaranteed for the queue. The maximum threshold is the maximum queue memory before packets are dropped.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/software/15-0_2_EX/qos/command_reference/b_qos_152ex_2960-x_cr/b_qos_152ex_2960x_cr_chapter_011.html#wp5865016930

QUESTION 155

Which of the following features does IGMPv3 support that IGMPv2 does not? (Select the best answer.)

- A. querier elections
- B. leave group messages
- C. groupspecific queries
- D. host membership report suppression
- E. multicast source filtering

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Internet Group Management Protocol version 3 (IGMPv3) supports multicast sourcefiltering. Multicast source filtering allows a host to specify the source addresses from which it will receive multicast traffic; it also allows a host to specify the source addresses from which it will not receive multicast traffic. IGMP version 1 (IGMPv1) and IGMPv2 do not provide support for multicast source filtering.

Although IGMPv3 supports querier elections, this feature was introduced in IGMPv2. The router with the lowest IP address on the subnet is elected as the querier. The querier is responsible for periodically sending out membership query messages to determine whether any hosts want to receive multicast packets for the multicast group. If at least one host responds with a membership report message, the querier will continue to send those multicast packets on that network segment. By default, membership query messages are sent every 60 seconds.

Although IGMPv3 supports leave group messages, this feature was introduced in IGMPv2.

In IGMPv1, a host leaves a multicast group silently. In IGMPv2, a host sends an IGMP leave message when it wants to leave a multicast group. IGMP routers maintain the IP address of the last reporter, which is the last host that sent a membership report message for that multicast group. If the last reporter sends a leave message, the IGMP router will wait an amount of time configured in the last member query response interval before sending a response and deleting the group. By default, the last member query response interval is one second.

Although IGMPv3 supports groupspecific queries, this feature was introduced in IGMPv2.

IGMPv1 queries are general queries sent to the 224.0.0.1 all hosts multicast address.

IGMPv2 queries are either general queries, which are sent to 224.0.0.1, or groupspecific queries, which are sent only to members of a particular multicast group. When IGMPv2 is used, the Max Response Time field in membership query messages contains a nonzero value. In IGMPv1 messages, the field is set to a value of 0, which is interpreted to mean 100 deciseconds, or 10 seconds. The IGMPv2 membership query message is the only message that contains a nonzero value in the Max Response Time field; all other message types set the field to a value of 0. IGMPv3 membership query messages, on the other hand, use a Max Resp Code field from which the Max Response Time value is derived.

IGMPv3 does not support host membership report suppression; in fact, IGMPv3 removed support for host membership report suppression. This feature, which is supported in IGMPv1 and IGMPv2, prevents the sending of a membership report if a similar report is detected from another host on the network. IGMPv3 removes this restriction.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfmulti.html#wp1066001
<https://www.ietf.org/rfc/rfc3376.txt>

QUESTION 156

Which of the following statements is correct regarding ISAKMP preshared keys that are stored in secure type 6 format? (Select the best answer.)

- A. The master key is stored in the router configuration and is encrypted with AES.
- B. The master key can be changed after it has been created.
- C. Deletion of the master key will unencrypt all of the encrypted passwords.
- D. Keys are encrypted as soon as you issue the key configkey password encryption masterkey command.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The master key can be changed after it has been created. Internet Security Association and Key Management Protocol (ISAKMP) preshared key encryption can be used to encrypt and store keys in secure type 6 format. To enable ISAKMP preshared key encryption, issue the following commands:

```
key config-key password-encryption master-key  
password encryption aes
```

The master key encrypts all of the other keys that are stored in the router configuration by using Advanced Encryption Standard (AES). Passwords are not encrypted until the password encryption aes command has been issued. The master key is not stored anywhere in the router configuration, nor can the master key be displayed.

To change the master key, issue the key configkey passwordencryption command. You will be prompted once for the old master key and twice for the new master key. If you successfully authenticate the old key, the existing encrypted preshared keys will be encrypted with the new master key.

You can delete the master key by issuing the no key configkey passwordencryptioncommand. However, the existing encrypted preshared keys will not be unencrypted, and they cannot be used by the router. Issuing the no password encryption aes command will also not unencrypt the existing preshared keys; once they are encrypted with secure type 6 encryption, they cannot be unencrypted.

Reference:

<https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/46420-pre-sh-keys-ios-rtr-cfg.html>

QUESTION 157

IP routing is not enabled on a Cisco router that you administer. You want to create a gateway of last resort on the router.

Which of the following commands should you issue to accomplish your goal? (Select the best answer.)

- A. ip route 0.0.0.0 0.0.0.0 s0/0
- B. ip defaultnetwork 10.10.100.0
- C. ip route 0.0.0.0 0.0.0.0 10.10.100.1
- D. ip defaultgateway 10.10.100.1

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the ip defaultgateway 10.10.100.1 command to create a gateway of last resort on the router in this scenario. The gateway of last resort is used when no route can be identified for routing packets. If a route cannot be determined and the gateway of last resort has been defined, packets are routed to the gateway's address. The ip defaultgateway command should be used to configure a gateway of last resort when IP routing is disabled on a Cisco router or switch. You should not issue the ip defaultnetwork 10.10.100.0 command to create a gateway of last resort on the router in this scenario, because the ip defaultnetworkcommand can be used only on devices that have IP routing enabled. Because IP routing is disabled on the router in this scenario, you cannot use the ip defaultnetwork command to create a gateway of last resort on the router. The ip defaultnetwork command is issued to flag routes in the routing table as candidates for the default route. If the ip defaultnetwork command is issued with a network address that matches an entry appearing in the routing table, that route becomes the default route. If there is no matching entry in the routing table, the route will not become a default route. You can issue multiple ip default network commands on a router; the router will use administrative distances (ADs) and metrics to determine the best default route.

You should not issue the ip route command to create a gateway of last resort on the router in this scenario, because the ip route command can be used only on devices that have IP routing enabled. Because IP routing is disabled on the router in this scenario, you cannot use the ip route command to create a gateway of last resort on the router. You can create a static gateway of last resort on a router by issuing the ip route command with a static route followed by an IP address. For example, you could issue the ip route 0.0.0.0 0.0.0.0 10.10.100.1 command to create a static gateway of last resort to the nexthop router at 10.10.100.1. Alternatively, you could specify a physical interface on the router as the gateway. For example, you could issue the ip route 0.0.0.0 0.0.0.0 s0/0 command to specify the Serial 0/0 interface on the local router as the gateway of last resort? packets with destinations not found in the routing table will be forwarded to the Serial 0/0 interface.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/routing-information-protocol-rip/16448-default.html>

QUESTION 158

Which of the following best describes how to create a MAC address that is mapped from a multicast IPv6 address? (Select the best answer.)

- A. Prepend FF to the last five octets of the IPv6 address.

- B. Prepend 3333 to the last four octets of the IPv6 address.
- C. Prepend FFFF to the last four octets of the IPv6 address.
- D. Prepend 0100.0CCC to the last two octets of the IPv6 address.
- E. Prepend 0180.C200 to the last two octets of the IPv6 address.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

To create a Media Access Control (MAC) address that is mapped from a multicast IPv6 address, prepend 3333 to the last four octets of the IPv6 address. This procedure is described in Internet Engineering Task Force (IETF) Request for Comments (RFC) 2464. The IPv6 prefix FF00::/8 is generally used for multicast addresses. IPv6 addresses in the

FF00::/8 range begin with the characters FF00 through FFFF. However, you would not prepend FF to the last five octets of the IPv6 address, nor would you prepend FFFF to the last four octets of the IPv6 address in order to create a MAC address that is mapped from a multicast IPv6 address.

You would not prepend 0100.0CCC to the last two octets of the IPv6 address. Multicast MAC address 0100.0CCC.CCCC is used by Cisco Discovery Protocol (CDP), Dynamic Trunking Protocol (DTP), Port Aggregation Protocol (PAgP), UniDirectional Link Detection (UDLD), and VLAN Trunking Protocol (VTP). CDP is a Layer 2 Cisco proprietary protocol that is used to advertise and discover only directly connected Cisco devices on a local network. DTP is a point-to-point protocol that is used to negotiate trunking. PAgP is an EtherChannel aggregation protocol. UDLD monitors a link to verify that both ends of the link are functioning.

VTP is used to centrally manage virtual LAN (VLAN) changes and to propagate those changes over trunk ports. Multicast MAC address 0100.0CCC.CCCD is used by 802.1D Spanning Tree Protocol (STP) to send nonnative VLAN bridge protocol data units (BPDUs).

You would not prepend 0180.C200 to the last two octets of the IPv6 address. Multicast MAC address 0180.C200.0000 is used by 802.1D STP to send native VLAN BPDUs. Multicast MAC address 0180.C200.0003 is used by 802.1X. Multicast MAC address 0180.C200.000E is used by Link Layer Discovery Protocol (LLDP).

Reference:

<https://tools.ietf.org/html/rfc2464>

QUESTION 159

Which of the following best defines an RD? (Select the best answer.)

- A. a value that indicates membership in an RFC 4364 VPN
- B. a path that labeled packets take through an MPLS network
- C. a value that enables RFC 4364 VPN customers to use overlapping IP address ranges
- D. a routing table instance for a VPN

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

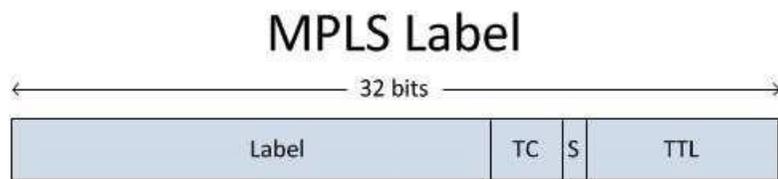
A route distinguisher (RD) is a value that enables Multiprotocol Label Switching (MPLS) virtual private network (VPN) customers to use overlapping IP address ranges; MPLS VPNs are described in Request for Comments (RFC) 4364. An ingress label switch router (LSR) creates a globally unique VPN version 4 (VPNv4) address by adding the RD to the beginning of an IP address. The LSR then assigns a label to the VPNv4 address prefix and stores the inbound to outbound label mapping in the Label Forwarding Information Base (LFIB). Authentication to the MPLS VPN is provided based on logical port and RD information. To create an RD, you should issue the rd value command, where the value parameter uses one of the following formats: - AS:nn, where AS is a 16bit autonomous system number (ASN) and nn is a 32bit decimal number

- A.B.C.D:nn, where A.B.C.D is a 32bit IP address and nn is a 16bit decimal number

There are three types of RDs: Type 0, Type 1, and Type 2. The type of RD configuration you create depends on how you issue the value parameter of the rd command and whether you are configuring a multicast VPN environment. Type 0 and Type 1 RDs are used in unicast configurations. A Type 0 RD is configured by issuing the value parameter of the rd command with the 16bit ASN in front of the 32bit decimal number. A Type 1 RD is configured by issuing the value parameter of the rd command with the 32bit decimal number in front of the 16bit ASN. A Type 2 RD is configured similarly to a Type 1 RD but only applies to multicast VPN configurations.

A route target (RT) is a value that is appended to a VPNv4 Border Gateway Protocol (BGP) route to indicate membership in an RFC 4364 MPLS VPN. Export RTs associate each route with one or more VPNs, and import RTs are associated with each VPN routing and forwarding (VRF) table to determine the routes that should be imported into the VRF? a VRF is a routing table instance for a VPN. By configuring import and export RTs, you can configure which sites can reach each other. For example, you can configure RTs so that CustomerA and CustomerB can communicate with ProviderZ, but CustomerA and CustomerB cannot communicate with one another. To configure RTs, you should issue the routetarget {import | export | both} value command, where the value parameter uses the same formats as the value parameter in the rd command.

A label switched path (LSP) is a path that labeled packets take through an MPLS network from one LSR to another. The 32bit MPLS label is used by LSRs to make forwarding decisions along the LSP. The MPLS label is placed between the Layer 2 header and the Layer 3 header. The structure of an MPLS label is shown below:



Reference:

CCIE Routing and Switching v5.0 Certification Guide, Volume 2, Chapter 11, MPBGP and Route Distinguishers, pp. 541-543

QUESTION 160

Which of the following statements is correct about external routes received by an NSSA? (Select the best answer.)

- A. External routes from an ASBR are converted to Type 3 LSAs and tunneled through the NSSA to the ABR, where they are converted to Type 5 LSAs.
- B. External routes from an ASBR are converted to Type 3 LSAs and tunneled through the NSSA to the ABR, where they are converted to Type 7 LSAs.
- C. External routes from an ASBR are converted to Type 5 LSAs and tunneled through the NSSA to the ABR, where they are advertised as Type 5 LSAs.
- D. External routes from an ASBR are converted to Type 5 LSAs and tunneled through the NSSA to the ABR, where they are converted to Type 7 LSAs.
- E. External routes from an ASBR are converted to Type 7 LSAs and tunneled through the NSSA to the ABR, where they are converted to Type 5 LSAs.
- F. External routes from an ASBR are converted to Type 7 LSAs and tunneled through the NSSA to the ABR, where they are advertised as Type 7 LSAs.

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

External routes from an autonomous system boundary router (ASBR) are converted to Type 7 linkstate advertisements (LSAs) and tunneled through the not-so-stubby area (NSSA) to the area border router (ABR), where they are converted to Type 5 LSAs. An NSSA is basically a stub area that contains one or more ASBRs. Type 7 LSAs are used to advertise external routes that are injected into an Open Shortest Path First (OSPF) NSSA.

External routes from an ASBR into an NSSA are not converted to Type 5 LSAs and tunneled through the NSSA to the ABR. Type 5 LSAs are used to advertise external routes that are injected into an OSPF backbone or standard area. When an ASBR in a backbone area or a standard area receives an external route, the ASBR creates a Type 5 LSA to advertise the external route. Like stub areas, NSSAs do not accept or create Type 5 LSAs.

External routes from an ASBR into an NSSA are not converted to Type 3 LSAs and tunneled through the NSSA to the ABR. Type 3 LSAs are used to advertise the area's subnets to another area. NSSAs accept Type 3 LSAs. However, Type 3 LSAs are not created by ASBRs; they are created by ABRs. Totally stubby areas and totally NSSAs do not accept Type 3, 4, or 5 summary LSAs. These LSAs are replaced by a default route at the ABR. As a result, routing tables are kept small within the totally stubby area.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html#definensstub>

QUESTION 161

Which of the following are true of both traffic policing and traffic shaping? (Select 2 choices.)

- A. Both buffer excess traffic.
- B. Both remark excess traffic.
- C. Both limit bandwidth utilization.
- D. Both use a token bucket.
- E. Both smooth traffic.

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:

Explanation:

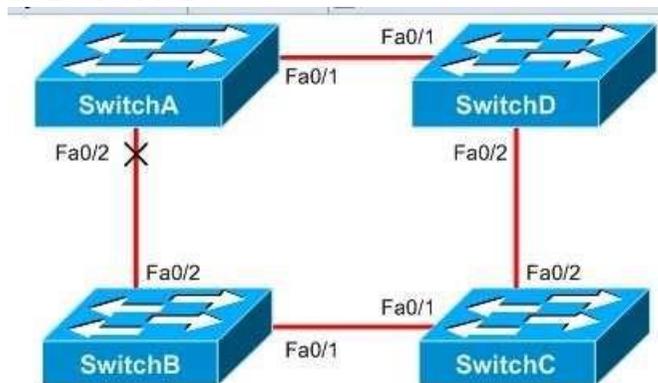
Traffic policing and traffic shaping both limit bandwidth utilization, and both use a token bucket. Traffic policing is used to slow down traffic to a value that the medium can support, to monitor bandwidth utilization, to enforce bandwidth limitations at the service provider edge, and to remark traffic that exceeds the Service Level Agreement (SLA). Traffic shaping is used to slow down traffic due to congestion, to enforce bandwidth rates, and to send traffic classes at different rates. To control the rate at which an interface sends packets, traffic policing and traffic shaping use a token bucket. Tokens are put into the token bucket at a specified rate, and tokens are removed from the bucket as bits are sent through the interface. If there are not enough tokens to send a packet, traffic policing drops or remarks the packet. As a result, traffic policing can cause traffic to be bursty. By contrast, traffic shaping queues packets when there are not enough tokens to send them. This generates a "leaky bucket" effect, which smooths traffic into a constant flow rather than a variable, bursty flow. The shaping parameters can also be configured so that packets can be sent in excess of the committed information rate (CIR) for a short period of time.

Traffic shaping does not remark excess traffic. Instead, traffic shaping buffers excess traffic and outofprofile packets in memory until the queue is full and drops traffic only if the queue is full. By contrast, traffic policing drops or remarks excess traffic and out-of-profile packets.

Reference:

<https://www.cisco.com/c/en/us/support/docs/quality-of-service-qos/qos-policing/19645-policevsshape.html>

QUESTION 162



You administer the switched network shown above. All switches are configured to use STP.

Which of the following could cause a switching loop? (Select the best answer.)

- A. configuring Fa0/1 on SwitchA and SwitchD for halfduplex mode
- B. configuring Fa0/1 on SwitchA and SwitchD for fullduplex mode
- C. configuring Fa0/1 on SwitchA and SwitchD to autonegotiate duplex settings
- D. configuring Fa0/1 on SwitchA for halfduplex mode and configuring Fa0/1 on SwitchD to autonegotiate duplex settings
- E. configuring Fa0/1 on SwitchA for fullduplex mode and configuring Fa0/1 on SwitchD to autonegotiate duplex settings

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Configuring Fa0/1 on SwitchA for fullduplex mode and configuring Fa0/1 on SwitchD to autonegotiate duplex settings could cause a switching loop. A switch port that has been manually configured to use fullduplex or halfduplex mode does not respond to a port that is attempting to autonegotiate duplex settings. When the autonegotiating port receives no reply, it will use the default duplex settings. A port configured to transmit at 100 Mbps defaults to halfduplex mode, and a port configured to transmit at 1000 Mbps defaults to fullduplex mode. Therefore, Fa0/1 on SwitchD will use halfduplex mode, causing a duplex mismatch with Fa0/1 on SwitchA.

Duplex mismatches can cause collisions, alignment errors, and intermittent connectivity. You can detect a duplex mismatch by monitoring a switch for %CDP4DUPLEXMISMATCH error messages. Additionally, you can issue the `show interfaces interface` command, which displays interface counter information. If you see an abnormal increase in frame check sequence (FCS) errors and alignment errors on a halfduplex port, you should suspect a duplex mismatch. An abnormal increase in FCS errors and runts on a fullduplex port is also an indicator of a duplex mismatch.

When SwitchD is in halfduplex mode, it performs carrier sense to determine whether the link is clear before sending packets. However, SwitchA does not perform carrier sense, because it is configured for fullduplex mode; this is what causes intermittent connectivity problems with a duplex mismatch. When SwitchA sends a high volume of traffic to SwitchD, the Spanning Tree Protocol (STP) bridge protocol data units (BPDUs) from SwitchD could be lost due to collisions. If SwitchA does not receive those BPDUs from SwitchD, SwitchA will assume that there is a loss of connectivity to SwitchD, unblock port Fa0/2, and forward all packets to SwitchB, which could cause a switching loop.

Manually configuring Fa0/1 on SwitchA and SwitchD to use the same duplex mode would not cause a switching loop. Configuring both switch ports for fullduplex mode would enable both ports to send and receive data simultaneously. Configuring both switch ports for halfduplex mode would enable only one port to send data at a time; however, communication could still occur, albeit slowly.

Configuring Fa0/1 on SwitchA for halfduplex mode and configuring Fa0/1 on SwitchD to autonegotiate duplex settings would not cause a switching loop. Fa0/1 on SwitchD would not receive a duplex autonegotiation response from SwitchA, so SwitchD would default to halfduplex mode. Both switch ports would then be using the same duplex mode, thereby enabling communication between the two ports.

Configuring Fa0/1 on SwitchA and SwitchD to autonegotiate duplex settings would not cause a switching loop. If both sides of a link were configured to autonegotiate duplex settings, they would negotiate fullduplex mode if both ports support fullduplex operation.

If either side of the link does not support fullduplex operation, the ports would negotiate halfduplex mode.

Reference:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10556-16.html#duplex>
https://www.cisco.com/c/en/us/support/docs/switches/catalyst-6500-series-switches/17053-46.html#auto_neg_valid

QUESTION 163

How long does it take for a PIM network to converge after the failure or addition of an Anycast RP? (Select the best answer.)

- A. about a second
- B. between one and two minutes
- C. as quickly as unicast routing converges
- D. as soon as an administrator manually reconfigures the RPs
- E. as soon as a multicast source or multicast listener attempts to contact the RP

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A Protocol Independent Multicast (PIM) network will converge as quickly as unicast routing converges after the failure or addition of an Anycast rendezvous point (RP). Convergence does not require a certain period of time, it does not require the presence or absence of multicast sources or receivers, and it does not require that an administrator manually reconfigure the RP.

Anycast RP enables multiple RPs to provide redundancy and load-sharing capabilities. Each multicast receiver will use the closest RP. Each of the Anycast RPs must be configured as Multicast Source Discovery Protocol (MSDP) peers of one another because they use MSDP to share information about multicast sources. All the Anycast RPs must have the same IP address on a loopback interface. Downstream routers must be configured with the shared loopback address of the Anycast RPs, either statically by using the `ip pim rp-address` command or dynamically by using AutoRP or Bootstrap Router (BSR).

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/solutions_docs/ip_multicast/White_papers/anycast.html

<https://www.ietf.org/rfc/rfc4610.txt>

QUESTION 164

Which of the following requires a physical RP? (Select the best answer.)

- A. PIM-DM
- B. PIM-SM
- C. PIM-SDM
- D. PIM-SSM
- E. Bidirectional PIM

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Protocol Independent Multicast sparse mode (PIMSM) requires a physical rendezvous point (RP). An RP is a well-connected, centrally located router that is responsible for keeping track of multicast group membership information. When a host wants to join a multicast group, it sends an Internet Group Management Protocol (IGMP) membership report message to its local router. The local router adds the interface to the multicast tree and forwards the message to the RP. This process creates a branch of the multicast tree from the host to the RP. A branch is not pruned until the group member leaves the group.

Bidirectional PIM (bidirPIM) can use a physical RP, but the RP does not have to be a physical device. Instead, bidirPIM can use a phantom RP, which is an address that is used as the RP address but is not assigned to a physical device. A physical RP is not required with bidir-PIM, because bidirPIM designated forwarders (DFs) can forward traffic up the shared tree directly to multicast receivers.

PIM dense mode (PIMDM) does not require an RP to keep track of multicast group membership information. Instead, PIMDM routers assume that all interfaces contain group members, so they periodically flood multicast traffic out all available interfaces, which causes a traffic spike. Each router in the network determines whether any hosts are interested in receiving the multicast traffic. If so, the router forwards the multicast traffic. If not, the router sends a prune message back to the multicast source and that branch of the multicast tree is pruned for a short period of time.

PIM sparse-dense mode (PIMSDM) does not require an RP. PIMSDM uses a combination of sparse mode and dense mode. The mode is determined on a per group basis. PIMSDM routers use sparse mode if an RP exists for a multicast group and use dense mode if no RP exists for a multicast group.

PIM Source-Specific Multicast (PIMSSM) does not require an RP. PIMSSM is best suited for onetomany applications, which are also called broadcast applications. When PIMSSM is used, a multicast host can specify the source addresses from which it will accept multicast traffic. Like PIMDM, PIMSSM uses sourcebased distribution trees, which are built from the multicast source to the multicast receivers.

Reference:

http://docwiki.cisco.com/wiki/Internet_Protocol_Multicast#PIM_Sparse_Mode

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

QUESTION 165

You have issued the following commands on Router1:

```
router eigrp 86 network
192.168.100.0 eigrp
stub
```

Which of the following commands or command sets must you issue to enable Router1 to advertise connected and summary routes? (Select the best answer.)

- A. No commands are necessary; the eigrp stub command enables the stub router to advertise connected and summary routes.
- B. eigrp stub connected summaryC. eigrp stub connected eigrp stub summary

D. eigrp stub static connected summary
E. redistribute connected redistribute static

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

No commands are necessary; the eigrp stub command enables the stub router to advertise connected and summary routes. The eigrp stub command configures a router as a stub router. When the eigrp stub command is issued without parameters, summary routes and directly connected routes are advertised by default. The following options can be issued with the eigrp stub command:

- receive-only-configures the router to receive routes but not advertise routes
- connected -configures the router to advertise directly connected networks
- redistributed -configures the router to advertise routes learned from another protocol
- static -configures the router to advertise static routes
- summary-configures the router to advertise summary routes
- leak-map map-name-configures the router to advertise specific dynamically learned prefixes

With the exception of the receive only option, all of the options can be included together in the eigrp stub command. For example, to configure a stub router to advertise connected, static, and summary routes, you should issue the eigrp stub connected static summary command.

When you issued the eigrp stub command in the scenario, you enabled Router1 to advertise connected and summary routes. Although issuing the eigrp stub connected summary command would also enable Router1 to advertise connected and summary routes, you are not required to advertise these routes.

Issuing the eigrp stub connected and eigrp stub summary commands in sequence would not enable Router1 to advertise connected and summary routes; only the last command issued defines the routes that are advertised by Router1. Thus, issuing these two commands in sequence would configure Router1 to advertise only summary routes.

Issuing the eigrp stub static connected summary command would enable Router1 to advertise connected, summary, and static routes. However, you are not required to advertise static routes, and the eigrp stub command already configures Router1 to advertise connected and summary routes.

The redistribute static command enables a router to redistribute static routes into Enhanced Interior Gateway Routing Protocol (EIGRP) but does not configure the router to advertise those static routes. Similarly, the redistribute connected command enables a router to redistribute directly connected routes into EIGRP but does not configure the router to advertise those connected routes.

Reference:

https://www.cisco.com/en/US/technologies/tk648/tk365/technologies_white_paper0900aecd8023df6f.html

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/command/ire-cr-book/ire-a1.html#wp1217649486

QUESTION 166

Which of the following statements are true regarding the differences between TACACS+ and RADIUS? (Select 2 choices.)

- A. TACACS+ encrypts the entire body of a packet, whereas RADIUS encrypts only the password.
- B. TACACS+ combines authorization and authentication functions, whereas RADIUS separates authentication, authorization, and accounting functions.
- C. TACACS+ provides router command authorization capabilities, whereas RADIUS does not provide router command authorization capabilities.
- D. TACACS+ uses UDP, whereas RADIUS uses TCP.
- E. TACACS+ is an IETF standard protocol, whereas RADIUS was developed by Cisco.

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Terminal Access Controller Access Control System Plus (TACACS+) encrypts the entire body of a packet, whereas Remote Authentication Dial-In User Server (RADIUS) encrypts only the password; also, TACACS+ provides router command authorization capabilities, whereas RADIUS does not provide router command authorization capabilities. TACACS+ is a Cisco proprietary protocol used during Authentication, Authorization, and Accounting (AAA) operations. TACACS+ provides more security and flexibility than RADIUS; because TACACS+ can be used to encrypt the entire body of a packet, users who intercept the encrypted packet cannot view the user name or contents of the packet. TACACS+ provides more flexibility by separating the authentication, authorization, and accounting functions of AAA. This enables more granular control of access to resources. TACACS+ gives administrators more control over access to configuration commands; users can be permitted or denied access to specific configuration commands. Because of this flexibility, TACACS+ is used with Cisco Secure Access Control Server (ACS), which is a software tool that is used to manage user authorization for router access.

RADIUS, not TACACS+, is an Internet Engineering Task Force (IETF) standard protocol. Like TACACS+, RADIUS is a protocol used with AAA operations. However, RADIUS is less secure and less flexible than TACACS+. RADIUS encrypts only the password of a packet; the rest of the packet would be viewable if the packet was intercepted by a malicious user. With RADIUS, the authentication and authorization functions of AAA are combined into a single function, which limits the flexibility that administrators have when configuring these functions. Furthermore, RADIUS does not provide router command authorization capabilities. TACACS+ uses Transmission Control Protocol (TCP) for transport. By contrast, RADIUS uses User Datagram Protocol (UDP) for packet delivery.

Reference:

<https://www.cisco.com/c/en/us/support/docs/security-vpn/remote-authentication-dial-user-service-radius/13838-10.html#comparing>

QUESTION 167

Which of the following addresses are not used by EIGRPv6 to form neighbor relationships? (Select the best answer.)

- A. addresses in the same subnet
- B. addresses in different subnets
- C. link-local addresses
- D. global addresses

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Global addresses are not used by Enhanced Interior Gateway Routing Protocol version 6 (EIGRPv6) to form neighbor relationships. Only linklocal addresses are used by EIGRPv6 to form neighbor relationships.

EIGRPv6 is also referred to as EIGRP for IPv6. To enable EIGRPv6 on a router, you should issue the `ipv6 router eigrp asnumber` command in global configuration mode, where `asnumber` is the autonomous system (AS) number, and then issue the `no shutdown` command in router configuration mode to start the routing process. If no IPv4 or IPv6 addresses are configured on the router, you must also issue the `routerid id` command in router configuration mode to manually configure a router ID, where `id` is a 32bit value similar to an IPv4 address.

For a neighbor relationship to form between two routers running EIGRP for IPv4, the primary IP address of each router must be on the same subnet; EIGRP will not form a neighbor relationship over a secondary IP address. However, EIGRPv6 does not require that neighbors be in the same subnet to form a neighbor relationship. Linklocal addresses are significant only on the local link, so EIGRPv6 routers must share a common medium. Therefore, it does not matter whether the linklocal addresses are on the same subnet or not.

Reference:

<https://learningnetwork.cisco.com/docs/DOC-11783> <https://learningnetwork.cisco.com/servlet/JiveServlet/downloadBody/8347-102-3-41650/CCNP%2520Route.pdf>

QUESTION 168

In which of the following situations does a router use AD values to determine route selection? (Select the best answer.)

- A. when multiple routes to the same destination network are received, and all of these routes are received from the same routing protocol
- B. when multiple routes to the same destination network are received, and each of these routes is received from a different routing protocol
- C. when multiple routes to different destination networks are received, and all of these routes are received from the same routing protocol
- D. when multiple routes to different destination networks are received, and each of these routes is received from a different routing protocol

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A router uses administrative distance (AD) values to determine route selection when multiple routes to the same destination network are received, and each of these routes is received from a different routing protocol. Lower ADs are preferred over higher ADs. The following list contains the most commonly used ADs:

Route Source	Distance
Connected route	0
Static route	1
EIGRP summary route	5
eBGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
iBGP	200
Unknown	255

A router uses metrics to determine route selection when multiple routes to the same destination network are received, and all of these routes are received from the same routing protocol. Each routing protocol uses different metrics. For example, Routing Information Protocol (RIP) uses hop count as a metric, Open Shortest Path First (OSPF) uses cost as a metric, and Enhanced Interior Gateway Routing Protocol (EIGRP) uses bandwidth and delay by default. When a routing protocol contains multiple routes to the same destination network, a router prefers the route with the lowest metric.

