

400-101

Number: 400-101
Passing Score: 900
Time Limit: 120 min
File Version: 18.0

400-101

CCIE Routing and Switching (v5.0)

By Marcio from Brazil

Sections

1. Network Principles
2. Layer 2 Technologies
3. Layer 3 Technologies
4. VPN Technologies
5. Infrastructure Security
6. Infrastructure Services
7. Mix Questions

Exam A

QUESTION 1

Refer to the exhibit.

```
#show interface FastEthernet0/0
FastEthernet0/0 is up, line protocol is up
  Hardware is PQII_PRO_UEC, address is 0024.14ac.0d3c (bia 001f.9e3c.a5c2)
  Internet address is 1.1.1.1/24
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full Duplex, 100Mbps, media type is RJ45
  output flow-control is XON, input flow-control is XON
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:00, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/1000/0/0 (size/max/drops/flushes); Total output drops: 10000
  Queueing strategy: Class-based queueing
  Output queue: 100/1000/10000 (size/max total/drops)
  30 second input rate 361000 bits/sec, 204 packets/sec
  30 second output rate 711000000 bits/sec, 223000 packets/sec
    1221583901 packets input, 3044421428 bytes, 0 no buffer
    Received 91124750 broadcasts (0 IP multicasts)
    0 runs, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
    0 watchdog, 0 multicast, 0 pause input
    1090847722 packets output, 796667418 bytes, 0 underruns
    0 output errors, 0 collisions, 0 interface resets
    0 unknown protocol drops
    0 babbles, 0 late collision, 0 deferred
    0 lost carrier, 0 no carrier, 0 pause output
    0 output buffer failures, 0 output buffers swapped out
```

Which two are causes of output queue drops on FastEthernet0/0? (Choose two.)

- A. an oversubscribed input service policy on FastEthernet0/0
- B. a duplex mismatch on FastEthernet0/0

- C. a bad cable connected to FastEthernet0/0
- D. an oversubscribed output service policy on FastEthernet0/0
- E. The router trying to send more than 100 Mb/s out of FastEthernet0/0

Correct Answer: DE

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Output drops are caused by a congested interface. For example, the traffic rate on the outgoing interface cannot accept all packets that should be sent out, or a service policy is applied that is oversubscribed. The ultimate solution to resolve the problem is to increase the line speed. However, there are ways to prevent, decrease, or control output drops when you do not want to increase the line speed. You can prevent output drops only if output drops are a consequence of short bursts of data. If output drops are caused by a constant high-rate flow, you cannot prevent the drops. However, you can control them.

Reference. <http://www.cisco.com/c/en/us/support/docs/routers/10000-series-routers/6343-queue-drops.html>

QUESTION 2

Refer to the exhibit.

```
Router#show ip cache flow
[...]
```

SrcIf	SrcIPaddress	DstIf	DstIPaddress	Pr	SrcP	DstP	Pkts
Vl1	144.254.10.206	Local	10.48.77.208	06	C363	01BB	2

Which statement about the output is true?

- A. The flow is an HTTPS connection to the router, which is initiated by 144.254.10.206.
- B. The flow is an HTTP connection to the router, which is initiated by 144.254.10.206.
- C. The flow is an HTTPS connection that is initiated by the router and that goes to 144.254.10.206.
- D. The flow is an HTTP connection that is initiated by the router and that goes to 144.254.10.206.

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

We can see that the connection is initiated by the Source IP address shown as 144.254.10.206. We also see that the destination protocol (DstP) shows

01BB, which is in hex and translates to 443 in decimal. SSL/HTTPS uses port 443.

QUESTION 3

Refer to the exhibit.

```
R101#show ip cache verbose flow
[...]
SrcIf          SrcIPaddress  DstIf          DstIPaddress   Pr TOS Flgs  Pkts
Port Msk AS    Port Msk AS    NextHop        B/Pk  Active
Et0/0          10.0.0.1       Et1/0*         14.0.0.2       01 80  10      1
0000 /0  0       0800 /0  0       0.0.0.0       100      0.0
```

What is the PHB class on this flow?

- A. EF
- B. none
- C. AF21
- D. CS4

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

This command shows the TOS value in hex, which is 80 in this case. The following chart shows some common DSCP/PHB Class values:

Service	DSCP value	TOS value	Juniper Alias	TOS hexadecimal	DSCP - TOS Binary
Premium IP	46	184	ef	B8	101110 - 101110xx
LBE	8	32	cs1	20	001000 - 001000xx
DWS	32	128	cs4	80	100000 - 100000xx
Network control	48	192	cs6	c0	110000 - 110000xx
Network control 2	56	224	cs7	e0	111000 - 111000xx

Reference. <http://www.tucny.com/Home/dscp-tos>

QUESTION 4

Refer to the exhibit.

```
R101#show ip cache flow
```

[...]							
SrcIf	SrcIPaddress	DstIf	DstIPaddress	Pr	SrcP	DstP	Pkts
Et0/0	10.0.0.1	Et1/0*	14.0.0.2	01	0000	0800	34
Et0/0	10.0.0.1	Et1/0	14.0.0.2	01	0000	0800	100
Et0/0	10.0.0.1	Se3/0*	14.0.0.2	01	0000	0800	33
Et0/0	10.0.0.1	Se2/0*	14.0.0.2	01	0000	0800	33
Et0/0	10.0.0.1	Null	224.0.0.5	59	0000	0000	26

What kind of load balancing is done on this router?

- A. per-packet load balancing
- B. per-flow load balancing
- C. per-label load balancing
- D. star round-robin load balancing

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Here we can see that for the same traffic source/destination pair of 10.0.0.1 to 14.0.0.2 there were a total of 100 packets (shown by second entry without the *) and that the packets were distributed evenly across the three different outgoing interfaces (34, 33, 33 packets, respectively).

QUESTION 5

What is a cause for unicast flooding?

- A. Unicast flooding occurs when multicast traffic arrives on a Layer 2 switch that has directly connected multicast receivers.
- B. When PIM snooping is not enabled, unicast flooding occurs on the switch that interconnects the PIM-enabled routers.
- C. A man-in-the-middle attack can cause the ARP cache of an end host to have the wrong MAC address. Instead of having the MAC address of the default gateway, it has a MAC address of the man-in-the-middle. This causes all traffic to be unicast flooded through the man-in-the-middle, which can then sniff all packets.
- D. Forwarding table overflow prevents new MAC addresses from being learned, and packets destined to those MAC addresses are flooded until space becomes available in the forwarding table.

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Causes of Flooding

The very cause of flooding is that destination MAC address of the packet is not in the L2 forwarding table of the switch. In this case the packet will be flooded out of all forwarding ports in its VLAN (except the port it was received on). Below case studies display most common reasons for destination MAC address not being known to the switch.

Cause 1: Asymmetric Routing

Large amounts of flooded traffic might saturate low-bandwidth links causing network performance issues or complete connectivity outage to devices connected across such low-bandwidth links

Cause 2: Spanning-Tree Protocol Topology Changes

Another common issue caused by flooding is Spanning-Tree Protocol (STP) Topology Change Notification (TCN). TCN is designed to correct forwarding tables after the forwarding topology has changed. This is necessary to avoid a connectivity outage, as after a topology change some destinations previously accessible via particular ports might become accessible via different ports. TCN operates by shortening the forwarding table aging time, such that if the address is not relearned, it will age out and flooding will occur

Cause 3: Forwarding Table Overflow

Another possible cause of flooding can be overflow of the switch forwarding table. In this case, new addresses cannot be learned and packets destined to such addresses are flooded until some space becomes available in the forwarding table. New addresses will then be learned. This is possible but rare, since most modern switches have large enough forwarding tables to accommodate MAC addresses for most designs.

Reference:

<http://www.cisco.com/c/en/us/support/docs/switches/catalyst-6000-series-switches/23563-143.html>

QUESTION 6

Which Cisco IOS XE process administers routing and forwarding?

- A. Forwarding manager
- B. Interface manager
- C. Cisco IOS
- D. Host manager

Correct Answer: C

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Some of the processes are listed in the table below:

Process	Purpose	Affected FRUs	SubPackage Mapping
Host Manager	Provides an interface between the IOS process and many of the information-gathering functions of the underlying platform kernel and operating system.	RP (one instance per RP) SIP (one instance per SIP) ESP (one instance per ESP)	RPControl SIPBase ESPBase
Interface Manager	Provides an interface between the IOS process and the per-SPA interface processes on the SIP.	RP (one instance per RP) SIP (one instance per SIP)	RPControl SIPBase
IOS	The IOS process implements all forwarding and routing features for the router.	RP (one per software redundancy instance per RP). Maximum of two instances per RP.	RPiOS
Forwarding Manager	Manages the downloading of configuration to each of the ESPs and the communication of forwarding plane information, such as statistics, to the IOS process.	RP (one per software redundancy instance per RP). Maximum of two instances per RP. ESP (one per ESP)	RPControl ESPBase

Reference.

http://www.cisco.com/c/en/us/td/docs/routers/asr1000/configuration/guide/chassis/asrswcfg/Software_Packaging_Architecture.html

QUESTION 7

Which circumstance can cause packet loss due to a microburst?

- A. slow convergence
- B. a blocked spanning-tree port
- C. process switching
- D. insufficient buffers

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Micro-bursting is a phenomenon where rapid bursts of data packets are sent in quick succession, leading to periods of full line-rate transmission that can overflow packet buffers of the network stack, both in network endpoints and routers and switches inside the network. Symptoms of micro bursts will manifest in the form of ignores and/ or overruns (also shown as accumulated in "input error" counter within show interface output). This is indicative of

receive ring and corresponding packet buffer being overwhelmed due to data bursts coming in over extremely short period of time (microseconds).

Reference. <http://ccieordie.com/?tag=micro-burst>

QUESTION 8

Which two statements about proxy ARP are true? (Choose two.)

- A. It is supported on networks without ARP.
- B. It allows machines to spoof packets.
- C. It requires larger ARP tables
- D. It reduces the amount of ARP traffic.

Correct Answer: BC

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Disadvantages of Proxy ARP

Hosts have no idea of the physical details of their network and assume it to be a flat network in which they can reach any destination simply by sending an ARP request. But using ARP for everything has disadvantages. These are some of the disadvantages:

- It increases the amount of ARP traffic on your segment.
- Hosts need larger ARP tables in order to handle IP-to-MAC address mappings.
- Security can be undermined. A machine can claim to be another in order to intercept packets, an act called "spoofing."
- It does not work for networks that do not use ARP for address resolution.
- It does not generalize to all network topologies. For example, more than one router that connects two physical networks.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation- resolution/13718-5.html>

QUESTION 9

Which service is disabled by the no service tcp-small-servers command?

- A. the finger service
- B. the Telnet service
- C. the Maintenance Operation Protocol service
- D. the chargen service

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

The TCP small servers are:

- **Echo:** Echoes back whatever you type through the **telnet x.x.x.x echo** command.
- **Chargen:** Generates a stream of ASCII data. Use the **telnet x.x.x.x chargen** command.
- **Discard:** Throws away whatever you type. Use the **telnet x.x.x.x discard** command.
- **Daytime:** Returns system date and time, if it is correct. It is correct if you run Network Time Protocol (NTP), or have set the date and time manually from the exec level. Use the **telnet x.x.x.x daytime** command.

Reference. <http://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/ios-software-releases-110/12815-23.html>

QUESTION 10

Which option is the most effective action to avoid packet loss due to microbursts?

- A. Implement larger buffers.
- B. Install a faster CPU.
- C. Install a faster network interface.
- D. Configure a larger tx-ring size.

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

You can't avoid or prevent them as such without modifying the sending host's application/network stack so it smoothes out the bursts. However, you can manage microbursts by tuning the size of receive buffers / rings to absorb occasional microbursts.

QUESTION 11

Which two statements about packet fragmentation on an IPv6 network are true? (Choose two.)

- A. The fragment header is 64 bits long.
- B. The identification field is 32 bits long.
- C. The fragment header is 32 bits long.
- D. The identification field is 64 bits long.
- E. The MTU must be a minimum of 1280 bytes.
- F. The fragment header is 48 bits long.

Correct Answer: AB

Section: Network Principles

Explanation


Explanation/Reference:

Explanation:

The fragment header is shown below, being 64 bits total with a 32 bit identification field:

Fragment Header

- **Offset:** the offset, in 8-octet units, of the data following this header, relative to the start of the Fragmentable Part of the packet
- **M flag:** 1 – more fragments, 0 – last fragment



25

Reference. <http://www.openwall.com/presentations/IPv6/img24.html>

QUESTION 12

You are backing up a server with a 1 Gbps link and a latency of 2 ms. Which two statements about the backup are true? (Choose two.)

- A. The bandwidth delay product is 2 Mb.
- B. The default TCP send window size is the limiting factor.
- C. The default TCP receive window size is the limiting factor.
- D. The bandwidth delay product is 500 Mb.
- E. The bandwidth delay product is 50 Mb.

Correct Answer: AC

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

1 Gbps is the same as 1000 Mbps, and $1000\text{Mb} \times .0002 = 2\text{ Mbps}$. With TCP based data transfers, the receive window is always the limiting factor, as the sender is generally able to send traffic at line rate, but then must wait for the acknowledgements to send more data.

QUESTION 13

Which two pieces of information does RTCP use to inform endpoint devices about the RTP flow? (Choose two.)

- A. the transmitted octet
- B. the lost packet count
- C. session control function provisioning information
- D. the CNAME for session participants
- E. the authentication method
- F. MTU size changes in the path of the flow

Correct Answer: AB

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

RTCP transports statistics for a media connection and information such as transmitted octet and packet counts, packet loss, packet delay variation, and round-trip delay time. An application may use this information to control quality of service parameters, perhaps by limiting flow, or using a different codec.

Reference. http://en.wikipedia.org/wiki/RTP_Control_Protocol

QUESTION 14

DRAG DROP

Drag and drop the argument of the `ip cef load-sharing algorithm` command on the left to the function it performs on the right.

Select and Place:

Drag and drop the argument of the <code>ip cef load-sharing algorithm</code> command on the left to the function it performs on the right.	
original	sets the load-balancing algorithm to use a source, a destination, and an ID hash
universal	sets the load-balancing algorithm for environments with a small number of source destination IP address pairs
tunnel	sets the load-balancing algorithm to use Layer 4 information
include-ports source destination	sets the load-balancing algorithm to use a source and destination hash

Correct Answer:

Drag and drop the argument of the <code>ip cef load-sharing algorithm</code> command on the left to the function it performs on the right.	
	universal
	tunnel
	include-ports source destination
	original

Section: Network Principles
Explanation

Explanation/Reference:

QUESTION 15

DRAG DROP

Drag and drop the Cisco IOX XE subpackage on the left to the function it performs on the right.

Select and Place:

Drag and drop the Cisco IOX XE subpackage on the left to the function it performs on the right.	
RPIOS	provisions the Cisco IOS Software kernel from which the IOS software features are housed and run
ESPBase	produces the ESP software, ESP operating system, and control processes
SIPBase	manages the Cisco IOS Software and the rest of the platform via the control plane
RPCControl	manages the Session Initiation Protocol carrier card operating system and control processes

Correct Answer:

Drag and drop the Cisco IOX XE subpackage on the left to the function it performs on the right.

	RPIOS
	ESPBase
	RPCControl
	SIPBase

Section: Network Principles

Explanation

Explanation/Reference:

QUESTION 16

DRAG DROP

Drag and drop each description of IPv6 transition technology on the left to the matching IPv6 transition technology category on the right.

Select and Place:

Drag and drop each description of IPv6 transition technology on the left to the matching IPv6 transition technology category on the right.

encapsulates IPv6 packets within IPv4 packets

supports translation between IPv4 and IPv6 by using algorithms to map addresses

supports stateful translation between IPv4 and IPv6 with static and manual mappings

requires IPv6-capable infrastructure

uses routing protocols to maintain IPv4 and IPv6 routing adjacencies

encapsulates IPv4 packets within IPv6 packets

Dual-Stack Network

Tunneling

NAT64

Correct Answer:

Drag and drop each description of IPv6 transition technology on the left to the matching IPv6 transition technology category on the right.

Dual-Stack Network

requires IPv6-capable infrastructure

uses routing protocols to maintain IPv4 and IPv6 routing adjacencies

Tunneling

encapsulates IPv6 packets within IPv4 packets

encapsulates IPv4 packets within IPv6 packets

NAT64

supports translation between IPv4 and IPv6 by using algorithms to map addresses

supports stateful translation between IPv4 and IPv6 with static and manual mappings

Section: Network Principles
Explanation

Explanation/Reference:

QUESTION 17

How many hash buckets does Cisco Express Forwarding use for load balancing?

- A. 8
- B. 16
- C. 24
- D. 32

Correct Answer: B

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

In order to understand how the load balance takes place, you must first see how the tables relate. The Cisco Express Forwarding table points to 16 hash buckets (load share table), which point to the adjacency table for parallel paths. Each packet to be switched is broken up into the source and destination address pair and checked against the loadshare table.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/express-forwarding-cef/18285-loadbal-cef.html>

QUESTION 18

Which three features require Cisco Express Forwarding? (Choose three.)

- A. NBAR
- B. AutoQoS
- C. fragmentation
- D. MPLS
- E. UplinkFast
- F. BackboneFast

Correct Answer: ABD

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

QoS Features That Require CEF

These class-based QoS features are supported only on routers that run CEF.

- Network Based Application Recognition (NBAR) provides intelligent network classification. For more information, refer to Network Based Application Recognition.

- The AutoQoS -VoIP feature simplifies and speeds up the implementation and provisioning of QoS for VoIP traffic. This feature is enabled with the help of the auto qos voip command. CEF must be enabled at the interface or ATM PVC before the auto qos command can be used. For more information about this feature and its prerequisites, refer to AutoQoS - VoIP.

From MPLS Fundamentals - Luc De Ghein

Why Is CEF Needed in MPLS Networks?

Concerning MPLS, CEF is special for a certain reason; otherwise, this book would not explicitly cover it. Labeled packets that enter the router are switched according to the label forwarding information base (LFIB) on the router. IP packets that enter the router are switched according to the CEF table on the router. Regardless of whether the packet is switched according to the LFIB or the CEF table, the outgoing packet can be a labeled packet or an IP packet

Reference. <http://www.cisco.com/c/en/us/support/docs/asynchronous-transfer-mode-atm/ip-to-atm-class-of-service/4800-cefreq.html>

QUESTION 19

Which two options are interface requirements for turbo flooding? (Choose two.)