A router uses prefix lengths to determine route selection when multiple routes to different destination networks are received, regardless of the routing protocol. When multiple routes to overlapping networks exist, a router will prefer the most specific route, which is the route with the longest prefix match. For example, if a router has a packet destined to 10.1.1.1, it will prefer a route to 10.1.1.0/24 over a route to 10.1.1.0/16, and it will prefer a route to 10.1.1.0/30 over a route to 10.1.1.0/24.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8651-21.html>

QUESTION 169

Which of the following EEM policy components is optional and contains code libraries? (Select the best answer.)

- A. event register keyword
- B. namespace import
- C. body
- D. entry status

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The namespace import component of the Embedded Event Manager (EEM) policy is optional and contains code libraries. An EEM policy can be written as an applet in the commandline interface (CLI) or in Tool Command Language (Tcl). EEM policies contain instructions on what action should take place if a defined event occurs. An EEM policy can consist of the following six components:

- Event register keyword
- Environmental must defines
- Namespace import
- Entry status
- Body
- Exit status

The event register keyword and the body are both required components of an EEM policy? the remaining four components are all optional. The event register keyword describes, registers, and schedules the event that is to be detected by the policy. The body contains the instructions regarding the actions to be carried out. The environmental must defines component determines whether required environmental variables have been defined before recovery actions are taken. The entry status determines whether another policy has been previously run for the defined event. The exit status determines whether the default action will be performed.

Reference:

<https://search.cisco.com/search?query=Cisco%20IOS%20Network%20Management%20Configuration%20Guide&locale=enUS&tab=Cisco>

QUESTION 170

Which of the following is most likely related to microbursts occurring on a network? (Select the best answer.)



<https://vceplus.com/>

- A. The network administrator reduces the size of the buffer to prevent packet loss.
- B. More than one device is sending traffic to a single destination at the same time.
- C. A gradual increase in traffic has occurred over a long period of time.
- D. An interface has shut down.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Of the available choices, more than one device sending traffic to a single destination at the same time is most likely related to microbursts occurring on a network.

A microburst is a significant increase in traffic over a very short period of time that can result in packet loss.

If the buffer size is inadequate and the buffer limit is exceeded, packets are dropped. Reducing the buffer size would increase, not decrease, the negative effects of a microburst; increasing the size of the buffer is the most effective way of avoiding packet loss due to a microburst. However, even if the network administrator reduces the size of the buffer, that action in itself will not cause a microburst to occur. Although a shutdown interface might cause some packets to not reach their destination, a shutdown interface would not cause microbursts to occur on a network.

Reference:

https://www.cisco.com/c/en/us/products/collateral/switches/nexus-5000-series-switches/white-paper-c11-733020.html#_Toc401129774

QUESTION 171

Which of the following statements is true regarding LISP? (Select the best answer.)

- A. LISP requires preconfigured tunnel endpoints.
- B. LISP must be running on both ends of the tunnel.
- C. The MR stores the registered EID prefixes and a mapping database.
- D. RLOC addresses are the IP addresses and prefixes that identify different routers in the IP network.
- E. LISP is used as the control plane protocol for EIGRP OTP.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Route Locator (RLOC) addresses are the IP addresses and prefixes that identify different routers in the IP network. Locator Identity Separation Protocol (LISP) splits the device identity and its location into separate numbering spaces. The Endpoint Identifier (EID) contains the locally relevant device identity and is used for endsite addressing. The RLOC contains the globally routed location of the device and is used to forward traffic between different networks.

The MapServer (MS), not the MapResolver (MR), stores the registered EID prefixes; the MS contains the mapping database of EID to RLOC mappings. The MR receives map request queries from LISP site Ingress Tunnel Routers (ITRs) when they attempt to populate the local mapcache of resolved EIDtoRLOC mappings.

An ITR receives packets from internal hosts and forwards them to external sites. Egress Tunnel Routers (ETRs) receive packets from external sites and forward them to internal hosts. If an edge device is both an ITR and an ETR, it is often called an xTR.

LISP tunnels are dynamically configured and do not require preconfigured endpoints. One advantage of LISP is its ability to offer mobility and scalability to a network. Endpoints can be relocated within a network and retain their configurations, including IP addressing, easing management tasks related to mobile endpoint devices.

LISP does not have to be running on both ends of a tunnel. LISP is designed to communicate with networks that are not using LISP.

LISP is not used as the control plane protocol for Enhanced Interior Gateway Routing Protocol (EIGRP) Over the Top (OTP); however, LISP is used as the data plane protocol for EIGRP OTP. EIGRP OTP is used to create a single contiguous EIGRP routing domain between sites over a service provider network. All Customer Edge (CE) routers must be configured for EIGRP OTP, and EIGRP neighbors must be manually configured with the neighbor command.

Reference:

https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/DCI/5-0/LISPMobility/DCI_LISP_Host_Mobility/LISPMobile_2.html#wp1277848

QUESTION 172

You want to ensure that a route in a VRF named ce1 uses the default gateway address of 172.16.0.1, which is commonly accessible from all other addresses on the router but does not exist in the ce1 table.

Which of the following commands should you issue? (Select the best answer.)

- A. ip route 0.0.0.0 0.0.0.0 172.16.0.1 global
- B. ip route 0.0.0.0 0.0.0.0 172.16.0.1 permanent
- C. ip route vrf 0.0.0.0 0.0.0.0 172.16.0.1 permanent
- D. ip route vrf 0.0.0.0 0.0.0.0 172.16.0.1 global
- E. ip route vrf ce1 0.0.0.0 0.0.0.0 172.16.0.1 global
- F. ip route vrf ce1 0.0.0.0 0.0.0.0 172.16.0.1 permanent



Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the ip route vrf ce1 0.0.0.0 0.0.0.0 172.16.0.1 global command to ensure that the route in a VPN routing and forwarding (VRF) named ce1 uses the default gateway address of 172.16.0.1, which is commonly accessible from all other addresses on the router but does not exist in the ce1 table. The global keyword configures VRF ce1 to use the 172.16.0.1 gateway that is present in the global routing table instead of attempting to look it up in the VRF ce1 routing table.

The global routing table stores paths that can be accessed by using any of the addresses on the router, not just the addresses associated with a given VRF. However, the global keyword applies only to the gateway address in the command, not to the entire static route. To configure a static route to apply to a given VRF, you should issue the vrf keyword along with the name of the VRF to which you want the route to apply. In this scenario, the ip route vrf ce1 0.0.0.0 0.0.0.0 172.16.0.1 global command configures a default route for the VRF named ce1.

You should not issue the ip route vrf ce1 0.0.0.0 0.0.0.0 172.16.0.1 permanent command in this scenario. The permanent keyword ensures that a route will not be removed from the associated VRF table even if the interface associated with the route is shut down. There are no conditions in this scenario that require you to issue the command with the permanent keyword.

You should not issue either the ip route vrf 0.0.0.0 0.0.0.0 172.16.0.1 global command or the ip route vrf 0.0.0.0 0.0.0.0 172.16.0.1 permanent command in this scenario. Neither of those commands contain valid syntax.

Reference: <https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/24508-internet-access-mpls-vpn.html#conf> https://www.cisco.com/c/en/us/td/docs/ios/12_2/switch/command/reference/fswtch_r/xfscmd2.html#wp1029959

QUESTION 173

Which of the following observations about potential BGP enhancements were documented in RFC 6774? (Select the best answer.)

- A. possible modifications to the bestpath algorithm
- B. possible software upgrades for PE routers
- C. possible addition of a session between a route reflector and its client
- D. possible addition of a fouroctet Path Identifier

Correct Answer: C

Section: (none)

Explanation



Explanation/Reference:

Explanation:

Observations about the possible addition of a session between a router reflector and its client were documented in Request for Comments (RFC) 6774, which discusses the distribution of diverse Border Gateway Protocol (BGP) paths. Specifically, RFC 6774 observed that BGP as it is typically deployed has no mechanism for distributing paths that are not considered the best path between speakers. However, the possible addition of a session between a route reflector and its client could enable a BGP router to distribute alternate paths.

RFC 6774 does not document possible modifications to the bestpath algorithm, nor does it document possible software upgrades for provider edge (PE) routers that are acting as route reflector clients. Although the document does discuss a possible means of distributing paths other than the best path, the means by which BGP determines the best path to a destination were not changed. Therefore, no software upgrade is required.

The BGP AddPaths proposal, not RFC 6774, proposed the possible addition of a fouroctet Path Identifier to Network Layer Reachability Information (NLRI) in order to enable BGP to distribute multiple paths.

Reference:

<https://tools.ietf.org/html/rfc6774>

QUESTION 174

Which of the following statements best describes poison reverse? (Select the best answer.)

- A. Poison reverse prevents switching loops.
- B. Poison reverse prevents routing loops by advertising a route as unreachable to all devices.
- C. Poison reverse prevents routing loops by advertising a route as unreachable to the interface from which the route was received.
- D. Poison reverse prevents routers from advertising a route through the same interface from which the route was learned.
- E. Poison reverse suppresses information regarding a better path to a route for a specified period of time.
- F. Poison reverse synchronizes VLAN configuration information between switches.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference:

Poison reverse prevents routing loops by advertising a route as unreachable to the interface from which the route was received. Split horizon is similar to poison reverse in that both methods prevent routing loops. However, poison reverse advertises a route as unreachable to the source interface, whereas split horizon does not.

Route poisoning is similar to poison reverse in that both methods prevent routing loops by advertising a route as unreachable. However, route poisoning immediately sends the advertisements to all interfaces, not just to the source interface.

Split horizon prevents routers from advertising a route through the same interface from which the route was learned. Thus split horizon prevents routing loops. By default, split horizon is enabled on all interfaces except those on which Frame Relay encapsulation or Switched Multimegabit Data Service (SMDS) encapsulation is enabled.

Holddown timers suppress information regarding a better path to a route for a specified period of time. When a router receives a routing update stating that a route is unreachable, the router waits a specified amount of time before accepting routes advertised by other sources.

Spanning Tree Protocol (STP) prevents switching loops on a network. Switching loops can occur when there is more than one switched path to a destination. The spanning tree algorithm determines the best path through a switched network, and any ports that create redundant paths are blocked. If the best path becomes unavailable, the network topology is recalculated and the port connected to the next best path is unblocked.

VLAN Trunking Protocol (VTP) is used to synchronize VTP and virtual LAN (VLAN) configuration information between switches. For switches to synchronize information over VTP, the following configuration parameters must match on all switches:

- VTP domain name
- VTP password
- VTP version

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html#splithorizon>

QUESTION 175

You are considering moving your company's software development to a public cloudbased solution. Which of the following are least likely to increase? (Select 2 choices.)

- A. availability
- B. redundancy
- C. security
- D. mobility
- E. control
- F. scalability

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Of the choices provided, security and control are least likely to increase. With a public cloudbased solution, the service provider, not the customer, controls the cloud infrastructure and devices. Therefore, physical security of the data and hardware is no longer in the customer's control. In addition, resources stored in the public cloud are typically accessed over the Internet. Care must be taken so that the data can be accessed securely.

Availability, redundancy, mobility, and scalability are all likely to increase by moving to a public cloudbased solution. Cloudbased resources are typically spread over several devices, sometimes even in multiple geographic areas, thereby ensuring availability. If one device or location becomes unavailable, other devices and locations can handle the workload. Data stored on cloudbased resources can be copied or moved to other devices or locations, thereby increasing redundancy and mobility. As usage increases, additional devices can be brought online, thereby providing scalability.

Reference:

<https://www.cisco.com/c/en/us/about/press/internet-protocol-journal/back-issues/table-contents-45/123-cloud1.html>

QUESTION 176

In a threenode OpenStack architecture, which services are part of the compute node? (Select 2 choices.)

- A. Ceilometer Agent
- B. Ceilometer Core
- C. Neutron DHCP Agent
- D. Neutron Server
- E. Nova Hypervisor
- F. Nova Management
- G. Correct

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In a three-node OpenStack architecture, the Ceilometer Agent and the Nova Hypervisor services are part of the compute node. OpenStack is an open source cloud computing platform. Each OpenStack modular component is responsible for a particular function, and each component has a code name. The following list contains several of the most popular OpenStack components:

- Nova -OpenStack Compute: manages pools of computer resources
- Neutron -OpenStack Networking: manages networking and addressing
- Cinder -OpenStack Block Storage: manages block level storage devices
- Glance -OpenStack Image: manages disk and server images
- Swift -OpenStack Object Storage: manages redundant storage systems
- Keystone -OpenStack Identity: is responsible for authentication
- Horizon -OpenStack Dashboard: provides a graphical user interface (GUI)
- Ceilometer -OpenStack Telemetry: provides counter based tracking that can be used for customer usage billing

A three node OpenStack architecture consists of the compute node, the controller node, and the network node. The compute node consists of the following services:

- Nova Hypervisor
- Kernel based Virtual Machine (KVM) or Quick Emulator (QEMU)
- Neutron Modular Layer 2 (ML2) PlugIn
- Neutron Layer 2 Agent
- Ceilometer Agent
- The controller node consists of the following services:

- Keystone
- Glance
- Nova Management
- Neutron Server
- Neutron ML2 PlugIn
- Horizon
- Cinder
- Swift
- Ceilometer Core

The network node consists of several Neutron services:

- Neutron ML2 PlugIn
- Neutron Layer 2 Agent
- Neutron Layer 3 Agent
- Neutron Dynamic Host Configuration Protocol (DHCP) Agent

Reference: <https://www.redhat.com/archives/rdo-list/2014-November/pdfzGvyHATdWc.pdf#page=12>

QUESTION 177

Which of the following is the mandatory transport protocol for NETCONF? (Select the best answer.)

- A. SSH
- B. SNMP
- C. SOAP
- D. YANG

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Secure Shell (SSH) is the mandatory transport protocol for Network Configuration Protocol (NETCONF). NETCONF, which is described in Request for Comments (RFC) 6241, provides the ability to automate the configuration of network devices. Protocol messages are encoded by using an Extensible Markup Language (XML) based format.

Simple Network Management Protocol (SNMP) is not the mandatory transport protocol for NETCONF. Although SNMP is used to monitor network devices, it is not typically used to manage network devices. NETCONF was created to address this lack of standardized functionality.

Simple Object Access Protocol (SOAP) is not the mandatory transport protocol for NETCONF. However, SOAP can be used to transport NETCONF. YANG is not the mandatory transport protocol for NETCONF. YANG, which is defined in RFC 6020, is a hierarchical data modeling language that can model configuration and state data for NETCONF. The YANG data can be encoded in an XML format.

Reference:

<https://tools.ietf.org/html/rfc6241>

<https://tools.ietf.org/html/rfc6020>

QUESTION 178

Which of the following hypervisors operates as a Type2 hypervisor? (Select the best answer.)

- A. HyperV
- B. KVM
- C. QEMU
- D. Xen

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Quick Emulator (QEMU) operates as a Type2 hypervisor. A hypervisor is used to create and run virtual machines (VMs). A Type2 hypervisor, which is also called a hosted hypervisor, runs within an operating system on the host computer. Other Type2 hypervisors include VMware Workstation, VirtualBox, and Parallels Desktop for Mac.

A Type1 hypervisor, which is also called a native hypervisor or a baremetal hypervisor, runs directly on the host computer's hardware. KVM, Xen, HyperV, and VMware ESX/ESXi operate as Type1 hypervisors.

Reference: <https://www.ibm.com/developerworks/library/l-hypervisor/>

https://www.ibm.com/developerworks/community/blogs/ibmvirtualization/entry/kvm_myths_uncovering_the_truth_about_the_open_source_hypervisor?lang=en

QUESTION 179

Which of the following statements are generally true of IoT devices? (Select 3 choices.)

- A. They are numerous.
- B. They are reliable.
- C. They consume a lot of power.
- D. They do not have much memory.
- E. They collectively produce a lot of data.

Correct Answer: ADE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Internet of Things (IoT) devices are numerous, do not have much memory, and collectively produce a lot of data. IoT devices, which are often called embedded devices or smart objects, are typically lowpower, lowmemory devices with limited processing capabilities. These devices are used in a variety of applications, such

as environmental monitoring, healthcare monitoring, process automation, and location tracking. Many embedded devices can transmit data wirelessly, and some are capable of transmitting over a wired connection. However, connectivity is generally unreliable and bandwidth is often constrained.

In 2003, there were only 500 million Internet-connected devices worldwide. By 2010, that number had grown to 12.5 billion devices, or 1.84 devices per person.

Cisco estimates that 50 billion IoT devices will exist by 2020 and more than 500 billion IoT devices will exist by 2030.

IoT devices collectively and individually produce a lot of data. For example, an airplane generates 10 terabytes (TB) of data for every 30 minutes of flight, and a tagged cow can generate an average of 200 megabytes (MB) of data per year. However, IoT devices often do not have the processing power to analyze the data, nor do they have the power or bandwidth to transmit a lot of data.

Reference:

https://www.cisco.com/web/AP/loEWebinarSeries/docs/the_internet_of_everythings_relevance_to_cloud_and_mobility_applications.pdf

https://www.cisco.com/c/dam/en_us/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf <https://developer.cisco.com/site/iox/documents/developer-guide/?ref=fog>

QUESTION 180

Which of the following statements is accurate regarding Salt? (Select the best answer.)

- A. Salt requires SSH.
- B. Salt requires installation of a master.
- C. Salt requires installation of a minion client.
- D. Salt requires Ruby programming knowledge.
- E. Salt requires Python programming knowledge.



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Salt requires installation of a master. Salt is a configuration management tool that is used to automate the installation, configuration, and maintenance of multiple computer systems, including the software that runs on those systems. Other configuration management tools include Puppet, Chef, and Ansible.

Salt is written in Python and operates on Linux distributions, UNIX-like systems, and Microsoft Windows.

Salt can use a client/server architecture by installing Salt master software on the server and Salt minion software on managed nodes. Masters and minions communicate by using ZeroMQ. Salt can also be used without installing Salt minion software by using Salt Secure Shell (SSH). However, Salt SSH is much slower than ZeroMQ.

Knowledge of Ruby or Python is not required to use Salt. Configuration information is stored primarily in state modules that are typically written in YAML? however, Python or Python Domain Specific Language (PyDSL) can also be used for complex configuration scripts.

Like Salt, Ansible is written in Python and operates on Linux distributions, UNIXlike systems, and Microsoft Windows. However, unlike the other configuration management software packages, Ansible does not use agent software on managed nodes. Configurations are stored on the Ansible server in playbooks that are written in YAML. Managed nodes can download scripted modules from an Ansible server by using SSH.

Puppet is written in Ruby and operates on Linux distributions, UNIXlike systems, and Microsoft Windows. Puppet uses a client/server architecture? managed nodes running the Puppet Agent application can receive configurations from a master server running Puppet Server. Modules are written in Ruby or by using a Rubylike Puppet language.

Like Puppet, Chef is written in Ruby and operates on Linux distributions, UNIXlike systems, and Microsoft Windows. Chef can use a client/server architecture or a standalone client configuration. Configuration information is contained within cookbooks that are written in Ruby and are stored on a Chef Server.

Managed nodes running the Chef Client can pull cookbooks from the server. Standalone clients that do not have access to a server can run chefsolo and pull cookbooks from a local directory or from a tar.gz archive on the Internet.

Reference:

<https://docs.saltstack.com/en/latest/topics/installation/index.html>

<https://docs.saltstack.com/en/latest/topics/ssh/index.html> <https://www.infoworld.com/article/2609482/data-center/data-center-review-puppet-vs-chef-vs-ansible-vs-salt.html?page=4>

QUESTION 181

Which of the following must match for two routers running OSPFv3 to establish a neighbor adjacency? (Select 2 choices.)

- A. area IDs
- B. router IDs
- C. process IDs
- D. instance IDs

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The area IDs must match and the instance IDs must match in order for two routers running Open Shortest Path First version 3 (OSPFv3) to establish a neighbor adjacency? OSPFv3 is also called OSPF for IPv6. Like its IPv4 counterpart, OSPFv2, OSPFv3 requires that routers have identically configured area IDs, hello timers, and dead timers in order to establish neighbor adjacencies. In addition, OSPFv3 requires that instance IDs also match; instance ID do not exist in OSPFv2. OSPFv3 allows multiple OSPF instances to run on a router. To keep track of each instance, OSPFv3 includes an instance ID field in the packet header. If no instance ID is specified on a link, the default value of 0 is used. When a router receives an OSPFv3 packet, it checks the instance ID in the packet header. If the instance ID in the header does not match the instance ID on the receiving interface, the router discards the packet even if the packet has a matching area ID. Similar to OSPFv2, OSPFv3 requires that hello timers and dead timers match in order for routers to establish a neighbor adjacency. Hello timers are used to

specify the amount of time between hello packets, which are used for neighbor discovery and maintaining neighbor relationships. By default, the hello timer is set to 10 seconds on point-to-point and broadcast links and 30 seconds on nonbroadcast multiaccess (NBMA) links. The dead timer is used to specify the amount of time to wait before declaring a neighbor to be down. By default, the dead timer is set to four times the hello timer value.

Router IDs should not match between two routers running OSPFv3. The router ID is a 32-bit value used to uniquely identify an OSPF router. By default, the router ID is the highest IPv4 loopback address configured on a router. If no loopback address is configured, the router ID is the highest IPv4 address among configured interfaces on the router. If no IPv4 addresses are configured on the router, the router ID must be manually configured before the OSPFv3 process will start. To manually configure the router ID, you should issue the `router-id` command in router configuration mode.

Process IDs do not have to match in order for two routers running OSPFv3 to establish a neighbor adjacency. Process IDs are used to identify an OSPF process on a router. However, unlike instance IDs, process IDs are only locally significant to the router.

Reference:

<https://tools.ietf.org/html/rfc5340#page-48>

https://www.cisco.com/c/en/us/td/docs/ios/ipv6/command/reference/ipv6_book/ipv6_07.html#wp2364384

QUESTION 182

You administer an OSPF network that contains a mixture of Ethernet, FastEthernet, GigabitEthernet, and TenGigabitEthernet links. The reference bandwidth is set to the default value of 100.

Which of the following will occur? (Select the best answer.)

- A. All links will have the same OSPF cost.
- B. FastEthernet, GigabitEthernet, and TenGigabitEthernet links will have the same OSPF cost.
- C. GigabitEthernet and TenGigabitEthernet links will have the same OSPF cost.
- D. Ethernet and FastEthernet links will have the same OSPF cost.
- E. All links will have different OSPF costs.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

FastEthernet, GigabitEthernet, and TenGigabitEthernet links will have the same Open Shortest Path First (OSPF) cost. An OSPF routing process uses a cost metric that is based on the bandwidth of an interface relative to a reference bandwidth. The formula to determine the cost of an interface is as follows: $\text{cost} = \frac{\text{reference bandwidth}}{\text{interface bandwidth}}$

The default reference bandwidth is 100 Mbps. You can issue the `autocost` command from router configuration mode to change the reference bandwidth for an OSPF routing process. The syntax for the `autocost` command is `autocost referencebandwidth refbw`, where `refbw` is the reference bandwidth expressed as an integer value in megabits per second between 1 and 4294967. Therefore, the default value of the `refbw` parameter is 100.

The minimum supported cost for an OSPF interface is 1, and any values that calculate to less than 1 are rounded up to 1. Therefore, any link with an interface bandwidth greater than or equal to 100 Mbps will result in a cost of 1 by default. As a result, the 100Mbps FastEthernet links, the 1Gbps GigabitEthernet links, and the 10Gbps TenGigabitEthernetlinks in this scenario will all have a cost of 1; the 10Mbps Ethernet links will have a cost of 10.

If the reference bandwidth is less than the fastest routed link on the network, a situation can arise where the cost of two interfaces is the same even though their link speeds are different. When an OSPF routing process is presented with multiple routes of the same cost, equal cost load balancing is used to distribute packets evenly among the available paths. This distribution will cause some packets in this scenario to take suboptimal routes to their destinations. To prevent this from occurring, the reference bandwidth should be a value greater than or equal to the bandwidth of the fastest routed link in the administrative domain. Alternatively, you can manually configure an OSPF cost for each interface by issuing the ip ospf cost command from interface configuration mode.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html#t6> https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/ospf-a1.html#wp3271966058 https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/ospf-a1.html#wp4045850100

QUESTION 183

Which of the following statements is true about the PortFast feature? (Select the best answer.)

- A. PortFast permanently places a switch port in the STP forwarding state.
- B. PortFast should not be enabled for ports that are connected to servers.
- C. PortFast prevents switching loops from occurring.
- D. PortFast can be configured only as a global default.
- E. PortFast can be configured only on a specific port.
- F. PortFast effectively disables STP on a port.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

PortFast permanently places a switch port in the Spanning Tree Protocol (STP) forwarding state, bypassing the listening and learning states. PortFast is a feature available on Catalyst switches that enables faster connectivity for hosts connected to an access layer switch port. If PortFast is not enabled, a switch port transitions through the STP listening and learning states before it enters the forwarding state. This process can take as long as 30 seconds if the default STP timers are used. Additionally, port initialization could take as long as 50 seconds if Port Aggregation Protocol (PAgP) is enabled. Since host computers or IP phones do not typically perform bridging functions, it is not necessary to make the switch port transition through the normal STP states, because the port should not encounter a switching loop. Thus STP skips the listening and learning states and places the port into the forwarding state so that the end host has immediate network connectivity. Although PortFast does accelerate the STP process, PortFast does not disable STP on the port.

Host ports that are not enabled for PortFast can cause a high number of STP topology changes to flood throughout the network, thereby causing high CPU utilization on network switches. Therefore, you should enable PortFast on ports that are connected to end hosts, such as IP phones, client workstations, or servers. Typically, servers and client workstations do not perform routing or switching, so there is no need to delay network connectivity while STP cycles through the listening and learning states. Conversely, PortFast should not be enabled on a port that is connected to a switch or other networking device. If you enable PortFast on such a port, you risk creating switching loops because the port is permanently in the STP forwarding state. PortFast can be enabled as a global default as well as on a specific port. If you enable PortFast as a global default, each port that is configured as an access port is enabled with PortFast. You can enable PortFast as a global default by issuing the spanningtree portfast default global configuration command. You can also enable PortFast on a perport basis by issuing the spanningtree portfast interface configuration command.

Reference:

<https://www.cisco.com/c/en/us/support/docs/switches/catalyst-6500-series-switches/10553-12.html#sptree>

https://www.cisco.com/en/US/products/hw/switches/ps679/products_configuration_guide_chapter09186a008007d779.html#xtocid24321

https://www.cisco.com/en/US/products/hw/switches/ps679/products_configuration_guide_chapter09186a008007d779.html#xtocid24323

<https://www.cisco.com/c/en/us/support/docs/switches/catalyst-4000-series-switches/65591-cat4500-high-cpu.html#host>

QUESTION 184

A nonroot switch receives several BPDUs from multiple forwarding switches. Each hello packet has the same root bridge ID and the same STP path cost to the root bridge.

Which of the following BPDU criteria is used next to determine the root port? (Select the best answer.)

- A. the lowest bridge ID of the forwarding switch
- B. the lowest port priority of the forwarding switch
- C. the lowest port number of the forwarding switch
- D. the highest bridge ID of the forwarding switch
- E. the highest port priority of the forwarding switch
- F. the highest port number of the forwarding switch

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The lowest bridge ID of the forwarding switch is used next to determine the root port. The root port on a switch is the port that receives the best Spanning Tree Protocol (STP) bridge protocol data unit (BPDU), which indicates the best path to the root bridge based on the best path cost. A root port is always in the forwarding state. Because there is only one best path to the root bridge, a switch cannot have more than one root port. Only the root bridge does not have a root port. The bridge ID is composed of a 2byte bridge priority and a 6byte Media Access Control (MAC) address. For example, a switch with a bridge priority of 32768

and a MAC address of 1234.5678.9abc would have a bridge ID of 32768.1234.5678.9abc. A switch with a lower priority value would also have a lower bridge ID. If priority values are equal, the switch with the lower MAC address is preferred; in MAC addresses, numbers are lower than letters and the hexadecimal value A is lower than the hexadecimal value F.

The root bridge sends hello packets every two seconds by default. When a switch receives a hello packet, the receiving switch modifies the forwarding switch's bridge ID, port priority, port number, and cost to reach the root bridge before forwarding the hello packet to neighboring switches. The interface that receives the hello packet with the lowest path cost will become the root port. When a switch receives multiple hello packets with the same path cost, it will choose the interface connected to the forwarding switch with the lowest bridge ID. When multiple equalcost paths to a forwarding switch exist, the receiving switch will choose the lowest port priority of the forwarding switch. If all port priorities are equal, the receiving switch will choose the lowest port number of the forwarding switch.

Reference:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/5234-5.html>

QUESTION 185

Which of the following statements best describe why WRED is useful for networks where the majority of traffic uses TCP? (Select 2 choices.)

- A. TCP packets that are dropped must be retransmitted.
- B. TCP packets cannot arrive out of sequence.
- C. TCP packets have large header sizes.
- D. TCP sources reduce traffic flow when congestion occurs.
- E. TCP packets must have priority over UDP packets.



Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Weighted random early detection (WRED) is useful for networks where the majority of traffic uses Transmission Control Protocol (TCP) because TCP packets that are dropped must be retransmitted. Additionally, TCP sources reduce traffic flow when congestion occurs, thereby further slowing down the network.

WRED is a congestion avoidance mechanism that addresses packet loss caused by tail drop, which occurs when new incoming packets are dropped because a router's queues are too full to accept them. Tail drop causes a problem called global TCP synchronization, whereby all of the TCP sources on a network reduce traffic flow during periods of congestion and then the TCP sources increase traffic flow when the congestion is reduced, which again causes congestion and dropped packets. When WRED is implemented, you can configure different tail drop thresholds for each IP precedence or Differentiated Services Code Point (DSCP) value so that lowerpriority traffic is more likely to be dropped than higher priority traffic, thereby avoiding global TCP synchronization.

WRED does not address header size. To compress the header of TCP packets, you should implement TCP header compression. Because TCP header compression compresses only the header, not the entire packet, TCP header compression works best for packets with small payloads, such as those carrying interactive data.

WRED does not address the order in which TCP packets arrive. TCP packets can arrive in any order because each packet is numbered with a sequence number. When the TCP packets arrive at their destination, TCP rearranges the packets into the correct order.

Although it is possible for TCP packets to require a higher priority than User Datagram Protocol (UDP) packets, it is also possible for UDP packets to require a higher priority than TCP packets. UDP traffic that requires a high priority includes Voice over IP (VoIP) traffic and realtime multimedia traffic. You should avoid placing TCP and UDP traffic in the same traffic class, because doing so can cause TCP starvation. UDP traffic is not aware of packet loss due to congestion control mechanisms, so devices sending UDP traffic might not reduce their transmission rates. This behavior causes the UDP traffic to dominate the queue and prevent TCP traffic from resuming a normal flow.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/qos/configuration/guide/fqos_c/qcfcnav.html

QUESTION 186

Which of the following features should not be enabled on a host port? (Select the best answer.)

- A. PortFast
- B. loop guard
- C. root guard
- D. BPDU guard

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Loop guard should not be enabled on a host port? it should be enabled on trunk ports to prevent Layer 2 switching loops from occurring. Loop guard is only used on interfaces that Spanning Tree Protocol (STP) considers to be point-to-point links. When a trunk port configured with loop guard stops receiving bridge protocol data units (BPDUs), loop guard will put the port into the loop-inconsistent state instead of allowing the port to transition to the forwarding state. If you were to enable loop guard on a port connected to a host computer, the port would transition to the loop-inconsistent state because a host does not send BPDUs.

PortFast can be enabled on a host port. PortFast enables a port to immediately access the network by transitioning the port into the STP forwarding state without passing through the STP listening and learning states. Because the ports are not expected to receive BPDUs, they are not required to learn the network topology. Host ports that are not enabled for PortFast can cause a high number of STP topology changes to flood throughout the network, thereby causing high CPU utilization on network switches. However, care should be taken to ensure that PortFast is not enabled on a port that is connected to a switch or other networking device. If you enable PortFast on such a port, you risk creating switching loops because the port is permanently in the STP forwarding state.

BPDU guard can be enabled on a host port to ensure that the port cannot receive BPDUs, thereby defining the edge of the STP domain. When a port that is configured with BPDU guard receives a BPDU, BPDU guard immediately puts the port into the err-disabled state and shuts down the port. The port must be manually reenabled, or it can be recovered automatically through the err-disabled timeout function.

Root guard can be enabled on a host port? however, it is more useful to enable PortFast and BPDU guard on a host port instead. Root guard is typically used to prevent a designated port from becoming a root port, thereby influencing which bridge will become the root bridge on the network. When root guard is applied to a port, the port is permanently configured as a designated port. A port that receives a superior BPDU will normally attempt to become a root port. However, if a designated port configured with root guard receives a superior BPDU, the port will be put into the rootinconsistent state and no data will flow through that port until it stops receiving superior BPDUs.

Reference:

https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10596-84.html#loop_guard_description

<https://www.cisco.com/c/en/us/support/docs/switches/catalyst-6500-series-switches/10553-12.html#sptree> <https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10588-74.html#diff>

QUESTION 187

Which of the following multicast addresses is used for Auto-RP announcement messages? (Select the best answer.)

- A. 224.0.0.2
- B. 224.0.0.13
- C. 224.0.0.102
- D. 224.0.1.39
- E. 224.0.1.40

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The multicast address 224.0.1.39 is used by AutoRP for RPAnnounce messages, which are sent by each candidate rendezvous point (RP) to advertise its eligibility to become an RP. The RPAnnounce messages are received by the mapping agent, which maps the candidate RPs to multicast groups. If multiple routers are advertised as candidate RPs for a multicast group, the router with the highest IP address is used as the RP for that group.

The multicast address 224.0.1.40 is used by AutoRP for RPDiscovery messages, which are sent by mapping agents to advertise the authoritative RP for a multicast group. AutoRP dynamically determines the RP for a multicast group so that RPs need not be manually configured. AutoRP uses a mapping agent to learn which routers are advertised as candidate RPs for each multicast group. The candidate list is then advertised to client routers.

The multicast address 224.0.0.2 is the allrouters address. This address is used by Protocol Independent Multicast version 1 (PIMv1) to send status messages, such as querymessages. The all routers address is also used by Internet Group Management Protocol (IGMP) and Hot Standby Router Protocol (HSRP).

The multicast address 224.0.0.13 is the allPIMrouters address. This address is used by PIMv2 to send status messages, such as hello messages, prune messages, and assert messages. The allPIMrouters address is also used by the Bootstrap Router (BSR) feature to dynamically assign RPs to multicast groups.

Other PIMv2 message types include the Register message, the RegisterStop message, and the Join/Prune message.

The multicast address 224.0.0.102 is used for Gateway Load Balancing Protocol (GLBP) hello messages. GLBP is a Cisco proprietary protocol that was developed to resolve some of the shortcomings of other router redundancy protocols, such as HSRP and Virtual Router Redundancy Protocol (VRRP). By default, hello messages are sent among GLBP group members every three seconds.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/solutions_docs/ip_multicast/White_papers/rps.html#wp1029236

QUESTION 188

Which of the following IPv6 address prefixes are not routable? (Select 2 choices.)

- A. 2000::/3
- B. FC00::/8
- C. FD00::/8
- D. FE80::/10
- E. FF02::/16
- F. FF05::/16

Correct Answer: DE

Section: (none)

Explanation



Explanation/Reference:

Explanation:

The IPv6 address prefixes FE80::/10 and FF02::/16 are not routable. The IPv6 prefix FE80::/10 is used for linklocal unicast addresses. IPv6 addresses in the FE80::/10 range begin with the characters FE80 through FEBF. Unicast packets are used for one-to-one communication. Linklocal addresses are unique only on the local segment. Therefore, linklocal addresses are not routable. An IPv6 capable host typically creates a linklocal unicast address automatically at startup. Linklocal unicast addresses are used for next-hop neighbor discovery and for environments in which no router is present to provide a routable IPv6 prefix.

The IPv6 prefix FF02::/16 is used for linklocal multicast addresses. Like linklocal unicast addresses, linklocal multicast addresses are not routable.

The IPv6 prefix 2000::/3 is used for global aggregatable unicast addresses. IPv6 addresses in the 2000::/3 range begin with the characters 2000 through 3FFF. Global aggregatable unicast address prefixes are distributed by the Internet Assigned Numbers Authority (IANA) and are globally routable over the Internet. The IPv6 prefixes FC00::/8 and FD00::/8 are used for unique local unicast addresses? together, these prefixes can be summarized as FC00::/7. IPv6 addresses in these ranges begin with the characters FC00 through FFFF. Unique local unicast addresses are not globally routable, but they are routable within an organization. All IPv6 addresses beginning with FF are multicast addresses, which are used for one-to-many communication. The following IPv6 multicast scopes are defined:

-FF01::/16 -node local

-FF02::/16 -link local

- FF05::/16 -sitelocal
- FF08::/16 -organizationlocal
- FF0E::/16 -global

The FF01::/16 prefix is used for nodelocal multicast addresses. These addresses are used only on the interface itself, much like a loopback address. Therefore, they are not routable.

The FF05::/16 prefix is used for sitelocal multicast addresses, and the FF08::/16 prefix is used for organization local multicast addresses. Like unicast addresses, sitelocal multicast and organization local multicast addresses are not globally routable, but they are routable within an organization. The FF0E::/16 prefix is used for globally routable multicast addresses.

IPv6 hosts use the multicasting capabilities of the Neighbor Discovery (ND) protocol to discover the link layer addresses of neighbor hosts. The Hop Limit field is typically set to 255 in ND packets that are sent to neighbors. Routers decrement the Hop Limit value as a packet is forwarded from hop to hop. Therefore, a router that receives an ND packet with a Hop Limit value of 255 considers the source of the ND packet to be a neighbor. If a router receives an ND packet with a Hop Limit that is less than 255, the packet is ignored, thereby protecting the router from threats that could result from the ND protocol's lack of neighbor authentication.

Reference:

<https://www.cisco.com/c/en/us/support/security/asa-5500-series-next-generation-firewalls/products-installation-and-configuration-guides-list.html#wp1010923>

https://www.cisco.com/c/en/us/products/collateral/physical-security/video-surveillance-manager/prod_white_paper0900aecd8073c232.html

QUESTION 189

You issue the show ip route command on RouterA and receive the following partial output:

```
O 10.20.0.0/28 [110/64] via 192.168.10.1, 00:02:38, Serial0/1
D 10.20.0.0/26 [90/2809856] via 192.168.10.4, 00:02:14, Serial0/4
R 10.20.0.0/24 [120/3] via 192.168.10.3, 00:33:38, Serial0/3
S 10.20.0.0/22 [1/0] via 192.168.10.2
```

RouterA receives a packet that is destined for 10.20.0.14.

Which of the following routes will RouterA use to send the packet? (Select the best answer.)



<https://vceplus.com/>

- A. the static route, because static routes are preferred over dynamic routes
- B. the EIGRP route, because it has the lowest administrative distance

- C. the RIP route, because it has the highest administrative distance
- D. the OSPF route, because it is the route with the longest prefix match

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RouterA will use the Open Shortest Path First (OSPF) route, because it is the route with the longest prefix match. When a packet is sent to a router, the router checks the routing table to see whether the nexthop address for the destination network is known. If multiple routes to a destination are known, the most specific route is used. Therefore, the following rules apply on RouterA:

-Packets sent to the 10.20.0.0/28 network use the OSPF route. This includes destination addresses from 10.20.0.0 through 10.20.0.15.

-Packets sent to the 10.20.0.0/26 network, except those sent to the 10.20.0.0/28 network, use the Enhanced Interior Gateway Routing Protocol (EIGRP) route. This includes destination addresses from 10.20.0.16 through 10.20.0.63.

-Packets sent to the 10.20.0.0/24 network, except those sent to the 10.20.0.0/26 network, use the Routing Information Protocol (RIP) route. This includes destination addresses from 10.20.0.64 through 10.20.0.255.

-Packets sent to the 10.20.0.0/22 network, except those sent to the 10.20.0.0/24 network, use the static route. This includes destination addresses from 10.20.1.0 through 10.20.3.255.

-Packets sent to any destination not listed in the routing table are forwarded to the default gateway, if one is configured.

Because the most specific route to 10.20.0.14 is the route toward the 10.20.0.0/28 network, RouterA will forward a packet destined for 10.20.0.14 to the Serial0/1 interface.

RouterA will not use the EIGRP route to send a packet that is destined for 10.20.0.14. Administrative distance (AD) values are used only to determine which route is placed in the routing table when multiple routes to a destination are known. However, a router considers routes with different prefix lengths as separate routes. If OSPF, EIGRP, and RIP had each advertised routes to 10.20.0.0/28, the EIGRP route would have been selected because EIGRP has the lowest AD. The following list contains the most commonly used ADs:

Route Source	Distance
Connected route	0
Static route	1
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
Internal BGP	200
Unknown	255

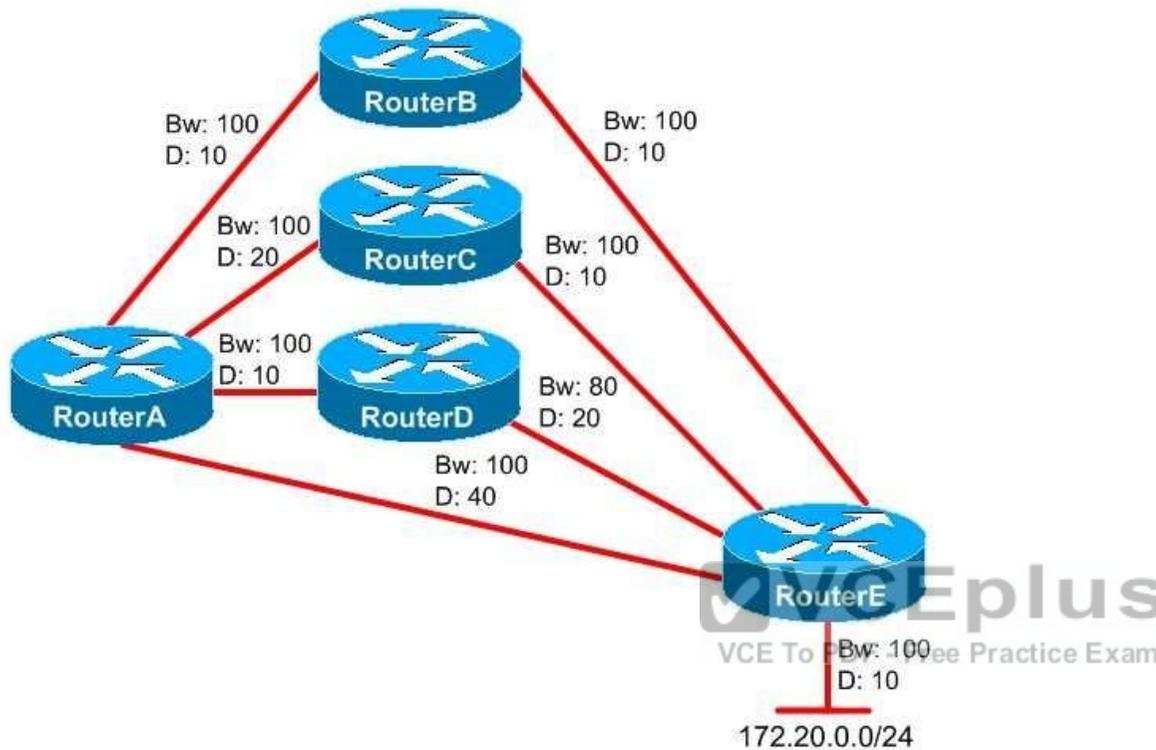
RouterA will not use the RIP route to send a packet that is destined for 10.20.0.14. Routes with longer prefix lengths are preferred over routes with shorter prefix lengths, and routes with lower ADs are preferred over routes with higher ADs.

RouterA will not use the static route to send a packet that is destined for 10.20.0.14. If the static route were configured so that the destination network was 10.20.0.0/28, the static route would be preferred over the OSPF route because static routes have a lower AD than dynamic routes.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8651-21.html>

QUESTION 190



You administer the EIGRP network shown above. The bandwidth (Bw) and delay (D) values are displayed above each link. Which of the following are feasible successors on RouterA for the 172.20.0.0/24 network? (Select the best answer.)

- A. only the route through RouterB
- B. only the route through RouterE
- C. only the routes through RouterB and RouterE
- D. only the routes through RouterC and RouterE
- E. only the routes through RouterB, RouterD, and RouterE
- F. the routes through RouterB, RouterC, RouterD, and RouterE

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Only the routes through RouterC and RouterE are feasible successors on RouterA for the 172.20.0.0/24 network. Feasible successors are backup routes that can be used immediately if the successor route is lost? the successor route is the best route to a destination. To qualify as a feasible successor, the advertising router must be closer to the destination than the router to which the route is advertised. To ensure that the advertising router meets this feasibility condition, compare the advertised distance (AD) to the feasible distance (FD) of the successor. The AD, which is also called the reported distance (RD), is the cost that the nexthop router has calculated for the route, and the FD is the cost that the local router has calculated for the route. If the AD of a route is lower than the FD of the successor, the route is a feasible successor.

By default, the Enhanced Interior Gateway Routing Protocol (EIGRP) composite metric is calculated by bandwidth and delay. The route with the lowest metric is the best route to the destination. EIGRP uses the lowest bandwidth along the path to a destination to calculate the bandwidth portion of the metric. Higher bandwidth values create lower metric values. By contrast, EIGRP uses the sum of the delays along the path to a destination to calculate the delay portion of the metric. Lower delay values create lower metric values. Therefore, higher bandwidth values and lower delay values are preferred over lower bandwidth values and higher delay values.

The following table displays the FD values for each route from RouterA to the 172.20.0.0/24 network:

Path	Bandwidth	Total Delay
A-B-E	100	30
A-C-E	100	40
A-D-E	80	40
A-E	100	50



The table shows that the route through RouterB has the highest bandwidth and lowest total delay. Therefore, the route through RouterB is the successor. For a route to qualify as a feasible successor, the metrics at the nexthop router should be lower than the metrics of the best route on the local router. The following table displays the AD values for each nexthop router to the 172.20.0.0/24 network:

Next Hop	Bandwidth	Total Delay
RouterB	100	20
RouterC	100	20
RouterD	80	30
RouterE	100	10

The route through RouterC is a feasible successor because the AD from RouterC is lower than the FD through RouterB. The bandwidth of the AD from RouterC is the same as the bandwidth of the FD through RouterB. However, because the total delay of the AD from RouterC is less than the total delay of the FD through RouterB, the route through RouterC is a feasible successor.

The route through RouterE is a feasible successor because the AD from RouterE is lower than the FD through RouterB. The bandwidth of the AD from RouterE is the same as the bandwidth of the FD through RouterB. However, because the total delay of the AD from RouterE is less than the total delay of the FD through RouterB, the route through RouterE is a feasible successor.

The route through RouterD is not a feasible successor, because the AD from RouterD is higher than the FD through RouterB. The total delay of the AD from RouterD is the same as the total delay of the FD through RouterB. However, because the bandwidth of the AD from RouterD is lower than the bandwidth of the FD through RouterB, the route through RouterD is not a feasible successor.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html#feasibleandreported>

QUESTION 191

Which of the following functions is not performed by LCP? (Select the best answer.)

- A. establishment of the PPP link
- B. termination of the PPP link
- C. detection of looped links
- D. negotiation of Network layer protocols
- E. negotiation of authentication parameters

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Network Control Protocol (NCP), not Link Control Protocol (LCP), negotiates Network layer protocols. NCP, a subprotocol in the PointtoPoint Protocol (PPP) suite, establishes and configures the Network layer protocols, such as IP and Internetwork Packet Exchange (IPX), that are used over the PPP link. When IP is used over the PPP link, either IP Control Protocol (IPCP) or IPv6 Control Protocol (IPv6CP) is the NCP. When IPX is used over the PPP link, IPX Control Protocol (IPXCP) is the NCP.

LCP is the most important subprotocol in the PPP suite. LCP establishes, configures, tests, maintains, and terminates PPP connections. The LCP phase must be complete and in an open state in order for a PPP link to establish. If the LCP phase fails, the output of the debug ppp negotiation command will indicate that LCP is in the closed state.

LCP uses a magic number parameter to determine whether a link is looped. A PPP router transmits the magic number within an LCP message. If the router receives an LCP message with that same magic number, the router recognizes that the line is looped and shuts down the interface.

LCP also negotiates authentication parameters. If PPP authentication is implemented, LCP negotiates whether to use Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).

Reference:

http://docwiki.cisco.com/wiki/Point-to-Point_Protocol <https://www.cisco.com/c/en/us/support/docs/wan/point-to-point-protocol-ppp/25440-debug-ppp-negotiation.html>

QUESTION 192

You have configured an IS-IS node with the NET 49.1741.c867.5309.af89.00.

Which of the following is the system ID? (Select the best answer.)

- A. 49
- B. 49.1741
- C. 1741.c867.5309
- D. c867.5309.af89
- E. af89.00
- F. 00

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The system ID of the Intermediate System to Intermediate System (ISIS) node is c867.5309.af89. A network entity title (NET) is a hexadecimal address that consists of the following three parts:

- The area ID
- The system ID
- The network service access point (NSAP) selector, or NSEL

The NET has a minimum length of 8 bytes and a maximum length of 20 bytes. Each byte consists of two hexadecimal characters. The NSEL is the last byte in the address and is typically set to 00.

The system ID is always 6 bytes long and precedes the NSEL. Level 1 routers must have a system ID that is unique within the area, and Level 2 routers must have a system ID that is unique within the domain. ISIS will not establish an adjacency between two routers with the same system ID.

The area ID is of variable length and precedes the system ID. The first byte, which is part of the area ID, is called the authority and format identifier (AFI) and is typically set to a value of 49 on privately addressed networks. Routers that share the same area address can form an adjacency.

In the NET 49.1741.c867.5309.af89.00, the area ID is 49.1741, the system ID is c867.5309.af89, and the NSEL is 00. You can configure a NET for a router by issuing the net command in ISIS router configuration mode. For example, to configure the NET on the node in this scenario, you would issue the command net 49.1741.c867.5309.af89.00.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/apollo/configuration/guide/fapolo_c/3cfcfns.html#wp1012582

https://www.cisco.com/c/en/us/td/docs/ios/12_2/iproute/command/reference/fiprrp_r/1rfisis.html#wp1018178

QUESTION 193

Which of the following DiffServ classes is most likely to be dropped? (Select the best answer.)

- A. AF11
- B. AF23
- C. AF31
- D. AF42

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Of the choices provided, the DiffServ class AF23 is most likely to be dropped. AF23 is a Differentiated Services Code Point (DSCP) value, which is a 6bit header value that identifies the Quality of Service (QoS) traffic class that is assigned to the packet. DSCP values beginning with AF are called Assured Forwarding (AF) perhop behaviors (PHBs), which are defined in Request for Comments (RFC) 2597. AF separates packets into four queue classes and three drop probabilities. The AF values are specified in the format AFxy, where x is the queue class and y is the drop probability. The following table displays the AF values with their queue classes and drop rates:

Queue Class (x)	Low Drop (y = 1)			Medium Drop (y = 2)			High Drop (y = 3)		
	DSCP	Binary	Decimal	DSCP	Binary	Decimal	DSCP	Binary	Decimal
1	AF11	001010	10	AF12	001100	12	AF13	001110	14
2	AF21	010010	18	AF22	010100	20	AF23	010110	22
3	AF31	011010	26	AF32	011100	28	AF33	011110	30
4	AF41	100010	34	AF42	100100	36	AF43	100110	38

The first three DSCP bits correspond to the queue class, the fourth and fifth DSCP bits correspond to the drop probability, and the sixth bit is always set to 0. To quickly convert AF values to decimal values, you should use the formula $8x + 2y$. For example, AF42 converts to a decimal value of 36, because $(8 \times 4) + (2 \times 2) = 32 + 4 = 36$.

Packets with higher AF values are not necessarily given preference over packets with lower AF values. Packets with a higher queue class value are given queuing priority over packets with a lower queue class, but packets with a higher drop rate value are dropped more often than packets with a lower drop rate value. Packets with DSCP values of AF13, AF23, AF33, and AF43 all have high drop probabilities; packets with DSCP values of AF11, AF21, AF31, and AF41 all have low drop probabilities.

Reference:

<https://www.cisco.com/c/en/us/support/docs/quality-of-service-qos/qos-packet-marking/10103-dscpvalues.html#assured>

QUESTION 194

What is the first step in a successful DHCP lease process? (Select the best answer.)

- A. A DHCP server sends a DHCPACK unicast.
- B. A DHCP client sends a DHCPDECLINE broadcast.
- C. A DHCP client sends a DHCPDISCOVER broadcast.
- D. A DHCP server sends a DHCPNAK broadcast.
- E. A DHCP server sends a DHCPOFFER unicast.
- F. A DHCP client sends a DHCPREQUEST broadcast.

Correct Answer: C

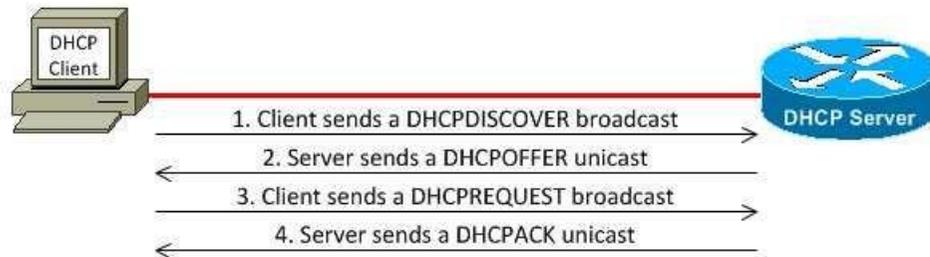
Section: (none)

Explanation

Explanation/Reference:

Explanation:

The first step in a successful Dynamic Host Configuration Protocol (DHCP) lease process occurs when a DHCP client sends a DHCPDISCOVER broadcast. The following graphic displays the steps in a successful DHCP lease process:



A DHCPDISCOVER packet is used to locate a DHCP server. If no DHCP server is available, the DHCP client will not be able to dynamically receive IP configuration information and, thus, will not be able to communicate on the network.

The second step in a successful DHCP lease process occurs when one or more DHCP servers send a DHCPOFFER unicast to the DHCP client. A DHCPOFFER packet contains IP configuration information, such as the IP address, subnet mask, default gateway, and Domain Name System (DNS) server addresses that a client should use.

The third step in a successful DHCP lease process occurs when the DHCP client sends a DHCPREQUEST broadcast. A DHCPREQUEST packet formally requests the IP address from the DHCP server. The DHCPREQUEST packet is broadcast to the entire network rather than unicast to the specific DHCP server so that the other DHCP servers can reallocate the IP addresses they offered to the DHCP client.

The fourth step in a successful DHCP lease process occurs when the DHCP server sends a DHCPACK unicast to the DHCP client. A DHCPACK packet confirms that the IP address has been officially assigned to the client for the duration of the lease.

Some packets are sent only during an unsuccessful DHCP lease process. A DHCPDECLINE packet is the opposite of a DHCPREQUEST packet.

A DHCPDECLINE packet is a broadcast packet that a DHCP client sends to formally reject a DHCP OFFER from a DHCP server. A DHCP client usually sends this kind of packet when the IP configuration is not valid for the client.

A DHCPNAK packet is the opposite of a DHCPACK packet. A DHCPNAK packet is a broadcast packet sent by a DHCP server to inform a DHCP client that the IP address in the DHCPREQUEST is no longer valid for the client to use. A DHCP server usually sends this kind of packet when the DHCP client is slow to respond to the DHCP server.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfdhcp.html

<https://www.ietf.org/rfc/rfc2131.txt>

QUESTION 195

You issue the show ip cache flow command and receive the following partial output:

SrcIF	SrcIPaddress	DstIf	DstIPaddress	Pr	SrcP	DstP	Pkts
Et0/0	10.1.1.36	Et1/0	10.2.1.74	11	C486	0045	1

Which of the following protocols is indicated in this flow? (Select the best answer.)

- A. DNS
- B. FTP
- C. HTTP
- D. HTTPS
- E. Telnet
- F. TFTP

Correct Answer: F

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Trivial File Transfer Protocol (TFTP) is indicated in this flow. The device at 10.1.1.36 has sent a User Datagram Protocol (UDP) TFTP packet to the device at 10.2.1.74.

The show ip cache flow command is used to display a summary of NetFlow statistics. The DstP field indicates the destination port field and is displayed in hexadecimal. The hexadecimal value 45 converts to the decimal value 69, which is the port number used by TFTP.

The Pr field is used to indicate the IP protocol number and is displayed in hexadecimal. The Pr field is set to a hexadecimal value of 06 for Transmission Control Protocol (TCP) and to a hexadecimal value of 11 for UDP.

Domain Name System (DNS) communicates over TCP and UDP port 53. The decimal value 53 converts to a hexadecimal value of 35. Therefore, the destination port field in the output of the show ip cache flow command would display a value of 0035 for DNS traffic.

File Transfer Protocol (FTP) communicates over TCP ports 20 and 21. The decimal values 20 and 21 convert to hexadecimal values of 14 and 15, respectively. Therefore, the destination port field in the output of the show ip cache flow command would display a value of 0014 or 0015 for FTP traffic.

Hypertext Transfer Protocol (HTTP) communicates over TCP port 80. The decimal value 80 converts to a hexadecimal value of 50. Therefore, the destination port field in the output of the show ip cache flow command would display a value of 0050 for HTTP traffic.

HTTP Secure (HTTPS) communicates over TCP port 443. The decimal value 443 converts to a hexadecimal value of 1BB. Therefore, the destination port field in the output of the show ip cache flow command would display a value of 01BB for HTTPS traffic.

Telnet communicates over TCP port 23. The decimal value 23 converts to a hexadecimal value of 17. Therefore, the destination port field in the output of the show ip cache flow command would display a value of 0017 for Telnet traffic.

Reference:

https://www.cisco.com/en/US/docs/ios/12_3t/netflow/command/reference/nfl_a1gt_ps5207_TSD_Products_Command_Reference_Chapter.html#wp1187159
https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/ios-netflow/prod_white_paper0900aecd80406232.html

QUESTION 196

Which of the following statements are correct regarding the differences between IGMPv1, IGMPv2, and IGMPv3? (Select 2 choices.)

- A. IGMPv2 introduced support for SSM.
- B. IGMPv2 introduced support for groupspecific queries.
- C. IGMPv2 introduced support for querier elections.
- D. IGMPv3 introduced support for leave group messages.
- E. IGMPv3 introduced support for membership report suppression.
- F. IGMPv1 has a default query interval of 120 seconds? IGMPv2 and IGMPv3 have a default queryinterval of 60 seconds.
- G. IGMPv1 has a default querier timeout of 120 seconds? IGMPv2 and IGMPv3 have a default querier timeout of 60 seconds.

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Internet Group Management Protocol version 2 (IGMPv2) introduced support for groupspecific queries and querier elections. IGMPv1 queries are general queries sent to the 224.0.0.1 allhosts multicast address. IGMPv2 and IGMPv3 queries are either general queries, which are sent to 224.0.0.1, or groupspecific queries, which are sent only to members of a particular multicast group.

Although IGMPv3 supports querier elections, this feature was introduced in IGMPv2. The router with the lowest IP address on the subnet is elected as the querier. The querier is responsible for periodically sending out membership query messages to determine whether any hosts want to receive multicast packets for the multicast group. If at least one host responds with a membership report message, the querier will continue to send those multicast packets on that network segment.

Although the Request for Comments (RFC) standard for IGMP query messages is 125 seconds, Cisco uses a default query interval of 60 seconds for all IGMP versions. The query interval determines how often the querier sends out membership query messages. If no member has responded to the query message within three times the query interval, the interface is pruned.

Cisco uses a default querier timeout of two times the query interval, or 120 seconds, for

IGMPv2 and IGMPv3? IGMPv1 does not support querier elections. The querier timeout is used to trigger querier elections. If an IGMP device has not received a query message from the querier within the querier timeout period, a querier election is triggered and a new querier is elected.

IGMPv3 introduced support for Source Specific Multicast (SSM). SSM enables IGMPv3 hosts to specify the source addresses from which they will accept multicast traffic. To enable SSM, you should issue the `ip pim ssm` command from global configuration mode, the `ip pim {sparsemode | sparsedensemode}` command from interface configuration mode, and the `ip igmp version 3` command from interface configuration mode.

Although IGMPv3 supports leave group messages, this feature was introduced in IGMPv2.

In IGMPv1, a host leaves a multicast group silently. In IGMPv2 and IGMPv3, a host sends an IGMP leave message when it wants to leave a multicast group. IGMP routers maintain the IP address of the last reporter, which is the last host that sent a membership report message for that multicast group. If the last reporter leaves a multicast group, the IGMP router immediately sends a membership query message to determine whether any interested hosts remain.

IGMPv3 does not support host membership report suppression; in fact, IGMPv3 removed support for host membership report suppression. This feature, which is supported in

IGMPv1 and IGMPv2, prevents the sending of a membership report if a similar report is detected from another host on the network. IGMPv3 removes this restriction.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfmulti.html#wp1066001

<https://www.ietf.org/rfc/rfc3376.txt>

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/command/imc-cr-book/imc_i1.html#wp4034771958 https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/command/imc-cr-book/imc_i1.html#wp1157094318

QUESTION 197

You issue the `show vlan private-vlan` command on SwitchA and receive the following partial output:

```
SwitchA#show vlan private-vlan
```

Primary	Secondary	Type	Ports
30	10	isolated	Fa0/1, Fa0/2, Fa0/6
30	20	community	Fa0/1, Fa0/3, Fa0/4
30	40	community	Fa0/1, Fa0/5

You are going to create secondary VLAN 50 and associate it with primary VLAN 30.

Which of the following commands should you issue in VLAN configuration mode for VLAN 50? (Select the best answer.)

- A. private-vlan community
- B. private-vlan isolated
- C. private-vlan primary
- D. private-vlan secondary
- E. private-vlan association 30
- F. private-vlan association add 30

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the privatevlan community command in virtual LAN (VLAN) configuration mode for VLAN 50. The privatevlan community command configures a VLAN to be a community secondary VLAN.