- A. The interface is Ethernet.
- B. The interface is configured for ARPA encapsulation.
- C. The interface is PPP.
- D. The interface is configured for GRE encapsulation.
- E. The interface is configured for 802.1Q encapsulation.

Correct Answer: AB

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

In the switch, the majority of packets are forwarded in hardware; most packets do not go through the switch CPU. For those packets that do go to the CPU, you can speed up spanning tree-based UDP flooding by a factor of about four to five times by using turbo-flooding. This feature is supported over Ethernet interfaces configured for ARPA encapsulation.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/metro/me3400/software/release/12-2_50_se/configuration/guide/scg/swiprout.html

QUESTION 20

Which option describes a limitation of Embedded Packet Capture?

- A. It can capture data only on physical interfaces and subinterfaces.
- B. It can store only packet data.
- C. It can capture multicast packets only on ingress.

D. It can capture multicast packets only on egress.

Correct Answer: C

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Restrictions for Embedded Packet Capture

- In Cisco IOS Release 12.2(33)SRE, EPC is supported only on 7200 platform.
- **EPC only captures multicast packets on ingress and does not capture the replicated packets on egress.**
- Currently, the capture file can only be exported off the device; for example, TFTP or FTP servers and local disk.

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/epc/configuration/15-mt/epc-15-mt-book/nm-packet-capture.html>

QUESTION 21

Refer to the exhibit.

```
switch#show mls cef exception status
Current IPv4 FIB exception state = TRUE
Current IPv6 FIB exception state = FALSE
Current MPLS FIB exception state = FALSE
```

A Cisco Catalyst 6500 Series Switch experiences high CPU utilization. What can be the cause of this issue, and how can it be prevented?

- A. The hardware routing table is full. Redistribute from BGP into IGP.
- B. The software routing table is full. Redistribute from BGP into IGP.
- C. The hardware routing table is full. Reduce the number of routes in the routing table.
- D. The software routing table is full. Reduce the number of routes in the routing table.

Correct Answer: C

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

FIB TCAM Exception - If you try to install more routes than are possible into the FIB TCAM you will see the following error message in the logs:

CFIB-SP-STBY-7-CFIB_EXCEPTION : FIB TCAM exception, Some entries will be software switched

%CFIB-SP-7-CFIB_EXCEPTION : FIB TCAM exception, Some entries will be software switched. %CFIB-SP-STBY-7-CFIB_EXCEPTION : FIB TCAM

exception, Some entries will be software switched.

This error message is received when the amount of available space in the TCAM is exceeded. This results in high CPU. This is a FIB TCAM limitation. Once TCAM is full, a flag will be set and FIB TCAM exception is received. This stops from adding new routes to the TCAM. Therefore, everything will be software switched. The removal of routes does not help resume hardware switching. Once the TCAM enters the exception state, the system must be reloaded to get out of that state. You can view if you have hit a FIB TCAM exception with the following command:

```
6500-2#sh mls cef exception status
Current IPv4 FIB exception state = TRUE
Current IPv6 FIB exception state = FALSE
Current MPLS FIB exception state = FALSE
```

When the exception state is TRUE, the FIB TCAM has hit an exception. The maximum routes that can be installed in TCAM is increased by the mls cef maximum-routes command.

Reference. <https://supportforums.cisco.com/document/59926/troubleshooting-high-cpu-6500-sup720>

QUESTION 22

Refer to the exhibit.

```
access-switch-1#show interface fastethernet0/9
FastEthernet0/9 is up, line protocol is up (connected)
  Hardware is Fast Ethernet, address is 04da.d237.9f09 (bia 04da.d237.9f09)
  Auto-duplex, Auto-speed, media type is 10/100BaseTX
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 59137853

access-switch-1#show mls qos interface fastethernet0/9 statistics
Queueset: 1
output queues dropped:
  queue:      threshold1  threshold2  threshold3
  -----
  queue 0:           0           0          48252
  queue 1:    23164955    35924645           1
  queue 2:           0           0           0
  queue 3:           0           0           0
```

Your network is suffering excessive output drops. Which two actions can you take to resolve the problem? (Choose two.)

- A. Install a switch with larger buffers.
- B. Configure a different queue set.
- C. Reconfigure the switch buffers.
- D. Configure the server application to use TCP.
- E. Update the server operating system.

Correct Answer: AB

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Installing a switch with larger buffers and correctly configuring the buffers can solve output queue problems.

For each queue we need to configure the assigned buffers. The buffer is like the 'storage' space for the interface and we have to divide it among the different queues. This is how to do it:

```
mls qos queue-set output <queue set> buffers Q1 Q2 Q3 Q4
```

In this example, there is nothing hitting queue 2 or queue 3 so they are not being utilized.

QUESTION 23


DRAG DROP

Drag and drop the Cisco IOS XE subpackage on the left to the function it performs on the right.

Select and Place:

RPBase	administers the shared port adaptor driver and related field-programmable device images
RPCControl	provisions the software needed to access the router
SIPSPA	manages the Cisco IOS Software and the rest of the platform via the control plane
RPAccess	provisions the operating system software route processor

Correct Answer:



Section: Network Principles

Explanation

Explanation/Reference:

QUESTION 24

Which two Cisco IOS XE commands can install a subpackage onto a router? (Choose two.)

- A. request platform software package install rp rpSlotNumber file fileURL
- B. boot system flash bootflash:filename
- C. copy sourceUrl destinationUrl
- D. license install file storedLocationUrl
- E. issu loadversion rp identifier file diskType imageFilename
- F. config-register value

Correct Answer: AC

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Managing and Configuring a Consolidated Package Using the request platform software package install Command

In the following example, the **request platform software package install** command is used to upgrade a consolidated package running on RP 0. The

force option, which forces the upgrade past any prompt (such as already having the same consolidated package installed), is used in this example.

Router# **request platform software package install rp 0 file bootflash:asr1000rp1-adventerprisek9.02.01.00.122-33.XNA.bin force**

To upgrade a consolidated package on the Cisco ASR 1000 Series Routers using the **copy** command, copy the consolidated package into the bootflash: directory on the router using the **copy** command as you would on most other Cisco routers. After making this copy, configure the router to boot using the consolidated package file.

In the following example, the consolidated package file is copied onto the bootflash: file system from TFTP. The config-register is then set to boot using **boot system** commands, and the **boot system** commands instruct the router to boot using the consolidated package stored in the bootflash: file system. The new configuration is then saved using the **copy running-config startup-config** command, and the system is then reloaded to complete the process.

Router# **dir bootflash:**

Directory of bootflash:/

```
11 drwx 16384 Dec 4 2007 04:32:46 -08:00 lost+found
86401 drwx 4096 Dec 4 2007 06:06:24 -08:00.ssh
14401 drwx 4096 Dec 4 2007 06:06:36 -08:00.rollback_timer
28801 drwx 4096 Mar 18 2008 17:31:17 -07:00.prst_sync
43201 drwx 4096 Dec 4 2007 04:34:45 -08:00.installer
13 -rw- 45977 Apr 9 2008 16:48:46 -07:00 target_support_output.tgz.tgz
```

928862208 bytes total (712273920 bytes free)

Router# **copy tftp bootflash:**

Address or name of remote host []? **172.17.16.81**

Source filename []? **/auto/tftp-users/user/asr1000rp1-adventerprisek9.02.01.00.122-33.XNA.bin** Destination filename [asr1000rp1-adventerprisek9.02.01.00.122-33.XNA.bin]?

Reference:

http://www.cisco.com/c/en/us/td/docs/routers/asr1000/configuration/guide/chassis/asrswcfg/Packag e_Management.html#78189

QUESTION 25

Which two statements about Cisco Express Forwarding are true? (Choose two.)

- A. Cisco Express Forwarding tables contain reachability information and adjacency tables contain forwarding information.
- B. Cisco Express Forwarding tables contain forwarding information and adjacency tables contain reachability information.
- C. Changing MAC header rewrite strings requires cache validation.
- D. Adjacency tables and Cisco Express Forwarding tables can be built separately.
- E. Adjacency tables and Cisco Express Forwarding tables require packet process-switching.

Correct Answer: AD

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Main Components of CEF

Information conventionally stored in a route cache is stored in several data structures for Cisco Express Forwarding switching. The data structures provide optimized lookup for efficient packet forwarding. The two main components of Cisco Express Forwarding operation are the forwarding information base (FIB) and the adjacency tables.

The FIB is conceptually similar to a routing table or information base. A router uses this lookup table to make destination-based switching decisions during Cisco Express Forwarding operation. The FIB is updated when changes occur in the network and contains all routes known at the time.

Adjacency tables maintain Layer 2 next-hop addresses for all FIB entries. This separation of the reachability information (in the Cisco Express Forwarding table) and the forwarding information (in the adjacency table), provides a number of benefits:

- The adjacency table can be built separately from the Cisco Express Forwarding table, allowing both to be built without any packets being process-switched.
- The MAC header rewrite used to forward a packet is not stored in cache entries, so changes in a MAC header rewrite string do not require validation of cache entries.

Reference: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipswitch_cef/configuration/15-mt/isw-cef-15-mt-book/isw-cef-overview.html

QUESTION 26

Refer to the exhibit.

```
Switch#show spanning-tree
VLAN0001

Spanning tree enabled protocol ieee

Root ID    Priority    32769
Address     001a.6d4b.c500
This bridge is the root
Hello Time  2 sec    Max Age 20 sec    Forward Delay 15 sec

Bridge ID   Priority    32769 (priority 32768 sys-id-ext 1)
Address     001a.6d4b.c500
Hello Time  2 sec    Max Age 20 sec    Forward Delay 15 sec
Aging Time  15

Interface   Role Sts Cost      Prio.Nbr Type
-----
Fa0/1       Desg FWD 19       128.1     P2p
```

If you change the Spanning Tree Protocol from pvst to rapid-pvst, what is the effect on the interface Fa0/1 port state?

- A. It transitions to the listening state, and then the forwarding state.
- B. It transitions to the learning state and then the forwarding state.
- C. It transitions to the blocking state, then the learning state, and then the forwarding state.
- D. It transitions to the blocking state and then the forwarding state.

Correct Answer: C

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

First, the port will transition to the blocking state, immediately upon the change, then it will transition to the new RSTP states of learning and forwarding.

Port States

There are only three port states left in RSTP that correspond to the three possible operational states. The 802.1D disabled, blocking, and listening states are merged into a unique 802.1w discarding state.

STP (802.1D) Port State	RSTP (802.1w) Port State	Is Port Included in Active Topology?	Is Port Learning MAC Addresses?
Disabled	Discarding	No	No
Blocking	Discarding	No	No
Listening	Discarding	Yes	No
Learning	Learning	Yes	Yes
Forwarding	Forwarding	Yes	Yes

QUESTION 27

Which type of port would have root guard enabled on it?

- A. A root port
- B. An alternate port
- C. A blocked port
- D. A designated port

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:


The root guard feature provides a way to enforce the root bridge placement in the network.

The root guard ensures that the port on which root guard is enabled is the designated port. Normally, root bridge ports are all designated ports, unless two or more ports of the root bridge are connected together. If the bridge receives superior STP Bridge Protocol Data Units (BPDUs) on a root guard-enabled port, root guard moves this port to a root-inconsistent STP state. This root-inconsistent state is effectively equal to a listening state. No traffic is forwarded across this port. In this way, the root guard enforces the position of the root bridge.

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

QUESTION 28

Refer to the exhibit.



All switches have default bridge priorities, and originate BPDUs with MAC addresses as indicated. The numbers shown are STP link metrics. Which two ports are forwarding traffic after STP converges? (Choose two.)

- A. The port connecting switch SWD with switch SWE
- B. The port connecting switch SWG with switch SWF
- C. The port connecting switch SWC with switch SWE

D. The port connecting switch SWB with switch SWC

Correct Answer: CD

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Here, we know SWB to SWC are forwarding because we already identified the blocking port. So for the last correct answer let's consider what must be done to prevent a switch loop between SWC/SWD/SWE. SWE to SWD will be blocked because SWC has a lower MAC address so it wins the forwarding port. And to look at it further, you could try to further understand what would happen with ports on SWG. Would the ports on SWG try to go through SWE or SWF? SWE has the lower MAC address so the port from SWG to SWE would win the forwarding election. Therefore, answer B could never be correct.

QUESTION 29

Refer to the exhibit.

```
Switch# show ip igmp snooping mrouter
Vlan      ports
-----
 10       Gi2/0/1(dynamic), Router
 20       Gi2/0/1(dynamic), Router
```

Which three statements about the output are true? (Choose three.)

- A. An mrouter port can be learned by receiving a PIM hello packet from a multicast router.
- B. This switch is configured as a multicast router.
- C. Gi2/0/1 is a trunk link that connects to a multicast router.
- D. An mrouter port is learned when a multicast data stream is received on that port from a multicast router.
- E. This switch is not configured as a multicast router. It is configured only for IGMP snooping.
- F. IGMP reports are received only on Gi2/0/1 and are never transmitted out Gi2/0/1 for VLANs 10 and 20.

Correct Answer: ABC

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

In this example, the switch has been configured as a multicast router since IGMP snooping has been enabled. All mrouter can learn about other

multicast routers by receiving a PIM hello packet from another multicast router. Also, since two different VLANs are being used by the same port of gi 2/0/1, it must be a trunk link that connects to another multicast router.

QUESTION 30

Refer to the exhibit.

```
Switch#show int fastEthernet0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: dynamic auto
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: On
Access Mode VLAN: 2 (VLAN0002)
Trunking Native Mode VLAN: 3 (VLAN0003)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk private VLANs: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL

Protected: false
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled
Appliance trust: none
```

If a port is configured as shown and receives an untagged frame, of which VLAN will the untagged frame be a member?

- A. VLAN 1
- B. VLAN 2

- C. VLAN 3
- D. VLAN 4

Correct Answer: B

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

When typing:

Switch(config-if)#**switchport mode** ?

access Set trunking mode to ACCESS unconditionally

dynamic Set trunking mode to dynamically negotiate access or trunk mode

trunk Set trunking mode to TRUNK unconditionally

and

Switch(config-if)#**switchport mode dynamic** ?

auto Set trunking mode dynamic negotiation parameter to AUTO

desirable Set trunking mode dynamic negotiation parameter to DESIRABLE

So if we configure Fa0/1 as **dynamic auto** mode, it will not initiate any negotiation but waiting for the other end negotiate to be a trunk with DTP. If the other end does not ask it to become a trunk then it will become an access port. Therefore when using the "show interface fastEthernet0/1 switchport" command we will see two output lines "**Administrative Mode. dynamic auto**" and "**Operational Mode. static access**"

Note. To set this port to VLAN 2 as the output above just use one additional command. "switchport access vlan 2".

Now back to our question, from the output we see that Fa0/1 is operating as an access port on VLAN 2 so if it receive untagged frame it will suppose that frame is coming from VLAN 2.

QUESTION 31

Refer to the exhibit.

```
Switch#show interfaces switchport backup detail
```

```
Switch Backup Interface Pairs:
```

Active Interface	Backup Interface	State
FastEthernet0/1	FastEthernet0/2	Active Up/Backup Standby

```
Interface Pair : Fa0/1, Fa0/2
```

```
Preemption Mode : off
```

```
Bandwidth : 100000 Kbit (Fa0/1), 10000 Kbit (Fa0/2)
```

```
Mac Address Move Update Vlan : auto
```

Which statement describes the effect on the network if FastEthernet0/1 goes down temporarily?

- A. FastEthernet0/2 forwards traffic only until FastEthernet0/1 comes back up.
- B. FastEthernet0/2 stops forwarding traffic until FastEthernet0/1 comes back up.
- C. FastEthernet0/2 forwards traffic indefinitely.
- D. FastEthernet0/1 goes into standby.

Correct Answer: C

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:


Use the switchport backup interface interface configuration command on a Layer 2 interface to configure Flex Links, a pair of interfaces that provide backup to each other. Use the no form of this command to remove the Flex Links configuration.

With Flex Links configured, one link acts as the primary interface and forwards traffic, while the other interface is in standby mode, ready to begin forwarding traffic if the primary link shuts down. The interface being configured is referred to as the active link; the specified interface is identified as the backup link. The feature provides an alternative to the Spanning Tree Protocol (STP), allowing users to turn off STP and still retain basic link redundancy.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960/software/release/12-2_53_se/command/reference/2960ComRef/cli3.html#wp3269214

QUESTION 32

Refer to the exhibit.



Which technology does the use of bi-directional BPDUs on all ports in the topology support?

- A. RSTP
- B. MST
- C. Bridge Assurance
- D. Loop Guard
- E. Root Guard
- F. UDLD

Correct Answer: C

Section: Layer 2 Technologies

Explanation


Explanation/Reference:

Explanation:


Spanning Tree Bridge Assurance

- Turns STP into a bidirectional protocol
- Ensures spanning tree fails "closed" rather than "open"
- If port type is "network" send BPDU regardless of state
- If network port stops receiving BPDU it's put in BA-inconsistent state

Without Bridge Assurance



With Bridge Assurance



Bridge Assurance (BA) can help protect against bridging loops where a port becomes designated because it has stopped receiving BPDUs. This is similar to the function of loop guard.

Reference. <http://lostintransit.se/tag/convergence/>

QUESTION 33

Which two statements are true about an EPL? (Choose two.)

- A. It is a point-to-point Ethernet connection between a pair of NNIs.
- B. It allows for service multiplexing.
- C. It has a high degree of transparency.
- D. The EPL service is also referred to as E-line.

Correct Answer: CD

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Ethernet private line (EPL) and Ethernet virtual private line (EVPL) are carrier Ethernet data services defined by the Metro Ethernet Forum. EPL provides a point-to-point Ethernet virtual connection (EVC) between a pair of dedicated user network interfaces (UNIs), with a high degree of transparency. EVPL

provides a point-to-point or point-to-multipoint connection between a pair of UNIs.

The services are categorized as an E-Line service type, with an expectation of low frame delay, frame delay variation and frame loss ratio. EPL is implemented using a point-to-point (EVC) with no Service Multiplexing at each UNI (physical interface), i.e., all service frames at the UNI are mapped to a single EVC (a.k.a. all-to-one bundling).