A private VLAN (PVLAN) provides separation between ports that belong to the same VLAN. Because the separation exists at Layer 2, the hosts can exist on the same IP subnet. The VLAN to which the hosts belong is called the primary VLAN. To configure a VLAN as a primary VLAN, you should issue the privatevlan primary command. You should not issue the privatevlan primary command for VLAN 50, because doing so would make VLAN 50 a primary VLAN, not a secondary VLAN. To create a PVLAN, you must create secondary VLANs and associate them with the primary VLAN. There are two types of secondary VLANs: community VLANs and isolated VLANs. Ports that belong to a community VLAN can communicate with promiscuous ports and with other ports that belong to the same community. However, they cannot communicate with isolated ports or with ports that belong to other communities. To configure a VLAN as a community VLAN, you should issue the privatevlan community command.

Ports that belong to an isolated VLAN can communicate with only promiscuous ports. To configure a VLAN as an isolated VLAN, you should issue the privatevlan isolated command. Only one isolated VLAN can be associated with a primary VLAN. Therefore, you should not issue the privatevlan isolated command for VLAN 50, because VLAN 10 is configured as an isolated VLAN and is associated with primary VLAN 30. If VLAN 10 did not exist, you could configure VLAN 50 as an isolated VLAN and associate it with primary VLAN 30.

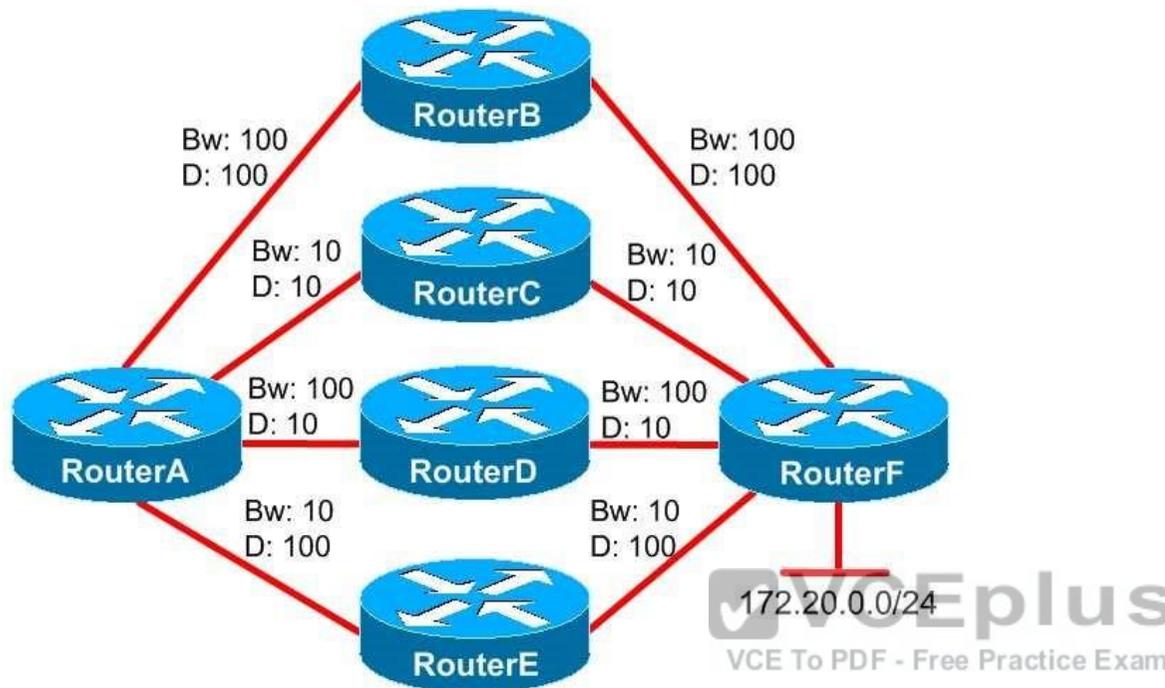
You should not issue the `privatevlan secondary` command for VLAN 50. The `privatevlan secondary` command uses incorrect syntax; therefore, issuing this command will generate an error.

You should not issue the `privatevlan association 30` command or the `privatevlan association add 30` command for VLAN 50. The `privatevlan association` command associates the primary VLAN with one or more secondary VLANs; this command should be issued in VLAN configuration mode for the primary VLAN. Therefore, the `privatevlan association` command should not be issued for VLAN 50; it should be issued for VLAN 30. Issuing the `privatevlan association 50` command for VLAN 30 would associate only secondary VLAN 50 with primary VLAN 30, but doing so would remove the existing secondary VLAN associations for VLAN 30. Issuing the `privatevlan association add 50` command for VLAN 30 would add secondary VLAN 50 to the list of existing secondary VLANs that are associated with VLAN 30.

Reference:

<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/CLIConfigurationGuide/PrivateVLANs.html>

QUESTION 198



You administer the EIGRP network shown above. The bandwidth (Bw) and delay (D) values are displayed above each link. RouterA uses the default variance value.

Which of the following statements are correct? (Select 2 choices.)

- A. RouterB is a successor for the 172.20.0.0/24 network.
- B. RouterC is a successor for the 172.20.0.0/24 network.
- C. RouterD is a successor for the 172.20.0.0/24 network.
- D. RouterE is a successor for the 172.20.0.0/24 network.
- E. RouterA will not perform load balancing.
- F. RouterA will perform unequalcost load balancing.

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RouterD is a successor for the 172.20.0.0/24 network. Additionally, RouterA will not perform load balancing.

A successor is a router with the lowest metric to a destination network. By default, the Enhanced Interior Gateway Routing Protocol (EIGRP) composite metric is calculated by bandwidth and delay. Higher bandwidth values create lower metric values. By contrast, lower delay values create lower metric values. Therefore, higher bandwidth values and lower delay values are preferred over lower bandwidth values and higher delay values. RouterD is a successor for the 172.20.0.0/24 network because the path through RouterD has the highest bandwidth and the lowest delay. Therefore, RouterD is the only successor for the 172.20.0.0/24 network.

RouterB is not a successor for the 172.20.0.0/24 network. Although the bandwidth through RouterB is the same as the bandwidth through RouterD, the delay through RouterB is higher than the delay through RouterD.

RouterC is not a successor for the 172.20.0.0/24 network. Although the delay through RouterC is the same as the delay through RouterD, the bandwidth through RouterC is lower than the bandwidth through RouterD.

RouterE is not a successor for the 172.20.0.0/24 network. The bandwidth through RouterE is lower than the bandwidth through RouterD. Additionally, the delay through RouterE is higher than the delay through RouterD.

RouterA will not perform load balancing. By default, an EIGRP router is configured with a variance value of 1, which enables the router to perform equalcost load balancing. However, there are no routes with a metric value as low as the route through RouterD. If there were another path to the 172.20.0.0/24 network with a bandwidth of 100 and a delay of 10, RouterA would load balance traffic between the two equalcost paths.

Although EIGRP is capable of unequalcost load balancing, RouterA is not configured to perform unequalcost load balancing, because RouterA is using the default variance value of 1. A variance value of 2 or higher means that a feasible successor can be used for unequalcost load balancing. However, voice traffic and other delay sensitive traffic will be negatively affected if those packets are routed over paths with lower bandwidths and higher delays.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/command/ire-cr-book/ire-s1.html#wp3209991315 <https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html#anc6>

QUESTION 199

You want to configure SSH for incoming VTY connections. The router has not been configured with a host name, a domain name, or an RSA key pair. Additionally, the VTY lines are not yet configured to accept incoming SSH connections.

You issue the ip ssh timeout 60 command from global configuration mode to configure the router with a 60second timeout.

Which of the following messages will you most likely receive? (Select the best answer.)

- A. Invalid input detected at '^' marker.
- B. Please define a hostname other than Router.
- C. Please define a domainname first.
- D. Please create RSA keys to enable SSH.
- E. Please enable SSH as a transport mode.

Explanation:

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

You will most likely receive the Please create RSA keys to enable SSH message when you issue the ip ssh timeout 60 command from global configuration mode.

To enable Secure Shell (SSH) for virtual terminal (VTY) lines on a Cisco router, you should complete the following steps:

1. Configure the router with a host name other than Router by issuing the hostname command.
2. Configure the router with a domain name by issuing the ip domain name command.
3. Generate an RSA key pair for the router by issuing the crypto key generate rsa command.
4. Configure the VTY lines to use SSH by issuing the transport input ssh command from line configuration mode.

You will not receive the Invalid input detected at '^' marker message when you issue the ip ssh timeout 60 command in this scenario. You would receive the Invalid input detected at '^' marker message if you were to mistype the timeout keyword or if you were to try to configure the SSH timeout with a value greater than 120 seconds. Although SSH is not yet enabled in this scenario, the router will accept the ip ssh timeout 60 command as a valid configuration. The ip ssh timeout 60 command would appear in the configuration if you were to issue the show runningconfig command.

You will not receive the Please define a hostname other than Router message when you issue the ip ssh timeout 60 command in this scenario. However, because you have not configured the router with a host name other than the default name of Router, you would receive the Please define a hostname other than Router message if you were to issue the crypto key generate rsa command. To configure a router with a host name other than the default, you should issue the hostname hostname command from global configuration mode.

You will not receive the Please define a domainname first message when you issue the ip ssh timeout 60 command in this scenario. However, if you had configured the router with a valid host name but had not configured the router with a domain name, you would receive the Please define a domainname first message if you were to issue the crypto key generate rsa command. In this scenario, you have configured neither the domain name nor the host name. To configure a router with a domain name, you should issue the ip domainname domainname command from global configuration mode.

You will not receive the Please enable SSH as a transport mode message when you issue the ip ssh timeout 60 command in this scenario. The Please enable SSH as a transport mode message is not a warning message that is displayed on Cisco routers. You can issue the transport input ssh command to configure SSH as the transport mode for VTY lines.

Reference:

<https://www.cisco.com/c/en/us/support/docs/security-vpn/secure-shell-ssh/4145-ssh.html#settingupanosrouterasssh>

QUESTION 200

Which of the following statements are incorrect? (Select 2 choices.)

- A. NBAR will not work without CEF.
- B. NBAR can classify IP and IPX traffic.
- C. NBAR can classify TCP and UDP traffic.

Explanation:

- D. NBAR can classify HTTP and FTP traffic.
- E. NBAR can classify unicast and multicast traffic.
- F. NBAR can classify inbound and outbound traffic.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Although Network Based Application Recognition (NBAR) can classify IP traffic, it cannot classify Internetwork Packet Exchange (IPX) traffic. Additionally, NBAR can classify unicast traffic, but it cannot classify multicast traffic.

NBAR enables a router to perform deep packet inspection for all packets that pass through an NBAR-enabled interface. With deep packet inspection, an NBAR-enabled router can classify traffic based on the content of a packet, not just the network header information.

Additionally, NBAR provides statistical reporting relative to each recognized application. For example, NBAR can be used to track bandwidth usage for each protocol type.

NBAR can classify traffic that uses Transmission Control Protocol (TCP), such as Hypertext Transfer Protocol (HTTP) and File Transfer Protocol (FTP) traffic, and traffic that uses User

Datagram Protocol (UDP), such as Dynamic Host Configuration Protocol (DHCP) and Trivial

File Transfer Protocol (TFTP) traffic. Additionally, NBAR can classify IP traffic that does not use TCP or UDP, such as Generic Routing Encapsulation (GRE) and IP Security (IPSec) traffic. Not only can NBAR classify traffic that uses static port numbers, it can also classify traffic that uses dynamically assigned port numbers.

Before NBAR can classify any traffic, Cisco Express Forwarding (CEF) must be enabled on the router. CEF is enabled by default on Cisco routers. If CEF has been disabled by the no ip cef command, you can reenable CEF by issuing the ip cef command.

You can configure NBAR to classify inbound traffic on an interface by issuing the service-policy input command. Alternatively, you can configure NBAR to classify outbound traffic by issuing the service-policy output command.

Reference:

https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/network-based-application-recognition-nbar/prod_case_study09186a00800ad0ca.html

<https://www.cisco.com/c/en/us/products/ios-nx-os-software/index.html>

QUESTION 201

DRAG DROP

Select the default STP timer values from the left, and place them over the corresponding STP timers on the right. Each timer value can be used only once. Not all timer values will be used.

Select and Place:

<http://www.gratisexam.com/>

1 second	hello
2 seconds	max_age
6 seconds	forward_delay
10 seconds	
15 seconds	
20 seconds	
30 seconds	
60 seconds	

Correct Answer:



Section: (none)
Explanation

Explanation/Reference:

Explanation:

Spanning Tree Protocol (STP) uses three timer values: the hello timer value, the forward_delay timer value, and the max_age timer value. The hello timer value is the time between the sending of bridge protocol data units (BPDUs). Switches send BPDUs to determine the path cost to the root bridge. A switch assumes that it has lost connectivity with a neighbor root bridge or neighbor designated bridge after it misses three BPDUs, which are sent every two seconds by default.

Therefore, the STP information is removed from a switch after six seconds pass without the switch receiving a BPDU.

STP information is also removed after a set period of time. The max_age timer value is the maximum length of time before old BPDU information is removed. By default, the max_age timer is set to 20 seconds.

STP interfaces exist in one of five states: blocking, listening, learning, forwarding, or disabled. The forward_delay timer value is the time a port spends in the listening state and learning state. By default, the forward_delay timer is set to 15 seconds.

Switches use the STP timer values on BPDUs that they receive, not on the local STP timer values configured on the switch itself. Therefore, only the timers configured on the root bridge will be used throughout the network. If you want to modify STP timer values, you should change them on the root bridge and the backup root bridge, at a minimum.

Reference:

https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/19120-122.html#stp_timers

QUESTION 202

Which of the following TOS values is used for the PHB value EF in the output of the show ip cache flow command? (Select the best answer.)

- A. E0
- B. 0
- C. A0
- D. B8.
- E. 80

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Type of Service (ToS) value B8 is used for the perhop behavior (PHB) value EF in the output of the show ip cache flow command. PHBs identify the Quality of Service (QoS) traffic class that is assigned to the packet. Expedited Forwarding (EF), which is defined in Request for Comments (RFC) 2598, indicates a highpriority packet that should be given queuing priority over other packets but should not be allowed to completely monopolize the interface. Voice over IP (VoIP) traffic is often assigned a value of EF.

The TOS value in the output of the show ip cache flow command is an 8bit hexadecimal value. EF has a binary value of 10111000. To convert an eightdigit binary value to a two character hexadecimal value, split the binary value in half and convert each section individually. The binary value 1011 converts to the hexadecimal character B, and the binary value 1000 converts to the hexadecimal character 8. Therefore, the value EF converts to the hexadecimal value B8, which is the TOS value that is displayed in the output of the show ip cache flow command.

Class Selector (CS) PHBs, which are defined in RFC 2475, only use the first three QoS bits; therefore, CS PHBs are backward compatible with 3bit IP precedence values. Packets with higher CS values are given queuing priority over packets with lower CS values. The following table displays the CS values with their binary values, hexadecimal TOS values, and IP precedence category names:



DSCP Value	Binary	Hexadecimal	IP Precedence
CS0	0000 0000	0	Routine
CS1	0010 0000	20	Priority
CS2	0100 0000	40	Immediate
CS3	0110 0000	60	Flash
CS4	1000 0000	80	Flash Override
CS5	1010 0000	A0	Critical
CS6	1100 0000	C0	Internetwork Control
CS7	1110 0000	E0	Network Control

Reference: https://www.cisco.com/en/US/docs/voice_ip_comm/bts/5.0/command/reference/clipc5.pdf#page=793
<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/netflow/command/nf-cr-book/nf-02.html#GUID-E3881A9E-7FD9-4BCE-83E5-603E55AE72DC>

QUESTION 203

Which of the following PfR modes use IP SLA probes? (Select the best answer.)

- A. only fast mode
- B. only active mode
- C. only passive mode
- D. only fast mode and active mode
- E. only fast mode and passive mode
- F. only active mode and passive mode
- G. fast mode, active mode, and passive mode

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Only fast mode and active mode use IP Service Level Agreement (SLA) probes. Cisco Performance Routing (PfR) enhances traditional routing methods by dynamically selecting the best path for applications based on network performance. The following three monitoring modes are used by PfR:

-Passive mode

-Active mode

-Fast mode

Active mode relies on IP SLA probes that generate traffic to capture performance metrics. Metrics used by active mode include delay, jitter, mean opinion score (MOS), and reachability. Shortterm monitoring uses the last five probe results; longterm monitoring uses the last 60 probe results.

Fast mode is similar to active mode. Active mode generates IP SLA probes only for the active exit path. By contrast, fast mode continuously generates IP SLA probes for all possible exit paths, not just the active exit path. Fast mode allows route changes to be made within three seconds. However, the performance benefits of fast mode require significant processor overhead; therefore, Cisco recommends that you use fast mode only for performancesensitive traffic, such as Voice over IP (VoIP) or video traffic.

Passive mode does not use IP SLA probes. Instead, passive mode relies on NetFlow to capture performance metrics. Metrics used by passive mode include delay, packet loss, reachability, and throughput. Throughput can be measured for all traffic flows. Delay, packet loss, and reachability can be measured only for Transmission Control Protocol (TCP) flows.

Reference:

https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/performance-routing-pfr/product_data_sheet0900aecd806c4ee4.html

http://docwiki.cisco.com/wiki/Performance_Routing_FAQs

QUESTION 204

You issue the show runningconfig command on RouterA and receive the following partial output:

```
class-map voip
  match ip dscp ef
class-map video
  match ip dscp 41
class-map ftp
  match protocol ftp
policy-map boson
  class voip
    priority percent 20
  class video
    bandwidth percent 40
  class ftp
    bandwidth remaining-percent 50
  class class-default
    bandwidth remaining-percent 25
!
interface FastEthernet0/1
  ip address 10.20.30.1 255.255.255.0
  max-reserved-bandwidth 100
  service-policy output boson
```

How much web traffic can RouterA send out the FastEthernet0/1 interface during periods of heavy voice and video traffic? (Select the best answer.)



<https://vceplus.com/>

- A. 10 Mbps
- B. 15 Mbps
- C. 20 Mbps
- D. 25 Mbps
- E. 40 Mbps

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RouterA can send 10 Mbps of web traffic out the FastEthernet0/1 interface during periods of heavy voice and video traffic. To create a Quality of Service (QoS) policy, you must perform the following steps:

1. Define one or more class maps by issuing the classmap name command.
2. Define the traffic that matches the class map by issuing one or more match commands.
3. Define one or more policy maps by issuing the policymap name command.
4. Link the class maps to the policy maps by issuing the classname command.
5. Define one or more actions that should be taken for that traffic class.
6. Link the policy map to an interface by issuing the servicepolicy {input | output} name command.

Bandwidth guarantees are set in policymap class configuration mode. You can specify the bandwidth as a rate or as a percentage with the bandwidth and priority commands. The syntax of the priority command is priority {bandwidth | percentpercentage} [burst], where bandwidth is specified in Kbps and burst is specified in bytes. The priority command creates a strictpriority queue where packets are dequeued before packets from other queues are dequeued. The strictpriority queue is given priority over all other traffic.

If no priority traffic is being sent, the other traffic classes can share the remaining bandwidth based on their configured values.

The bandwidth command specifies a guaranteed amount of bandwidth for a particular traffic class. The syntax of the bandwidth command is bandwidth {kbps | remaining percentpercentage | percentpercentage}, where kbps is the amount of bandwidth that is guaranteed to a particular traffic class.

<http://www.gratisexam.com/>

In this scenario, Voice over IP (VoIP) traffic is given a guaranteed 20 percent of the interface's bandwidth. Video traffic is given a guaranteed 40 percent of the interface's bandwidth. Voice and video traffic can exceed these bandwidth percentages if any unused bandwidth remains.

The remaining 40 percent, or 40 Mbps, of the interface's bandwidth can be used by other traffic. If traffic does not match any traffic class, it will become part of the classdefault class. In this scenario, web traffic belongs to the classdefault class. Therefore, web traffic can consume 25 percent of the remaining bandwidth. If no other traffic is being sent on the interface, web traffic can consume 25 percent of the interface's bandwidth. However, when voice and video traffic are heavy, web traffic can consume 25 percent of the remaining 40 Mbps, which is equal to 10 Mbps.

Even less web traffic can be sent if File Transfer Protocol (FTP) traffic or other unclassified traffic is heavy. FTP traffic can consume 50 percent of the remaining bandwidth on the interface. If no other traffic is being sent on the interface, FTP traffic can consume 50 percent of the interface's bandwidth. During periods of heavy voice and video usage, FTP traffic can consume 50 percent of the remaining 40 Mbps, which is equal to 20 Mbps.

Reference:

<https://www.cisco.com/c/en/us/support/docs/quality-of-service-qos/qos-packet-marking/10100-priorityvsbw.html>

QUESTION 205

Which of the following is appended to a VPNv4 BGP route to indicate membership in an RFC 4364 MPLS VPN? (Select the best answer.)

- A. a label
- B. an RT
- C. an RD
- D. an LSP
- E. a VRF



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A route target (RT) is appended to a virtual private network version 4 (VPNv4) Border

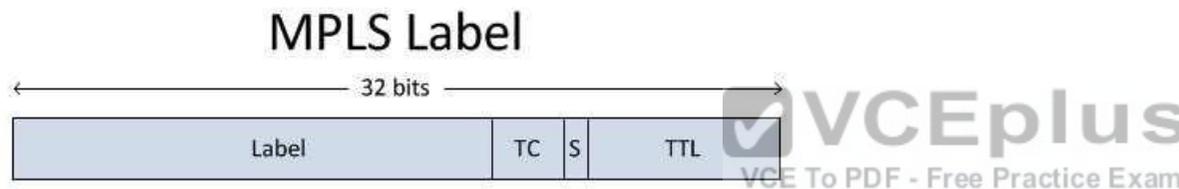
Gateway Protocol (BGP) route to indicate membership in a Request for Comments (RFC)4364 Multiprotocol Label Switching (MPLS) VPN. Export RTs associate each route with one or more VPNs, and import RTs are associated with each VPN routing and forwarding (VRF) table to determine the routes that should be imported into the VRF; a VRF is a routing table instance for a VPN. A label is assigned to each VPNv4 address prefix, and the inboundtooutbound label mapping is stored in the Label Forwarding Information Base (LFIB). By configuring import and export RTs, you can configure which sites can reach each other. For example, you can configure RTs so that CustomerA and CustomerB can communicate with ProviderZ, but CustomerA and CustomerB cannot communicate with one another. To configure RTs, you should issue the `routetarget {import | export | both} valuecommand`. The import, export, and both keywords specify whether extended community attributes should be imported, exported, or both. The value parameter should use one of the following formats:

- AS:nn, where AS is a 16bit autonomous system number (ASN) and nn is a 32bit decimal number
- A.B.C.D:nn, where A.B.C.D is a 32bit IP address and nn is a 16bit decimal number

A route distinguisher (RD) is a value that is added to the beginning of an IP address to create a globally unique VPNv4 address. RDs enable customers to use the same or overlapping IP address ranges on their internal networks. To create an RD, you should issue the rd value command, where the value parameter uses the same formats as the value parameter in the route-target command.

There are three types of RDs: Type 0, Type 1, and Type 2. The type of RD configuration you create depends on how you issue the value parameter of the rd command and whether you are configuring a multicast VPN environment. Type 0 and Type 1 RDs are used in unicast configurations. A Type 0 RD is configured by issuing the value parameter of the rd command with the 16bit ASN in front of the 32bit decimal number. A Type 1 RD is configured by issuing the value parameter of the rd command with the 32bit decimal number in front of the 16bit ASN. A Type 2 RD is configured similarly to a Type 1 RD but only applies to multicast VPN configurations.

A label switched path (LSP) is the path that labeled packets take through an MPLS network from one label switch router (LSR) to another. The 32bit MPLS label is used by LSRs to make forwarding decisions along the LSP. The MPLS label is placed between the Layer 2 header and the Layer 3 header. The structure of an MPLS label is shown below:



Reference:

<https://tools.ietf.org/html/rfc4364>

QUESTION 206

On which of the following interfaces can a port ACL be applied? (Select 3 choices.)

- A. an SVI
- B. a trunk port
- C. an EtherChannel interface
- D. a routed port
- E. a Layer 2 port

Correct Answer: BCE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A port access control list (PACL) can be applied to a trunk port, a Layer 2 port, or an EtherChannel interface. PACLs filter inbound Layer 2 traffic on a switch port interface; PACLs cannot filter outbound traffic. When PACLs are applied on a switch, packets are filtered based on several criteria, including IP addresses, port numbers, or upperlayer protocol information. If a PACL is applied to a trunk port, it will filter all virtual LAN (VLAN) traffic traversing the trunk, including voice and data VLAN traffic. A PACL can be used with an EtherChannel configuration, but the PACL must be applied to the logical EtherChannel interface? physical ports within the EtherChannel group cannot have a PACL applied to them.

PACLs cannot be applied to a switch virtual interface (SVI) or to a routed port. An SVI is a virtual interface that is used as a gateway on a multilayer switch. SVIs can be used to route traffic across Layer 3 interfaces. However, PACLs can only be applied to Layer 2 switching interfaces. Furthermore, because PACLs operate at Layer 2, they cannot be applied to routed ports, which operate at Layer 3.

Reference:

<https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SX/configuration/guide/book/vacl.html#wp1119764>

QUESTION 207

Which of the following attacks results in unicast traffic being sent out every port on a switch, regardless of the intended destination of the traffic? (Select the best answer.)

- A. an ARP poisoning attack
- B. a VLAN hopping attack
- C. a MAC flooding attack
- D. a DHCP spoofing attack
- E. an STP attack



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A Media Access Control (MAC) flooding attack results in traffic being sent out every port on a switch, regardless of the intended destination of the traffic. Switches and bridges store any learned MAC addresses in a MAC address table. When the MAC address table becomes full, no more MAC addresses can be learned. If a switch receives traffic destined for a MAC address that is not in its MAC address table, the switch floods the traffic out every port except the port that originated the traffic. Consequently, in a MAC flooding attack, an attacker attempts to fill the MAC address table so that any further traffic will be sent to all hosts on the network, causing excessive unicast flooding. As a result, the attacker can access any traffic that is sent to the switch.

In a Dynamic Host Configuration Protocol (DHCP) spoofing attack, a rogue DHCP server is attached to the network in an attempt to intercept DHCP requests. The rogue DHCP server can then respond to the DHCP requests with its own IP address as the default gateway address so that all traffic is routed through the rogue

DHCP server. As a result, a host that has obtained an IP address from a rogue DHCP server could become the victim of a man-in-the-middle attack in which a malicious individual eavesdrops on a network conversation between two hosts.

In an Address Resolution Protocol (ARP) poisoning attack, which is also known as an ARPspoofing attack, the attacker intercepts an ARP request packet and replies with its own MAC address, rather than the address of the intended recipient. Subsequently, the attacker is able to intercept any traffic intended for the original recipient.

In a virtual LAN (VLAN) hopping attack, an attacker attempts to inject packets into other VLANs by accessing the VLAN trunk and doubletagging 802.1Q frames. A successful VLAN hopping attack enables an attacker to send traffic to other VLANs without using a router.

In a Spanning Tree Protocol (STP) attack, an attacker listens for STP frames to determine the port ID of the interface that is transmitting the STP frames. The attacker can then send bridge protocol data units (BPDUs) in an attempt to become the root bridge for the network.

Reference:

https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-6500-series-switches/prod_white_paper0900aecd805457cc.html

QUESTION 208

Which of the following error messages might appear when a unidirectional link failure exists between two switches? (Select 2 choices.)

- A. %STP-2-BLOCK_BPDUGUARD
- B. %STP-2-BLOCK_PVID_LOCAL
- C. %STP-2-BLOCK_PVID_PEER
- D. %STP-2-BRIDGE_ASSURANCE_BLOCK
- E. %STP-2-DISPUTE_DETECTED



Correct Answer: DE

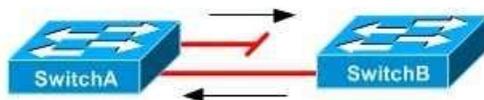
Section: (none)

Explanation

Explanation/Reference:

Explanation:

Of the choices provided, the %STP2BRIDGE_ASSURANCE_BLOCK error message or the %STP2DISPUTE_DETECTED error message might appear when a unidirectional link failure exists between two switches. A unidirectional link failure exists when a defective cable causes one device to not receive what the other device sends, as illustrated by the following exhibit:



The %STP2BRIDGE_ASSURANCE_BLOCK error message appears when a Spanning Tree Protocol (STP) bridge protocol data unit (BPDU) has not been received on an interface with Bridge Assurance enabled. Bridge Assurance ensures that BPDUs are sent bidirectionally on all network interfaces. If an interface with Bridge Assurance does not receive a BPDU, or if the connected interface does not have Bridge Assurance enabled, the interface is put into an inconsistent state

and is blocked. Bridge Assurance is supported only with Rapid PerVLAN Spanning Tree Plus (RPVST+) and Multiple Spanning Tree (MST) and only on point-to-point links. The %STP2DISPUTE_DETECTED error message appears when a switch receives inferior BPDUs that are marked as designated and learning or forwarding. This indicates that the connected switch is not receiving superior BPDUs from the local switch. As a result, the local switch will record the %STP2DISPUTE_DETECTED error and shut down the interface to prevent a bridging loop.

The %STP2BLOCK_BPDUGUARD error message appears not when a unidirectional link failure exists, but when a BPDU has been received on an interface with BPDU guard enabled. When an interface that is configured with BPDU guard receives a BPDU, BPDU guard immediately puts the interface into the errdisable state and shuts down the interface. Afterward, the interface must be manually reenabled, or it can be recovered automatically through the errdisable timeout function. The %STP2BLOCK_PVID_LOCAL and %STP2BLOCK_PVID_PEER error messages appear not when a unidirectional link failure exists, but when an interface has received a BPDU that is tagged with the same virtual LAN (VLAN) ID as the interface's native VLAN. Native VLAN BPDUs are sent untagged, so if a switch receives BPDUs that are tagged with the native VLAN for that interface, the local switch will record the % STP2BLOCK_PVID_LOCAL error message and block the interface? the remote switch will record the % STP2BLOCK_PVID_PEER error message.