Reference. http://en.wikipedia.org/wiki/Ethernet_Private_Line

QUESTION 34

Which two statements describe characteristics of HDLC on Cisco routers? (Choose two.)

- A. It supports multiple Layer 3 protocols.
- B. It supports multiplexing.
- C. It supports only synchronous interfaces.
- D. It supports authentication.

Correct Answer: AC

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Cisco High-Level Data Link Controller (HDLC) is the Cisco proprietary protocol for sending data over synchronous serial links using HDLC. Cisco HDLC also provides a simple control protocol called Serial Line Address Resolution Protocol (SLARP) to maintain serial link keepalives. Cisco HDLC is the default for data encapsulation at Layer 2 (data link) of the Open System Interconnection (OSI) stack for efficient packet delineation and error control. The absence of a protocol type field in the HDLC header posed a problem for links that carried traffic from more than one Layer 3 protocol. Cisco, therefore, added an extra Type field to the HDLC header, creating a Cisco-specific version of HDLC. Cisco routers can support multiple network layer protocols on the same HDLC link. For example an HDLC link between two Cisco routers can forward both IPv4 and IPv6 packets because the Type field can identify which type of packet is carried inside each HDLC frame.

Reference.

http://www.cisco.com/c/en/us/td/docs/routers/access/800/819/software/configuration/Guide/819_S_CG/6ser_conf.html#pgfId-1073734

QUESTION 35

Which mechanism can be used on Layer 2 switches so that only multicast packets with downstream receivers are sent on the multicast router-connected ports?

- A. IGMP snooping
- B. Router Guard
- C. PIM snooping
- D. multicast filtering

Correct Answer: C

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Ideally, the Layer 2 device should forward the multicast transmission only out ports to which receivers are connected and also out any ports that are connected to downstream multicast routers. This configuration requires a Layer 2 device to be able to determine the ports on which multicast routers and receivers for each separate (S,G) or (*,G) multicast group are located. To facilitate intelligent forwarding of multicast traffic on the LAN, Cisco Catalyst switches support two mechanisms:

- **IGMP snooping**-- The switch listens in or "snoops" IGMP communications between receivers and multicast routers. This snooping enables the switch to determine which ports are connected to receivers for each multicast group and which ports are connected to multicast routers.
- **Cisco Group Management Protocol (CGMP)**-- The switch communicates with multicasts routers, with multicast routers relaying group membership information to switches.

Reference. https://www.informit.com/library/content.aspx?b=CCNP_Studies_Switching&seqNum=59

QUESTION 36

Which statement is true about Fast Link Pulses in Ethernet?

- A. They are used during collision detection.
- B. They are used only if the media type is optical.
- C. They are part of UniDirectional Link Detection.
- D. They are used during autonegotiation.

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

To make sure that your connection is operating properly, IEEE 802.3 Ethernet employs normal link pulses (NLPs), which are used for verifying link integrity in a 10BaseT system. This signaling gives you the link indication when you attach to the hub and is performed between two directly connected link interfaces (hub-to-station or station-to-station). NLPs are helpful in determining that a link has been established between devices, but they are not a good indicator that your cabling is free of problems.

An extension of NLPs is fast link pulses. These do not perform link tests, but instead are employed in the autonegotiation process to advertise a device's capabilities.

Reference. <http://www.cisco.com/en/US/docs/internetworking/troubleshooting/guide/tr1904.html>

QUESTION 37

Which statement is true regarding UDLD and STP timers?

- A. The UDLD message timer should be two times the STP forward delay to prevent loops.
- B. UDLD and STP are unrelated features, and there is no relation between the timers.
- C. The timers need to be synced by using the spanning-tree uddl-sync command.
- D. The timers should be set in such a way that UDLD is detected before the STP forward delay expires.

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

UDLD is designed to be a helper for STP. Therefore, UDLD should be able to detect an unidirectional link before STP would unblock the port due to missed BPDUs. Thus, when you configure UDLD timers, make sure your values are set so that unidirectional link is detected before "STP MaxAge + 2xForwardDelay" expires.

Reference. <http://blog.ine.com/tag/stp/>

QUESTION 38

Which switching technology can be used to solve reliability problems in a switched network?

- A. fragment-free mode
- B. cut-through mode
- C. check mode
- D. store-and-forward mode

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Characteristics of Store-and-Forward Ethernet Switching

This section provides an overview of the functions and features of store-and-forward Ethernet switches.

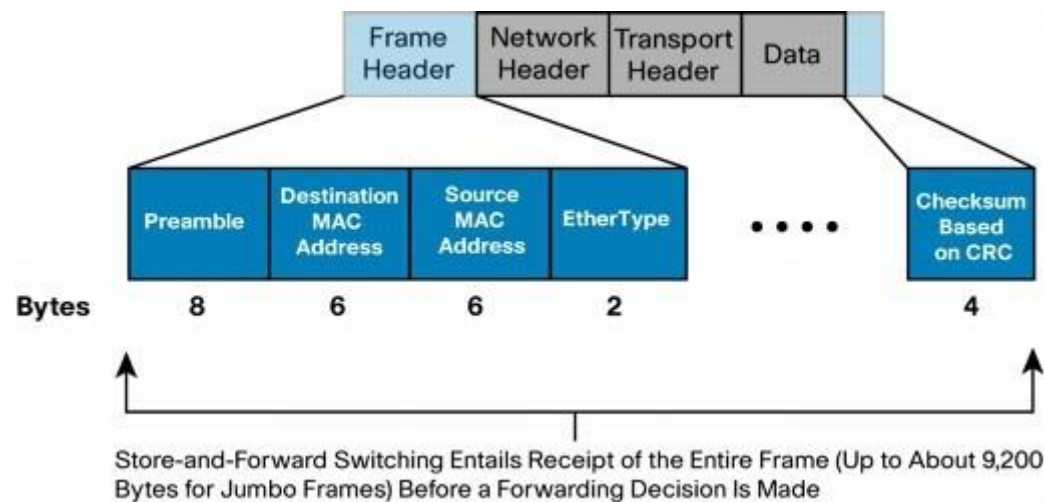
Error Checking

Figure 1 shows a store-and-forward switch receiving an Ethernet frame in its entirety. At the end of that frame, the switch will compare the last field of the datagram against its own frame-check- sequence (FCS) calculations, to help ensure that the packet is free of physical and data-link errors.

The switch then performs the forwarding process.

Whereas a store-and-forward switch solves reliability issues by dropping invalid packets, cut- through devices forward them because they do not get a chance to evaluate the FCS before transmitting the packet.


Figure 1. Ethernet Frame Entering a Store-and-Forward Bridge or Switch (from Left to Right)



Reference. http://www.cisco.com/c/en/us/products/collateral/switches/nexus-5020-switch/white_paper_c11-465436.html

QUESTION 39

Refer to the exhibit.



All switches have default bridge priorities, and originate BPDUs with MAC addresses as indicated. The numbers shown are STP link metrics. Which two ports are in blocking state after STP converges? (Choose two.)

- A. the port on switch SWD that connects to switch SWE
- B. the port on switch SWF that connects to switch SWG
- C. the port on switch SWD that connects to switch SWC

D. the port on switch SWB that connects to switch SWD

Correct Answer: CD

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

This is a scenario that wants you to demonstrate understanding of the Root switch and Root port election process. So, it's best to start with where the root switch will be and work down from there. It's setup nicely because the lowest MAC address switch starts at the top and then the lower priority/higher mac addresses move down the architecture. SWA wins the root election and of course all ports in SWA are forwarding. SWB introduces the possibility for a switching loop so it's important to understand which ports will be put into the blocking state. Since SWD is a higher MAC address it will end up with a blocked port connected to SWB to prevent a loop; and this is one of the correct answers. To prevent the possibility of another potential switching loop, SWD again ends up with the higher MAC address so blocking the link between D and C prevents a B/C/D switching loop.

QUESTION 40

Which statement is true about IGMP?

- A. Multicast sources send IGMP messages to their first-hop router, which then generates a PIM join message that is then sent to the RP.
- B. Multicast receivers send IGMP messages to their first-hop router, which then forwards the IGMP messages to the RP.
- C. IGMP messages are encapsulated in PIM register messages and sent to the RP.
- D. Multicast receivers send IGMP messages to signal their interest to receive traffic for specific multicast groups.


Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:



In the example shown above, the receivers (the designated multicast group) are interested in receiving the video data stream from the source. The receivers indicate their interest by sending an Internet Group Management Protocol (IGMP) host report to the routers in the network. The routers are then responsible for delivering the data from the source to the receivers.

Reference.

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

QUESTION 41

Which two statements are true about RSTP? (Choose two.)

- A. By default, RTSP uses a separate TCN BPDU when interoperating with 802.1D switches.
- B. By default, RTSP does not use a separate TCN BPDU when interoperating with 802.1D switches.
- C. If a designated port receives an inferior BPDU, it immediately triggers a reconfiguration.
- D. By default, RTSP uses the topology change TC flag.
- E. If a port receives a superior BPDU, it immediately replies with its own information, and no reconfiguration is triggered.

Correct Answer: BD

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

The RSTP does not have a separate topology change notification (TCN) BPDU. It uses the topology change (TC) flag to show the topology changes.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2950/software/release/12-1_9_ea1/configuration/guide/scg/swmstp.html

QUESTION 42

Refer to the exhibit.

```
switch#show spanning-tree detail

MST0 is executing the mstp compatible Spanning Tree protocol
Bridge Identifier has priority 32768, sysid 0, address f4ac.c1c4.2b80
Configured hello time 2, max age 20, forward delay 15, transmit hold-count 6
Current root has priority 24576, address 0019.07aa.9ac0
Root port is 56 (Port-channel1), cost of root path is 0
Topology change flag not set, detected flag not set
Number of topology changes 296 last change occurred 00:01:17 ago
      from GigabitEthernet0/15
```

Which two statements are true about the displayed STP state? (Choose two.)

- A. The STP version configured on the switch is IEEE 802.1w.
- B. Port-channel 1 is flapping and the last flap occurred 1 minute and 17 seconds ago.
- C. The switch does not have PortFast configured on Gi0/15.
- D. BPDUs with the TCN bit set are transmitted over port channel 1.

Correct Answer: CD

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

A port enabled with portfast will not send topology changes when a port goes up or down, but here we see that 296 TCN's were sent so we know that Gi 0/15 does not have portfast enabled. TCN's are sent using BPDU's over the root port, which we see is port channel 1.

QUESTION 43

DRAG DROP

Drag and drop the multicast protocol definition on the left to the correct default time interval on the right.

Select and Place:

Drag and drop the multicast protocol definition on the left to the correct default time interval on the right.

IGMPv2 query interval

30 seconds

IGMPv2 querier timeout

IGMPv1 query interval

60 seconds

PIMv1 query interval

IGMPv3 query interval

120 seconds

Correct Answer:

Drag and drop the multicast protocol definition on the left to the correct default time interval on the right.

	30 seconds
	PIMv1 query interval
	60 seconds
	IGMPv2 query interval
	IGMPv1 query interval
	IGMPv3 query interval
	120 seconds
	IGMPv2 querier timeout

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

QUESTION 44

When you migrate a network from PVST+ to rapid-PVST+, which two features become inactive? (Choose two.)

- A. Root guard
- B. Loop guard
- C. UplinkFast
- D. UDLD
- E. BackboneFast

F. Bridge Assurance

Correct Answer: CE

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

It is good to know the UplinkFast and BackboneFast behavior before you start the migration process.

Here, the Access1 switch runs Cisco IOS. This output is taken before migration to the rapid-PVST+ mode:

```
Access1#show spanning-tree vlan 10
```

```
VLAN0010
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    24586
           Address    0015.63f6.b700
           Cost      3019
           Port      107 (FastEthernet3/0/1)
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID  Priority    49162 (priority 49152 sys-id-ext 10)
           Address    000f.f794.3d00
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
           Aging Time 300
```

```
Uplinkfast enabled
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
-----	----	---	-----	-----	-----
Fa3/0/1	Root	FWD	3019	128.107	P2p
Fa3/0/2	Altn	BLK	3019	128.108	P2p

```
Access1#show spanning-tree summary
```

```
Switch is in pvst mode
```

```
Root bridge for: none
```

```
Extended system ID      is enabled
Portfast Default        is disabled
PortFast BPDU Guard Default is enabled
Portfast BPDU Filter Default is disabled
Loopguard Default       is disabled
EtherChannel misconfig guard is enabled
UplinkFast              is enabled
BackboneFast            is enabled
Configured Pathcost method used is short
```

Name	Blocking	Listening	Learning	Forwarding	STP Active
-----	-----	-----	-----	-----	-----
VLAN0010	1	0	0	1	2
VLAN0020	1	0	0	1	2
-----	-----	-----	-----	-----	-----
2 vlans	2	0	0	2	4

```
This output is taken after the mode is changed to rapid-PVST+:
```

```
Access1#show spanning-tree vlan 10
```

You can see in the show spanning-tree summary command output that UplinkFast and BackboneFast are enabled, but are inactive in rapid-PVST mode.

Reference. <http://www.cisco.com/c/en/us/support/docs/switches/catalyst-6500-series-switches/72836-rapidpvst-mig-config.html#upback1>

QUESTION 45

Which statement is true about LLDP?

- A. LLDP provides VTP support.
- B. LLDP does not use a multicast address to communicate.
- C. LLDP can indicate only the duplex setting of a link, and not the speed capabilities.
- D. LLDP does not support native VLAN indication.

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Cisco Discovery Protocol Versus LLDP-MED TLV Comparison

TLV Function	LLDP TLV	Cisco Discovery Protocol TLV
Native VLAN support-Indicates the native VLAN	No	Native VLAN TLV

Reference.

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

QUESTION 46

Which statement is true when using a VLAN ID from the extended VLAN range (10064094)?

- A. VLANs in the extended VLAN range can be used with VTPv2 in either client or server mode.
- B. VLANs in the extended VLAN range can only be used as private VLANs.
- C. STP is disabled by default on extended-range VLANs.
- D. VLANs in the extended VLAN range cannot be pruned.

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Enabling VTP pruning on a VTP server enables pruning for the entire management domain. Making VLANs pruning-eligible or pruning-ineligible affects pruning eligibility for those VLANs on that device only (not on all switches in the VTP domain). VTP pruning takes effect several seconds after you enable it. VTP pruning does not prune traffic from VLANs that are pruning-ineligible. VLAN 1 and VLANs 1002 to 1005 are always pruning-ineligible; traffic from these VLANs cannot be pruned. Extended-range VLANs (VLAN IDs higher than 1005) are also pruning-ineligible.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3550/software/release/12-1_13_ea1/configuration/guide/3550scg/swvtp.html#wpxref48156

QUESTION 47

Which statement is true about trunking?

- A. Cisco switches that run PVST+ do not transmit BPDUs on nonnative VLANs when using a dot1q trunk.
- B. When removing VLAN 1 from a trunk, management traffic such as CDP is no longer passed in that VLAN.
- C. DTP only supports autonegotiation on 802.1q and does not support autonegotiation for ISL.
- D. DTP is a point-to-point protocol.

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Ethernet trunk interfaces support different trunking modes. You can set an interface as trunking or nontrunking or to negotiate trunking with the neighboring interface. To autonegotiate trunking, the interfaces must be in the same VTP domain.

Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which is a Point-to-Point Protocol. However, some internetworking devices might forward DTP frames improperly, which could cause misconfigurations.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750/software/release/12-2_55_se/configuration/guide/scg3750/swvlan.html

QUESTION 48

Which three statements are true about an EtherChannel? (Choose three.)

- A. PAGP and LACP can be configured on the same switch if the switch is not in the same EtherChannel.
- B. EtherChannel ports in suspended state can receive BPDUs but cannot send them.
- C. An EtherChannel forms between trunks that are using different native VLANs.
- D. LACP can operate in both half duplex and full duplex, if the duplex setting is the same on both ends.

E. Ports with different spanning-tree path costs can form an EtherChannel.

Correct Answer: ABE

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Answer A. EtherChannel groups running PAgP and LACP can coexist on the same switch or on different switches in the stack. Individual EtherChannel groups can run either PAgP or LACP, but they cannot interoperate.

Answer B:


EtherChannel Member Port States	
Port States	Description
bundled	The port is part of an EtherChannel and can send and receive BPDUs and data traffic.
suspended	The port is not part of an EtherChannel. The port can receive BPDUs but cannot send them. Data traffic is blocked.
standalone	The port is not bundled in an EtherChannel. The port functions as a standalone data port. The port can send and receive BPDUs and data traffic.

Answer E. Ports with different spanning-tree path costs can form an EtherChannel if they are otherwise compatibly configured. Setting different spanning-tree path costs does not, by itself, make ports incompatible for the formation of an EtherChannel.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/software/15-0_2_EX/layer2/configuration_guide/b_lay2_152ex_2960-x_cg/b_lay2_152ex_2960-x_cg_chapter_010.html

QUESTION 49

Refer to the exhibit.



Which statement about configuring the switch to manage traffic is true?

- A. The switchport priority extend cos command on interface FastEthernet0/0 prevents traffic to and from the PC from taking advantage of the high-priority data queue that is assigned to the IP phone.
- B. The switchport priority extend cos command on interface FastEthernet0/0 enables traffic to and from the PC to use the high priority data queue that is assigned to the IP phone.
- C. When the switch is configured to trust the CoS label of incoming traffic, the trusted boundary feature is disabled automatically.
- D. The mls qos cos override command on interface FastEthernet0/0 configures the port to trust the CoS label of traffic to and from the PC.

Correct Answer: A

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

In some situations, you can prevent a PC connected to the Cisco IP Phone from taking advantage of a high-priority data queue. You can use the switchport priority extend cos interface configuration command to configure the telephone through the switch CLI to override the priority of the traffic received from the PC.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2950/software/release/12-1_22_ea2/configuration/guide/2950scg/swqos.html

QUESTION 50

DRAG DROP

Drag and drop the PPPoE packet type on the left to the corresponding description on the right.

Select and Place:

Drag and drop the PPPoE packet type on the left to the corresponding description on the right.

PADR	A packet that is sent with the destination_addr set to the broadcast address. The p indicates the type of service requested.
PADT	A packet that is sent with the destination_addr set to the unicast address of the PP client. The packet contains an offer for the client.
PADO	A packet that is sent from the PPPoE client with the destination_addr set to the ch access concentrator. The packet contains a session request from the client.
PADI	A packet that is sent as confirmation to the client. The packet contains the unique PPPoE session ID.
PADS	A packet that is sent to terminate the PPPoE session.

Correct Answer:

Drag and drop the PPPoE packet type on the left to the corresponding description on the right.

PADI

PADO

PADR

PADS

PADT

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

QUESTION 51

What is the destination multicast MAC address for BPDUs on the native VLAN, for a switch that is running 802.1D?

- A. 0185.C400.0000
- B. 0100.0CCC.CCCC
- C. 0100.0CCC.CCCD
- D. 0180.C200.0000

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

If the native vlan is 1:

A STP BPDU for VLAN 1 will be sent untagged to MAC 0180.c200.0000 (this is the common spanning tree)

A PVST+ BPDU for VLAN 1 will be sent untagged to MAC 0100.0ccc.cccd

A PVST+ BPDU for all other vlans will be sent with a 802.1Q tag to MAC 0100.0ccc.cccd (with a PVID = to the VLAN)

If the native vlan is not 1:

A STP BPDU for VLAN 1 will be sent untagged (on the native vlan) to MAC 0180.c200.0000 (this is the common spanning tree)


A PVST+ BPDU for VLAN1 will be sent with a 802.1Q tag to MAC 0100.0ccc.cccd (with a PVID=1)

A PVST+ BPDU for the native vlan will be sent untagged to MAC 0100.0ccc.cccd (with a PVID=native vlan)

A PVST+ BPDU for all other vlans will be sent with a 802.1Q tag to MAC 0100.0ccc.cccd (with a PVID = to the VLAN)

QUESTION 52

Refer to the exhibit.



All switches have default bridge priorities, and originate BPDUs with MAC addresses as indicated. The numbers shown are STP link metrics.

After STP converges, you discover that traffic from switch SWG toward switch SWD takes a less optimal path. What can you do to optimize the STP tree in this switched network?

A. Change the priority of switch SWA to a lower value than the default value.

- B. Change the priority of switch SWB to a higher value than the default value.
- C. Change the priority of switch SWG to a higher value than the default value.
- D. Change the priority of switch SWD to a lower value than the default value.

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

In this topology, we see that all port paths and priorities are the same, so the lowest MAC address will be used to determine the best STP path. From SWG, SWE will be chosen as the next switch in the path because it has a lower MAC address than SWF. From SWE, traffic will go to SWC because it has a lower MAC address, and then to SWD, instead of going from SWE directly to SWD. If we lower the priority of SWD (lower means better with STP) then traffic will be sent directly to SWD.

QUESTION 53

Which three statements are true about VSS? (Choose three.)

- A. VSS separates the control planes of the active and the standby chassis.
- B. Configuration changes can be made on both active and standby chassis.
- C. When the VSS active chassis recovers after a failure, it initiates a switchover and takes on the active role again.
- D. VSS unifies the control planes of the active and the standby chassis.
- E. HSRP configuration is not required to run VSS.
- F. The VSS standby chassis monitors the VSS active chassis using the VSL.

Correct Answer: DEF

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

VSS operates on a unified control plane with a distributed forwarding architecture in which the active supervisor (or switch) is responsible for actively participating with the rest of the network and for managing and maintaining control plane information.

VSS actually removes the need for a next-hop redundancy protocol like HSRP or VRRP. These first-hop redundancy protocols are usually heavily tied to a fast-converging routing protocol like EIGRP, and still require that each device maintain its own control plane.

The standby chassis monitors the active chassis using the VSL. If it detects failure, the standby chassis initiates a switchover and takes on the active role. When the failed chassis recovers, it takes on the standby role.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/15-1SY/config_guide/sup2T/15_1_sy_swcg_2T/virtual_switching_systems.pdf

QUESTION 54

Which flag in a configuration BPDU instructs all switches to shorten their bridge table aging process from the default 300 seconds to the current forward delay value?

- A. topology change bit
- B. topology change acknowledgment bit
- C. priority bit
- D. max-age bit

Correct Answer: A

Section: Layer 2 Technologies

Explanation

Explanation/Reference:


Explanation:

The Root Bridge continues to set the Topology Change flag (TCN bit) in all Configuration BPDUs that it sends out for a total of Forward Delay + Max Age seconds (default = 35 (20+15) seconds). This flag instructs all bridges to shorten their MAC address table (Bridge table) aging process from the default value of 300 seconds to the current Forward Delay value of the bridge (default=15 seconds).

The TCA flag is set by the upstream bridge to tell the downstream bridges to stop sending TCN BPDUs. The TC flag is set in configuration BPDU by the Root Bridge to shorten the bridge table age-out period from default 300 seconds to Forward Delay seconds.

QUESTION 55

Refer to the exhibit.



Which technology can be used on the switch to enable host A to receive multicast packets for 239.2.2.2 but prevent host B from receiving them?

- A. IGMP filtering
- B. MLD snooping
- C. IGMP snooping
- D. MLD filtering

Correct Answer: C

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

IGMP snooping is the process of listening to Internet Group Management Protocol (IGMP) network traffic. The feature allows a network switch to listen in on the IGMP conversation between hosts and routers. By listening to these conversations the switch maintains a map of which links need which IP multicast streams. Multicasts may be filtered from the links which do not need them and thus controls which ports receive specific multicast traffic.

QUESTION 56

Which option describes the purpose of the PPP endpoint discriminator?

- A. It identifies the maximum payload packet.

- B. It notifies the peer that it prefers 12-bit sequence numbers.
- C. It identifies the system attached to the link.
- D. It determines whether a loopback is on the link.

Correct Answer: C

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

In situations in which many clients use the same username to initiate an MP connection, or when interoperating with non-Cisco routers, you need to control the order in which the bundle name is created. It is necessary to configure the access server to create a bundle name based on the endpoint discriminator first, the username second, or both. The endpoint discriminator identifies the system transmitting the packet and advises the network access server (NAS) that the peer on this link could be the same as the peer on another existing link. Because every client has a unique endpoint discriminator, only multiple links from the same client are bundled into a single unique MP connection. For example, consider when two PC clients initiate a multilink connection to an access server using the same username. If the multilink bundle name is established based on the endpoint discriminator first, then on the username or on both, the NAS can accurately bundle the links from each client using the endpoint discriminator as a bundle name. This bundle name is unique to the peer system transmitting the packet.

Reference. <http://www.cisco.com/c/en/us/support/docs/wan/point-to-point-protocol-ppp/10238-mppp-bundle-name.html>

QUESTION 57

Which option describes how a VTPv3 device responds when it detects a VTPv2 device on a trunk port?

- A. It sends VTPv3 packets only.
- B. It sends VTPv2 packets only.
- C. It sends VTPv3 and VTPv2 packets.
- D. It sends a special packet that contains VTPv3 and VTPv2 packet information.

Correct Answer: C

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

When a VTP version 3 device on a trunk port receives messages from a VTP version 2 device, the VTP version 3 device sends a scaled-down version of the VLAN database on that particular trunk in a VTP version 2 format. A VTP version 3 device does not send out VTP version 2-formatted packets on a trunk port unless it first receives VTP version 2 packets on that trunk. If the VTP version 3 device does not receive VTP version 2 packets for an interval of time on the trunk port, the VTP version 3 device stops transmitting VTP version 2 packets on that trunk port.

Even when a VTP version 3 device detects a VTP version 2 device on a trunk port, the VTP version 3 device continues to send VTP version 3 packets in addition to VTP version 3 device 2 packets, to allow two kinds of neighbors to coexist on the trunk. VTP version 3 sends VTP version 3 and VTP version 2 updates on VTP version 2-detected trunks.

Reference: <http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SX/configuration/guide/book/vtp.html>

QUESTION 58

Which three statements about bridge assurance are true? (Choose three.)

- A. Bridge assurance must be enabled on both ends of a link.
- B. Bridge assurance can be enabled on one end of a link or on both ends.
- C. Bridge assurance is enabled on STP point-to-point links only.
- D. Bridge assurance is enabled on STP multipoint links only.
- E. If a bridge assurance port fails to receive a BPDU after a timeout, the port is put into a blocking state.
- F. If a bridge assurance port fails to receive a BPDU after a timeout, the port is put into an error disabled state.

Correct Answer: ACE

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Bridge Assurance is enabled by default and can only be disabled globally. Also, Bridge Assurance can be enabled only on spanning tree network ports that are point-to-point links. Finally, both ends of the link must have Bridge Assurance enabled.

With Bridge Assurance enabled, BPDUs are sent out on all operational network ports, including alternate and backup ports, for each hello time period. If the port does not receive a BPDU for a specified period, the port moves into the blocking state and is not used in the root port calculation. Once that port receives a BPDU, it resumes the normal spanning tree transitions.

Reference:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/C_LIConfigurationGuide/SpanningEnhanced.html

QUESTION 59

What is the hop limit for an MLD message?

- A. 1
- B. 2
- C. 15
- D. 255

Correct Answer: A

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

MLD uses the Internet Control Message Protocol (ICMP) to carry its messages. All MLD messages are link-local with a hop limit of 1, and they all have the alert option set. The alert option implies an implementation of the hop-by-hop option header.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti_lsm/configuration/xr-3s/imc-lsm-xr-3s-book/ipv6-mcast-mld-xr.html

QUESTION 60

DRAG DROP

Drag and drop the LACP elements on the left into the correct priority order in the hot-standby port-selection process on the right.

Select and Place:

Drag and drop the LACP elements on the left into the correct priority order in the hot-standby port-selection process on the right.	
switch MAC address	1
port number	2
LACP system priority	3
LACP port priority	4

Correct Answer:

Drag and drop the LACP elements on the left into the correct priority order in the hot-standby port-selection process on the right.	
	LACP system priority
	switch MAC address
	LACP port priority
	port number

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

QUESTION 61

Which statement about Cisco Discovery Protocol is true?

- A. The multicast address 0100.0cdd.dddd is used as the destination address for periodic advertisements.
- B. An inactive VLAN that is configured on an access port passes periodic Cisco Discovery Protocol advertisements.
- C. The multicast address 0100.0ccc.ccd is used as the destination address for periodic advertisements.
- D. A VLAN must be active on an access port before periodic Cisco Discovery Protocol advertisements are passed.

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

All CDP packets include a VLAN ID. If you configure CDP on a Layer 2 access port, the CDP packets sent from that access port include the access port VLAN ID. If you configure CDP on a Layer 2 trunk port, the CDP packets sent from that trunk port include the lowest configured VLAN ID allowed on that trunk port.

CDP messages on the active physical interfaces (Ethernet NIC) to a well-known multicast address (0100.0CCC.CCCC).

QUESTION 62

Which three TLVs does LLDP use to discover network devices? (Choose three.)

- A. Management address
- B. Port description
- C. Network policy
- D. System name
- E. Location information
- F. Power management

Correct Answer: ABD

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Basic Management TLV Set

This set includes the following five TLVs used in LLDP:

1. *Port description TLV*: Provides a description of the port in an alpha-numeric format. The value equals the ifDescr object, if the LAN device supports RFC 2863.
2. *System name TLV*: Provides the system's assigned name in an alpha-numeric format. The value equals the sysName object, if the LAN device supports RFC 3418.
3. *System description TLV*: Provides a description of the network entity in an alpha-numeric format. This includes system's name and versions of hardware, operating system and networking software supported in the device. The value equals the sysDescr object, if the LAN device supports RFC 3418.
4. *System capabilities TLV*: Indicates the primary function(s) of the device and whether or not these functions are enabled in the device. The capabilities are indicated by two octets. Bits 0 through 7 indicate Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device and Station respectively. Bits 8 through 15 are reserved.
5. *Management address TLV*: Indicates the addresses of the local LLDP agent. Other remote managers can use this address to obtain information related to the local device.

Reference. http://www.eetimes.com/document.asp?doc_id=1272069

QUESTION 63

Which command enables L2 QoS support in all VLANs (including the native VLAN)?

- A. switchport priority extend cos
- B. mls qos trust dscp
- C. mls qos rewrite ip dscp
- D. switchport trunk native vlan tag

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

You can enter the switchport trunk native vlan tag command to enable the tagging of native VLAN traffic on a per-port basis. When tagging is enabled, all the packets on the native VLAN are tagged and all incoming untagged data packets are dropped, but untagged control packets are accepted. When tagging is enabled, it will allow for L2 QoS support in all VLANs, including the native VLAN.

QUESTION 64

Which two fields reside in the initial CHAP challenge packet? (Choose two.)

- A. the authentication name of the challenger
- B. a random hash value generated by the device
- C. the hashed packet type ID
- D. the packet type ID in clear text

Correct Answer: AD
Section: Layer 2 Technologies
Explanation

Explanation/Reference:

Explanation:

When a caller A dials in to an access server B, The Access server sends across the link an initial Type 1 authentication packet called a Challenge. This Challenge packet contains a randomly generated number, an ID sequence number to identify the challenge (sent in clear text) and the authentication name of the challenger.

Reference. <http://www.rhyshaden.com/ppp.htm>

QUESTION 65

Which statement about WAN Ethernet Services is true?

- A. Rate-limiting can be configured per EVC.
- B. Point-to-point processing and encapsulation are performed on the customer network.
- C. Ethernet multipoint services function as a multipoint-to-multipoint VLAN-based connection.
- D. UNIs can perform service multiplexing and all-in-one bundling.

Correct Answer: A
Section: Layer 2 Technologies
Explanation

Explanation/Reference:

Explanation:

The MEF has defined a set of bandwidth profiles that can be applied at the UNI or to an EVC. A bandwidth profile is a limit on the rate at which Ethernet frames can traverse the UNI or the EVC.

Reference. <http://www.ciscopress.com/articles/article.asp?p=101367&seqNum=2>

QUESTION 66

DRAG DROP

Drag and drop each STP port role on the left to the matching statement on the right.

Select and Place:

alternate port	the port whose path cost deems it closest to the root bridge
backup port	the port that sends the best BPDUs on its segment
designated port	a blocked port that receives more useful BPDUs from a different bridge
root port	a blocked port that receives more useful BPDUs from its own bridge

Correct Answer:

	root port
	designated port
	alternate port
	backup port

Section: Layer 2 Technologies
Explanation

Explanation/Reference:

QUESTION 67**DRAG DROP**

Drag and drop the VLAN number on the left to the corresponding default VLAN name on the right.

Select and Place:

1001	fddi-default
1002	fddinet-default
1003	trnet-default
1004	ethernet
1005	token-ring-default

Correct Answer:

1002

1004

1005

1001

1003

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

QUESTION 68

DRAG DROP

Drag and drop the StackWise stack master election rule on the left into the correct priority order on the right.

Select and Place:

the switch with the highest software priority	1
the switch with the lowest MAC address	2
the current stack master	3
the switch with a defined interface-level configuration	4
the switch with the highest priority value	5
the switch with the longest up time	6

Correct Answer:



the current stack master

the switch with the highest priority value

the switch with a defined interface-level configuration

the switch with the highest software priority

the switch with the longest up time

the switch with the lowest MAC address

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

QUESTION 69

DRAG DROP

Drag and drop the IGMPv2 timer on the left to its default value on the right.

Select and Place:

Group Membership Interval	1 second
Last Member Query Interval	10 seconds
Query Interval	60 seconds
Query Response Interval	255 seconds
Other Querier Present Interval	260 seconds
Version 1 Router Present Timeout	400 seconds

Correct Answer:

Last Member Query Interval

Query Response Interval

Query Interval

Other Querier Present Interval

Group Membership Interval

Version 1 Router Present Timeout

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

QUESTION 70

DRAG DROP

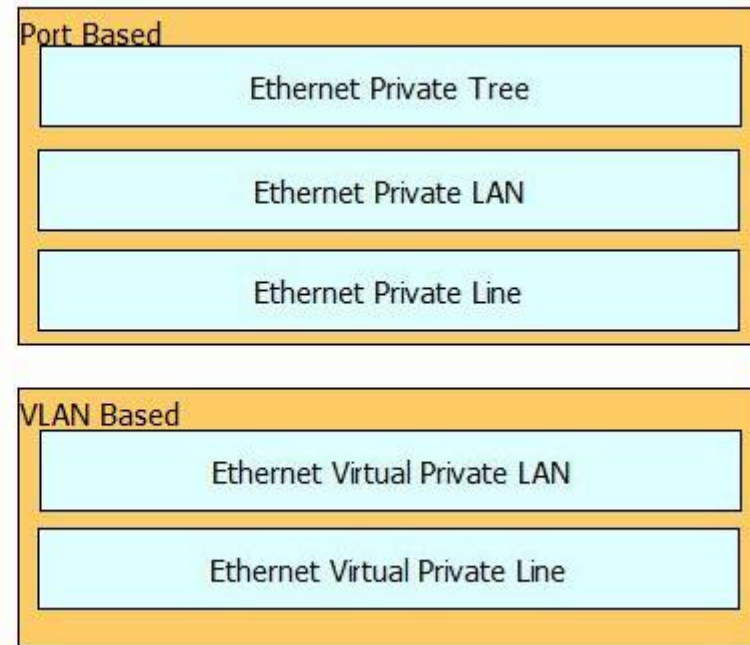
Drag and drop the Metro Ethernet circuit on the left to the corresponding Service Type category on the right.

Select and Place:

Ethernet Virtual Private Line
Ethernet Private Tree
Ethernet Private LAN
Ethernet Virtual Private LAN
Ethernet Private Line

Port Based
1
2
3
VLAN Based
1
2

Correct Answer:



Section: Layer 2 Technologies
Explanation

Explanation/Reference:

QUESTION 71

You are configuring Wireshark on a Cisco Catalyst 4500E Switch with a Supervisor 8. Which three actions can you take to prevent the capture from overloading the CPU? (Choose three.)

- A. Attach the specific ports that are part of the data path.
- B. Use an in-line filter.
- C. Use an appropriate ACL.
- D. Add memory to the Supervisor.
- E. Reconfigure the buffers to accommodate the additional traffic.
- F. Configure a policy map, class map, and an access list to express the match conditions.

Correct Answer: ABC

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Because packet forwarding typically occurs in hardware, packets are not copied to the CPU for software processing. For Wireshark packet capture, packets are copied and delivered to the CPU, which causes an increase in CPU usage. To avoid high CPU, do the following:

- Attach only relevant ports.
- Use a class map, and secondarily, an access list to express match conditions. If neither is viable, use an explicit, in-line filter.
- Adhere closely to the filter rules. Restrict the traffic type (such as, IPv4 only) with a restrictive, rather than relaxed ACL, which elicits unwanted traffic.

Reference: http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/15-1-2/XE_340/configuration/guide/config/wireshrk.pdf

QUESTION 72

Which three statements about VTP version 3 are true? (Choose three.)

- A. It supports other databases in addition to VLAN.
- B. It supports VLANs up to 4095.
- C. It supports the synchronization of switch configuration templates between switches in the domain.
- D. It supports the transfer of information about private VLAN structures.
- E. It supports the transfer of PVST+ configuration information.
- F. It supports RSTP.

Correct Answer: ABD

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Much work has gone into improving the usability of VTP version 3 in three major areas:

- The new version of VTP offers better administrative control over which device is allowed to update other devices' view of the VLAN topology. The chance of unintended and disruptive changes is significantly reduced, and availability is increased. The reduced risk of unintended changes will ease the change process and help speed deployment.
- Functionality for the VLAN environment has been significantly expanded. Two enhancements are most beneficial for today's networks:
 - In addition to supporting the earlier ISL VLAN range from 1 to 1001, the new version supports the whole IEEE 802.1Q VLAN range up to 4095.
 - In addition to supporting the concept of normal VLANs, VTP version 3 can transfer information regarding Private VLAN (PVLAN) structures.
- The third area of major improvement is support for databases other than VLAN (for example, MST).

Reference: http://www.cisco.com/c/en/us/products/collateral/switches/catalyst-6500-series-switches/solution_guide_c78_508010.html

QUESTION 73

In an STP domain, which two statements are true for a nonroot switch, when it receives a configuration BPDU from the root bridge with the TC bit set? (Choose two.)

- A. It sets the MAC table aging time to max_age + forward_delay time.
- B. It sets the MAC table aging time to forward_delay time.
- C. It recalculates the STP topology upon receiving topology change notification from the root switch.
- D. It receives the topology change BPDU on both forwarding and blocking ports.

Correct Answer: BD

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

When the TC bit is received, every bridge is then notified and reduces the aging time to forward_delay (15 seconds by default) for a certain period of time (max_age + forward_delay). It is more beneficial to reduce the aging time instead of clearing the table because currently active hosts, that effectively transmit traffic, are not cleared from the table.

Once the root is aware that there has been a topology change event in the network, it starts to send out its configuration BPDUs with the topology change (TC) bit set. These BPDUs are relayed by every bridge in the network with this bit set. As a result all bridges become aware of the topology change situation and it can reduce its aging time to forward_delay. Bridges receive topology change BPDUs on both forwarding and blocking ports.

An important point to consider here is that a TCN does not start a STP recalculation. This fear comes from the fact that TCNs are often associated with unstable STP environments; TCNs are a consequence of this, not a cause. The TCN only has an impact on the aging time. It does not change the topology nor create a loop.

Reference: <http://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/12013-17.html#topic1>

QUESTION 74

Which two statements about RSTP and MSTP BPDUs are true? (Choose two.)

- A. MSTP switches can detect boundary ports when they receive RSTP version 2 BPDUs.
- B. MSTP switches can detect boundary ports when they receive RSTP version 1 BPDUs.
- C. RSTP switches can process MSTP version 3 BPDUs.
- D. When all boundary switches are running RSTP, MST sends only version 0 configuration BPDUs.

Correct Answer: AC

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

A switch running both MSTP and RSTP supports a built-in protocol migration mechanism that enables it to interoperate with legacy 802.1D switches. If

this switch receives a legacy 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only 802.1D BPDUs on that port. An MST switch can also detect that a port is at the boundary of a region when it receives a legacy BPDU, an MST BPDU (version 3) associated with a different region, or an RST BPDU (version 2).

Reference: http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2950/software/release/12-1_9_ea1/configuration/guide/scg/swmstp.html

QUESTION 75

Which three options are sources from which a SPAN session can copy traffic? (Choose three.)

- A. ports
- B. EtherChannels
- C. VLANs
- D. subnets
- E. primary IP addresses
- F. secondary IP addresses

Correct Answer: ABC

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

SPAN Sources

The interfaces from which traffic can be monitored are called SPAN sources. Sources designate the traffic to monitor and whether to copy ingress, egress, or both directions of traffic. SPAN sources include the following:

- Ethernet ports
 - Port channels
 - The inband interface to the control plane CPU -- You can monitor the inband interface only from the default VDC. Inband traffic from all VDCs is monitored.
 - VLANs -- When a VLAN is specified as a SPAN source, all supported interfaces in the VLAN are SPAN sources.
 - Remote SPAN (RSPAN) VLANs
 - Fabric port channels connected to the Cisco Nexus 2000 Series Fabric Extender
 - Satellite ports and host interface port channels on the Cisco Nexus 2000 Series Fabric Extender
- These interfaces are supported in Layer 2 access mode, Layer 2 trunk mode, and Layer 3 mode.

Reference: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/system_management/configuration/guide/sm_nx_os_cg/sm_14span.html#wp1239492

QUESTION 76

Which three capabilities are provided by MLD snooping? (Choose three.)

- A. dynamic port learning

- B. IPv6 multicast router discovery
- C. user-configured ports age out automatically
- D. a 5-minute aging timer
- E. flooding control packets to the egress VLAN
- F. a 60-second aging timer

Correct Answer: ABD

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Like IGMP snooping, MLD snooping performs multicast router discovery, with these characteristics:

- · Ports configured by a user never age out.
- · Dynamic port learning results from MLDv1 snooping queries and IPv6 PIMv2 packets.
- · If there are multiple routers on the same Layer 2 interface, MLD snooping tracks a single multicast router on the port (the router that most recently sent a router control packet).
- · Dynamic multicast router port aging is based on a default timer of 5 minutes; the multicast router is deleted from the router port list if no control packet is received on the port for 5 minutes.
- · IPv6 multicast router discovery only takes place when MLD snooping is enabled on the switch.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750/software/release/12-2_55_se/configuration/guide/scg3750/swv6mld.pdf

QUESTION 77

Refer to the exhibit.