Reference:

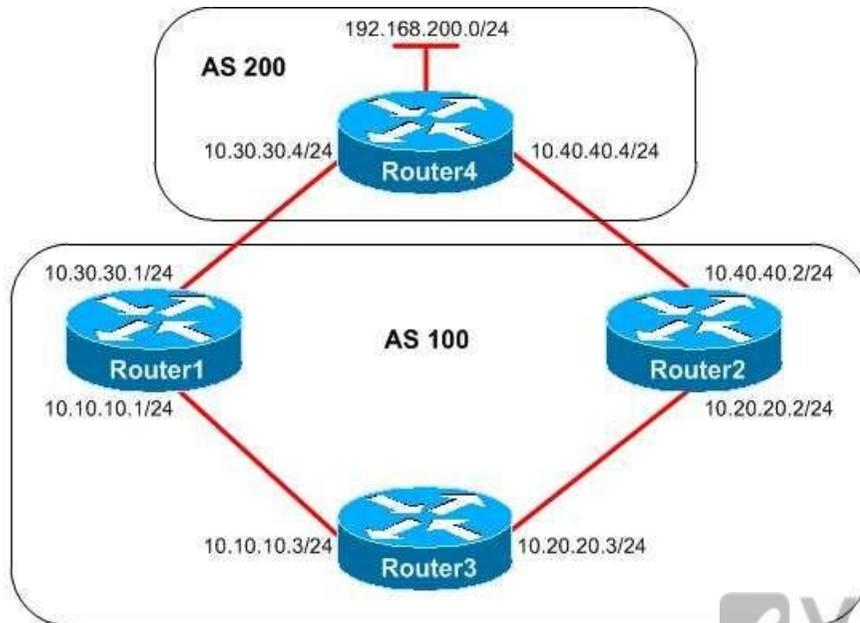
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/system_messages/reference/sys_Book.html#wp1400041

[https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_2/nx-](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_2/nx-os/layer2/configuration/guide/Cisco_Nexus_7000_Series_NXOS_Layer_2_Switching_Configuration_Guide_Release_4-2_chapter6.html#con_1490082)

[os/layer2/configuration/guide/Cisco_Nexus_7000_Series_NXOS_Layer_2_Switching_Configuration_Guide_Release_4-2_chapter6.html#con_1490082](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_2/nx-os/layer2/configuration/guide/Cisco_Nexus_7000_Series_NXOS_Layer_2_Switching_Configuration_Guide_Release_4-2_chapter6.html#con_1490082)

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/24063-pvid-inconsistency-24063.html#topic1>

QUESTION 209



You administer the network shown in the diagram above. All routers are running BGP, and all attributes are set to the default values. You want to ensure that Router3 sends packets to the 192.168.200.0/24 network through Router2.

Which of the following command sets can you issue to accomplish your goal? (Select 2 choices.)

- A. Router2(config)#router bgp 100
Router2(configrouter)#neighbor 10.20.20.3 weight 45000
- B. Router2(config)#router bgp 100
Router2(configrouter)#neighbor 10.40.40.4 weight 45000
- C. Router2(config)#router bgp 100
Router2(configrouter)#bgp default localpreference 400
- D. Router2(config)#routemap map1 permit 10
Router2(configroutemap)#set localpreference 500
Router2(configroutemap)#exit
Router2(config)#router bgp 100
Router2(configrouter)#neighbor 10.20.20.3 routemap map1 in
- E. Router2(config)#routemap map1 permit 10
Router2(configroutemap)#set localpreference 300

```
Router2(config routemap)#exit
Router2(config)#router bgp 100
Router2(config router)#neighbor 10.40.40.4 routemap map1 in
```

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You can issue the following command set to ensure that Router3 sends packets to the 192.168.200.0/24 network through Router2:

```
Router2(config)#router bgp 100
Router2(config router)#bgp default localpreference 400
```

Alternatively, you can issue the following command set to accomplish your goal:

```
Router2(config)#routemap map1 permit 10
Router2(config routemap)#set localpreference 300
Router2(config routemap)#exit
Router2(config)#router bgp 100
Router2(config router)#neighbor 10.40.40.4 routemap map1 in
```



When a router has multiple paths to a destination and weight values of the routes are equal, the route with the highest local preference is preferred. Typically, the multiexit discriminator (MED) value is used to indicate a preferred path into an AS with multiple entry points and is advertised to external Border Gateway Protocol (eBGP) routers.

However, local preference is considered before the MED value, so you can configure the local preference to prefer one route over another.

Issuing the `bgp default localpreference 400` command configures Router2 to advertise a local preference value of 400 to Router3. By default, a local preference value of 100 is advertised. Therefore, Router1 will advertise a local preference of 100 to Router3. Because the local preference of the route through Router2 is higher than the local preference of the route through Router1, Router3 will prefer the route through Router2.

The local preference value can also be specified by using a route map. The `routemap map1 permit 10` command configures a route map named map1. The permit keyword indicates that the conditions specified in the set command will be processed, and the 10 keyword is a sequence number that specifies the order in which route maps should be processed. Since no match command is specified, the route map will apply to all packets. The `set localpreference 300` command configures a local preference value of 300 for routes affected by the route map. Finally, the `neighbor 10.40.40.4 routemap map1 in` command applies the route map named map1 to incoming routes from 10.40.40.4. Router2 will then advertise this route to Router3 with a local preference value of 300. Because the local preference of the route through Router2 is higher than the default local preference of the route through Router1, Router3 will prefer the route through Router2. You cannot accomplish your goal by issuing the following command set:

```
Router2(config)#routemap map1 permit 10
Router2(config)routemap)#set localpreference 500
Router2(config)routemap)#exit
Router2(config)#router bgp 100
Router2(config)router)#neighbor 10.20.20.3 routemap map1 in
```

In this command set, the neighbor command incorrectly specifies Router3, not Router4, as the neighbor router. Therefore, routes received by Router2 from Router3 would be assigned a local preference value of 500 and advertised to Router4.

You cannot accomplish your goal by issuing the following command set:

```
Router2(config)#router bgp 100
Router2(config)router)#neighbor 10.40.40.4 weight 45000
```

When determining the best path, a BGP router first chooses the route with the highest weight. By default, routes generated by the local router are assigned a weight of 32768 and routes learned from another BGP router are assigned a weight of 0. Issuing the neighbor 10.40.40.4 weight 45000 command on Router2 would configure the path toward Router4 with a weight value of 45000. However, this weight value is significant only to Router2; it would not be advertised to Router3. Therefore, issuing the neighbor 10.40.40.4 weight 45000 command on Router2 would not influence routing decisions on Router3. Issuing the neighbor 10.20.20.2 weight 45000 command on Router3 would ensure that Router3 preferred the route through Router2.

You cannot accomplish your goal by issuing the following command set, because the neighbor command would configure Router2 with a weight value for the path toward Router3:

```
Router2(config)#router bgp 100
Router2(config)router)#neighbor 10.20.20.3 weight 45000
```

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-a1.html#wp9538078130 https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/command/iri-cr-book/iri-cr-s1.html#wp4033207811 https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-m1.html#wp2222404444 <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html> https://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfbgp.html#wp1040707

QUESTION 210

Which of the following statements is accurate regarding 6to4 tunneling? (Select the best answer.)

- A. IPv6 addresses from the 2002::/16 prefix must be used.
- B. IPv6 packets are encapsulated inside an IPv4 packet that has a protocol type of 42.
- C. IPv6 and CLNS can be used.
- D. Any IPv6 unicast address can be used.

E. A tunnel destination is required.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

IPv6 addresses from the 2002::/16 prefix must be used with 6to4 tunneling. The 32 bits following the 2002::/16 prefix correspond to the IPv4 address of the tunnel source. For example, if the IPv4 address of the tunnel source is 10.14.177.1, you could use 2002:0A0E:B101::/128 for the IPv6 address of the tunnel interface; the hexadecimal value 0A corresponds to the decimal value 10, the hexadecimal value 0E corresponds to the decimal value 14, the hexadecimal value B1 corresponds to the decimal value 177, and the hexadecimal value 01 corresponds to the decimal value 1.

Although an IPv6 unicast address can be used with the IntraSite Automatic Tunnel Addressing Protocol (ISATAP) tunneling method, only IPv6 addresses with the 2002::/16 prefix can be used with 6to4 tunneling.

IPv6 packets are encapsulated inside an IPv4 packet that has a protocol type of 41, not a protocol type of 42. When an IPv4 packet with a protocol type of 41 arrives on the router interface, the IPv6 packet is decapsulated and mapped to the appropriate IPv6 tunnel interface based on the IPv4 address.

Connectionless Network Service (CLNS) cannot be used with 6to4 tunneling. CLNS packets can be encapsulated using Generic Routing Encapsulation (GRE) and IPv4-compatible tunneling. GRE and IPv4-compatible tunneling can encapsulate a variety of Network layer packets, including IPv6.

Because 6to4 tunneling is an automatic, point-to-multipoint tunneling method, a tunnel destination is not required. Other point-to-multipoint tunneling methods include ISATAP and IPv4-compatible tunneling. GRE and manual tunneling are point-to-point tunneling methods that require an IPv4 address for the tunnel destination.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ipv6-15-2mt-book/ip6-tunnel.html>

QUESTION 211

Which of the following planes is responsible for exchanging routing and label information? (Select the best answer.)

- A. the data plane
- B. the forwarding plane
- C. the control plane
- D. the management plane
- E. the routing plane

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The control plane is responsible for exchanging routing and label information. Routing information is exchanged by using a routing protocol, such as Border Gateway Protocol (BGP), Enhanced Interior Gateway Routing Protocol (EIGRP), Intermediate System to Intermediate System (ISIS), or Open Shortest Path First (OSPF). Label information is exchanged by using a label exchange protocol, such as Resource Reservation Protocol (RSVP), Tag Distribution Protocol (TDP), or Label Distribution Protocol (LDP). LDP is a newer standard that includes features of the Cisco proprietary TDP. RSVP is used by Multiprotocol Label Switching Traffic Engineering (MPLS TE) to also reserve network bandwidth. Bandwidth is reserved on demand based on destination address or traffic type so that enough bandwidth is available for the traffic.

Cisco routers are separated into three planes: the management plane, the control plane, and the data plane. Cisco does not define a routing plane. The data plane, which is also called the forwarding plane, is responsible for forwarding packets. Packets are forwarded based on destination address or label information. The Forwarding Information Base (FIB), which is part of the data plane, is built from information in the routing table. When the routing table is updated, the next hop information in the FIB is also updated. Multiprotocol Label Switching (MPLS) label information is stored in the Label Forwarding Information Base (LFIB) table? the data plane then uses that information to forward the packets to the correct destination.

The management plane is responsible for device management and coordination between the three planes. Protocols used by the management plane include Hypertext Transfer Protocol (HTTP), Secure HTTP (HTTPS), Simple Network Management Protocol (SNMP), Secure Shell (SSH), and Telnet.

Reference:

https://www.cisco.com/c/en/us/products/collateral/security/ios-network-foundation-protection-nfp/prod_white_paper0900aecd805ffde8.html

QUESTION 212

Which of the following statements best describes RTCP? (Select the best answer.)

- A. RTCP is used to monitor and transport audio and video packets.
- B. RTCP is used to monitor and report statistics about an RTP session.
- C. RTCP is used to monitor and control packet loss for an RTP session.
- D. RTCP is used to monitor and control jitter for an RTP session.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RealTime Transport Control Protocol (RTCP) is used to monitor and report statistics about a RealTime Transport Protocol (RTP) session. RTCP monitors the RTP session and collects statistical information about transmitted octet and packet counts, packet delay, packet loss, and jitter. An RTCP packet is sent from the originating device to the remote device every five seconds.

A two way audio session, such as a telephone conversation, requires two RTP streams and two RTCP streams: one pair of streams originating from each device. An RTP session is established on an even numbered User Datagram Protocol (UDP) port ranging from 16384 through 32767. The corresponding RTCP session is

established on the port immediately following the UDP port established by RTP. Once a UDP port pair is established by an IP phone, the phone uses those ports for the duration of the session.

RTP, not RTCP, is used to monitor and transport audio and video packets. RTP adds three pieces of information to an audio or video data packet header: the payload type, the sequence number, and a time stamp. The payload type indicates whether the data is audio or video data. Sequence numbers are used to determine how to order incoming audio or video packets to reconstruct the data. Time stamps are useful for creating a buffer to mitigate delays between packets. RTCP is not used to control packet loss for an RTP session. Although RTCP monitors and reports on packet loss for an RTP stream, RTCP does not prevent packets from being lost during transfer.

RTCP is not used to control jitter, which is a variation in delay, for an RTP session.

Although RTCP monitors and reports on packet delay and jitter, RTCP cannot actively control packet delay and jitter.

Reference: https://www.cisco.com/c/en/us/td/docs/interfaces_modules/shared_port_adapters/configuration/ASR1000/asr1000-sip-spa-book.html#pgfId-1296118 **QUESTION 213**

You administer a network that contains only nonCisco routers. You purchase a new Cisco router and connect the S0/0 interface to one of the nonCisco routers. During the initial system configuration dialog, you choose yes when you are asked whether you want to configure S0/0. You configure the IP address and subnet mask, and you choose the router defaults for the other options.

After you complete the configuration, you discover that you are unable to ping the serial interface on the nonCisco router, which is configured with an IP address of 10.10.10.2/30. You issue the show interfaces serial 0/0 command on the Cisco router and receive the following partial output:

```
Serial0/0 is up, line protocol is down
Hardware is HD64570
Internet address is 10.10.10.1 255.255.255.252
MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec, rely 255/255, load 1/255
Encapsulation HDLC, loopback not set, keepalive set (10 sec)
```

Why are you unable to ping the serial interface of the nonCisco router? (Select the best answer.)

- A. The subnet mask used during the initial configuration is incorrect.
- B. There is another device with the IP address 10.10.10.1 on the network.
- C. The serial interface on the Cisco router is administratively down.
- D. The encapsulation type should be set to PPP.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You are unable to ping the serial interface of the nonCisco router because the encapsulation type on the Cisco router should be set to PointtoPoint Protocol (PPP).

The status line Serial0/0 is up, line protocol is down indicates a Layer 2 problem? Layer 2 problems are often caused by mismatched encapsulation modes. On Cisco routers, HighLevel Data Link Control (HDLC) is the default serial interface encapsulation protocol. Cisco's implementation of HDLC is proprietary and should be used only with other Cisco routers. Because the existing network comprises a mixture of Cisco and nonCisco equipment, PPP encapsulation is a more viable choice than HDLC.

The subnet mask used in the initial configuration is correct. The 255.255.255.252 subnet mask allows for two host addresses on the 10.10.10.0 network: 10.10.10.1 and 10.10.10.2.

If there were another device with the IP address 10.10.10.1 on the network, the Ciscorouter would show a status of Serial0/0 is administratively down, line protocol is down. In the output, S0/0 is up? thus no IP address conflict exists.

Ha the interface been shut down administratively, the interface status would have been Serial0/0 is administratively down, line protocol is down. In the output, S0/0 is up? thus the interface has not been shut down administratively with the shutdown command.

Reference:

<https://www.cisco.com/en/US/docs/internetworking/troubleshooting/guide/tr1915.html>

<https://www.cisco.com/c/en/us/td/docs/routers/access/800/805/software/configuration/guide/805swcfg/overview.html>

QUESTION 214

What AD is assigned to external BGP routes by default? (Select the best answer.)

- A. 1
- B. 5
- C. 20
- D. 90
- E. 170
- F. 200



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

External Border Gateway Protocol (eBGP) routes are assigned an administrative distance(AD) of 20 by default. AD values are used to determine the routing protocol that should be preferred when multiple routes to a destination network exist. A routing protocol with a lower AD will be preferred over a route with a higher AD. The following list contains the most commonly used ADs:

Route Source	Distance
Connected route	0
Static route	1
EIGRP summary route	5
eBGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
iBGP	200
Unknown	255

Internal BGP (iBGP) routes are assigned an AD of 200 by default. Therefore, eBGP routes are preferred over iBGP routes.

Directly connected routes have an AD of 0. Therefore, directly connected routes are trusted over routes from any other source.

Static routes have an AD of 1. Therefore, static routes are more trusted than routes from any routing protocol. Static routes are optimal for routing networks that do not change often. To create a static route, you should issue the ip route command.

Enhanced Interior Gateway Routing Protocol (EIGRP) summary routes are assigned an AD of 5 by default. Routes that are learned by EIGRP are called internal EIGRP routes and have an AD of 90. Routes that are redistributed into EIGRP are called external EIGRP routes and have an AD of 170. To modify the AD values used by EIGRP, you should issue the distance eigrp internal external command, where internal is the AD used for internal EIGRP routes and external is the AD used for external EIGRP routes.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/15986-admin-distance.html>

QUESTION 215

How does SNMPv3 use HMAC-SHA or HMAC-MD5? (Select 2 choices.)

- A. as authentication hashes at the noAuthNoPriv security level
- B. as authentication hashes at the authNoPriv security level
- C. as authentication hashes at the authPriv security level
- D. as encryption hashes at the noAuthNoPriv security level
- E. as encryption hashes at the authNoPriv security level
- F. as encryption hashes at the authPriv security level

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Simple Network Management Protocol version 3 (SNMPv3) uses Hashbased Message Authentication CodeSecure Hash Algorithm (HMACSHA) or HMACMessage Digest 5(HMACMD5) as authentication hashes at the authNoPriv security level and at the authPriv security level. The difference between the authNoPriv security level and the authPriv security level is that the authPriv security level also encrypts the authentication process by using either Data Encryption Standard (DES), Triple DES (3DES), or Advanced Encryption Standard (AES).

There are three SNMPv3 security levels: noAuthNoPriv, authNoPriv, and authPriv. Both the authPriv security level and the authNoPriv security level authenticate by matching a hash of the user name. The noAuthNoPriv security level, on the other hand, authenticates by matching the user name in clear text.

The noAuthNoPriv security level operates differently in SNMPv3 than it does in SNMPv1 and SNMPv2C. Both SNMPv1 and SNMPv2C match a clear-text community string, not a user name, to authenticate.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/software/15-0_2_EX/network_management/configuration_guide/b_nm_15ex_2960-x_cg/b_nm_15ex_2960-x_cg_chapter_0100.html#reference_160326642C03413B92A68E856426EABA

QUESTION 216

Which of the following is an optional, nontransitive BGP path attribute? (Select the best answer.)

- A. aggregator
- B. AS-path
- C. origin
- D. originator ID
- E. next hop
- F. weight

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Originator ID is an optional, nontransitive Border Gateway Protocol (BGP) path attribute. An optional BGP path attribute can be supported by a particular BGP implementation but is not required. A nontransitive BGP path attribute is not forwarded to BGP peers if the attribute is not supported on a particular BGP implementation. Multiexit discriminator (MED) and cluster list are also optional, nontransitive BGP path attributes. Internet Engineering Task Force (IETF)standard BGP path attributes can be broken down into the following categories:

<http://www.gratisexam.com/>

Well-Known

Mandatory	Discretionary
AS-path	atomic aggregate
origin	local preference
next hop	

Optional

Transitive	Nontransitive
aggregator	MED
community	originator ID
	cluster list

Aggregator and community are both optional, transitive BGP path attributes. A transitive BGP path attribute must be passed to BGP peers, even if the attribute is not supported.

AS-path, origin, and next hop are all wellknown, mandatory BGP path attributes. All BGP implementations are required to recognize wellknown BGP path attributes. A mandatory BGP attribute must be included in every BGP update. A discretionary BGP attribute, on the other hand, can be included in a BGP update under specific sets of circumstances. Atomic aggregate and local preference are wellknown, discretionary BGP attributes.

Weight is a Cisco proprietary BGP path attribute. The value assigned to the weight attribute is not passed to BGP peers. The route that has been assigned the highest BGP weight value is considered the best route.

Reference:

<https://tools.ietf.org/html/rfc4271#section-5.1.4>

CCIE Routing and Switching v5.0 Certification Guide, Volume 2, Chapter 2, Generic Terms and Characteristics of BGP Path Attributes, pp. 93-95

QUESTION 217

DRAG DROP

Select the routing methods on the left, and place them in order in which PfR will search their respective databases in order to find a parent route.

Select and Place:

BGP	1
RIB	2
EIGRP	3
static	4

Correct Answer:

 **VCEplus**
VCE To PDF - Free Practice Exam

	BGP
	EIGRP
	static
	RIB

Section: (none)
Explanation

Explanation/Reference:

Explanation:

Performance Routing (PfR) enhances traditional routing methods by dynamically selecting the best path for traffic classes based on network performance. The path selection procedure can be influenced by several factors, including delay, packet loss, reachability, throughput, jitter, and mean opinion score (MOS). When PfR wants to modify a path for a traffic class, it will search for a parent route, which is an exact matching route or a less specific route. PfR will search for a parent route in the following locations, in order:

1. Border Gateway Protocol (BGP) routing database
2. Enhanced Interior Gateway Routing Protocol (EIGRP) routing database
3. Static route database
4. Routing Information Base (RIB)

PfR can directly control path selection for BGP, EIGRP, and static routes. Protocol

Independent Routing Optimization (PIRO) enables PfR to also search for a parent route within the IP RIB, extending support for PfR to routing protocols such as Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS).

Reference: https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/performance-routing-pfr/product_data_sheet0900aecd806c4ee4.html http://docwiki.cisco.com/wiki/Performance_Routing_FAQs
<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/pfr/configuration/xr-3s/pfr-xr-3s-book/pfr-simple-ph1.html>
https://www.cisco.com/c/dam/global/bg_bg/assets/ciscoexpo2011/pdf/Next_Generation_Routing_Architectures-Gerd_Pflueger.pdf

QUESTION 218

You issue the ip sla schedule 20 start-time 9:00:00 command.

How long will the IP SLA operation run? (Select the best answer.)

- A. until the operation has completed once
- B. for 20 minutes
- C. for one hour
- D. for nine hours
- E. forever

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The IP Service Level Agreement (SLA) operation will run for one hour. IP SLA operations are a suite of tools on Cisco devices that enable an administrator to analyze and troubleshoot IP networks. For example, the following command set configures IP SLA to regularly test and verify the reachability of IP address 10.10.10.2:

```
ip sla 1
  type echo protocol ipIcmpEcho 10.10.10.2
  timeout 1000
  threshold 2
  frequency 3
ip sla schedule 1 life forever start-time now
```

The syntax of the ip sla schedule command is ip sla schedule operationnumber [life{forever | seconds}] [starttime {hh:mm[:ss] [monthday | daymonth] | pending | now | afterhh:mm:ss | random milliseconds}] [ageoutseconds] [recurring]. The ip sla schedule command has replaced the ip sla monitor schedule command, which you might see on older IOS versions.

The life keyword specifies how long the operation should run. If the life keyword is not specified, such as in the ip sla schedule 20 starttime 9:00:00 command, the operation will run for 3,600 seconds, or one hour. The life keyword is not specified in the ip sla schedule 20 starttime 9:00:00 command in this scenario? therefore, the operation will run for one hour.

The operationnumber variable indicates the number of the IP SLA operation that is to be scheduled. The ip sla schedule 20 starttime 9:00:00 command in this scenario specifies that IP SLA operation 20 is to be scheduled.

The starttime keyword indicates when the IP SLA operation should start. If the starttime keyword is not specified, the operation is placed in a pending state and will not run automatically? issuing the starttime pending keywords also places the operation in a pending state. The ip sla schedule 20 starttime 9:00:00 command in this scenario specifies that IP SLA operation 20 should start at 9 a.m.

The ip sla schedule command does not influence how often an IP SLA operation is repeated. To change how often an IP SLA operation is repeated, you can issue the frequency command from an IP SLA configuration submenu. If the frequency command is not configured, the IP SLA operation will repeat every 60 seconds. The variable for the frequency command is specified in seconds; therefore, the frequency 60 command has the same effect as the default frequency of 60 seconds.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/command/sla-cr-book/sla_i1.html#wp3201991432

QUESTION 219

Which of the following can you configure on a router to gather userlevel network resource utilization statistics? (Select the best answer.)

- A. CoPP
- B. AutoQoS
- C. NetFlow
- D. traffic policing
- E. traffic shaping

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You can configure NetFlow on a router to gather userlevel network resource utilization statistics. NetFlow is a Cisco IOS feature that can be used to silently monitor traffic flows. A traffic flow is defined as a series of packets with the same source IP address, destination IP address, protocol, and Layer 4 information. NetFlow caches flowbased statistics, such as packet counts, byte counts, and protocol distribution. The data gathered by NetFlow is typically exported to flow collectors or other management software. You can then analyze the data to facilitate network planning, customer billing, and traffic engineering.

You can configure Control Plane Policing (CoPP) to protect the management and control planes on a Cisco router, not to gather userlevel network utilization statistics. The control plane is one of the four logical components that collectively define a router? the remaining components are the data plane, the management plane, and the services plane. The control plane contains the route processor; routing protocol operation, network management, and processbased switching are handled by the control plane. CoPP filters the types of packets that enter or exit the control plane and controls the rate at which permitted packets enter or exit the control plane. Because traffic must pass through the control plane to reach the management plane, CoPP protects the management plane as well.

You can configure AutoQoS on a router to help automate the configuration of Quality of Service (QoS), not to gather userlevel network utilization statistics. For example, you can implement AutoQoS to generate the commands necessary to configure QoS for a specific network configuration, thereby simplifying QoS configuration.

You can configure traffic policing on a router to limit the rate of traffic that passes through an interface, not to gather userlevel network utilization statistics. With traffic policing, packet flows that exceed the configured thresholds are typically dropped. Alternatively, traffic can be remarked with a lower priority before being transmitted.

Similarly, you can configure traffic shaping on a router to limit the rate of traffic that passes through an interface. However, with traffic shaping, flows that exceed the configured thresholds are typically buffered and not dropped. Because traffic is buffered to maintain a desired packet rate, erratic patterns are smoothed into a uniform flow. You cannot configure traffic shaping on a router to gather userlevel network utilization statistics.

Reference:

https://www.cisco.com/en/US/tech/tk812/technologies_white_paper09186a008022bde8.shtml#wp1002404

QUESTION 220

Which of the following statements is true regarding HSRPv2 and HSRPv6? (Select the best answer.)

- A. Both support 4,096 groups.
- B. Both use the same UDP port number.
- C. Both use the same multicast address.
- D. Both use the same virtual MAC address range.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Both Hot Standby Router Protocol version 2 (HSRPv2) and HSRPv6 support 4,096 groups. HSRP is a First Hop Redundancy Protocol (FHRP) that enables multiple routers to act as a single gateway. The HSRP virtual IP address can be configured as the default gateway address for client devices.

Multiple routers can be assigned to an HSRP group, but each group has only one active router and one standby router. The active router is the router in the group with the highest HSRP priority value, and the standby router is the router with the second highest HSRP priority value. Other routers in the HSRP group are in the listen state. If the active router fails, the standby router assumes the active router role and a new standby router is elected.

HSRPv1 supports only 256 groups, numbered from 0 through 255. HSRPv2 improves upon HSRPv1 by increasing the number of groups to 4,096. HSRPv6, which is also called HSRP for IPv6, further improves HSRPv2 by adding support for IPv6. The default HSRP group value for all HSRP versions is 0.

HSRPv2 and HSRPv6 do not use the same User Datagram Protocol (UDP) port number. HSRPv1 and HSRPv2 use UDP port number 1985. HSRPv6 uses UDP port number 2029.

HSRPv2 and HSRPv6 do not use the same multicast address. Multicast addresses are used to send Hello packets to group members. By default, Hello packets are sent by the active router every three seconds. Only the standby router monitors the active router's Hello packets. If the standby router does not receive a Hello packet from the active router for the duration configured in the Hold time, the standby router takes over the role of the active router. By default, the Hold time is set to 10 seconds. HSRPv1 uses multicast address 224.0.0.2. HSRPv2 uses multicast address 224.0.0.102. HSRPv6 uses multicast address FF02::66.

HSRPv2 and HSRPv6 do not share the same virtual Media Access Control (MAC) address range. Although each router has a unique MAC address, the routers in an HSRP group will share a virtual MAC address. HSRPv1 uses the virtual MAC address 0000.0c07.acxx, where xx is the group number in hexadecimal format.

Therefore, an HSRPv1 router could have a virtual MAC address from 0000.0c07.ac00 through 0000.0c07.acff. HSRPv2 uses the virtual MAC address 0000.0c9f.fxxx, where xxx is the group number in hexadecimal format. Therefore, an HSRPv2 router could have a virtual MAC address from 0000.0c9f.f000 through 0000.0c9f.ffff. HSRPv6 uses the virtual MAC address 0005.73a0.0xxx, where xxx is the group number in hexadecimal format. Therefore, an HSRPv6 router could have a virtual MAC address from 0005.73a0.0000 through 0005.73a0.0fff.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750e_3560e/software/release/12-2_55_se/configuration/guide/3750escg/swhsrp.html

<https://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol-hsrp/9281-3.html#q34> https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp_fhrp/configuration/15-sy/fhp-15-sy-book/HSRP-for-IPv6.html

QUESTION 221

You have just configured an NHRP hub router and an NHRP spoke router. You issue the show runningconfig command on the hub and receive the following partial output:

```
interface Tunnel0
  ip nhrp network-id 1
  ip nhrp authentication Boson
  ip nhrp map multicast dynamic
```

You issue the show running-config command on the spoke and receive the following partial output:

<http://www.gratisexam.com/>

```
interface Tunnel1
 ip nhrp network-id 2
 ip nhrp authentication Boson
 ip nhrp nhs 10.1.1.1
 ip nhrp map multicast 192.168.1.1
 ip nhrp map 10.1.1.1 192.168.1.1
```

Which of the following statements is true? (Select the best answer.)



<https://vceplus.com/>

- A. The tunnel should establish normally.
- B. The tunnel will not establish, because the network IDs are different.
- C. The tunnel will not establish, because the authentication keys cannot be fewer than eight characters in length.
- D. The tunnel will not establish, because the tunnel interface IDs are not the same.
- E. The tunnel will not establish, because the hub is missing the ip nhrp nhs command.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The tunnel should establish normally. Next Hop Resolution Protocol (NHRP) dynamically learns IP addresses of spoke routers in a hubandspoke Dynamic Multipoint virtual private network (DMVPN) environment, which uses Generic Routing Encapsulation (GRE) tunneling.

The authentication key must match on an NHRP hub router and an NHRP spoke router. If the authentication key does not match on the hub and the spoke, the tunnel will not establish. To configure the authentication key, issue the ip nhrp authentication keycommand in interface configuration mode for the tunnel interface.

The authentication key can be any value up to eight characters in length? therefore, having a fivecharacter authentication key on the hub and spoke will allow the tunnel to establish as long as the authentication keys match. The authentication key is casesensitive.

The NHRP network ID need not match on the hub and spokes. The NHRP network ID is used to identify the NHRP domain for an interface when two or more NHRP domains are configured on the same device. The network ID is a locally significant value and is not sent out in any NHRP packets. When NHRP packets arrive on an interface, those packets are assigned to the network ID that is configured on that interface. Although it is easier to keep track of NHRP domains if all of

<http://www.gratisexam.com/>

the devices in the NHRP domain are configured with the same NHRP network ID, it is not required. To configure the network ID, issue the ip nhrp networkid networkid command in interface configuration mode for the tunnel interface. The network key can be any value from 1 through 4294967295. The tunnel interface ID need not match on the hub and spokes. The tunnel interface ID number is a locally significant value. As long as the tunnel is configured properly, the tunnel between the hub and the spoke will establish regardless of the interface ID that is used for the tunnel. The ip nhrp nhs command configures a spoke router with the tunnel address of the hub. Hu routers need not be configured with the ip nhrp nhs command? as spoke routers register with the hub, their IP addresses are automatically discovered.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_4/ip_addr/configuration/guide/hadnhrp.html

QUESTION 222

Which of the following best defines a BGP confederation? (Select the best answer.)