```
Interface Serial0/1
  ppp multilink
  multilink-group 2
  ppp multilink interleave
  ppp multilink multiclass
```

Which two statements about the implementation are true? (Choose two.)

- A. The PPP multilink protocol header is omitted on delay-sensitive packets.
- B. The maximum number of fragments is 1.
- C. Small real-time packets are multilink-encapsulated.
- D. A transmit queue is provided for smaller packets.

Correct Answer: AD
Section: Layer 2 Technologies
Explanation

Explanation/Reference:

Explanation:

Previous implementations of Cisco IOS Multilink PPP (MLP) include support for Link Fragmentation Interleaving (LFI). This feature allows the delivery of delay-sensitive packets, such as the packets of a Voice call, to be expedited by omitting the PPP Multilink Protocol header and sending the packets as raw PPP packets in between the fragments of larger data packets. This feature works well on bundles consisting of a single link. However, when the bundle contains multiple links there is no way to keep the interleaved packets in sequence with respect to each other.

The Multiclass Multilink PPP (MCMP) feature in Cisco IOS Release 12.2(13)T addresses the limitations of MLP LFI on bundles containing multiple links by introducing multiple data classes. With multiclass multilink PPP interleaving, large packets can be multilink-encapsulated and fragmented into smaller packets to satisfy the delay requirements of real-time voice traffic; small real-time packets, which are not multilink encapsulated, are transmitted between fragments of the large packets. The interleaving feature also provides a special transmit queue for the smaller, delay-sensitive packets, enabling them to be transmitted earlier than other flows. Interleaving provides the delay bounds for delay-sensitive voice packets on a slow link that is used for other best-effort traffic.

References:

http://www.cisco.com/c/en/us/td/docs/ios/dial/configuration/guide/12_4t/dia_12_4t_book/dia_multi_class_link_ppp.pdf

http://www.cisco.com/c/en/us/td/docs/routers/access/500/520/software/configuration/guide/520_SC_G_Book/520scg_concepts.html

QUESTION 78

Which two statements are characteristics of Ethernet private LAN circuits? (Choose two.)

- A. They support communication between two or more customer endpoints.
- B. They utilize more than one bridge domain.
- C. They support point-to-multipoint EVC.
- D. They support multipoint-to-multipoint EVC.

Correct Answer: AD
Section: Layer 2 Technologies
Explanation

Explanation/Reference:

Explanation:

An Ethernet Private LAN (EPLAN) is a multipoint-to-multipoint EVC. EPLAN is an EVC that supports communication between two or more UNIs. In EPLAN, only one EVC can exist on a port and the port can have only one EFP.

Reference:

https://www.cisco.com/c/en/us/td/docs/optical/cpt/r9_3/configuration/guide/cpt93_configuration/cpt93_configuration_chapter_0100.pdf

QUESTION 79

Which two statements about Inverse ARP are true? (Choose two.)

- A. It uses the same operation code as ARP.
- B. It uses the same packet format as ARP.
- C. It uses ARP stuffing.
- D. It supports static mapping.
- E. It translates Layer 2 addresses to Layer 3 addresses.
- F. It translates Layer 3 addresses to Layer 2 addresses.

Correct Answer: BE

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Inverse Address Resolution Protocol (Inverse ARP or InARP) is used to obtain Network Layer addresses (for example, IP addresses) of other nodes from Data Link Layer (Layer 2) addresses. It is primarily used in Frame Relay (DLCI) and ATM networks, in which Layer 2 addresses of virtual circuits are sometimes obtained from Layer 2 signaling, and the corresponding Layer 3 addresses must be available before those virtual circuits can be used. Since ARP translates Layer 3 addresses to Layer 2 addresses, InARP may be described as its inverse. In addition, InARP is implemented as a protocol extension to ARP: it uses the same packet format as ARP, but different operation codes.

Reference: http://en.wikipedia.org/wiki/Address_Resolution_Protocol

QUESTION 80

Refer to the exhibit.

```
R1#show ip mroute 232.1.1.1
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, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(10.1.4.7, 232.1.1.1), 00:17:24/00:02:53, flags: sTI
  Incoming interface: Ethernet1/0, RPF nbr 10.1.5.6*
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:14:42/00:01:21
```

What is the meaning of the asterisk (*) in the output?

- A. PIM neighbor 10.1.5.6 is the RPF neighbor for the group 232.1.1.1 for the shared tree.
- B. PIM neighbor 10.1.5.6 is the one that is seen as the RPF neighbor when performing the command show ip rpf 10.1.4.7.
- C. PIM neighbor 10.1.5.6 is the winner of an assert mechanism.
- D. The RPF neighbor 10.1.5.6 is invalid.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

show ip mroute Field Descriptions

Field	Description
RPF neighbor or RPF nbr	IP address of the upstream router to the source. Tunneling indicates that this router is sending data to the RP encapsulated in register packets. The hexadecimal number in parentheses indicates to which RP it is registering. Each bit indicates a different RP if multiple RPs per group are used. If an asterisk (*) appears after the IP address in this field, the RPF neighbor has been learned through an assert.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_2/ipmulti/command/reference/fiprmc_r/rfmult3.html

QUESTION 81

Refer to the exhibit.


```
interface Ethernet0/1
 ip address 110.100.1.4 255.255.255.0
!
router ospf 100
 router-id 4.4.4.4
 redistribute static metric-type 1 subnets tag 704
 network 110.110.0.0 0.0.255.255 area 110
!
ip route 192.168.10.0 255.255.255.0 Ethernet0/1 110.100.1.1
!
```

External LSA:

OSPF Router with ID (4.4.4.4) (Process ID 100)

Type-5 AS External Link States

LS age: 101
Options: (No TOS-capability, DC, Upward)
LS Type: AS External Link
Link State ID: 192.168.10.0 (External Network Number)
Advertising Router: 4.4.4.4
LS Seq Number: 80000084
Checksum: 0x74E2
Length: 36
Network Mask: /24
Metric Type: 1 (Comparable directly to link state metric)
MTID: 0
Metric: 20
Forward Address: 0.0.0.0
External Route Tag: 704

Which option explains why the forwarding address is set to 0.0.0.0 instead of 110.100.1.1?

- A. The interface Ethernet0/1 is in down state.
- B. The next-hop ip address 110.100.1.1 is not directly attached to the redistributing router.
- C. The next-hop interface (Ethernet0/1) is specified as part of the static route command; therefore, the forwarding address is always set to 0.0.0.0.

D. OSPF is not enabled on the interface Ethernet0/1.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

From the output of the "show ip ospf database" command (although this command is not shown) we can conclude this is an ASBR (with Advertising Router is itself) and E0/1 is the ASBR's next hop interface for other routers to reach network 192.168.10.0.

The Forwarding Address is determined by these conditions:

* The forwarding address is set to 0.0.0.0 if the ASBR redistributes routes and OSPF is not enabled on the next hop interface for those routes.

* These conditions set the forwarding address field to a non-zero address:

+ OSPF is enabled on the ASBR's next hop interface AND

+ ASBR's next hop interface is non-passive under OSPF AND

+ ASBR's next hop interface is not point-to-point AND

+ ASBR's next hop interface is not point-to-multipoint AND + ASBR's next hop interface address falls under the network range specified in the router ospf command.

* Any other conditions besides these set the forwarding address to 0.0.0.0. -> We can see E0/1 interface is not running OSPF because it does not belong to network 110.110.0.0 0.0.255.255 which is declared under OSPF process -> F.A address is set to 0.0.0.0.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13682-10.html>

QUESTION 82

Refer to the exhibit.

```
Hub2#sh ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(123)
H   Address          Interface          Hold Uptime    SRTT   RTO  Q  Seq
                               (sec)          (ms)                Cnt  Num
0   192.168.0.2       Et0/3             11 01:49:56     1   3000  0   1
Hub2#sh ip ospf neighbor

Neighbor ID   Pri  State           Dead Time   Address        Interface
192.168.0.2   1    FULL/DR         00:00:31    192.168.0.2    Ethernet0/3
```

You have configured two routing protocols across this point-to-point link. How many BFD sessions will be established across this link?

A. three per interface

- B. one per multicast address
- C. one per routing protocol
- D. one per interface

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Cisco devices will use one Bidirectional Forwarding Detection (BFD) session for multiple client protocols in the Cisco implementation of BFD for Cisco IOS Releases 12.2(18)SXE, 12.0(31)S, and 12.4(4)T. For example, if a network is running OSPF and EIGRP across the same link to the same peer, only one BFD session will be established, and BFD will share session information with both routing protocols.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html#wp1053749

QUESTION 83

Refer to the exhibit.

```
R1#show ipv6 route

C   2001:DB8::/64 [0/0]
    via Ethernet0/0, directly connected
L   2001:DB8::1/128 [0/0]
    via Ethernet0/0, receive
```

Which statement is true?

- A. 2001:DB8::1/128 is a local host route, and it can be redistributed into a dynamic routing protocol.
- B. 2001:DB8::1/128 is a local host route, and it cannot be redistributed into a dynamic routing protocol.
- C. 2001:DB8::1/128 is a local host route that was created because ipv6 unicast-routing is not enabled on this router.
- D. 2001:DB8::1/128 is a route that was put in the IPv6 routing table because one of this router's loopback interfaces has the IPv6 address 2001:DB8::1/128.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

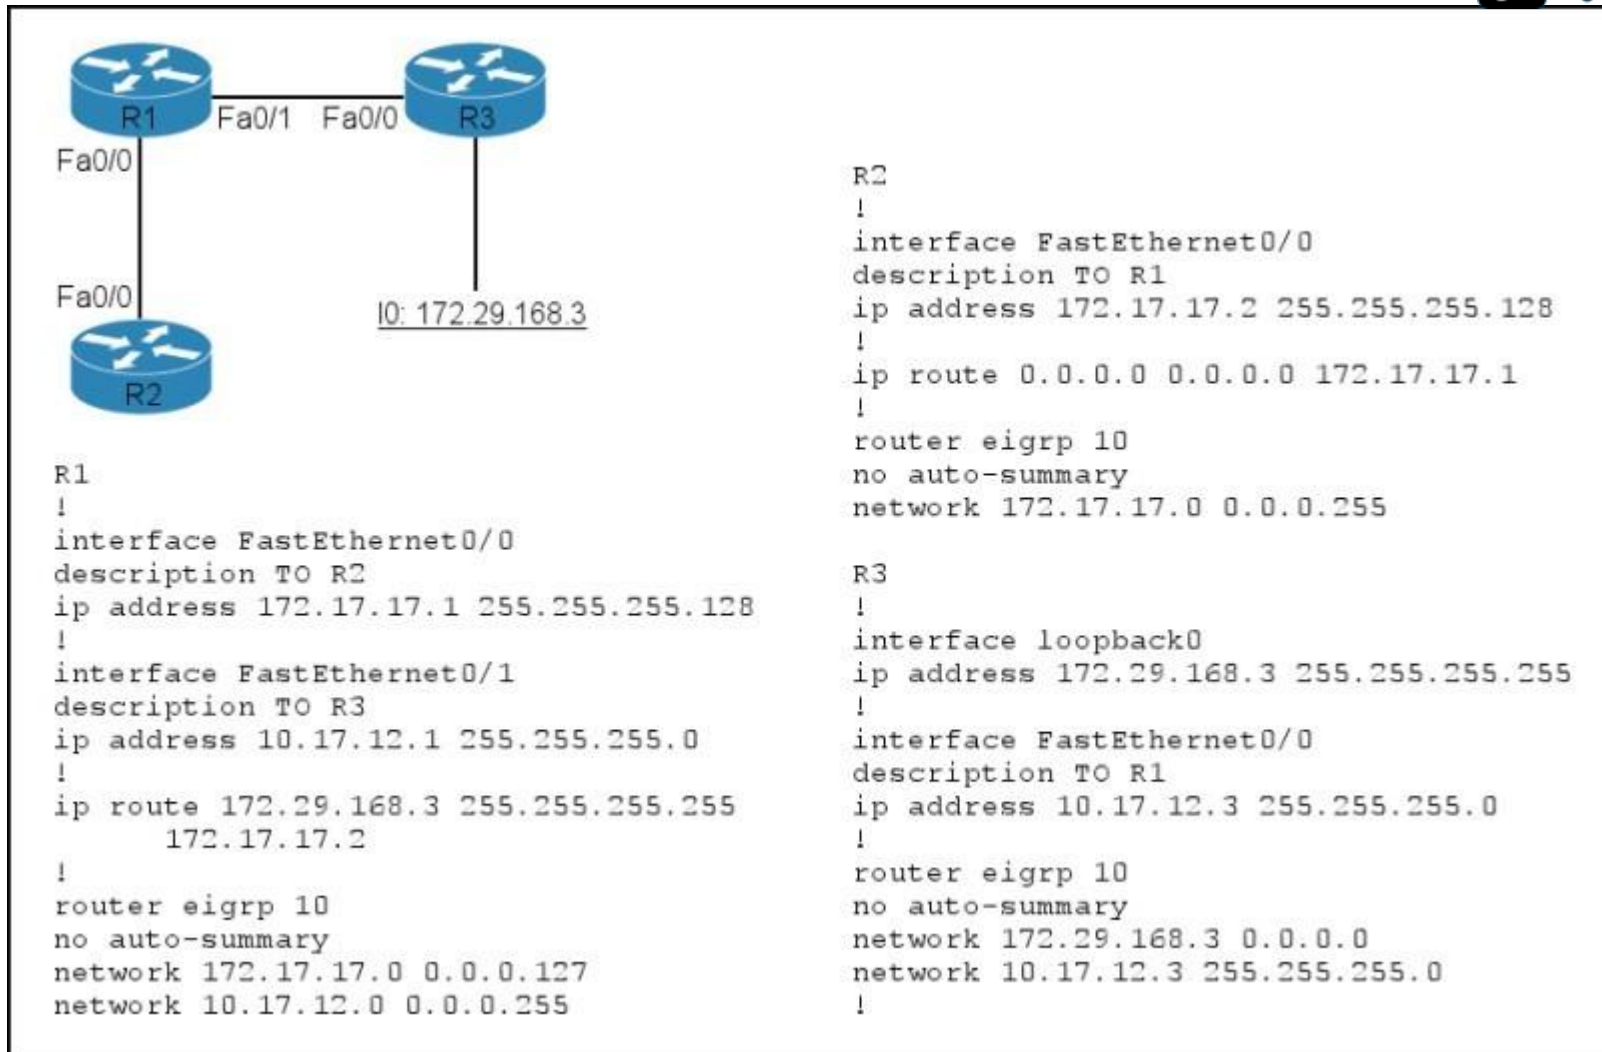
Explanation:

The local routes have the administrative distance of 0. This is the same administrative distance as connected routes. However, when you configure redistributed connected under any routing process, the connected routes are redistributed, but the local routes are not. This behavior allows the networks to not require a large number of host routes, because the networks of the interfaces are advertised with their proper masks. These host routes are only needed on the router that owns the IP address in order to process packets destined to that IP address. It is normal for local host routes to be listed in the IPv4 and IPv6 routing table for IP addresses of the router's interfaces. Their purpose is to create a corresponding CEF entry as a receive entry so that the packets destined to this IP address can be processed by the router itself. These routes cannot be redistributed into any routing protocol.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/ip-routing/116264-technote-ios-00.html>

QUESTION 84

Refer to the exhibit.



Routers R1, R2, and R3 are configured as shown, and traffic from R2 fails to reach 172.29.168.3. Which action can you take to correct the problem?

- A. Correct the static route on R1.
- B. Correct the default route on R2.
- C. Edit the EIGRP configuration of R3 to enable auto-summary.

D. Correct the network statement for 172.29.168.3 on R3.

Correct Answer: A

Section: Layer 3 Technologies

Explanation


Explanation/Reference:

Explanation:

On R1 we see there is a wrongly configured static route: ip route 172.29.168.3 255.255.255.255 172.17.17.2. It should be ip route 172.29.168.3 255.255.255.255 10.17.12.3.

QUESTION 85

Refer to the exhibit.



R3 prefers the path through R1 to reach host 10.1.1.1.

Which option describes the reason for this behavior?

- A. The OSPF reference bandwidth is too small to account for the higher speed links through R2.
- B. The default OSPF cost through R1 is less than the cost through R2.
- C. The default OSPF cost through R1 is more than the cost through R2.
- D. The link between R2 and R1 is congested.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The default formula to calculate OSPF bandwidth is $BW = \text{Bandwidth Reference} / \text{interface bandwidth [bps]} = 10^8 / \text{interface bandwidth [bps]}$

BW of the R1-R3 link = $10^8 / 100\text{Mbps} = 10^8 / 10^8 = 1$

BW of the R2-R3 link = $10^8 / 1\text{Gbps} = 10^8 / 10^9 = 1$ (round up)

Therefore OSPF considers the two above links have the same Bandwidth -> R3 will go to 10.1.1.1 via the R1-R3 link. The solution here is to increase the Bandwidth Reference to a higher value using the "auto-cost reference-bandwidth" command under OSPF router mode. For example:

```
Router(config)#router ospf 1
```

```
Router(config-router)#auto-cost reference-bandwidth 10000
```

This will increase the reference bandwidth to 10000 Mbps which increases the BW of the R2-R3 link to $10^{10} / 10^8 = 100$.

QUESTION 86

Refer to the exhibit.

```
*>172.21.95.0/22 172.17.192.1 0 120 0 65534 65535 65100 65235 ?
```

For which reason could a BGP-speaking device in autonomous system 65534 be prevented from installing the given route in its BGP table?

- A. The AS number of the BGP is specified in the given AS_PATH.
- B. The origin of the given route is unknown.
- C. BGP is designed only for publicly routed addresses.
- D. The AS_PATH for the specified prefix exceeds the maximum number of ASs allowed.
- E. BGP does not allow the AS number 65535.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

BGP is considered to be a 'Path Vector' routing protocol rather than a distance vector routing protocol since it utilises a list of AS numbers to describe the path that a packet should take. This list is called the AS_PATH. Loops are prevented because if a BGP speaking router sees it's own AS in the AS_PATH of a route it rejects the route.

QUESTION 87

Which statement about the feasibility condition in EIGRP is true?

- A. The prefix is reachable via an EIGRP peer that is in the routing domain of the router.
- B. The EIGRP peer that advertises the prefix to the router has multiple paths to the destination.
- C. The EIGRP peer that advertises the prefix to the router is closer to the destination than the router.
- D. The EIGRP peer that advertises the prefix cannot be used as a next hop to reach the destination.

Correct Answer: C

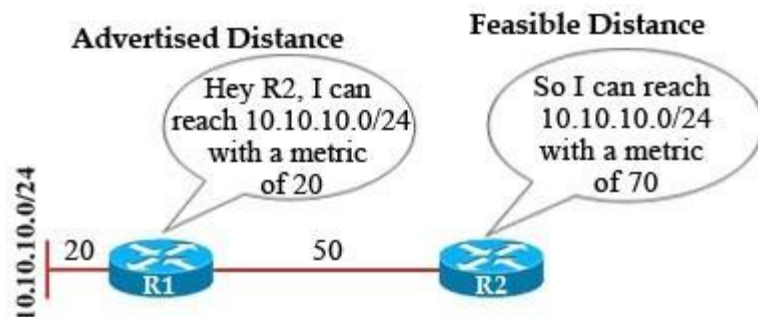
Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The advertised metric from an EIGRP neighbor (peer) to the local router is called Advertised Distance (or reported distance) while the metric from the local router to that network is called Feasible Distance. For example, R1 advertises network 10.10.10.0/24 with a metric of 20 to R2. For R2, this is the advertised distance. R2 calculates the feasible distance by adding the metric from the advertised router (R1) to itself. So in this case the feasible distance to network 10.10.10.0/24 is $20 + 50 = 70$.



Before a router can be considered a feasible successor, it must pass the feasibility condition rule. In short, the feasibility condition says that if we learn about a prefix from a neighbor, the advertised distance from that neighbor to the destination must be lower than our feasible distance to that same destination.

Therefore we see the Advertised Distance always smaller than the Feasible Distance to satisfy the feasibility condition.

QUESTION 88

Which two statements about the function of the stub feature in EIGRP are true? (Choose two.)

- A. It stops the stub router from sending queries to peers.
- B. It stops the hub router from sending queries to the stub router.
- C. It stops the stub router from propagating dynamically learned EIGRP prefixes to the hub routers.
- D. It stops the hub router from propagating dynamically learned EIGRP prefixes to the stub routers.

Correct Answer: BC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

When using the EIGRP Stub Routing feature, you need to configure the distribution and remote routers to use EIGRP, and to configure only the remote router as a stub. Only specified routes are propagated from the remote (stub) router. The router responds to queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message "inaccessible." A router that is configured as a stub will send a special peer information packet to all neighboring routers to report its status as a stub router.

Without the stub feature, even after the routes that are sent from the distribution router to the remote router have been filtered or summarized, a problem might occur. If a route is lost somewhere in the corporate network, EIGRP could send a query to the distribution router, which in turn will send a query to the remote router even if routes are being summarized. If there is a problem communicating over the WAN link between the distribution router and the remote router, an EIGRP stuck in active (SIA) condition could occur and cause instability elsewhere in the network. The EIGRP Stub Routing feature allows a network administrator to prevent queries from being sent to the remote router.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/eigrpstb.html

QUESTION 89

In which type of EIGRP configuration is EIGRP IPv6 VRF-Lite available?

- A. stub
- B. named mode
- C. classic mode
- D. passive

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The EIGRP IPv6 VRF Lite feature provides EIGRP IPv6 support for multiple VRFs. EIGRP for IPv6 can operate in the context of a VRF. The EIGRP IPv6 VRF Lite feature provides separation between routing and forwarding, providing an additional level of security because no communication between devices belonging to different VRFs is allowed unless it is explicitly configured. The EIGRP IPv6 VRF Lite feature simplifies the management and troubleshooting of traffic belonging to a specific VRF.

The EIGRP IPv6 VRF Lite feature is available only in EIGRP named configurations.

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ipv6-15-2mt-book/ip6-eigrp.html#GUID-92B4FF4F-2B68-41B0-93C8-AAA4F0EC1B1B>

QUESTION 90

Two routers are trying to establish an OSPFv3 adjacency over an Ethernet link, but the adjacency is not forming. Which two options are possible reasons that prevent OSPFv3 to form between these two routers? (Choose two.)

- A. mismatch of subnet masks
- B. mismatch of network types
- C. mismatch of authentication types
- D. mismatch of instance IDs
- E. mismatch of area types

Correct Answer: DE

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

An OSPFv3 interface must have a compatible configuration with a remote interface before the two can be considered neighbors. The two OSPFv3 interfaces must match the following criteria:

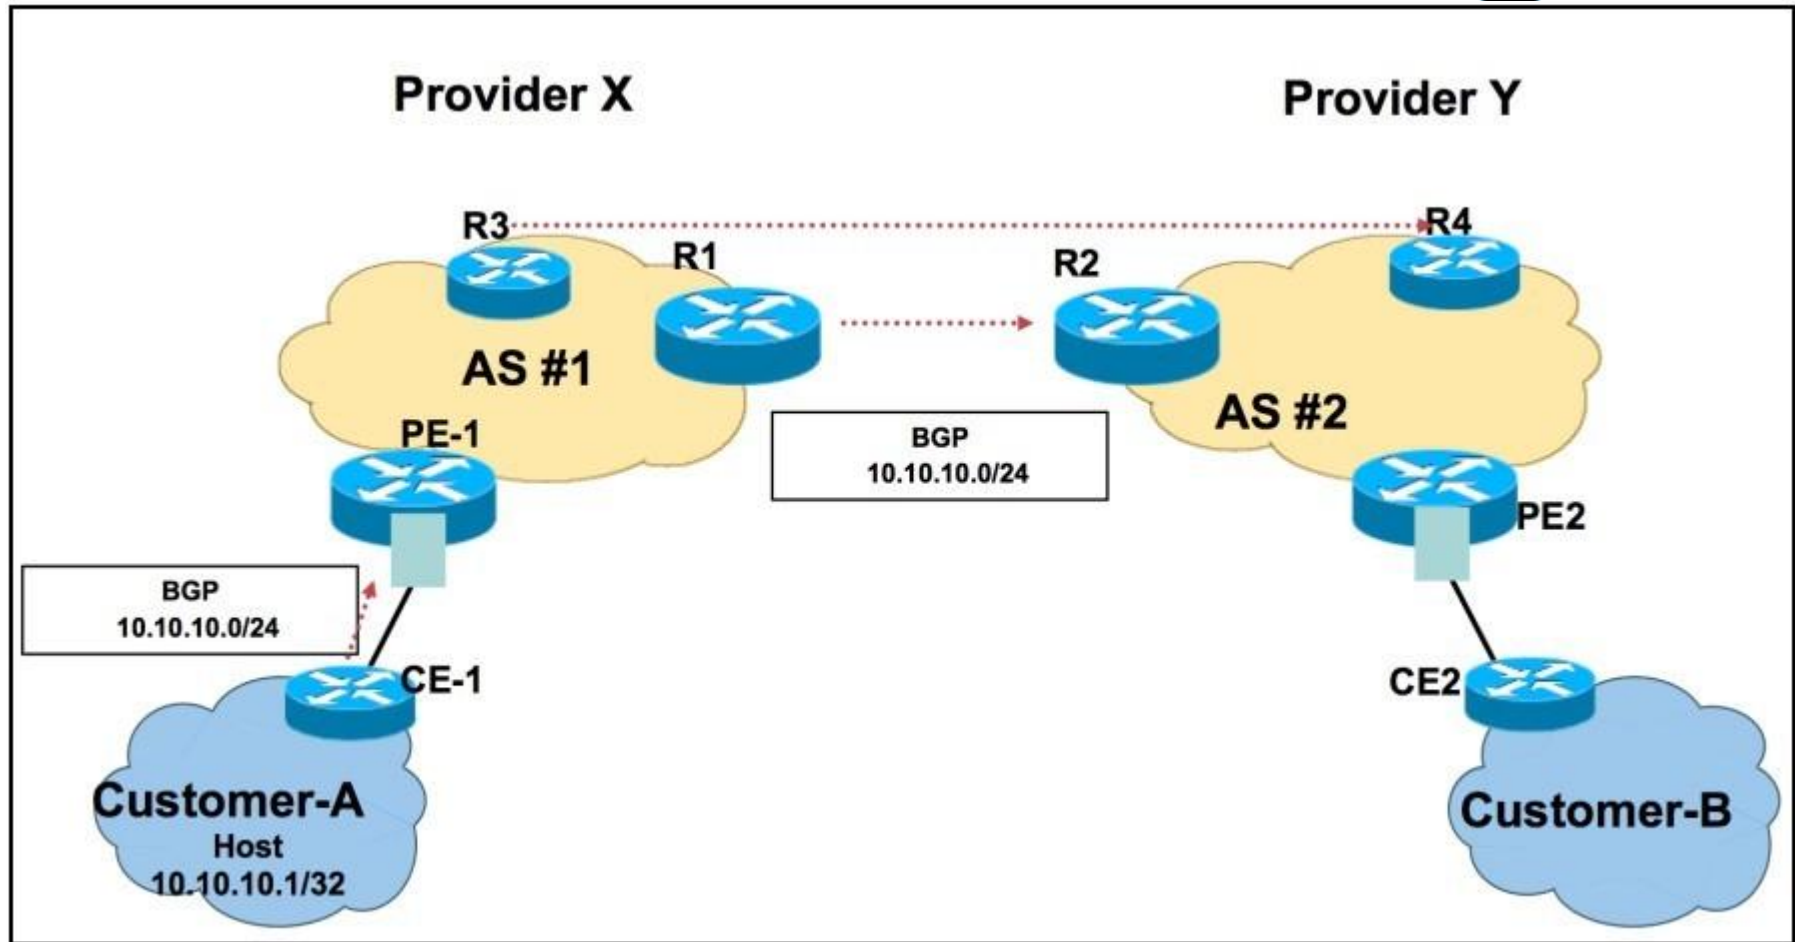
- . Hello interval
- . Dead interval
- . Area ID
- . Optional capabilities

The OSPFv3 header includes an instance ID field to identify that OSPFv3 packet for a particular OSPFv3 instance. You can assign the OSPFv3 instance. The interface drops all OSPFv3 packets that do not have a matching OSPFv3 instance ID in the packet header.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/unicast/configuration/guide/l3_cli_nxos/l3_ospfv3.html

QUESTION 91

Refer to the exhibit.



AS #1 and AS #2 have multiple EBGP connections with each other. AS #1 wants all return traffic that is destined to the prefix 10.10.10.1/32 to enter through the router R1 from AS #2. In order to achieve this routing policy, the AS #1 advertises a lower MED from R1, compared to a higher MED from R3, to their respective BGP neighbor for the prefix 10.10.10.0/24. Will this measure guarantee that the routing policy is always in effect?

- A. Yes, because MED plays a deterministic role in return traffic engineering in BGP.
- B. Yes, because a lower MED forces BGP best-path route selection in AS #2 to choose R1 as the best path for 10.10.10.0/24.
- C. Yes, because a lower MED in AS #2 is the highest BGP attribute in BGP best-path route selection.
- D. No, AS #2 can choose to alter the weight attribute in R2 for BGP neighbor R1, and this weight value is cascaded across AS #2 for BGP best-path route selection.

- E. No, AS #2 can choose to alter the local preference attribute to overwrite the best-path route selection over the lower MED advertisement from AS #1. This local preference attribute is cascaded across AS #2 for the BGP best-path route selection.

Correct Answer: E

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

MED and AS path prepending can both be used to influence the way incoming traffic from other Autonomous Systems get sent to the local AS, but they provide no guarantee as the other AS ultimately has the final word in how they send traffic. Since local preference is preferred over MED in the BGP decision process, the other AS can configure local preference to override the MED settings you have configured.

QUESTION 92

Refer to the exhibit.