- A. a division of an AS
- B. one or more route reflectors and their client peers
- C. a group of peers with the same update policies
- D. a group of destinations to which the same routing decisions should be applied

Correct Answer: A

Section: (none)

Explanation



Explanation/Reference:

Explanation:

A Border Gateway Protocol (BGP) confederation is a division of an autonomous system (AS). A confederation enables an AS to be divided into discrete units, each of which acts like a separate AS. The routers within each confederation must be fully meshed with one another unless a route reflector is established. BGP updates are passed within the confederation and between confederations. To configure a confederation, you should issue the bgp confederation identifier as-number and bgp confederation peers as-number commands.

A cluster consists of one or more route reflectors and their client peers. Internal BGP (iBGP) routes are not advertised to iBGP peers. In order to avoid having to create a full mesh configuration, you can configure one or more route reflectors to pass iBGP routes between iBGP routers. Each route reflector in the cluster should be fully meshed with one another. In addition, each route reflector in the cluster should be configured with the same 4-byte cluster ID so that the route reflectors in the cluster can recognize routing updates from each other. To configure a route reflector with a cluster ID, you should issue the bgp cluster-id cluster-id command from BGP router configuration mode.

A peer group is a group of peers with the same update policies. Peer groups can simplify administration by enabling an administrator to simultaneously configure a group of peers with the same update policies, such as route maps, filter lists, and distribute lists. Any configuration options that are configured with the specified peer group name will be applied to members of the peer group. To define a peer group, you should issue the neighbor peer-group-name peer-group command. A community is a group of destinations to which the same routing decisions should be applied. By default, Cisco routers do not pass community attributes to BGP

neighbors. To configure a router to send community attributes to a neighbor, you should issue the neighbor {ipaddress | peergroupname} sendcommunity [standard | extended | both] command. The community attribute can be modified within a route map by issuing the set community command with one of the following four keywords:

- no-advertise -prevents advertisements to any BGP peer
- no-export-prevents advertisements to eBGP peers
- local-as -prevents advertising outside the AS, or in confederation scenarios, outside the sub-AS
- internet-advertises the route to any router

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfbgp.html#wp1024370

<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/26634-bgp-toc.html#bgpconfed>

QUESTION 223

Which of the following routes from the show ip route command was learned from a Type 3 or Type 4 LSA? (Select the best answer.)

- A. O E1 172.17.1.0/24 [110/74] via 10.1.2.3, 00:00:23, FastEthernet1/0
- B. O 172.17.2.0/24 [110/74] via 10.1.2.3, 00:00:23, FastEthernet1/0
- C. O IA 172.17.3.0/24 [110/74] via 10.1.2.3, 00:00:23, FastEthernet1/0
- D. C 172.17.4.0/24 is directly connected, FastEthernet1/1
- E. S 172.17.5.0/24 is directly connected, FastEthernet0/1



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following route from the show ip route command was learned from a Type 3 or Type 4 linkstate advertisement (LSA):

O IA 172.17.3.0/24 [110/74] via 10.1.2.3, 00:00:23, FastEthernet1/0

A routing table entry that begins with O IA indicates an Open Shortest Path First (OSPF) interarea summary route from a Type 3 or Type 4 LSA. Interarea routes are routes that are advertised between areas. These routes are not propagated through totally stubby areas.

A routing table entry that begins with O E1 or O E2 indicates an OSPF external summary route from a Type 5 LSA. Type 5 LSAs are not propagated through stub areas, not so stubby areas (NSSAs), or totally stubby areas. By default, routes are redistributed into OSPF as Type 2 external routes, which are indicated by an E2 in the output of the show ip route command. Type 2 external routes have a metric that remains constant throughout the autonomous system (AS). Type 1 external routes, which are indicated by an E1 in the output of the show ip route command, have a metric that increases as the route is propagated throughout the AS. A

routing table entry that begins with O indicates an intraarea route from a Type 1 or Type2 LSA. Intraarea routes are advertised within an area. Type 1 and Type 2 LSAs are accepted by all OSPF area types.

A routing table entry that begins with C indicates a directly connected route. A routing table entry that begins with S indicates a static route, which is configured by issuing the ip route command. Neither of these routes is learned from a Type 3 or Type 4 LSA.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/47865-ospfdb6.html>

QUESTION 224

Which of the following statements are true regarding the ipv6 ospf authentication command and the ipv6 ospf encryption command? (Select 3 choices.)

- A. Both commands require AH.
- B. Both commands require ESP.
- C. Both commands enable encryption.
- D. Both commands enable authentication.
- E. Both commands require an SPI value.
- F. Both commands must be configured from interface configuration mode.

Correct Answer: DEF

Section: (none)

Explanation



Explanation/Reference:

Explanation:

Both commands enable authentication, require a Security Parameter Index (SPI) value, and must be configured from interface configuration mode. Open Shortest Path First version 3 (OSPFv3) uses IP Security (IPSec) to provide security. IPSec relies upon Authentication Header (AH) for authentication and Encapsulating Security Payload (ESP) for encryption. An IPSec security policy consists of an SPI and a key. The SPI value that is configured in the ipv6 ospf authentication command and the ipv6 ospf encryption command must be a value from 256 through 4294967295.

The ipv6 ospf authentication command enables only authentication, not encryption. Therefore, it requires AH. Either Message Digest 5 (MD5) or Secure Hash Algorithm 1 (SHA1) can be used as the authentication algorithm.

The ipv6 ospf encryption command enables both authentication and encryption. Therefore, it requires ESP.

ESP can be used by itself or in conjunction with AH. Either

Advanced Encryption Standard (AES), Data Encryption Standard (DES), or Triple DES(3DES) can be used as the encryption algorithm.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/15-sy/iro-15-sy-book/ip6-route-ospfv3-auth-ipsec.html

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/command/ipv6-cr-book/ipv6-i3.html#wp3695874190> <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/command/ipv6-cr-book/ipv6-i4.html#wp2881247299>

QUESTION 225

What restrictions apply to OSPF virtual link creation? (Select 3 choices.)

- A. The routers at each end of the virtual link must share a common area.
- B. The router at the far end of the virtual link cannot connect to a stub area.
- C. The transit area cannot be a stub area.
- D. One router must connect to the backbone area.
- E. The virtual link must pass through the backbone area.

Correct Answer: ACD

Section: (none)

Explanation

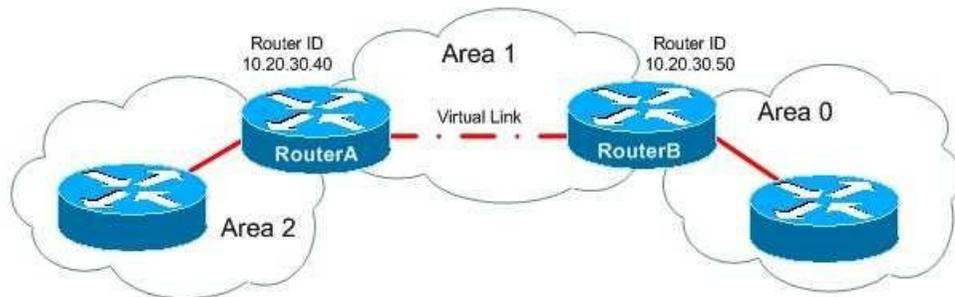
Explanation/Reference:

Explanation:

The following restrictions apply to Open Shortest Path First (OSPF) virtual link creation:

1. The routers at each end of the virtual link must share a common area.
2. The transit area cannot be a stub area.
3. The transit area cannot be the backbone area.
4. One router must connect to the backbone area.

All areas in an OSPF internetwork must be connected to the backbone area, Area 0. If a remote area has no direct connection to the backbone area, you can create a virtual link between two area border routers (ABRs) to connect the remote area to the backbone area through a transit area. One ABR is connected to Area 0 and the transit area, and the other ABR is connected to the transit area and the remote area. The following displays a virtual link between Area 2 and Area 0 through the transit area, Area 1:



The virtual link passes through the transit area, Area 1, not the backbone area, Area 0. Although Area 2 can be a stub area, the transit area, Area 1, cannot be a stub area.

To create a virtual link, you should issue the `area areaidvirtuallinkrouterid` command in router configuration mode, where `areaid` is the transit area ID and `routerid` is the router ID of the router at the other end of the virtual link. For example, to create a virtual link between RouterA and RouterB in the example, you should issue the `area 1 virtuallink 10.20.30.50` command on RouterA and the `area 1 virtuallink 10.20.30.40` command on RouterB.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html> <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html>

QUESTION 226

Which of the following ICMPv6 message types is used to query for the link-layer address of a host? (Select the best answer.)

- A. router solicitation
- B. router advertisement
- C. neighbor solicitation
- D. neighbor advertisement

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

An Internet Control Message Protocol version 6 (ICMPv6) neighbor solicitation message is used to query for the linklayer address of a host. Neighbor solicitation messages are sent with the sender's own linklayer address to the solicitednode multicast address. The solicitednode multicast address is created by adding the FF02::1:FF00/104 prefix to the last 24 bits of the destination host's IPv6 address. After a destination host's linklayer address is discovered, neighbor solicitations can be used to verify the reachability of a destination host.

When a host receives a neighbor solicitation message, it will reply with a neighbor advertisement message that contains the linklayer address of the host. The neighbor advertisement is sent directly to the host that sent the neighbor solicitation. A host will send an unsolicited neighbor advertisement whenever its address changes. Unsolicited neighbor advertisements are sent to the allnodes linklocal multicast address FF02::1.

A router solicitation message is sent by an IPv6capable host at startup. When IPv6 is enabled on a router interface, a linklocal address is created. Before the address is assigned to the interface, duplicate address detection (DAD) is performed to determine whether the IPv6 address is unique on the link. If DAD determines that the address is unique, the linklocal address is assigned to the interface and the router solicitation message is sent to the allrouters multicast address FF02::2. Hosts use router solicitation messages to request an immediate router advertisement.

A router advertisement that is sent in response to a router solicitation message is sentdirectly to the host that sent the router solicitation. Routers also periodically send unsolicited router advertisements to the allnodes multicast address FF02::1. Router advertisements contain the following information: -The IPv6 address of the router interface attached to the link

-One or more IPv6 prefixes for the local link

- The lifetime for each prefix
- Flags that specify whether stateless or stateful autoconfiguration can be used -
- The hop limit and maximum transmission unit (MTU) that the host should use
- Whether the router is a default router
- The amount of time that the router can be used as a default router

When a host receives a router advertisement, the IPv6 linklocal prefix is added to the host's interface identifier to create the host's full IPv6 address. The first three octets of the interface identifier are set to the Organizationally Unique Identifier (OUI) of the Media Access Control (MAC) address of the interface. The fourth and fifth octets are set to FFFE. The sixth, seventh, and eighth octets are equal to the last three octets of the MAC address.

Reference: <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ipv6-15-2mt-book/ip6-addrg-bsc-con.html#GUID-82127508-12DE-489C-B004-52CE8DB1415F>
<https://tools.ietf.org/html/rfc4861>

QUESTION 227

Which of the following is true regarding RTC? (Select the best answer.)

- A. RTC sends only the prefixes that the PE router wants.
- B. RTC finds route inconsistencies.
- C. RTC synchronizes peers without a hard reset.
- D. RTC works with only VPNv4.
- E. RTC makes the ABR an RR and sets the next hop to self.



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Route Target Constraint (RTC) sends only the prefixes that the Provider Edge (PE) router wants. In a normal Multiprotocol Label Switching (MPLS) virtual private network (VPN), the route reflector (RR) sends all of its VPN version 4 (VPNv4) and VPNv6 prefixes to the PE router. The PE router then drops the prefixes for which it does not have a matching VPN routing and forwarding (VRF). RTC allows a PE router to send its route target (RT) membership data to the RR within an address family named rfilter. The RR then uses rfilter to determine which prefixes to send to the PE. In order for RTC to work, both the RR and the PE need to support

RTC.

RTC does not find route inconsistencies, nor does it synchronize peers without a hard reset. This functionality is provided by Border Gateway Protocol (BGP) Enhanced Route Refresh.

BG Enhanced Route Refresh is enabled by default. If two BGP peers support EnhancedRoute Refresh, each peer will send a RouteRefresh StartofRIB (SOR) message and a RouteRefresh EndofRIB (EOR) message before and after an AdjRIBOut message, respectively. After a peer receives an EOR message, or after the EOR timer expires, the peer will check to see whether it has any routes that were not readvertised. If any stale routes remain, they are deleted and the route inconsistency is logged.

RTC does not make the area border router (ABR) an RR, nor does it set the next hop to self. This behavior is exhibited by Unified MPLS. Unified MPLS increases scalability for an MPLS network by extending the label switched path (LSP) from end to end, not by redistributing interior gateway protocols (IGPs) into one another, but by distributing some of the IGP prefixes into BGP. BGP then distributes those prefixes throughout the network.

Reference: <https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/116062-technologies-technote-restraint-00.html>
<https://search.cisco.com/search?query=Cisco%20IOS%20BGP%20Configuration%20Guide&locale=enUS&tab=Cisco>
<https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/116127-configure-technology-00.html>

QUESTION 228

DRAG DROP

Select each BGP command on the left, and drag it o its corresponding description on the right.

Select and Place:

Commands	Descriptions
aggregate-address	displays the number of keepalive messages exchanged between BGP peers
clear ip bgp	displays the session state
show ip bgp neighbors	indicates a summarized route that should be advertised
show ip bgp summary	rebuilds the BGP routing table

Help Reset Done

Correct Answer:

Commands	Descriptions
	show ip bgp neighbors
	show ip bgp summary
	aggregate-address
	clear ip bgp

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show ip bgp neighbors command displays the number of keepalive messages exchanged between Border Gateway Protocol (BGP) peers. The keepalive statistics are displayed within the Message statistics block of output. Additional information provided by the show ip bgp neighbors command includes detailed neighbor path, prefix, capability, and attribute information.

The following is sample output from the show ip bgp neighbors command:

```
BGP neighbor is 192.168.15.15, remote AS 1, internal link
BGP version 4, remote router ID 192.168.15.15
BGP state = Established, up for 00:14:21
Last read 00:00:14, last write 00:00:14, hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received(old & new)
  Graceful Restart Capability:advertised and received
  Address family IPv4 Unicast: advertised and received
Message statistics:
  InQ depth is 0
  OutQ depth is 0

      Sent      Rcvd
Opens:          5          5
Notifications: 0           0
Updates:        0           0
Keepalives:    245        244
Route Refresh: 0           0
Total:         250        249

Default minimum time between advertisement runs is 5 seconds
<output omitted>
Number of NLRI's in the update sent: max 0, min 0

Connections established 5; dropped 4
Last reset never
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled
Local host: 192.168.15.1, Local port: 29485
Foreign host: 192.168.15.15, Foreign port: 179
```



The clear ip bgp command rebuilds the BGP routing table. This command can be used to begin soft reconfiguration or a hard reset. Soft reconfiguration uses stored prefix information in order to rebuild BGP routing tables without breaking down any active peering sessions, whereas a hard reset breaks down the active peering sessions and then rebuilds the BGP routing tables.

The aggregateaddress command indicates a summarized route that should be advertised. The syntax of the aggregateaddress command is aggregateaddressipaddresssubnetmask [summaryonly] [asset]. Typically, the aggregateaddresscommand is issued with the optional summaryonly keyword, which prevents the advertisement of routes with longer prefixes within the summarized range. The optional asset keyword enables BGP to detect loops by generating an aggregate address mathematically from a set of autonomous systems (ASes). You can determine whether an aggregate address has been calculated from a set of ASes by examining the output of the show ip bgp command. For example, the following output displays an aggregate address that summarizes AS 600 by using paths through AS 500 and AS 400:

```

BGP table version is 2, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*> 170.0.0.0/16     2.2.2.1             0         600 {500,400} i
  
```

The show ip bgp summary command displays the session state. You can also use the show ip bgp summary command to determine neighbor path, prefix, capability, and attribute information. If the output indicates that network entries and path entries are consuming a lot of memory, the BGP database might be too large; this can occur when the router is attempting to store the entire global BGP routing table. The following is sample output from the show ip bgp summary command:

```

BGP router identifier 172.16.100.1, local AS number 100
BGP table version is 201, main routing table version 201
25 network entries using 2950 bytes of memory
45 path entries using 4849 bytes of memory
15 BGP path attribute entries using 875 bytes of memory
2 multipath network entries and 4 multipath paths
10 BGP AS-PATH entries using 240 bytes of memory
7 BGP community entries using 168 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
36 received paths for inbound soft reconfiguration
BGP using 33225 total bytes of memory
Dampening enabled. 4 history paths, 0 dampened paths
BGP activity 25/2949 prefixes, 60/1 paths, scan interval 15 secs
Neighbor    V    AS  MsgRcvd  MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
10.111.1.1   4   101     26      22     201    0     0  00:12:42  19
10.112.1.1   4   102     21      51     201    0     0  00:11:25   0
  
```

Reference: https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp5.html#wp1159860
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp2.html#wp1107408
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp1.html#wp1111300
<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5441-aggregation.html#aggregatingwiththeassetargument>
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp5.html#wp1162534

QUESTION 229
 DRAG DROP

Select the multicast MAC addresses on the left, and place them on the corresponding protocols on the right.

Select and Place:

0100.0CCC.CCCC	802.1X
0100.0CCC.CCCD	CGMP
0100.0CDD.DDDD	LLDP
0180.C200.0000	CDP, DTP, PAgP, UDLD, and VTP
0180.C200.0003	native VLAN STP BPDUs
0180.C200.000E	nonnative VLAN STP BPDUs

Help

Reset

Done

Correct Answer:

	0180.C200.0003
	0100.0CDD.DDDD
	0180.C200.000E
	0100.0CCC.CCCC
	0180.C200.0000
	0100.0CCC.CCCD

Help

Reset

Done

Section: (none)
Explanation

Explanation/Reference:

Explanation:

Multicast Media Access Control (MAC) address 0180.C200.0003 is used by 802.1X. The Institute of Electrical and Electronics Engineers (IEEE) 802.1X standard defines a method that uses Extensible Authentication Protocol (EAP) to establish portbased connections.

IEEE 802.1X is designed to require authentication before a client is allowed access to a network.

Multicast MAC address 0100.0CDD.DDDD is used by Cisco Group Management Protocol (CGMP). CGMP is used between routers and switches to manage IP multicast traffic at the Data Link layer.

Multicast MAC address 0180.C200.000E is used by Link Layer Discovery Protocol (LLDP). LLDP is a Layer 2 openstandard discovery protocol that is used to facilitate interoperability between Cisco devices and nonCisco devices. Attributes that can be learned from neighboring devices contain type, length, and value (TLV) information including port description, system description, system name, and management address.

Multicast MAC address 0100.0CCC.CCCC is used by Cisco Discovery Protocol (CDP), Dynamic Trunking Protocol (DTP), Port Aggregation Protocol (PAgP), UniDirectional LinkDetection (UDLD), and VLAN Trunking Protocol (VTP). CDP is a Layer 2 Cisco proprietary protocol that is used to advertise and discover only directly connected Cisco devices on a local network. DTP is a pointto point protocol that is used to negotiate trunking. PAgP is an EtherChannel aggregation protocol. UDLD monitors a link to verify that both ends of the link are functioning. VTP is used to centrally manage virtual LAN (VLAN) changes and to propagate those changes over trunk ports.

Multicast MAC address 0180.C200.0000 is used by 802.1D Spanning Tree Protocol (STP) to send native VLAN bridge protocol data units (BPDUs). Multicast MAC address 0100.0CCC.CCCD is used by 802.1D STP to send nonnative VLAN BPDUs.

Reference: <https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/24063-pvid-inconsistency-24063.html#topic1> https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/01xo/configuration/guide/config/cntl_pln.html#wp1151092

QUESTION 230

Which of the following packet types can be captured by EPC on egress? (Select the best answer.)

- A. only unicast packets
- B. only unicast and broadcast packets
- C. only broadcast and multicast packets
- D. only unicast and multicast packets
- E. unicast, broadcast, and multicast packets
- F. no packets

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Only unicast and broadcast packets can be captured by Embedded Packet Capture (EPC) on egress; unicast and broadcast packets can also be captured by EPC on ingress. Multicast packets can be captured by EPC, but only on ingress, not on egress.

Cisco IOS EPC is a feature that you can implement to assist with tracing packets and troubleshooting issues with packet flow in and out of Cisco devices. To implement EPC, you must perform the following steps:

1. Create a capture buffer.
2. Create a capture point.
3. Associate the capture point with the capture buffer.
4. Enable the capture point.

To create a capture buffer, you should issue the `monitor capture bufferbuffername[clear | exportexportlocation | filteraccesslistipaccesslist | limit {allownthpacket | duration seconds | packetcounttotalpackets | packetsper sec packets} | [maxsizeelements] [sizebuffersize] [circular | linear]]` command from global configuration mode. The capture buffer contains packet data and metadata. The packet data does not contain a timestamp indicating when the packet was added to the buffer; the timestamp is contained within the metadata. In addition, the metadata contains information regarding the direction of transmission of the packet, the switch path, and the encapsulation type.

To create a capture point, you should issue the `monitor capture point {ip | ipv6} {cefcapturepointname interfacename interfacetype {both | in | out} | processswitched capturepointname {both | fromus | in | out}}` command from global configuration mode. You can create multiple capture points with unique names and parameters on a single interface; however, you can associate each capture point with only one capture buffer.

To associate a capture point with a capture buffer, you should issue the `monitor capture point associatecapturepointname capturebuffername` command from global configuration mode. Each capture point can be associated with only one capture buffer.

Finally, to enable the capture point so that it can begin to capture packet data, you should issue the `monitor capture point start {capturepointname | all}` command.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/epc/configuration/15-mt/epc-15-mt-book/nm-packet-capture.html#GUID-CF6AFCF6-34DC-4B16-9406E73C806FBD2E>

QUESTION 231

You want to establish an EtherChannel between SwitchA and SwitchB using an IEEE standardsbased protocol. Which of the following channelgroup modes could you configure on the switches? (Select 2 choices.)

- A. SwitchA set to on and SwitchB set to on
- B. SwitchA set to passive and SwitchB set to active
- C. SwitchA set to active and SwitchB set to active

- D. SwitchA set to desirable and SwitchB set to auto
- E. SwitchA set to desirable and SwitchB set to desirable

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You could set one switch to active and the other switch to passive? alternatively, you could set both switches to active. Link Aggregation Control Protocol (LACP) is an Institute of Electrical and Electronics Engineers (IEEE) standardsbased protocol that groups individualphysical ports into a single logical link, called an EtherChannel. The ports that constitute an EtherChannel are grouped according to various parameters, such as hardware, port, and administrative limitations. Because LACP is a standardsbased protocol, it can be used between Cisco and nonCisco switches.

Port Aggregation Protocol (PAgP) is an older, Ciscoproprietary alternative to LACP. Because PAgP is a Ciscoproprietary protocol, it can be used only on Cisco switches. LikeLACP, PAgP identifies neighboring ports and their group capabilities? however, PAgP does not assign roles to the EtherChannel's endpoints like LACP does.

The following table displays the channelgroup configurations that will establish an EtherChannel:

SwitchA \ SwitchB	off	auto	desirable	passive	active	on
off	NO	NO	NO	NO	NO	NO
auto	NO	NO	PAgP	NO	NO	NO
desirable	NO	PAgP	PAgP	NO	NO	NO
passive	NO	NO	NO	NO	LACP	NO
active	NO	NO	NO	LACP	LACP	NO
on	NO	NO	NO	NO	NO	ON

The channelgroup command configures the EtherChannel mode. The syntax of the channelgroup command is channelgroup numbermode {on | active | passive | {auto | desirable} [nonsilent]}, where number is the port channel interface number. The on keyword configures the channel group to unconditionally create the channel with no LACP or PAgP negotiation. The active and passive keywords can be used only with LACP. The active keyword configures the channel group to actively negotiate LACP, and the passive keyword configures the channel group to listen for LACP negotiation to be offered. Either or both sides of the link must be set to active to establish an EtherChannel over LACP? setting both sides to passive will not establish an EtherChannel over LACP.

The auto, desirable, and nonsilent keywords can be used only with PAgP. The desirable keyword configures the channel group to actively negotiate PAgP, and the autokeyword configures the channel group to listen for PAgP negotiation to be offered. Either or both sides of the link must be set to desirable to establish an EtherChannel over PAgP? setting both sides to auto will not establish an EtherChannel over PAgP. The optional nonsilent keyword requires that a port receive PAgP packets before the port is added to the channel.

Reference: https://www.cisco.com/c/en/us/support/docs/switches/catalyst-4000-series-switches/23408-140.html#lACP_page
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750/software/release/12-2_52_se/command/reference/3750cr/cli1.html#wp11890010
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750/software/release/12-2_52_se/command/reference/3750cr/cli1.html#wp11890203

QUESTION 232

You issue the following commands on a switch:

```
SwitchA#configure terminal
SwitchA(config)#interface fastethernet 0/7
SwitchA(config-if)#switchport mode trunk
SwitchA(config-if)#switchport trunk encapsulation dot1q
SwitchA(config-if)#switchport trunk native vlan 44
SwitchA(config-if)#switchport trunk allowed vlan
remove 55
```

Which of the following statements is true regarding VLAN traffic when it is sent over port Fa0/7? (Select the best answer.)

- A. VLAN 1 traffic will be untagged.
- B. VLAN 44 traffic will be untagged.
- C. VLAN 55 traffic will be untagged.
- D. All VLAN traffic will be untagged.
- E. All VLAN traffic will be tagged.



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Traffic from virtual LAN (VLAN) 44 will be untagged when it is sent over port Fa0/7. VLAN 44 traffic is untagged because it has been configured as the native VLAN by the switchport trunk native vlan 44 command. By default, the native VLAN is VLAN 1. You can issue the switchport trunk native vlan vlanid command to change the native VLAN.

All VLAN traffic will not be untagged when it is sent over port Fa0/7; only native VLAN traffic will be untagged. Traffic from all other VLANs will be tagged when it is sent over port Fa0/7. To ensure that traffic from the native VLAN is tagged, you can issue the switchport trunk native vlan tag command. Tagging native VLAN traffic is necessary to enable Layer 2 Quality of Service (QoS) support on the native VLAN.

Traffic from VLAN 1 will not be untagged when it is sent over port Fa0/7? it will be tagged because VLAN 1 is no longer the native VLAN. To reconfigure VLAN 1 to be the native VLAN, you can issue the switchport trunk native vlan 1 command or the noswitchport trunk native vlan command.

Traffic from VLAN 55 cannot be sent over port Fa0/7. The switchport trunk allowed vlan remove 55 command removes VLAN 55 from the list of allowed VLANs that can be trunked over port Fa0/7. The switchport trunk allowed vlan {add | all | except | remove} vlanlist command is issued from interface configuration mode to manually prune VLANs. Manual pruning enables an administrator to strictly specify which VLANs are allowed or denied on a trunk port.

You can issue the show interfaces trunk command to display the list of ports that are configured for trunking, the native VLAN for each trunk port, and the list of currently allowed VLANs for each trunk port. The following displays the output of the show interfaces trunk command:

```
SwitchA#show interfaces trunk
Port    Mode          Encapsulation  Status    Native vlan
Fa0/7   desirable    802.1q         trunking  44

Port    Vlans allowed on trunk
Fa0/7   1-54,56-4094

Port    Vlans allowed and active in management domain
Fa0/7   1-10,44

Port    Vlans in spanning tree forwarding state and not pruned
Fa0/7   1-10,44
```



Reference:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/8021q/8758-43.html> <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/command/ir-cr-book/ir-s7.html#wp3698486980>

QUESTION 233

Which of the following commands configures PIM-SM interfaces to use dense mode to flood Auto-RP traffic to 224.0.1.39 and 224.0.1.40? (Select the best answer.)

- A. ip pim sparse-dense-mode
- B. ip pim send-rp-discovery
- C. ip pim send-rp-announce
- D. ip pim autorp listener

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The ip pim autorp listener command configures Protocol Independent Multicast sparse mode (PIMSM) interfaces to use dense mode to flood AutoRP traffic to 224.0.1.39 and 224.0.1.40. AutoRP candidate rendezvous points (RPs) use multicast address 224.0.1.39 to transmit RPAnnounce messages, which advertise that a router is eligible to become an RP.

AutoRP mapping agents use multicast address 224.0.1.40 to transmit RPDiscovery messages, which advertise the authoritative RP for a multicast group.

The ip pim sendrpdiscovery command does not configure PIMSM interfaces to use dense mode to flood AutoRP traffic to 224.0.1.39 and 224.0.1.40? it configures a router as an AutoRP mapping agent. AutoRP dynamically determines the RP for a multicast group so that RPs do not have to be manually configured. AutoRP uses a mapping agent to learn which routers are advertised as candidate RPs for each multicast group. The candidate list is then advertised to client routers.

The ip pim sendrpannounce command does not configure PIMSM interfaces to use dense mode to flood AutoRP traffic to 224.0.1.39 and 224.0.1.40? it configures a router as an AutoRP candidate RP. A candidate RP advertises itself to the mapping agent, and the mapping agent maps the candidate RPs to multicast groups. If multiple routers are advertised as candidate RPs for a multicast group, the router with the highest IP address is used as the RP for that group.

The ip pim sparsedensemode command does not configure PIMSM interfaces to use dense mode to flood AutoRP traffic to 224.0.1.39 and 224.0.1.40? it configures a PIM router to operate in PIM sparsedense mode (PIMSDM). PIMSDM uses PIMSM for groups that have an RP configured and PIM dense mode (PIMDM) for groups that do not have an RP configured.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/ipmulti/command/reference/fiprnc_r/1rfmult2.html#wp1090395

QUESTION 234

A switch receives a BPDU with the TC bit set.

By default, how long will it take for the switch to age out the MAC address table? (Select the best answer.)

- A. two seconds
- B. 15 seconds
- C. 20 seconds
- D. 35 seconds
- E. 300 seconds

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The switch will age out the Media Access Control (MAC) address table in 15 seconds by default. After the switch receives a bridge protocol data unit (BPDU) with the topology change (TC) bit set, the switch will reduce the default aging time of 300 seconds to the forward_delay value, which is 15 seconds by default.

When a switch needs to signal that a topology change has occurred, it will send topology change notification (TCN) BPDUs on its root port every two seconds, which is the default hello_time value. The designated bridge will forward the TCN BPDU to its root port? additionally, it will send a topology change acknowledgment (TCA) BPDU back to the switch that sent the TCN. This process will continue until the root bridge receives the TCN.

When the root bridge receives the TCN, it will send BPDUs with the TC bit set. By default, the root bridge will set the TC bit for 35 seconds, which is the default max_age timer of 20 seconds plus the default forward_delay value of 15 seconds. The TC BPDUs will be propagated throughout the spanningtree topology.

Reference:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/12013-17.html#anc6>

QUESTION 235

You issue the following commands on RouterA:

```
RouterA(config)#ipv6 router rip boson
```

```
RouterA(config-rtr)#distribute-list prefix-list bosonrip in FastEthernet0/0
```

Which of the following statements best describes what will occur? (Select the best answer.)