```
R1>sh ip bgp 10.1.1.1
BGP routing table entry for 10.1.0.0/16, version 182
Paths: (2 available, best #1, table default, not advertised to EBGp peer)
  Advertised to update-groups:
    2
  Refresh Epoch 1
  50811 65112
    172.28.1.5 from 172.28.1.5 (192.168.236.222)
      Origin incomplete, localpref 800, valid, external, best
      Community: no-export
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 1
  50811 65112, (received-only)
    172.28.1.5 from 172.28.1.5 (192.168.236.222)
      Origin incomplete, localpref 100, valid, external
      Community: 65112:21147 50811:11145
      rx pathid: 0, tx pathid: 0
R1>
```

What does "(received-only)" mean?

- A. The prefix 10.1.1.1 can not be advertised to any eBGP neighbor.
- B. The prefix 10.1.1.1 can not be advertised to any iBGP neighbor.
- C. BGP soft reconfiguration outbound is applied.

D. BGP soft reconfiguration inbound is applied.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

When you configure bgp soft-configuration-inbound, all the updates received from the neighbor will be stored unmodified, regardless of the inbound policy, and these routes appear as "(received- only)."

QUESTION 93

Which regular expression will only allow prefixes that originated from AS 65000 and that are learned through AS 65001?

A. ^65000_65001\$

B. 65000_65001\$

C. ^65000_65001

D. ^65001_65000\$

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

The following table lists the regular expressions and their meanings:

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

Some commonly used regular expressions include:

Expression	Meaning
.*	Anything
^\$	Locally originated routes
^100_	Learned from AS 100
_100\$	Originated in AS 100
100	Any instance of AS 100
^[0-9]+\$	Directly connected ASes

Reference. <http://blog.ine.com/2008/01/06/understanding-bgp-regular-expressions/>

QUESTION 94

Which statement describes the BGP add-path feature?

- A. It allows for installing multiple IBGP and EBGP routes in the routing table.
- B. It allows a network engineer to override the selected BGP path with an additional path created in the config.
- C. It allows BGP to provide backup paths to the routing table for quicker convergence.
- D. It allows multiple paths for the same prefix to be advertised.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

BGP routers and route reflectors (RRs) propagate only their best path over their sessions. The advertisement of a prefix replaces the previous announcement of that prefix (this behavior is known as an implicit withdraw). The implicit withdraw can achieve better scaling, but at the cost of path diversity.

Path hiding can prevent efficient use of BGP multipath, prevent hitless planned maintenance, and can lead to MED oscillations and suboptimal hot-potato routing. Upon nexthop failures, path hiding also inhibits fast and local recovery because the network has to wait for BGP control plane convergence to restore traffic. The BGP Additional Paths feature provides a generic way of offering path diversity; the Best External or Best Internal features offer path diversity only in limited scenarios.

The BGP Additional Paths feature provides a way for multiple paths for the same prefix to be advertised without the new paths implicitly replacing the previous paths. Thus, path diversity is achieved instead of path hiding.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xr-3s/irg-xr-3s-book/irg-additional-paths.html

QUESTION 95

Refer to the exhibit.


```
R1#show bgp ipv4 unicast 10.100.1.1/32
BGP routing table entry for 10.100.1.1/32, version 8
Paths: (2 available, best #1, table default, RIB-failure(17))
  Advertised to update-groups:
    2
  Refresh Epoch 2
  4
    10.1.3.4 from 10.1.3.4 (10.100.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 2
  5 4
    10.1.5.5 from 10.1.5.5 (10.1.5.5)
      Origin IGP, localpref 100, valid, external
      rx pathid: 0, tx pathid: 0
```

What is a reason for the RIB-failure?

- A. CEF is not enabled on this router.
- B. The route 10.100.1.1/32 is in the routing table, but not as a BGP route.
- C. The routing table has yet to be updated with the BGP route.
- D. The BGP route is filtered inbound and hence is not installed in the routing table.

Correct Answer: B

Section: Layer 3 Technologies

Explanation


Explanation/Reference:

Explanation:

A rib-failure occurs when BGP tries to install the bestpath prefix into the RIB, but the RIB rejects the BGP route because a route with better administrative distance already exists in the routing table. An inactive Border Gateway Protocol (BGP) route is a route that is not installed in the RIB, but is installed in the BGP table as rib-failure.

Example Topology

Router 1 (R1) and router 2 (R2) have two parallel links; one link runs BGP AS 65535 and the other link runs Enhanced Interior Gateway Routing Protocol (EIGRP) AS 1. Both BGP and EIGRP are advertising the network 10.1.1.1/32 on R1.



R2 learns about the 1.1.1.1/32 route through both EIGRP and BGP, but installs only the EIGRP route in the routing table because of the lower administrative distance. Since the BGP route is not installed in the R2 routing table, the route appears as a rib-failure in the R2 BGP table.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/116146-config-bgp-next-hop-00.html>

QUESTION 96

Refer to the exhibit.

```
R1#show bgp ipv4 unicast summary
BGP router identifier 10.1.3.1, local AS number 1
BGP table version is 2, main routing table version 2
1 network entries using 144 bytes of memory
1 path entries using 80 bytes of memory
1/1 BGP path/bestpath attribute entries using 144 bytes of memory
1 BGP AS-PATH entries using 24 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
BGP using 392 total bytes of memory
BGP activity 1/0 prefixes, 1/0 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.2	4	2	69	69	2	0	0	01:00:54	0
10.1.2.3	4	3	69	70	1	0	0	01:00:45	0
10.1.3.4	4	4	72	70	2	0	0	01:01:12	1

Which statement is true?

- A. BGP peer 10.1.2.3 is performing inbound filtering.
- B. BGP peer 10.1.2.3 is a route reflector.
- C. R1 is a route reflector, but BGP peer 10.1.2.3 is not a route reflector client.
- D. R1 still needs to send an update to the BGP peer 10.1.2.3.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

On R1 the routing table version (Tbl Ver) for 10.1.2.3 is 1, other routers have version 2, so it needs to send an update to the 10.1.2.3 peer.

QUESTION 97

Refer to the exhibit.

```
RouterA#  
conf t  
router isis  
net 49.5200.1580.3500.6002.00  
  
RouterB#  
conf t  
router isis 1  
net 49.5200.1580.3500.6002.00
```

Router A and router B are physically connected over an Ethernet interface, and ISIS is configured as shown. Which option explains why the ISIS neighborhood is not getting formed between router A and router B?

- A. same area ID
- B. same N selector
- C. same domain ID
- D. same system ID

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

With IS-IS, the LSP identifier is derived from the system ID (along with the pseudonode ID and LSP number). Each IS is usually configured with one NET and in one area; each system ID within an area must be unique.

The big difference between NSAP style addressing and IP style addressing is that, in general, there will be a single NSAP address for the entire router, whereas with IP there will be one IP address per interface. All ISs and ESs in a routing domain must have system IDs of the same length. All routers in an area must have the same area address. All Level 2 routers must have a unique system ID domain-wide, and all Level 1 routers must have a unique system ID area-wide.

Reference.

http://www.cisco.com/en/US/products/ps6599/products_white_paper09186a00800a3e6f.shtml

QUESTION 98

Refer to the exhibit.

```
R4#show isis database R4.00-00 detail
IS-IS Level-2 LSP R4.00-00
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
R4.00-00             * 0x000022BE  0xD36A        1194          0/0/0
  Area Address: 49.0001
  NLPID:         0x81 0xCC 0x8E
  Hostname: R4
  IP Address:    10.1.100.4
  IPv6 Address:  2001:100::1:4
  Metric: 10     IS-Extended R3.00
  Metric: 10     IS-Extended R5.03
  Metric: 10     IP 10.1.1.0/24
  Metric: 10     IP 10.1.2.0/24
  Metric: 10     IP 10.1.3.0/24
  Metric: 10     IP 10.1.100.4/32
  Metric: 50     IP 10.200.200.200/32
  Metric: 10     IPv6 2001:1::1:0/112
  Metric: 10     IPv6 2001:1::2:0/112
  Metric: 10     IPv6 2001:100::1:4/128
```

Which statement is true?

- A. IS-IS has been enabled on R4 for IPv6, single-topology.
- B. IS-IS has been enabled on R4 for IPv6, multitopology.
- C. IS-IS has been enabled on R4 for IPv6, single-topology and multitopology.
- D. R4 advertises IPv6 prefixes, but it does not forward IPv6 traffic, because the protocol has not been enabled under router IS-IS.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

When working with IPv6 prefixes in IS-IS, you can configure IS-IS to be in a single topology for both IPv4 and IPv6 or to run different topologies for IPv4 and IPv6.


By default, IS-IS works in single-topology mode when activating IPv4 and IPv6. This means that the IS-IS topology will be built based on IS Reachability TLVs. When the base topology is built, then IPv4 prefixes (IP Reachability TLV) and IPv6 prefixes (IPv6 Reachability TLV) are added to each node as

leaves, without checking if there is IPv6 connectivity between nodes.

Reference: <https://blog.initialdraft.com/archives/3381/>

QUESTION 99

Refer to the exhibit.



Why is the neighbor relationship between R2 and R4 shown as ES-IS?

- A. because there is an MTU mismatch between R2 and R4
- B. because interface S3/0 of R4 is configured as L1/L2
- C. because interface S3/0 of R2 is configured as L1
- D. because there is a hello interval mismatch between R2 and R4

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

With IS-IS we will see ES-IS when one of the following is true:

1. One side is configured to send only L2 and another side is configured to send L1. In this case both sides show each-other as ES-IS.
2. There is an MTU Mismatch so we see ES-IS in only one side.

So in this question because we do not know about the other side's "show CLNS neighbor" A must be the better choose.

QUESTION 100

Refer to the exhibit.

```
R4
interface FastEthernet0/1
ip address 192.168.2.1 255.255.255.0
ip pim sparse-dense-mode
duplex auto
speed auto
standby 1 ip 192.168.2.4
standby 1 priority 150
standby 1 preempt

R5
interface FastEthernet0/1
ip address 192.168.2.2 255.255.255.0
ip pim sparse-dense-mode
duplex auto
speed auto
standby 1 ip 192.168.2.4
```

The interface FastEthernet0/1 of both routers R4 and R5 is connected to the same Ethernet segment with a multicast receiver. Which two statements are true? (Choose two)

- A. Multicast traffic that is destined to a receiver with IP address 192.168.2.6 will flow through router R4.
- B. Both routers R4 and R5 will send PIM join messages to the RP.
- C. Only router R5 will send a multicast join message to the RP.
- D. Multicast traffic that is destined to a receiver with IP address 192.168.2.6 will flow through router R5.

Correct Answer: CD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Even though R4 is the active HSRP router, traffic will flow through R5 and only R5 will send the join messages. The Multicast DR is elected by the higher IP address or priority. R5 has 192.168.2.2 and R4 has 192.168.2.1. R5 is the DR which send all packets to the RP.

QUESTION 101

Refer to the exhibit.

```
router ospf 100
router-id 4.4.4.4
area 110 nssa
summary-address 192.168.0.0 255.255.0.0 nssa-only
redistribute static metric-type 1 subnets tag 704
network 110.110.0.0 0.0.255.255 area 110
```

This is the configuration of the ASBR of area 110. Which option explains why the remote ABR should not translate the type 7 LSA for the prefix 192.168.0.0/16 into a type 5 LSA?

- A. The remote ABR translates all type 7 LSA into type 5 LSA, regardless of any option configured in the ASBR.
- B. The ASBR sets the forwarding address to 0.0.0.0 which instructs the ABR not to translate the LSA into a type 5 LSA.
- C. The ASBR originates a type 7 LSA with age equal to MAXAGE 3600.
- D. The ABR clears the P bit in the header of the type 7 LSA for 192.168.0.0/16.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

When external routing information is imported into an NSSA, LSA Type 7 is generated by the ASBR and it is flooded within that area only. To further distribute the external information, type 7 LSA is translated into type 5 LSA at the NSSA border. The P-bit in LSA Type 7 field indicates whether the type 7 LSA should be translated. This P-bit is automatically set by the NSSA ABR (also the Forwarding Address (FA) is copied from Type 7 LSA). The P-bit is not set only when the NSSA ASBR and NSSA ABR are the same router for the area. If bit P = 0, then the NSSA ABR must not translate this LSA into Type 5.

The nssa-only keyword instructs the device to instigate Type-7 LSA with cleared P-bit, thereby, preventing LSA translation to Type 5 on NSSA ABR device.

Note. If a router is attached to another AS and is also an NSSA ABR, it may originate both a type-5 and a type-7 LSA for the same network. The type-5 LSA will be flooded to the backbone and the type-7 will be flooded into the NSSA. If this is the case, the P-bit must be reset (P=0) in the type-7 LSA so the type-7 LSA isn't again translated into a type-5 LSA by another NSSA ABR.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/15-e/iro-15-e-book/iro-ospfv3-nssa-cfg.html

QUESTION 102

What is the function of an EIGRP sequence TLV packet?

- A. to acknowledge a set of sequence numbers during the startup update process
- B. to list the peers that should listen to the next multicast packet during the reliable multicast process
- C. to list the peers that should not listen to the next multicast packet during the reliable multicast process
- D. to define the initial sequence number when bringing up a new peer

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

EIGRP sends updates and other information between routers using multicast packets to 224.0.0.10. For example in the topology below, R1 made a change in the topology and it needs to send updates to R2 & R3. It sends multicast packets to EIGRP multicast address 224.0.0.10. Both R2 & R3 can receive the updates and acknowledge back to R1 using unicast. Simple, right? But what if R1 sends out updates, only R2 replies but R3 never does? In the case a router sends out a multicast packet that must be reliably delivered (like in this case), an EIGRP process will wait until the RTO (retransmission timeout) period has passed before beginning a recovery action. This period is calculated from the SRTT (smooth round-trip time). After R1 sends out updates it will wait for this period to expire. Then it makes a list of all the neighbors from which it did not receive an Acknowledgement (ACK). Next it sends out a packet telling these routers stop listening to multicast until they are been notified that it is safe again. Finally the router will begin sending unicast packets with the information to the routers that didn't answer, continuing until they are caught up. In our example the process will be like this:

1. R1 sends out updates to 224.0.0.10
2. R2 responds but R3 does not
3. R1 waits for the RTO period to expire
4. R1 then sends out an unreliable-multicast packet, called a sequence TLV (Type-Length-Value) packet, which tells R3 not to listen to multicast packets any more
5. R1 continues sending any other multicast traffic it has and delivering all traffic, using unicast to R3, until it acknowledges all the packets
6. Once R3 has caught up, R1 will send another sequence TLV, telling R3 to begin listening to multicast again.

The sequence TLV packet contains a list of the nodes that should not listen to multicast packets while the recovery takes place. But notice that the TLV packet in step 6 does not contain any nodes in the list.

Note. In the case R3 still does not reply in step 4, R1 will attempt to retransmit the unicast 16 times or continue to retransmit until the hold time for the neighbor in question expires. After this time, R1 will declare a retransmission limit exceeded error and will reset the neighbor.

(Reference. EIGRP for IP: Basic Operation and Configuration)

QUESTION 103

What are two reasons to define static peers in EIGRP? (Choose two.)

- A. Security requirements do not allow dynamic learning of neighbors.
- B. The link between peers requires multicast packets.
- C. Back-level peers require static definition for successful connection.
- D. The link between peers requires unicast packets.

Correct Answer: AD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:


Explanation:

There are two ways we can create EIGRP neighbor relationship:

- + Use "network " command. This is the more popular way to create EIGRP neighbor relationship. That router will check which interfaces whose IP addresses belong to the and turn EIGRP on that interface. EIGRP messages are sent via multicast packets.
- + Use "neighbor" command. The interface(s) that have this command applied no longer send or receive EIGRP multicast packets. EIGRP messages are sent via unicast. The router only accepts EIGRP packets from peers that are explicitly configured with a neighbor statement. Consequently, any messages coming from routers without a corresponding neighbor statement are discarded. This helps prevent the insertion of unauthorized routing peers -> A and D are correct.

QUESTION 104

Refer to the exhibit.



R2 is mutually redistributing between EIGRP and BGP.

Which configuration is necessary to enable R1 to see routes from R3?

- A. The R3 configuration must include ebgp-multihop to the neighbor statement for R2.

- B. The R2 BGP configuration must include bgp redistribute-internal.
- C. R1 must be configured with next-hop-self for the neighbor going to R2.
- D. The AS numbers configured on R1 and R2 must match.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Whenever you redistribute from BGP to something else, BGP will only advertise externally learned routes. To allow the redistribution of iBGP routes into an interior gateway protocol such as EIGRP or OSPF, use the bgp redistribute-internal command in router configuration mode.

QUESTION 105

What is the purpose of EIGRP summary leaking?

- A. to allow a summary to be advertised conditionally on specific criteria
- B. to allow a component of a summary to be advertised in addition to the summary
- C. to allow overlapping summaries to exist on a single interface
- D. to modify the metric of the summary based on which components of the summary are operational

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

When you do manual summarization, and still you want to advertise some specific routes to the neighbor, you can do that using leak-map. Please read more about leaking routes here.

http://www.cisco.com/c/en/us/td/docs/ios/iproute_eigrp/command/reference/ire_book/ire_i1.html#wp1037685.

QUESTION 106

Refer to the exhibit.

```
*May20 12:16: BGP(4):10.1.1.2 rcvd UPDATE w/ attr:nexthop 10.1.1.2,origin ?, localpref 100,metric 0,extended community RT:999:999
*May20 12:16: BGP(4):10.1.1.2 rcvd 999:999:192.168.1.99/32,label 29--DENIED due to:extended community not supported
```

You have just created a new VRF on PE3. You have enabled debug ip bgp vpnv4 unicast updates on PE1, and you can see the route in the debug, but not in the BGP VPNv4 table. Which two statements are true? (Choose two.)

- A. VPNv4 is not configured between PE1 and PE3.
- B. address-family ipv4 vrf is not configured on PE3.
- C. After you configure route-target import 999:999 for a VRF on PE3, the route will be accepted.
- D. PE1 will reject the route due to automatic route filtering.
- E. After you configure route-target import 999:999 for a VRF on PE1, the route will be accepted.

Correct Answer: DE

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The route target extended community for VPLS auto-discovery defines the import and export policies that a VPLS instance uses. The export route target sets an extended community attribute number that is appended to all routes that are exported from the VPLS instance. The import route target value sets a filter that determines the routes that are accepted into the VPLS instance. Any route with a value in its import route target contained in its extended attributes field matching the value in the VPLS instance's import route target are accepted. Otherwise the route is rejected.

QUESTION 107

Which two DHCP messages are always sent as broadcast? (Choose two.)

- A. DHCP OFFER
- B. DHCP DECLINE
- C. DHCP RELEASE
- D. DHCP REQUEST
- E. DHCP DISCOVER

Correct Answer: DE

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

DHCP discovery

The client broadcasts messages DHCPDISCOVER on the network subnet using the destination address 255.255.255.255 or the specific subnet broadcast address.