- A. IPv6 routing updates matching bosonrip will not be accepted.
- B. IPv6 routing updates matching bosonrip and arriving on the FastEthernet0/0 interface will be accepted.
- C. IPv6 routing updates matching bosonrip and destined for the FastEthernet0/0 interface will be advertised.
- D. IPv6 routing updates matching bosonrip and arriving on the FastEthernet0/0 interface will not be advertised.
- E. IPv6 routing updates not matching bosonrip will be advertised.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

IPv6 routing updates matching bosonrip and arriving on the FastEthernet0/0 interface will be accepted if you issue the following commands on RouterA:

```
RouterA(config)#ipv6 router rip boson
```

```
RouterA(config-rtr)#distribute-list prefix-list bosonrip in FastEthernet0/0
```

The distributelist prefixlist command configures the Routing Information Protocol for IPv6 (RIPv6) process in this scenario to match IPv6 prefixes arriving on the FastEthernet0/0 interface to the IPv6 prefixes that are defined in the bosonrip prefix list. If the prefixes match, the route is accepted. If the prefixes do not match, the route is notaccepted. The distributelist prefixlist command can also be used with Enhanced Interior Gateway Routing Protocol (EIGRP) for IPv6.

IPv6 routing updates matching bosonrip and destined for the FastEthernet0/0 interfacewill not be advertised. To configure the RIPv6 process to match IPv6 prefixes that are destined for the FastEthernet0/0 interface, you should issue the distributelist prefixlistcommand with the out keyword. For example, in this

scenario, you would issue the `distribute list prefixlist bosonrip out FastEthernet0/0` command to match IPv6 prefixes in the bosonrip list that are destined for the FastEthernet0/0 interface.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/ipv6/command/reference/ipv6_book/ipv6_04.html#wp2510121

QUESTION 236

Which of the following flags in the output of the `show ip mroute` command indicates that a receiver is directly connected to the network segment that is connected to the interface? (Select the best answer.)

- A. A
- B. C
- C. D
- D. L
- E. S

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The C flag in the output of the `show ip mroute` command indicates that a receiver is directly connected to the network segment that is connected to the interface. You can view the IP multicast routing table by issuing the `show ip mroute` command, as shown in the following output:



```
RouterA#show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report, Z - Multicast Tunnel,
       Y - Joined MDT-data group, y - Sending to MDT-data group
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode
(*, 224.0.0.3), uptime 5:29:15, RP is 192.168.99.5, flags: SC
  Incoming interface: Tunnel0, RPF neighbor 192.168.50.1, Dvmrp
  Outgoing interface list:
    FastEthernet0, Forward/Sparse, 3:00:10/0:01:50
(192.168.55.0/24, 224.0.0.3), uptime 3:00:10, expires 0:01:50, flags: C
  Incoming interface: Tunnel0, RPF neighbor 192.168.50.1
  Outgoing interface list:
    FastEthernet0, Forward/Sparse, 3:00:10/0:01:50, A
```

The A flag in the flags field would indicate that the router is a candidate for Multicast Source Discovery Protocol (MSDP) advertisement. However, if the A flag is specified in the outgoing interface list, as shown in the previous output, the router is the winner of an assert mechanism and therefore becomes the forwarder. The D flag would indicate that the router is using dense mode. The L flag would indicate that the local router is a member of the multicast group. The S flag would indicate that the router is using sparse mode.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/command/imc-cr-book/imc_s1.html#wp3622554730

QUESTION 237

Which of the following is true regarding the structure of a VPN ID? (Select the best answer.)

- A. It begins with a 4-byte VPN index and ends with a 6-byte MAC address.
- B. It begins with an 8-byte RD and ends with a 4-byte IPv4 address.
- C. It begins with a 4-byte IPv4 address and ends with a 3-byte OUI.
- D. It begins with a 3-byte OUI and ends with a 4-byte VPN index.
- E. It begins with a 6-byte MAC address and ends with a 4-byte IPv4 address.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A virtual private network (VPN) ID begins with a 3byte Organizationally Unique Identifier(OUI) and ends with a 4byte VPN index. The VPN ID identifies a VPN routing and forwarding (VRF). To update a VPN ID for a VRF, issue the vpn id oui: vpn-index command from VRF configuration mode.

Although a Media Access Control (MAC) address contains an OUI, a VPN ID does not contain a MAC address. A VPN ID also does not contain a route distinguisher (RD) or an

IPv4 address. However, a multiprotocol Border Gateway Protocol (BGP) VPNIPv4 address begins with an 8byte RD and ends with a 4byte IPv4 address.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2b/12_2b4/feature/guide/12b_vpn.html

QUESTION 238

Which of the following statements best describes the purpose of holddown timers? (Select the best answer.)

- A. Holddown timers are used by RIP to specify the amount of time to wait before deleting a route from the routing table.
- B. Holddown timers are used by OSPF to specify the amount of time between sending hello packets.
- C. Holddown timers are used by OSPF to specify the amount of time to wait before declaring a neighbor to be down.
- D. Holddown timers are used by RIP to specify the amount of time to suppress information regarding a better path to a route.
- E. Holddown timers are used by RIP to specify the amount of time to wait between broadcasting routing table updates.
- F. Holddown timers are used by RIP to specify the amount of time to wait before declaring a route to be unreachable.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Holddown timers are used by Routing Information Protocol (RIP) to specify the amount of time to suppress information regarding a better path to a route. When a router receives a routing update stating that a route is unreachable, the router waits a specified amount of time before accepting routes advertised by other sources. By default, the holddown timer is set to 180 seconds. RIP uses holddown timers and split horizon to prevent routing loops. Split horizon prevents routers from sending routing information to the same interface from which it was received.

RIP uses four different network timers: update, invalid, holddown, and flush. The update timer is used to specify the amount of time to wait between broadcasting routing table updates. By default, the update timer is set to 30 seconds. The invalid timer is used to specify the amount of time to wait before declaring a route to be unreachable. By default, the invalid timer is set to 180 seconds, and it should always be set to at least three times the value of the update timer. The flush timer is used to specify the amount of time to wait before deleting a route from the routing table. By default, the flush timer is set to 240 seconds, and it should always be

set to a value greater than the invalid timer. To manually configure the four RIP network timers, you should issue the `timers basic update invalid hold down flush` command in RIP router configuration mode, where `update`, `invalid`, `hold down`, and `flush` are specified in seconds.

The Open Shortest Path First (OSPF) hello timer is used to specify the amount of time between sending hello packets. Hello packets are used for neighbor discovery and maintaining neighbor relationships. By default, the hello timer is set to 10 seconds on point-to-point and broadcast links and to 30 seconds on nonbroadcast multiaccess (NBMA) links. The OSPF dead timer is used to specify the amount of time to wait before declaring a neighbor to be down. By default, the dead timer is set to four times the hello timer value. In order for OSPF to work correctly, the hello and dead timers should be consistent across all OSPF routers. To manually configure the hello timer interval, you should issue the `ip ospf hello interval seconds` command in OSPF interface configuration mode. To manually configure the dead timer interval, you should issue the `ip ospf dead interval seconds` command in OSPF interface configuration mode.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/iproute/command/reference/fiprrp_r/1rfrip.html#wp1018019 <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html>

QUESTION 239

You issue the `snmpserver host auth version 3 dot1x` command on a Cisco switch. No previous SNMP commands have been issued on the switch. Which of the following statements are true? (Select 2 choices.)

- A. Notifications are sent as informs.
- B. Notifications are sent as traps.
- C. Only 802.1X notifications are sent.
- D. Only authentication notifications are sent.
- E. The `noAuthNoPriv` security level is applied to the host.
- F. The `AuthNoPriv` security level is applied to the host.



Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Notifications are sent as traps, and the `noAuthNoPriv` security level is applied to the host. Simple Network Management Protocol (SNMP) is used to remotely monitor and manage network devices.

The basic syntax of the `snmpserver host` command is `snmpserver host {hostname | ipaddress} [vrfvrfname | informs | traps | version {1 | 2c | 3 [auth | noauth | priv]]] communitystring [udpportport [notificationtype] | notificationtype]`. Therefore, the command `snmpserver host auth version 3 dot1x` specifies that the switch should send SNMP version 3 (SNMPv3) notifications to a device with the hostname `auth` using the community string `dot1x`.

By default, notifications are sent as traps. You can also explicitly specify that notifications be sent as traps by issuing the `traps` keyword in the `snmpserver host` command. To send notifications as informs, you should issue the `informs` keyword in the `snmpserver host` command.

There are three SNMPv3 security levels: noAuthNoPriv, authNoPriv, and authPriv. If no security level is specified in the snmpserver host command, the noAuthNoPriv security level is used. The noAuthNoPriv security level, which is also enabled by issuing the noauthkeyword, authenticates by matching the user name in clear text. The authNoPriv security level, which is enabled by issuing the auth keyword, matches an unencrypted Message Digest 5 (MD5) or Secure Hash Algorithm (SHA) hash of the user name in order to authenticate. The authPriv security level, which is enabled by issuing the priv keyword, authenticates by matching an MD5 or SHA hash of the user name that is also encrypted by using either Data Encryption Standard (DES), Triple DES (3DES), or Advanced Encryption Standard (AES). In this scenario, the placement of the auth keyword configures a host name, not a security level. If a notification type is not specified in the snmpserver host command, all notificationtypes are sent. Although the dot1x keyword can be used to specify that 802.1X notifications are sent, the placement of the dot1x keyword in this scenario configures a community string, not a notification type.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipswitch/command/isw-cr-book/isw-s3.html#wp3482803949>

QUESTION 240

Which of the following commands should you issue from interface configuration mode to associate an NHRP group with a QoS policy map? (Select the best answer.)

- A. ip nhrp group
- B. ip nhrp map
- C. ip nhrp map group
- D. ip nhrp map multicast
- E. ip nhrp map multicast dynamic
- F. ip nhrp responder



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the ip nhrp map group command from interface configuration mode to associate a Next Hop Resolution Protocol (NHRP) group with a Quality of Service (QoS) policy map. NHRP is used to create a database of tunnel addresses to real addresses. When a spoke router wants to send a packet to another spoke router by using an ondemand spoketospoke tunnel, the sending router queries the NHRP database to determine the receiving router's dynamic spoke address; the sending router then creates the ondemand tunnel between the spoke routers.

Before you can map an NHRP group with a QoS policy map, you must first create the NHRPgroup by issuing the ip nhrp group command from interface configuration mode. The following command set creates an NHRP group named boson and maps the group to a QoS policy map named exsim:

```
Router(config)#interface Tunnel 0
```

```
Router(config)#ip nhrp group boson
```

```
Router(config)#ip nhrp map group boson servicepolicy output exsim
```

The ip nhrp map command configures spoke routers with a static mapping that maps the hub router's tunnel IP address to the hub router's physical IP address. The syntax of the ip nhrp map command is ip nhrp map ipaddress nbmaaddress command, where ip address is the hub router's tunnel IP address and nbmaaddress is the hub router's physical IP address. Hub routers need not be configured with the ip nhrp map ipaddress nbmaaddress command; as spoke routers register with the hub, the mappings are dynamically created.

The ip nhrp map multicast dynamic command configures a hub router to allow spoke routers to register with the hub as multicast receivers. Spoke routers should not be configured with the ip nhrp map multicast dynamic command? instead, they should be configured with the ip nhrp map multicastipaddress command, where ipaddress is the physical IP address of the hub router. The ip nhrp map multicast command enables the spoke router to send broadcast and multicast packets over the tunnel.

The ip nhrp responder command specifies the IP address that the next hop server should use when replying to Responder Address queries. The syntax of the ip nhrp responder command is ip nhrp responder interface type interface number. The primary IP address of the interface is the IP address that the next hop server will use in the NHRP reply.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/command/ipaddr-cr-book/ipaddr-i4.html#wp2318545993>

QUESTION 241

Which of the following is the VLAN name for VLAN 1004? (Select the best answer.)

- A. fddi-default
- B. fddinet-default
- C. token-ring-default
- D. trnet-default

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The virtual LAN (VLAN) name for VLAN 1004 is fddinetdefault. VLANs 1002 through 1005 cannot be pruned, deleted, or used to send data over Ethernet. VLANs 1002 and 1004 are reserved for Fiber Distributed Data Interface (FDDI). VLANs 1003 and 1005 are reserved for Token Ring. The following table displays VLANs 1002 through 1005 along with their corresponding names:

VLAN	Name
1002	fddi-default
1003	token-ring-default
1004	fddinet-default
1005	trnet-default

You can see these names when issuing the show vlan command:

```
SwitchA#show vlan
VLAN Name                Status    Ports
-----
1    default                active    Fa0/2, Fa0/4, Fa0/5, Fa0/6
                                           Fa0/7, Fa0/9, Fa0/10, Fa0/11
                                           Fa0/12, Gi0/1, Gi0/2
2    VLAN0002              active
10   VLAN0010              active    Fa0/3
20   VLAN0020              active

1002 fddi-default          active
1003 token-ring-default  active
1004 fddinet-default     active
1005 trnet-default       active
```



Reference:

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

QUESTION 242

Which of the following will occur when you issue the ntp server 10.11.12.13 command on RouterA? (Select 2 choices.)

- A. RouterA will become an NTP server for a host at 10.11.12.13.
- B. RouterA will become an NTP client of a server at 10.11.12.13.
- C. RouterA will listen for NTP packets on the interface that is configured to use IP address 10.11.12.13.
- D. RouterA will use IP address 10.11.12.13 as the source IP address for all NTP packets leaving the interface.
- E. RouterA will use NTP version 3.
- F. RouterA will be set to stratum 8.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RouterA will become a Network Time Protocol (NTP) client of a server at 10.11.12.13 and will use NTP version 3. NTP is used to synchronize the time on network devices. Issuing the `ntp server` command from global configuration mode configures a Cisco router to operate as an NTP static client. An NTP static client receives its time from an NTP server.

The basic syntax of the `ntp server` command is `ntp server ipaddress [version number] [prefer]`, where `ipaddress` is the IP address of the NTP server that the client will use to receive its time. The optional `version` keyword can be issued to specify the NTP version? if no version is specified, NTP version 3 will be used by default. The optional `prefer` keyword can be used so that the specified NTP server will be preferred by the client over other NTP servers? if the `prefer` keyword is not used, the client will synchronize with the server that has the lowest stratum number.

RouterA will not become an NTP server for a host at 10.11.12.13. You should issue the `ntp master` command from global configuration mode to configure a Cisco router to operate as an NTP server. The syntax of the `ntp master` command is `ntp master [stratum]`, where `stratum` is an NTP stratum value from 1 through 15. By default, an NTP server is configured to use a stratum value of 8. Devices with higher stratum numbers receive time from devices with lower stratum numbers. For example, a stratum 2 device typically receives its time from a stratum 1 device, a stratum 3 device typically receives its time from a stratum 2 device, and so on.

RouterA will not be set to stratum 8 unless it receives its time from an NTP server at stratum 7. You cannot manually set the stratum for an NTP client. To configure RouterA so that it is set to stratum 8, you must issue the `ntp master` command or the `ntp master 8` command.

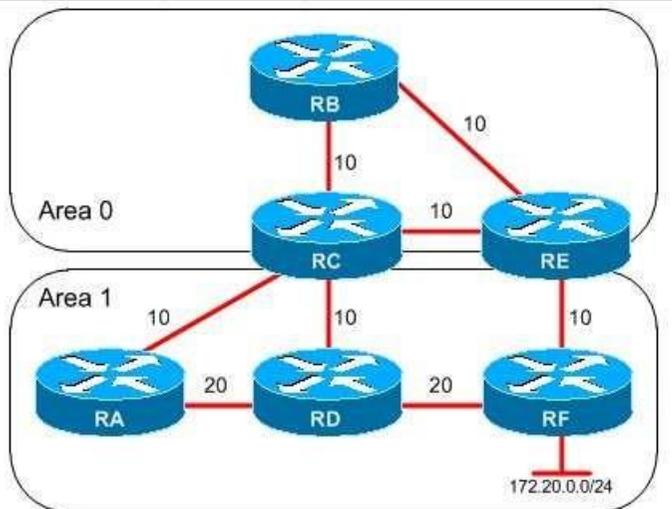
RouterA will not listen for NTP packets on the interface that is configured to use IP address 10.11.12.13, nor will RouterA use IP address 10.11.12.13 as the source IP address for all NTP packets leaving the interface. IP address 10.11.12.13 is the address of the NTP server, not an IP address on RouterA. Issuing the `ntp broadcast client` command from interface configuration mode configures a Cisco router to operate as an NTP broadcast client. An NTP broadcast client listens on the configured interface for NTP broadcasts from an NTP server, which the NTP client uses to adjust its time. The difference between a broadcast client and a static client is that a broadcast client can receive its time from any NTP server. By contrast, a static client receives its time from the NTP server specified in the `ntp server` command.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/configfun/command/reference/ffun_r/frf012.html#wp1123899

Cisco: Basic System Management Commands: `ntp server`

QUESTION 243



You administer the OSPF network shown above. The cost values are displayed next to each link. RA receives packets destined for the 172.20.0.0/24 network. Which path or paths will RA use to send the packets? (Select the best answer.)

- A. only R-ARD-RF
- B. only RA-RC-RE-RF
- C. only RA-RD-RF and RA-RC-RD-RF
- D. only RA-RD-RF, RA-RC-RD-RF, and RA-RC-RB-RE-RF

Correct Answer: C

Section: (none)

Explanation

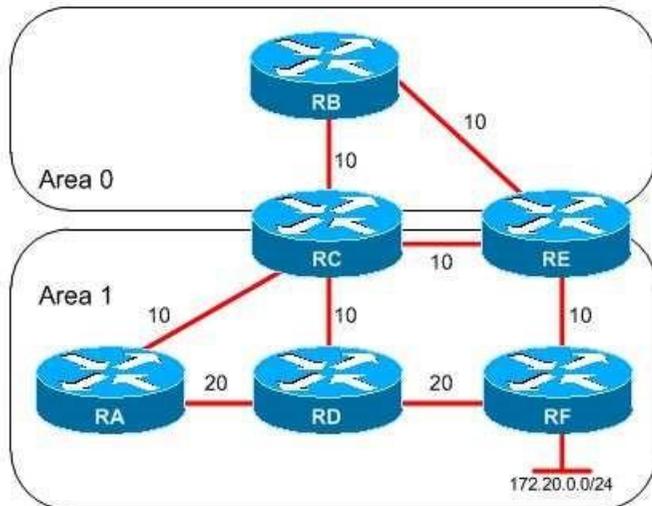
Explanation/Reference:

Explanation:

RA will use only the RARDRF path and the RARCRDRF path to send the packets to the 172.20.0.0/24 network. The total cost from RA through RD to RF is $20 + 20 = 40$, and the total cost from RA through RC and RD to RF is $10 + 10 + 20 = 40$. Open Shortest Path First (OSPF) can load balance traffic across equalcost paths? since both paths have a totalcost of 40, RA can use both paths to send the packets. RA will not always prefer the route through the least number of routers. Instead, RA prefers the intraarea route with the lowest total cost, regardless of the number of routers the packets must pass through. RA will not use the RARCREF path to send the packets. Although the total cost from RA through RC and RE to RF is $10 + 10 + 10 = 30$, OSPF prefers intraarea routes over interarea routes, regardless of the total path cost. OSPF uses the following preference order when selecting the best route to a destination:

1. Intraarea routes
2. Interarea routes
3. External Type 1 routes
4. External Type 2 routes

Therefore, RA prefers an intraarea route with a cost of 40 over an interarea route with a cost of 30. If the link between RC and RE were within Area 1, as shown in the following graphic, RA would prefer the route from RA through RC and RE to RF:



RA will not use the RA-RC-RB-RE-RF path to send the packets. Although the total cost of this path is $10 + 10 + 10 + 10 = 40$, the RARCRBRERF route is an interarea route; RA prefers the intraarea routes with a cost of 40.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html#7>

QUESTION 244

Which of the following command sets correctly configures and applies a MAC ACL? (Select the best answer.)

- A. Router(config)#accesslist 800 permit 0000.0c01.abcd
Router(config)#interface gigabitethernet 1/0
Router(config)#mac accessgroup 800 in
- B. Router(config)#accesslist 800 permit 0000.0c01.abcd
Router(config)#interface gigabitethernet 1/0

- Router(config)#mac accessgroup 800 out
- C. Router(config)#accesslist 800 permit 0000.0c01.abcd
Router(config)#interface gigabitethernet 1/0.1
Router(config)#mac accessgroup 800 in
- D. Router(config)#accesslist 700 permit 0000.0c01.abcd
Router(config)#interface gigabitethernet 1/0.1
Router(config)#mac accessgroup 700 out
- E. Router(config)#accesslist 700 deny 0000.0c01.abcd
Router(config)#accesslist 700 permit any
Router(config)#interface gigabitethernet 1/0.1
Router(config)#mac accessgroup 700 in
- F. Router(config)#accesslist 700 deny 0000.0c01.abcd
Router(config)#accesslist 700 permit any
Router(config)#interface gigabitethernet 1/0.1
Router(config)#mac accessgroup 700 out

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following command set correctly configures and applies a Media Access Control (MAC) access control list (ACL):

```
Router(config)#accesslist 700 deny 0000.0c01.abcd
Router(config)#accesslist 700 permit any
Router(config)#interface gigabitethernet 1/0.1
Router(config)#mac accessgroup 700 in
```

A MAC ACL filters inbound Ethernet packets based on the source MAC address. Because a MAC ACL filters at Layer 2, both IP and nonIP packets can be filtered. MAC ACLs support Ethernet, 802.1Q virtual LAN (VLAN), and 802.1QinQ packets.

To create a MAC ACL, you should issue the `accesslist accesslistnumber {deny | permit} {macaddress | any}` command. The accesslist number variable must be a number from 700 through 799. MAC ACLs that include the any keyword are automatically moved to the end of the ACL. Like normal ACLs, MAC ACLs also have an implicit deny any statement that is applied to any routes that have not been explicitly permitted or denied by previous accesslist statements.

You can then apply the MAC ACL to an interface or a subinterface by issuing the `mac accessgroup accesslistnumber in` command from interface or subinterface configuration mode.

The command sets that configure the accesslistnumber variable in the accesslistcommand to a value of 800 cannot be used to create a MAC ACL. Access lists numbered from 800 through 899 are used for Internetwork Packet Exchange (IPX) ACLs.

The command sets that attempt to configure the out keyword in the mac accessgroupcommand cannot be used to create a MAC ACL. MAC ACLs can be configured only for inbound filtering.

Reference:

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

QUESTION 245

Which of the following provides antireplay protection for GET VPN group members? (Select the best answer.)

- A. KEK
- B. SAR
- C. TEK
- D. TSK

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Synchronous Antireplay (SAR) provides antireplay protection for Group Encrypted Transport (GET) virtual private network (VPN) group members. GET VPN is a connectionless, nontunneling VPN technology based on the Group Domain of Interpretation (GDOI) standard proposed in Request for Comments (RFC) 3547. Nontunneling VPNs such as GET VPN can be used on a variety of networks, including IP, Frame Relay, Multiprotocol Label Switching (MPLS), and Asynchronous Transfer Mode (ATM) networks. Although GET VPN does not use tunneling, it does rely upon Internet Key Exchange (IKE) and IP Security (IPSec) security associations (SAs).

GE VPN requires a key server, which is responsible for maintaining the policy, creating and maintaining group keys, and servicing registration requests. The key server prevents replay attacks by maintaining a pseudotime clock to keep track of time. Group members regularly synchronize to the pseudotime on the key server. If an intercepted message is replayed, the replayed message will likely fall outside the pseudotime window. A group member will detect the pseudotime discrepancy and will therefore reject the replayed message.

A traffic encryption key (TEK) is used to encrypt data between GET VPN group members. When a group member registers with the key server, the group member downloads the IPSec policy and encryption keys from the key server. If a group member fails to register with a key server, all traffic is sent unencrypted through the group member unless the Fail Close feature is activated.

A key encryption key (KEK) is used to encrypt data between the key server and group members. Periodically, the key server will send rekey messages to group members in order to refresh the IPSec SA before it expires. The KEK protects the rekey message, which contains new encryption keys that the group members should use, thereby securing the control plane.

A transmission security key (TSK) is used by direct sequence spread spectrum (DSSS) or frequency hopping radios. TSKs are not used by GET VPN group members.

Reference:

https://www.cisco.com/en/US/docs/ios-xml/ios/sec_conn_getvpn/configuration/15-2mt/sec-get-vpn.html#GUID-AA25D5AD-BDD5-499E-AA2F-4F4F113C57D2

QUESTION 246

Which of the following commands configures the CEF loadbalancing algorithm to use only a source, a destination, and an ID hash? (Select the best answer.)

- A. ip cef load-sharing algorithm universal
- B. ip cef load-sharing algorithm original
- C. ip cef load-sharing algorithm include-ports source destination
- D. ip cef load-sharing algorithm include-ports source destination gtp

Correct Answer: A

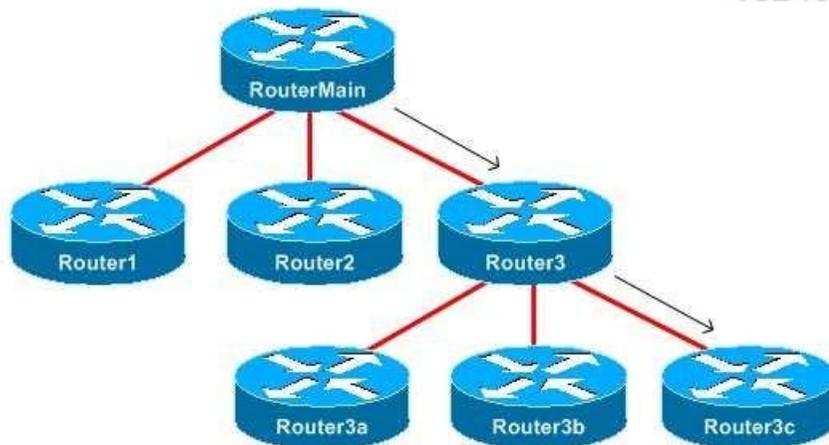
Section: (none)

Explanation

Explanation/Reference:

Explanation:

The ip cef loadsharing algorithm universal command configures the Cisco Express Forwarding (CEF) loadbalancing algorithm to use only a source, a destination, and an IDhash. The original CEF loadbalancing algorithm is prone to CEF polarization, which occurs when multiple routers in sequence use the same loadbalancing mechanism. To understand CE polarization, consider the following topology:



RouterMain will run the loadbalancing algorithm on a flow and, based on the hash result, will send the flow to Router1, Router2, or Router3. If Router1, Router2, and Router3 run the same loadbalancing algorithm as RouterMain uses, those routers will get the same hash result and will therefore no longer load balance. For

example, flows that are sent from RouterMain to Router3 will always be forwarded to Router3c because Router3 generates the same hash for each flow that RouterMain does. The ip cef loadsharing algorithm original command configures a router to use the original CEF loadbalancing algorithm, which uses only a source and destination hash and is prone to CEF polarization.

Universal mode is an improvement to CEF that causes each router to use a 32bit Universal ID as a hashing seed. Because each router uses a different Universal ID, each router will produce different hashing values, thereby avoiding CEF polarization by enabling each router to load balance differently. Universal mode is enabled by default or by issuing the ip cef loadsharing algorithm universal command.

The ip cef loadsharing algorithm includeports source destination command configures CEF to not only use the universal loadbalancing algorithm but also to consider Layer 4 source and destination port information. Because this command uses the Universal ID it also avoids CEF polarization. The ip cef loadsharing algorithm includeports source destination gtp command is similar to the ip cef loadsharing algorithm includeports source destination command. However, the ip cef loadsharing algorithm includeports source destination gtp command also considers the GPRS Tunneling Protocol (GTP) Tunnel Endpoint Identifier (TEID), when applicable.

Reference: <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipswitch/command/isw-cr-book/isw-i1.html#wp5609710740>

CCIE Routing and Switching v5.0 Certification Guide, Volume 1, Chapter 6, Load Sharing with CEF and Related Issues, pp. 282-285

QUESTION 247

Which of the following is a point-to-point protocol that can be configured on a switch to dynamically control the establishment of 802.1Q and ISL trunk links? (Select the best answer.)

A. DTP



<https://vceplus.com/>

B. LACP

C. PAgP

D. STPE

E. VTP

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Dynamic Trunking Protocol (DTP) is a Cisco proprietary point-to-point protocol that can be configured on a switch to dynamically control the establishment of 802.1Q and InterSwitch Link (ISL) trunk links. ISL is Cisco's proprietary encapsulation method; it can be used to configure trunks on Cisco switches only. The Institute of Electrical and Electronics Engineers (IEEE) 802.1Q standard is widely supported and should be used when you must incorporate third-party switches into an existing Cisco topology.

There are five DTP switch port modes: trunk, desirable, auto, access, and nonegotiate. The switchport mode {trunk | dynamic {desirable | auto} | access | nonegotiate} command configures the DTP mode for a switch port. Cisco recommends that you set both sides of a trunk link to desirable mode when using DTP. When connecting to a non-Cisco device, you should manually configure the switch port for trunk mode. The following table indicates whether a trunk link will be established between a pair of switches configured with the switchport mode command:

SwitchA \ SwitchB	access	nonegotiate	dynamic auto	dynamic desirable	trunk
access	NO	NO	NO	NO	NO
nonegotiate	NO	YES	NO	NO	YES
dynamic auto	NO	NO	NO	YES	YES
dynamic desirable	NO	NO	YES	YES	YES
trunk	NO	YES	YES	YES	YES

Port Aggregation Protocol (PAgP) is a Cisco proprietary protocol that groups individual physical PAgP-configured ports into a single logical link, called an EtherChannel. The ports that constitute an EtherChannel are grouped according to various parameters, such as hardware, port, and administrative limitations. Once PAgP has created an EtherChannel, it adds the EtherChannel to the spanning tree as a single switch port.

Link Aggregation Control Protocol (LACP) is a newer, standards-based alternative to PAgP that is defined by the IEEE 802.3ad standard. LACP is available on switches newer than the Catalyst 2950, which only offers PAgP. Like PAgP, LACP identifies neighboring ports and their group capabilities; however, LACP goes further by assigning roles to the EtherChannel's endpoints.

Spanning Tree Protocol (STP) is defined in the IEEE 802.1D standard. Layer 2 protocols use STP to determine the best path through a switched network. STP prevents switching loops on a network. Switching loops can occur when there is more than one switched path to a destination. The spanning tree algorithm determines the best path through a switched network, and any ports that create redundant paths are blocked. If the best path becomes unavailable, the network topology is recalculated and the port connected to the next best path is unblocked.

VLAN Trunking Protocol (VTP) is used to synchronize VTP and virtual LAN (VLAN) configuration information between switches. Changes that are made on one VTP server are propagated throughout the VTP domain. For switches to synchronize information over VTP, the following configuration parameters must match on all switches:

- VTP domain name
- VTP password
- VTP version

Reference:

<http://www.gratisexam.com/>

<https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/31sga/configuration/guide/config/layer2.html#wp1020498>
https://www.cisco.com/c/en/us/support/docs/switches/catalyst-4000-series-switches/23408-140.html#lcp_page

QUESTION 248

In which of the following situations will an EIGRP router transition a route from the passive state to the active state and send multicast query packets to its neighbors? (Select the best answer.)

- A. when the successor becomes unreachable and a feasible successor exists
- B. when the successor becomes unreachable and no feasible successors exist
- C. when the router sees its own router ID in the Neighbor field of an update packet
- D. when all the routes in the router's topology table are in the passive state
- E. when the DR fails and no BDR exists on the multiaccess segment

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference:

When the successor becomes unreachable and no feasible successors exist, an Enhanced Interior Gateway Routing Protocol (EIGRP) router transitions a route from the passive state to the active state and sends multicast query packets to its neighbors by using the multicast IP address 224.0.0.10. These multicast query packets are used to interrogate EIGRP neighbors to determine whether they have a route to a destination network. The route will remain in the active state until replies are received for each of the neighbor queries. You can display which routers have not yet replied to a query by issuing the `show ip eigrp topology active` command, as shown in the following output:

```
RouterA#show ip eigrp topology active
IP-EIGRP Topology Table for AS(1)/ID(192.168.99.2)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
A: 192.168.99.3/32, 1 successors, FD is Inaccessible, tag is 1
   1 replies, active 00:05:15, query-origin: Local origin
     via Redistributed (2297856/0)
     Remaining replies:
       via 192.168.99.3, r, Serial0/0.223
```

If feasible successors exist when the current successor becomes unreachable, the EIGRP router immediately uses a feasible successor. Because route recomputation is not required, the route will stay in the passive state.

EIGRP routers maintain a neighbor table, a topology table, and a routing table. The neighbor table lists all the adjacent routers configured to run EIGRP. The topology table lists the next hop for all the network destinations known to all the EIGRP neighbors that represent a loopfree path to that destination. The routing table lists only the best route to each destination network.

EIGRP calculates a feasible distance (FD) for each neighbor. The FD represents the cost of using that particular next hop to reach the destination. The neighbor with the lowest FD to a destination network is known as the successor, and the route to that successor is stored in the routing table. If a second neighbor can reach the same destination, but at a higher FD it is listed in the topology table as a feasible successor and is stored in the topology table as a backup in the event that the successor fails. In the topology table, EIGRP will list all the successors and feasible successors that can reach a given destination.

Routes are always in one of two states: passive or active. Routes are in the passive state when no route recomputation is necessary. Route recomputation occurs when the network topology changes due to a link failure or recovery. The routes will remain in the passive state as long as at least one feasible successor exists. When no feasible successors exist for a destination, the route enters the active state and EIGRP initiates a route recomputation by using the 224.0.0.10 multicast address to send query packets to all known EIGRP neighbors.

Routers running EIGRP will not check for their own router ID in the Neighbor field of an update packet. Only routers running Open Shortest Path First (OSPF) perform this action. Before entering a twoway state with an OSPF neighbor, a router running OSPF must see its own router ID in the Neighbor field of an update packet that it receives from that neighbor.

Routers running EIGRP do not elect designated routers (DRs) or backup designated routers (BDRs). Only routers running OSPF elect DRs and BDRs. OSPF elects one router to be the DR and another router to be the BDR in each multiaccess segment. The DR serves as a single point of contact for all OSPF routers on the multiaccess segment, and the BDR exists in case the DR fails. All OSPF routers on the segment exchange updates only with the DR and BDR.

Reference:

http://docwiki.cisco.com/wiki/Enhanced_Interior_Gateway_Routing_Protocol

QUESTION 249

You issue the `no service tcpssmallservers` command on a Cisco router.

Which of the following servers are disabled by the command? (Select 3 choices.)

- A. BOOTP
- B. chargen
- C. discard
- D. echo
- E. finger

Correct Answer: BCD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The discard, echo, and chargen servers are disabled by issuing the `no service tcpssmallservers` command. The discard server, which uses Transmission Control Protocol (TCP) port 9, discards any data that is sent to it. The echo server, which uses TCP port 7, receives data and echoes that data back to the sender. The chargen server, which uses TCP port 19, generates a stream of ASCII data back to the sender. These servers are used to test TCP functionality; however, they can be exploited. Therefore, it is recommended that these servers be disabled. The TCP small servers are disabled by default on Cisco IOS version 11.3 and later. The finger server is not disabled by issuing the `no service tcpssmallservers` command. The finger server, which uses TCP port 79, displays user information. By default, the finger server is disabled. If the finger server has been enabled, you can disable it by issuing the `no ip finger` command or the `no service finger` command.

The Bootstrap Protocol (BOOTP) server is not disabled by issuing the `no service tcpssmallservers` command. The BOOTP server, which uses User Datagram Protocol (UDP) port 67, is a predecessor of Dynamic Host Configuration Protocol (DHCP) and is used to assign IP addresses to client devices. By default, the BOOTP server is enabled. To disable the BOOTP server, you should issue the `no ip bootp server` command.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fundamentals/command/cf_command_ref/R_through_setup.html#wp2590195365
<https://www.cisco.com/image/gif/paws/12815/23.pdf> Cisco: TCP and UDP Small Servers (PDF)

QUESTION 250

DRAG DROP

Select the characteristics from the left, and place them underneath the corresponding line type on the right. Fill all boxes. Characteristics can be used more than once; some characteristics might not be used.

To complete this question, click Launch Simulator and follow the onscreen instructions.

Select and Place:

Select and Place:

	EPL	EVPL
uses a point-to-point EVC between two NNIs		
uses a point-to-point EVC between two UNIs		
provides full transparency		
does not provide full transparency		
allows for service multiplexing		
does not allow for service multiplexing		
is an E-Line service		
is an E-LAN service		

 **VCEplus**
VCE To PDF - Free Practice Exam

Correct Answer:

	EPL	EVPL
uses a point-to-point EVC between two NNIs		
uses a point-to-point EVC between two UNIs	uses a point-to-point EVC between two UNIs	uses a point-to-point EVC between two UNIs
provides full transparency	provides full transparency	does not provide full transparency
does not provide full transparency	does not allow for service multiplexing	allows for service multiplexing
allows for service multiplexing	is an E-Line service	is an E-Line service
does not allow for service multiplexing		
is an E-Line service		
is an E-LAN service		


VCE To PDF - Free Practice Exam

Section: (none)

Explanation

Explanation/Reference:

Explanation:

An Ethernet private line (EPL) is a Metro Ethernet ELine service that uses a point-to-point Ethernet virtual connection (EVC) between two User Network Interfaces (UNIs). An EVC associates two or more UNIs. A UNI is the demarcation point at which the service provider's responsibility ends and the customer's responsibility begins. Bandwidth profiles can be established per EVC or per UNI.

An EPL provides full transparency such that Layer 2 protocols are the same at the source and destination UNIs. However, an EPL does not allow for service multiplexing; only one EP is supported at the UNI. Instead, an EPL allows all-to-one bundling. Generally, if a UNI is configured for service multiplexing, all-to-one bundling must be disabled, and conversely, if a UNI is configured for all-to-one bundling, service multiplexing must be disabled.

<http://www.gratisexam.com/>

An Ethernet virtual private line (EVPL) is also a Metro Ethernet ELine service that uses a point-to-point EVC between two UNIs. Unlike an EPL, an EVPL does not provide full transparency, because Layer 2 control protocols are discarded at the UNI. However, an EVPL allows for service multiplexing so that more than one EVC is supported at the UNI.

Neither an EPL nor an EVPL use EVCs between Network to Network Interfaces (NNIs). Like UNIs, NNIs are demarcation points; however, NNIs are demarcation points between service provider networks.

Neither an EPL nor an EVPL are ELAN services. Whereas ELine services are point-to-point Ethernet services, ELAN services are multipoint-to-multipoint Ethernet services.

Reference:

https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/fulfillment/6-1/theory/operations/guide/theory/l2ce.html

QUESTION 251

DRAG DROP

Select the terms that are displayed in the output from the show ip eigrp neighbors command from the left, and drag them to the corresponding definitions on the right.

Select and Place:

Correct Answer:

Section: (none)
Explanation

Explanation/Reference:

Explanation:

Enhanced Interior Gateway Routing Protocol (EIGRP) maintains three tables for each Layer3 protocol: -

Neighbor table

-Topology table

-Routing table

The EIGRP neighbor table lists the directly connected neighbors that have an established

EIGRP adjacency with the router, as shown in the following output from the show ip eigrp neighbors command:

```
RouterA#show ip eigrp neighbors
IP-EIGRP neighbors for process 1
H   Address                Interface      Hold Uptime   SRTT  RTO  Q  Seq
   (sec)                    (ms)          (sec)         (ms)  Cnt  Num
0   192.168.1.2              Fa0/0         12 00:00:11   26    200  0   3
```

In the output above, the EIGRP process on RouterA has established an adjacency, or neighbor relationship, with another EIGRP router that has been assigned the IP address of 192.168.1.2. The SRTT column value, which is 26, indicates the smooth roundtrip time (SRTT) in milliseconds. The SRTT indicates how long it takes for an EIGRP packet to be sent and for an acknowledgment to be returned. The RTO column value, which is 200, indicates the retransmit interval in milliseconds. The retransmit interval, or retransmission timeout (RTO), is the amount of time an EIGRP router will wait before attempting to resend a packet that has been stored in the retransmission queue. The retransmission queue contains packets that a router needs to resend to a neighboring router. Related to the retransmit interval is the queue (Q) count, which is displayed in the Q Cnt column. The Q count is the number of EIGRP reply, query, and update packets that are waiting to be sent. The Q count in the output above is 0, which indicates that there are no packets that are waiting to be sent.

Reference:

https://www.cisco.com/c/en/us/td/docs/net_mgmt/ip_solution_center/6-0/infrastructure/reference/guide/infrastructure/iscglss1.pdf

https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/fulfillment/6-1/theory/operations/guide/theory/l2ce.html

QUESTION 252

You issue the mls ip cef loadsharing full simple command.

Which of the following statements is true? (Select the best answer.)

- A. CEF load balancing will use Layer 3 information with multiple adjacencies.
- B. CEF load balancing will use Layer 3 information without multiple adjacencies.
- C. CEF load balancing will use Layer 3 and Layer 4 information with multiple adjacencies.
- D. CEF load balancing will use Layer 3 and Layer 4 information without multiple adjacencies.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Cisco Express Forwarding (CEF) load balancing will use Layer 3 and Layer 4 information without multiple adjacencies. The syntax of the mls ip cef loadsharing command is mls ip cef loadsharing [full] [excludeport {destination | source}] [simple]. When the full keyword is used, CEF load balancing will use Layer 3 and Layer 4 information with multiple adjacencies. When the simple keyword is used, CEF load balancing will use Layer 3 information without multiple adjacencies. When the full and simple keywords are both used, CEF load balancing will use Layer 3 and Layer 4 information without multiple adjacencies.

The excludeport keyword configures CEF to exclude either source or destination Layer 4 ports from the load balancing algorithm. In addition, the excludeport keyword configures

CE to exclude both source and destination IP addresses from the load balancing algorithm, regardless of whether the source or destination keywords are used.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipswitch/command/isw-cr-book/isw-i1.html#wp4169023325>

QUESTION 253

In which of the following scenarios will a BGP route be advertised when conditional advertisements are configured? (Select 2 choices.)

- A. when a prefix appears in the advertise map and in the exist map
- B. when a prefix appears in the advertise map but not in the exist map
- C. when a prefix appears in the advertise map and in the nonexist map
- D. when a prefix appears in the advertise map but not in the nonexist map
- E. when a prefix does not appear in the advertise map but does appear in the exist map
- F. when a prefix appears in neither the advertise map nor the exist map
- G. when a prefix does not appear in the advertise map but does appear in the nonexist map
- H. when a prefix appears in neither the advertise map nor the nonexist map

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When conditional advertisements are configured, a Border Gateway Protocol (BGP) route will be advertised in either of these scenarios:

-When a prefix appears in the advertise map and in the exist map

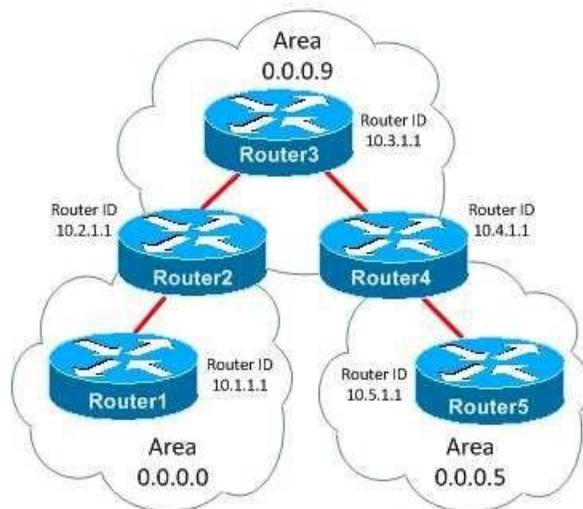
-When a prefix appears in the advertise map but not in the nonexistent map

A conditional advertisement is a route that is withheld from a neighbor until a condition is met. The prefixes are contained within two route maps: an advertise map and either an exist map or a nonexistent map. The advertise map indicates the prefixes that will be advertised when the condition is met, provided that the prefix exists in the BGP routing table. If an exist map is used, the prefix must appear within the advertise map and the exist map in order for the router to advertise the prefix. If a nonexistent map is used, the prefix must appear within the advertise map but not within the nonexistent map in order for the router to advertise the prefix. Conditional advertisements are created by issuing the neighbor advertise-map command from BGP router configuration mode. The syntax of the neighbor advertise-map command is neighbor ip address advertise-map mapname {exist-map mapname | non-exist-map mapname}. For example, the neighbor 192.168.1.1 advertise-map BOSON1 exist-map EXIST1 command creates a conditional advertisement where a prefix will be advertised to a neighbor at 192.168.1.1 if a prefix exists in BOSON1 as well as the exist map EXIST1. Conversely, the neighbor 192.168.1.1 advertise-map BOSON1 non-exist-map NONEXIST1 command creates a conditional advertisement where a prefix will be advertised to a neighbor at 192.168.1.1 if the prefix appears in BOSON1 but does not appear in the non-exist map NONEXIST1.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp3.html#wp1105432

QUESTION 254



Which of the following commands should you issue to create a virtual link between Area 0.0.0.0 and Area 0.0.0.5? (Select 2 choices.)

- A. Router1(configrouter)#area 0.0.0.5 virtuellink 10.5.1.1
- B. Router2(configrouter)#area 0.0.0.9 virtuellink 10.4.1.1
- C. Router2(configrouter)#area 0.0.0.0 virtuellink 10.3.1.1
- D. Router4(configrouter)#area 0.0.0.5 virtuellink 10.3.1.1
- E. Router4(configrouter)#area 0.0.0.9 virtuellink 10.2.1.1

F. Router5(configrouter)#area 0.0.0.0 virtuallink 10.1.1.1

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the area 0.0.0.9 virtuallink 10.4.1.1 command on Router2 and the area 0.0.0.9 virtuallink 10.2.1.1 command on Router4 to create a virtual link between Area 0.0.0.0 and Area 0.0.0.5. To create a virtual link, you should issue the area areaidvirtuallinkrouterid command in router configuration mode on each area border router (ABR), where areaid is the transit area ID, and routerid is the router ID of the router at the other end of the virtual link.

All areas in an Open Shortest Path First (OSPF) internetwork must be connected to the backbone area, Area 0.0.0.0, also known as simply Area 0. A virtual link must be created between two ABRs to connect a remote area to the backbone area through a transit area. The following restrictions apply to virtual links:

- The routers at each end of the virtual link must share a common area.
- The transit area cannot be a stub area.
- The transit area cannot be the backbone area.
- One router must connect to the backbone area.

In this scenario, Router2 and Router4 are ABRs that share a common area, Area 0.0.0.9. Router2 connects to the backbone area, and Router4 connects to the remote area. Because Area 0.0.0.9 will be used as the transit area for the virtual link, Area 0.0.0.9 cannot be configured as a stub area.

You cannot create a virtual link between Router1 and Router5, because they do not share a common area. Therefore, you should not issue the area 0.0.0.5 virtuallink 10.5.1.1 command on Router1 or the area 0.0.0.0 virtuallink 10.1.1.1 command on Router5.

You should not issue the area 0.0.0.0 virtuallink 10.3.1.1 command on Router2 or the area 0.0.0.5 virtuallink 10.3.1.1 command on Router4. The areaid parameter should be the transit area ID, not the backbone or remote area ID. Additionally, you should establish the virtual link directly between two routers, not between two routers and an intermediate router.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html>

QUESTION 255

Which of the following commands creates a single-rate, dual-bucket, three-color policer? (Select the best answer.)

- A. police 100000 5000 8000 conform-action transmit exceed-action set-dscp-transmit af21
- B. police 100000 5000 8000 conform-action transmit exceed-action set-dscp-transmit af21 violate-action drop
- C. police cir 50000 bc 50000 pir 100000 conform-action transmit exceed action se-tdscp-transmit af21
- D. police cir 50000 pir 100000 conform-action transmit exceed-action set dscp-transmit af21 violate-actiondrop

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Explanation:

The police 100000 5000 8000 conformaction transmit exceedaction setdscp transmit af21 violateaction drop command creates a singlerate, dualbucket, threecolor policer. Traffic policing is used to control the amount of traffic sent on an interface. The police command transmits, reclassifies, or drops traffic based on how much traffic is being sent.

The syntax of the singlerate police command is policebps [burstnormal] [burstmax] conformaction actionexceedaction action [violateaction action], where the optional burstnormal and burstmax parameters are specified in bytes. In the police 100000 5000 8000 conformaction transmit exceedaction setdscptransmit af21 violateaction drop command, the average rate is 100,000 bps, the normal burst size is 5,000 bytes, and the maximum burst size is 8,000 bytes. The conformactionkeyword specifies what happens to packets that conform to the bps rate? in this example, traffic up to 100,000 bps is transmitted. The exceedaction keyword specifies what happens to traffic that exceeds the rate limit but does not exceed the maximum burst size? in this example, traffic that exceeds the rate limit up to 8,000 bytes is transmitted with a

Differentiated Services Code Point (DSCP) value of AF21. The optional violateactionkeyword specifies what happens to packets that exceed the maximum burst size? in this example, packets that exceed the maximum burst size are dropped.

Issuing the singlerate police command with the conformaction, exceedaction, andviolateaction keywords creates a dualbucket, threecolor policer. Conforming traffic is considered green traffic, bursting traffic is considered yellow traffic, and traffic that violates the policy is considered red traffic. Traffic flowing into the first bucket is green traffic. When the first bucket is full, traffic flows into the second bucket. Traffic flowing into the second bucket is yellow traffic. When the second bucket is full, red traffic overflows the second bucket.

Issuing the singlerate police command with the conformaction and exceedactionkeywords but without the violateaction keyword creates a singlebucket, twocolor policer. Therefore, the police 100000 5000 8000 conformaction transmit exceedaction setdscptransmit af21 command creates a singlerate, singlebucket, twocolor policer. You can also issue the police command with two traffic rates: a committed information rate (CIR) and a peak information rate (PIR). The syntax of the dualrate police command is policecir [bc conformburst] [pirpir] [bepeakburst] [conformaction action[exceedaction action [violateaction action]]], where cir and pir are specified in bps. Therefore, the police cir 50000 bc 50000 pir 100000 conformaction transmit exceedaction setdscptransmit af21 command and the police cir 50000 pir 100000 conformaction transmit exceedaction setdscptransmit af21 violateaction drop command create a dualrate policer, not a singlerate policer. A dualrate policer always uses a dualbucket policer regardless of the number of colorspecified. Therefore, the police cir 50000 bc 50000 pir 100000 conformaction transmit exceedaction setdscptransmit af21 command creates a dualrate, dualbucket, twocolor policer, and the police cir 50000 pir 100000 conformaction transmit exceedaction setdscptransmit af21 violateaction drop command creates a dualrate, dualbucket, threecolor policer.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/qos/command/reference/qos_book/qos_n1.html#wp1060117

https://www.cisco.com/c/en/us/td/docs/ios/qos/command/reference/qos_book/qos_n1.html#wp1047550

<https://www.cisco.com/c/en/us/td/docs/routers/10000/10008/configuration/guides/qos/qoscf/10qpolce.html#wp1041572>

QUESTION 256

DRAG DROP

Select the pseudowire FEC element fields from the left, and place them on the corresponding descriptions on the right.

Select and Place:

C-Bit	is a 32-bit arbitrary value
Group ID	is also known as the VC ID
Interface Parameters	is always set to a value of 128
Pseudowire ID	indicates whether control words are used
Pseudowire ID FEC	is set to 0x0005 for Ethernet
Pseudowire Information Length	is set to 0 if the Pseudowire ID and Interface Parameters are not present
Pseudowire Type	is a variable-length field

VCEplus
VCE To PDF - Free Practice Exam

Help Reset Done

Correct Answer:

	Group ID
	Pseudowire ID
	Pseudowire ID FEC
	C-Bit
	Pseudowire Type
	Pseudowire Information Length
	Interface Parameters



Section: (none)
Explanation

Explanation/Reference:

Explanation:

A pseudowire FEC element consists of the following fields:

Pseudowire ID FEC

CBit

Pseudowire Type

Pseudowire Information Length

Group ID

Pseudowire ID

Interface Parameters

The Pseudowire ID Forwarding Equivalence Class (FEC) is an 8bit field that is always set to a value of 128. This value indicates that the packet is a pseudowire FEC element.

The CBit, or Control Word Bit, indicates whether a 4bit control word will be present in every pseudowire packet. If the CBit is set to a value of 1, the control word will be placed between the Multiprotocol Label Switching (MPLS) label stack and the Layer 2 payload.

The Pseudowire Type is a 15bit field that indicates the type of pseudowire. This field will be set to a value of 0x0005 if the pseudowire is an Ethernet pseudowire.

The Pseudowire Information Length is an 8bit field that indicates the octet length of the Pseudowire ID field and the Interface Parameters field. If the Pseudowire Information Length field is set to a value of 0, the Pseudowire ID and Interface Parameters fields are not present? the pseudowire FEC element applies to all pseudowires using the specified Group ID. The Group ID is a 32bit arbitrary value that represents a group of pseudowires. The Pseudowire ID field is a 32bit specific value that represents a particular pseudowire.

The Pseudowire ID is a 32bit value that identifies a particular pseudowire. Both endpoints must be configured with the same pseudowire type and ID.

The Interface Parameters field is the only variable length field. This field provides circuitspecific information, such as the maximum transmission unit (MTU) for the interface.

Reference:

<https://tools.ietf.org/html/rfc4447#section-5.2>

<http://www.ciscopress.com/articles/article.asp?p=386788&amp;amp;amp;seqNum=2>

QUESTION 257

Which of the following statements best describes the BGP split horizon rule? (Select the best answer.)

- A. Routes learned through eBGP are not advertised to iBGP peers.
- B. Routes learned through iBGP are not advertised to eBGP peers.
- C. Routes learned through iBGP are not advertised to iBGP peers.
- D. Routes learned through eBGP are not advertised to eBGP peers.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Border Gateway Protocol (BGP) split horizon rule states that routes learned through internal BGP (iBGP) are not advertised to iBGP peers, which are BGP routers that exist within the same autonomous system (AS). An iBGP peer advertises the following routes to another iBGP peer: -Routes learned through external BGP (eBGP)

-Routes learned through redistribution

-Routes originated by a network statement

Because iBGP routes are not advertised to iBGP peers, one of the following actions must be taken to enable routers running iBGP to communicate: -

Configure a full mesh.

-Configure a confederation.

-Configure a route reflector.

A full mesh configuration enables each router to learn each iBGP route independently without passing through a neighbor. However, a full mesh configuration requires the most administrative effort to configure. A confederation enables an AS to be divided into discrete units, each of which acts like a separate AS. Within each confederation, the routers must be fully meshed unless a route reflector is established. A route reflector can be used to pass iBGP routes between iBGP routers, eliminating the need for a full mesh configuration.

However, it is important to note that route reflectors advertise best paths only to routereflector clients. Additionally, if multiple paths exist, a route reflector will always advertise the exit point that is closest to the route reflector. eBGP peers are BGP routers that belong to different ASes. An eBGP peer advertises the following routes to another eBGP peer:

- Routes learned through iBGP
- Routes learned through eBGP
- Routes learned through redistribution -
- Routes originated by a network statement

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-m1.html#wp9015889000

https://www.cisco.com/en/US/docs/net_mgmt/ip_manager/1.0/user/guide/app_integrity_checks.html#wp4405

QUESTION 258

Which of the following values in an MPLS label indicates that this label is the last label in the stack? (Select the best answer.)

- A. a value of 0 in the TTL field
- B. a value of 255 in the TTL field
- C. a value of 0 in the TC field
- D. a value of 1 in the TC field
- E. a value of 0 in the S field
- F. a value of 1 in the S field

Correct Answer: F

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A value of 1 in the S field in a Multiprotocol Label Switching (MPLS) label, which is also known as a MPLS header, indicates that this label is the last label in the stack. The structure of a typical 4byte MPLS label is shown below:

MPLS Label



The S field, sometimes referred to as the Stack bit or the BottomofStack field, is a 1bit field that indicates whether the label is the last MPLS label in a packet. An S field set to 0 indicates that one or more MPLS labels follow this label. An S field set to 1 indicates that this label is the last label in the stack.

The TimeToLive (TTL) field is not used to indicate whether a label is the last label in the stack. Similar to an IP TTL field, the MPLS TTL field is an 8bit field that is used to control the propagation of packets through an MPLS network. When an IP packet enters an MPLS network, the ingress router decrements the IP TTL value by 1 and copies that value to the MPLS TTL field. Each MPLS router along the path decrements the MPLS TTL field by 1. When the packet reaches the egress router, the MPLS TTL value is decremented by 1 and copied to the IP TTL field. A TTL field set to 0 indicates that the packet should be discarded. If MPLS TTL propagation is disabled, the MPLS TTL field is set to 255 and decrements as the packet passes through the MPLS network; when the packet reaches the egress router, the MPLS TTL value is not copied to the IP TTL field.

The Traffic Class (TC) field is not used to indicate whether an MPLS label is the last label in the stack. Cisco routers use the 3bit TC field to carry the IP precedence value, which is used to classify and prioritize network traffic. The TC field was formerly designated as the Experimental (EXP) field in Request for Comments (RFC) 3032. However, RFC 3032 did not officially designate the use of the EXP field, so some nonCisco routers use this field for other purposes. RFC 5462 officially renames the EXP field as the TC field and designates it to carry traffic class information, such as IP precedence values. A TC field set to 0 indicates a packet that contains lowpriority traffic, and a TC field set to 1 indicates a packet with a slightly higher priority than a packet with a TC field set to 0. Highpriority traffic would have a TC field set to 7.

Reference: <https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/4649-mpls-faq-4649.html> https://www.cisco.com/en/US/tech/tk828/technologies_q_and_a_item09186a00800a43f5.shtml
<https://tools.ietf.org/html/rfc5462>

QUESTION 259

Which of the following OSPF areas does not accept Type 3, 4, and 5 summary LSAs? (Select the best answer.)

- A. stub area
- B. ordinary area
- C. backbone area
- D. notsostubby area
- E. totally stubby area

Correct Answer: E

Section: (none)
Explanation

Explanation/Reference:

Explanation:

An Open Shortest Path First (OSPF) totally stubby area does not accept Type 3, 4, and 5 summary linkstate advertisements (LSAs), which advertise routes outside the area. These LSAs are replaced by a default route at the area border router (ABR). As a result, routing tables are kept small within the totally stubby area. To create a totally stubby area, you should issue the `area areaid stub nosummary` command in router configuration mode.

The backbone area, Area 0, accepts all LSAs. All OSPF areas must directly connect to the backbone area or must traverse a virtual link to the backbone area. To configure a router to be part of the backbone area, you should issue the `area 0` command in router configuration mode.

An ordinary area, which is also called a standard area, accepts all LSAs. Every router in an ordinary area contains the same OSPF routing database. To configure an ordinary area, you should issue the `area areaid` command in router configuration mode.

A stub area does not accept Type 5 LSAs, which advertise external summary routes. Routers inside the stub area will send all packets destined for another area to the ABR. To configure a stub area, you should issue the `area areaid stub` command in router configuration mode.

A notso stubby area (NSSA) is basically a stub area that contains one or more autonomous system boundary routers (ASBRs). Like stub areas, NSSAs do not accept Type 5 LSAs.

External routes from the ASBR are converted to Type 7 LSAs and tunneled through the NSSA to the ABR, where they are converted back to Type 5 LSAs. To configure an NSSA, you should issue the `area areaid nssa` command in router configuration mode. To configure a totally NSSA, which does not accept summary routes, you should issue the `area areaid nssa nosummary` command in router configuration mode.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html>

QUESTION 260

For which of the following reasons would a multicast host send a packet to 235.77.34.2? (Select 2 choices.)

- A. to join a multicast group at 235.77.34.2
- B. to leave a multicast group at 235.77.34.2
- C. to send a general query
- D. to reply to a general query

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

<http://www.gratisexam.com/>

A multicast host would send a packet to 235.77.34.2 to join a multicast group at 235.77.34.2 and to reply to either a general membership query or a group-specific membership query. Internet Group Management Protocol (IGMP) has three message types: a membership report message, a membership query message, and a leave group message. When a host wants to join a multicast group, it sends an IGMP membership report message to that multicast group IP address. Packets for that multicast group are then sent on that network segment so that the host can receive the multicast traffic. When a host receives a membership query, it will send a membership report message to the multicast groups from which the host wants to receive traffic.

A multicast host would not send a packet to 235.77.34.2 to send a general query. The querier router on a network segment sends out general query messages to the 224.0.0.1 all hosts multicast address to determine whether any hosts on that network segment want to continue to receive multicast packets for any multicast group. If at least one host responds with a membership report message, the querier will continue to send those multicast packets on that network segment. If no host responds to three consecutive membership query messages, the router will stop forwarding the multicast traffic on that network segment. When IGMPv2 is used, the Max Response Time field in membership query messages contains a nonzero value. In IGMPv1 messages, the field is set to a value of 0, which is interpreted to mean 100 deciseconds, or 10 seconds. The IGMPv2 membership query message is the only message that contains a nonzero value in the Max Response Time field; all other message types set the field to a value of 0.

A multicast host would not send a packet to 235.77.34.2 to leave a multicast group at 235.77.34.2. When IGMP version 1 (IGMPv1) is used, hosts leave a multicast group without sending any notification. When IGMPv2 is used, hosts send a leave group message to 224.0.0.2 when leaving a multicast group; the 224.0.0.2 multicast address is used to send a message to all multicast-capable routers. When a multicast router receives a leave group message, the router will send a group-specific membership query to the multicast group to determine whether there are any hosts on the segment that want to continue to receive the multicast traffic from that group. If at least one host responds with a membership report message, the querier will continue to send those multicast packets on that network segment.

Reference:

CCIE Routing and Switching v5.0 Certification Guide, Volume 2, Chapter 7, IGMPv2 Host Membership Query Functions, pp. 285-286

CCIE Routing and Switching v5.0 Certification Guide, Volume 2, Chapter 7, IGMPv2 Leave Group and Group-Specific Query Messages, pp. 289-291



<https://vceplus.com/>

<http://www.gratisexam.com/>

www.vceplus.com - Download A+ VCE (latest) free Open VCE Exams - VCE to PDF Converter - VCE Exam Simulator - VCE Online