DHCP request

In response to the DHCP offer, the client replies with a DHCP request, broadcasts to the server, requesting the offered address. A client can receive DHCP offers from multiple servers, but it will accept only one DHCP offer.

Reference. http://en.wikipedia.org/wiki/Dynamic_Host_Configuration_Protocol

QUESTION 108

With which ISs will an ISIS Level 1 IS exchange routing information?

- A. Level 1 ISs
- B. Level 1 ISs in the same area
- C. Level 1 and Level 2 ISs
- D. Level 2 ISs

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:


Explanation:

IS-IS differs from OSPF in the way that "areas" are defined and routed between. IS-IS routers are designated as being: Level 1 (intra-area); Level 2 (inter area); or Level 1-2 (both). Level 2 routers are inter area routers that can only form relationships with other Level 2 routers. Routing information is exchanged between Level 1 routers and other Level 1 routers, and Level 2 routers only exchange information with other Level 2 routers. Level 1-2 routers exchange information with both levels and are used to connect the inter area routers with the intra area routers.

Reference. <http://en.wikipedia.org/wiki/IS-IS>

QUESTION 109

Refer to the exhibit.



Why is the neighbor relationship between R1 & R2 and R1 & R3 an L2-type neighborship?

- A. because the area ID on R1 is different as compared to the area ID of R2 and R3
- B. because the circuit type on those three routers is L1/L2
- C. because the network type between R1, R2, and R3 is point-to-point
- D. because the hello interval is not the same on those three routers

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

With IS-IS, an individual router is in only one area, and the border between areas is on the link that connects two routers that are in different areas. A Level 2 router may have neighbors in the same or in different areas, and it has a Level 2 link-state database with all information for inter-area routing. Level 2 routers know about other areas but will not have Level 1 information from its own area.

Reference.

http://www.cisco.com/en/US/products/ps6599/products_white_paper09186a00800a3e6f.shtml

QUESTION 110

Which three statements about the designated router election in IS-IS are true? (Choose three.)

- A. If the IS-IS DR fails, a new DR is elected.
- B. The IS-IS DR will preempt. If a new router with better priority is added, it just becomes active in the network.
- C. If there is a tie in DR priority, the router with a higher IP address wins.
- D. If there is a tie in DR priority, the router with a higher MAC address wins.
- E. If the DR fails, the BDR is promoted as the DR.
- F. The DR is optional in a point-to-point network.

Correct Answer: ABD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

DR/DIS election

- highest priority (0-127)
- highest mac address

Setting priority to 0 doesn't disable DIS election; use point-to-point to disable it.

There can be separate DRs for L1 and L2 adjacencies.

There is no backup DR. If the primary DR fails, a new DR is elected.

DR preemption is enabled by default.

Reference. <http://ccie-in-2-months.blogspot.com/2013/12/is-is-hints.html>

QUESTION 111

Which three elements compose a network entity title? (Choose three.)

- A. area ID
- B. domain ID
- C. system ID
- D. NSAP selector
- E. MAC address
- F. IP address

Correct Answer: ACD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

An IS (Intermediate system) is identified by an address known as a network access point (NASAP). The NSAP is divided up into three parts as specified

by ISO/IEC 10589:

Area address -- This field is of variable length, composed of high order octets, and it excludes the System ID and N-selector (NSEL) fields. This area address is associated with a single area within the routing domain.

System ID -- This field is 6 octets long and should be set to a unique value with Level 1 and Level

2. The system ID defines an end system (ES) or an IS in an area. You configure the area address and the system ID with the NET command. You can display the system ID with the show isis topology command.

NSEL -- This field is called the N-selector, also referred to as the NSAP, and it specifies the upper-layer protocol. The NSEL is the last byte of the NSAP and identifies a network service user. A network service user is a transport entity or the IS network entity itself. When the N-selector is set to zero, the entire NSAP is called a network entity title (NET).

A NET is an NSAP where the last byte is always the n-selector and is always zero. A NET can be from 8 to 20 bytes in length.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/iproute_isis/command/reference/irs_book/irs_is2.html

QUESTION 112

Which three statements about IS-IS are true? (Choose three.)

- A. IS-IS can be used only in the service provider network.
- B. IS-IS can be used to route both IP and CLNP.
- C. IS-IS has three different levels of authentication: interface level, process level, and domain level.
- D. IS-IS is an IETF standard.
- E. IS-IS has the capability to provide address summarization between areas.

Correct Answer: BCE

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Intermediate System to Intermediate System (IS-IS) was designed as the routing protocol for ISO's CLNP described in ISO 10589. IS-IS is a Link State routing protocol akin to OSPF and was developed by DEC for use with DECnet Phase V. It was originally thought that TCP/IP would gradually make way for the seven layer OSI architecture so an enhancement to IS-IS was developed called Integrated IS-IS also known as Dual IS-IS that could route both Connectionless-Mode Network Service (CLNS) as well as IP.

Cisco IOS supports IS-IS authentication on 3 different levels; between neighbors, area-wide, and domain-wide, where each can be used by themselves or together.

summary-address address mask {level-1 | level-1-2 | level-2} is used to configure IP address summarization.

References:

<http://www.rhyshaden.com/isis.htm>

http://mynetworkingwiki.com/index.php/Configuring_IS-IS

QUESTION 113

Which statement describes the function of the tracking object created by the track 10 ip route 192.168.99.0/24 reachability command?

- A. It tracks the reachability of route 192.168.99.0/24.
- B. It tracks the line protocol status of the interface on which route 192.168.99.0/24 is received.
- C. It tracks exactly 10 occurrences of route 192.168.99.0/24.
- D. It tracks the summary route 192.168.99.0/24 and all routes contained within.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	track object-number {interface interface-id line-protocol ip routing} ip route ip-address/prefix-length {metric threshold reachability} list {boolean {and or}} {threshold {weight percentage}}}	<p>(Optional) Create a tracking list to track the configured state and enter tracking configuration mode.</p> <ul style="list-style-type: none"> • ___ The <i>object-number</i> range is from 1 to 500. • ___ Enter interface interface-id to select an interface to track. • ___ Enter line-protocol to track the interface line protocol state or enter ip routing to track the interface IP routing state. • ___ Enter ip route ip-address/prefix-length to track the state of an IP route. • ___ Enter metric threshold to track the threshold metric or enter reachability to track if the route is reachable. The default up threshold is 254 and the default down threshold is 255. • ___ Enter list to track objects grouped in a list. Configure the list as described on the previous pages. <p>Note Repeat this step for each interface to be tracked.</p>

Reference:

http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3560/software/release/12-2_52_se/configuration/guide/3560scg/sweot.html

QUESTION 114

Which group of neighbors can be configured as a BGP peer group?

- A. a group of iBGP neighbors that have the same outbound route policies
- B. a group of iBGP and eBGP neighbors that have the same inbound distribute-list

- C. a group of eBGP neighbors in the same autonomous system that have different outbound route policies
- D. a group of iBGP neighbors that have different outbound route policies

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

You can group BGP neighbors who share the same outbound policies together in what is called a BGP peer group. Instead of configuring each neighbor with the same policy individually, a peer group allows you to group the policies which can be applied to individual peers thus making efficient update calculation along with simplified configuration.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13755-29.html>

QUESTION 115

Refer to the exhibit.

```
BGP(0): 10.1.3.4 rcvd UPDATE w/ attr: nexthop 10.1.3.4, origin i,
metric 0, merged path 4, AS_PATH
BGP(0): 10.1.3.4 rcvd 10.100.1.1/32...duplicate ignored
```

Notice that debug ip bgp updates have been enabled. What can you conclude from the debug output?

- A. This is the result of the clear ip bgp 10.1.3.4 in command.
- B. This is the result of the clear ip bgp 10.1.3.4 out command.
- C. BGP neighbor 10.1.3.4 performed a graceful restart.
- D. BGP neighbor 10.1.3.4 established a new BGP session.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

If you enter the **clear ip bgp out** command for a BGP peer, that router resends its BGP prefixes to that peer. This does not cause a change in the best path on the receiving BGP peer. Hence, there is no change in the Table Version on that peer.

When you run the **debug ip bgp updates** on the receiving router, you see:

```
BGP(0): 10.1.3.4 rcvd UPDATE w/ attr: nexthop 10.1.3.4, origin i, metric 0, merged path 4, AS_PATH
```

```
BGP(0): 10.1.3.4 rcvd 10.100.1.1/32...duplicate ignored
```

The received update is recognized as a duplicate, so it is ignored and no best path change occurs.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/116511-technote-tableversion-00.html>

QUESTION 116

Which set of commands conditionally advertises 172.16.0.0/24 as long as 10.10.10.10/32 is in the routing table?

A.

```
neighbor x.x.x.x advertise-map ADV exist-map EXT
route-map ADV
  match IP address prefix-list ADV
!
route-map EXT
  match IP address prefix-list EXT
!
ip prefix-list EXT permit 172.16.0.0/24
!
ip prefix-list ADV permit 10.10.10.10/32
```

B.

```
neighbor x.x.x.x advertise-map ADV exist-map EXT
route-map ADV
  match IP address prefix-list ADV
!
route-map EXT
  match IP address prefix-list EXT
!
ip prefix-list ADV permit 172.16.0.0/24
!
ip prefix-list EXT permit 10.10.10.10/32
```

C.

```
neighbor x.x.x.x advertise-map ADV
route-map ADV
  match IP address prefix-list ADV
  match IP address prefix-list EXT
!
ip prefix-list ADV permit 172.16.0.0/24
!
ip prefix-list EXT permit 10.10.10.10/32
```

D.

```
neighbor x.x.x.x exist-map EXT
route-map EXT
  match IP address prefix-list ADV
  match IP address prefix-list EXT
!
ip prefix-list ADV permit 172.16.0.0/24
!
ip prefix-list EXT permit 10.10.10.10/32
```

Correct Answer: B

Section: Layer 3 Technologies

Explanation


Explanation/Reference:

Explanation:

Advertise maps are used for conditional routing to advertise specified prefixes if something which is specified in exist map exists. In our question we need to advertise 172.16.0.0/24 if 10.10.10.10/32 exists in the routing table so we have to use command. "neighbor x.x.x.x advertise- map <prefix-list of 172.16.0.0/24> exist-map <prefix-list of 10.10.10.10/32>". Therefore B is correct.

QUESTION 117

Refer to the exhibit.



Why is R2 unable to ping the loopback interface of R4?

- A. The local preference is too high.
- B. The weight is too low.
- C. The next hop is not reachable from R2.
- D. The route originated from within the same AS.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Before a BGP speaker installs a route to a network in the main IP routing table, the router must know how to reach the next hop that is used to get to that network. Route reachability is verified by searching for a route to the next hop in the main IP routing table. Unlike IGP routing protocols, such as EIGRP and OSPF, which assume that a route is reachable if they learned it through a valid adjacency, BGP does not install routes that it cannot verify as reachable. If a route to the next hop for a BGP network is found in the main IP routing table, BGP assumes that the network is reachable, and that the particular BGP route might be stored in the main IP routing table. If the router receives a route to a network that is not reachable, that route continues to be stored in the incoming BGP table, adj-RIB-In, and might be seen using the show ip bgp command, but is not placed in the main IP routing table.

Reference. https://www.informit.com/library/content.aspx?b=CCIE_Practical_Studies_II&seqNum=75

QUESTION 118

Which statement about the BGP originator ID is true?

- A. The route reflector always sets the originator ID to its own router ID.
- B. The route reflector sets the originator ID to the router ID of the route reflector client that injects the route into the AS.
- C. The route reflector client that injects the route into the AS sets the originator ID to its own router ID.
- D. The originator ID is set to match the cluster ID.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

An RR reflecting the route received from a RR-Client adds:

1. **Originator ID**- a 4-byte BGP attribute that is created by the RR. This attribute carries the Router ID of the originator of the route in the local AS. If the update comes back to the originator, it ignores the update.
2. **Cluster List**- A Cluster List is a list of Cluster IDs that an update has traversed. When a route reflector sends a route received from a client to a non-client, it appends the local Cluster ID. If a route reflector receives a route whose Cluster List contains the local Cluster ID, it ignores the update.

Reference. <https://sites.google.com/site/amitsciscozone/home/bgp/bgp-route-reflectors>

QUESTION 119

Refer to the exhibit.


```
R5#show ip bgp
BGP table version is 24, local router ID is 10.100.1.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*> 10.100.1.1/32   10.1.1.1             0      100      0 65001 23456 2 i
r> 10.100.1.2/32   10.1.2.1             0      100      0 65001 23456 i
```

Which two statements are true? (Choose two.)

- A. This router is not 4-byte autonomous system aware.
- B. This router is 4-byte autonomous system aware.
- C. The prefix 10.100.1.1/32 was learned through an autonomous system number with a length of 4 bytes, and this router is 4-byte autonomous system aware.
- D. The prefix 10.100.1.1/32 was learned through an autonomous system number with a length of 4 bytes, and this router is not 4-byte autonomous system aware.
- E. The prefix 10.100.1.1/32 was originated from a 4-byte autonomous system.

Correct Answer: AD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Prior to January 2009, BGP autonomous system (AS) numbers that were allocated to companies were 2-octet numbers in the range from 1 to 65535 as described in RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*. Due to increased demand for AS numbers, the Internet Assigned Number Authority (IANA) started to allocate four-octet AS numbers in the range from 65536 to 4294967295. RFC 5396, *Textual Representation of Autonomous System (AS) Numbers*, documents three methods of representing AS numbers. Cisco has implemented the following two methods:

- Asplain -- Decimal value notation where both 2-byte and 4-byte AS numbers are represented by their decimal value. For example, 65526 is a 2-byte AS number and 234567 is a 4-byte AS number.
- Asdot -- Autonomous system dot notation where 2-byte AS numbers are represented by their decimal value and 4-byte AS numbers are represented by a dot notation. For example, 65526 is a 2-byte AS number and 1.169031 is a 4-byte AS number (this is dot notation for the 234567 decimal number).

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xr-3s/irg-3s-book/irg-4byte-asn.html

QUESTION 120

Refer to the exhibit.


```
R2# show bgp ipv4 unicast summary
```

```
BGP router identifier 10.100.1.2, local AS number 2
```

```
BGP table version is 1, main routing table version 1
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.100.1.1	4	1	0	0	1	0	0	6d20h	Idle (PfxCt)

Which command is configured on this router?

- A. bgp update-delay 60
- B. neighbor 10.100.1.1 maximum-prefix 200
- C. neighbor 10.100.1.1 maximum-path 2
- D. neighbor 10.100.1.1 ebgp-multihop 2

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The BGP Maximum-Prefix feature allows you to control how many prefixes can be received from a neighbor. By default, this feature allows a router to bring down a peer when the number of received prefixes from that peer exceeds the configured Maximum-Prefix limit. This feature is commonly used for external BGP peers, but can be applied to internal BGP peers also. When the maximum number of prefixes has been received, the BGP sessions closes into the IDLE state.

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

QUESTION 121

Refer to the exhibit.

```
R3#sho ip bgp 172.16.1.0
BGP routing table entry for 172.16.1.0/24, version 5
Paths: (1 available, no best path)
Not advertised to any peer
Refresh Epoch 1
1
192.168.1.1 (inaccessible) from 192.168.2.1 (192.168.3.1)
  Origin IGP, metric 0, localpref 100, valid, internal
  rx pathid: 0x0, tx pathid: 0
```

Why is network 172.16.1.0/24 not installed in the routing table?

- A. There is no ARP entry for 192.168.1.1.
- B. The router cannot ping 192.168.1.1.
- C. The neighbor 192.168.1.1 just timed out and BGP will flush this prefix the next time that the BGP scanner runs.
- D. There is no route for 192.168.1.1 in the routing table.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Here we see that the next hop IP address to reach the 172.16.1.0 network advertised by the BGP peer is 192.168.1.1. However, the 192.168.1.1 IP is not in the routing table of R3 so it adds the route to the BGP table but marks it as inaccessible, as shown.

QUESTION 122

Consider a network that mixes link bandwidths from 128 kb/s to 40 Gb/s. Which value should be set for the OSPF reference bandwidth?

- A. Set a value of 128.
- B. Set a value of 40000.
- C. Set a manual OSPF cost on each interface.
- D. Use the default value.
- E. Set a value of 40000000.
- F. Set a value of 65535.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Unlike the metric in RIP which is determined by hop count and EIGRP's crazy mathematical formulated metric, OSPF is a little more simple. The default formula to calculate the cost for the OSPF metric is $(10^8/BW)$.

By default the metrics reference cost is 100Mbps, so any link that is 100Mbps will have a metric of

1. a T1 interface will have a metric of 64 so in this case if a router is trying to get to a FastEthernet network on a router that is through a T1 the metric would be 65 ($64 + 1$). You do however have the ability to statically specify a metric on a per interface basis by using the `ip ospf cost #` where the cost is an integer between 1-65535.

So the big question is why would you want to statically configure a metric?

The biggest advantage of statically configuring an OSPF metric on an interface is to manipulate which route will be chosen dynamically via OSPF. In a nut shell it's like statically configuring a dynamic protocol to use a specific route. It should also be used when the interface bandwidths vary greatly (some very low bandwidth interfaces and some very high speed interfaces on the same router).

QUESTION 123

Which statement about a type 4 LSA in OSPF is true?

- A. It is an LSA that is originated by an ABR, that is flooded throughout the AS, and that describes a route to the ASBR.
- B. It is an LSA that is originated by an ASBR, that is flooded throughout the AS, and that describes a route to the ASBR.
- C. It is an LSA that is originated by an ASBR, that is flooded throughout the area, and that describes a route to the ASBR.
- D. It is an LSA that is originated by an ABR, that is flooded throughout the AS, and that describes a route to the ABR.
- E. It is an LSA that is originated by an ABR, that is flooded throughout the area, and that describes a route to the ASBR.

Correct Answer: E

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

LSA Type 4 (called Summary ASBR LSA) is generated by the ABR to describe an ASBR to routers in other areas so that routers in other areas know how to get to external routes through that ASBR.

QUESTION 124

Refer to the exhibit.

```
R1#show ip route 1.1.1.1
% Network not in table

R1#sho ip ospf database router 1.1.1.1

OSPF Router with ID (10.10.10.10) (Process ID 1)

Router Link States (Area 0)

  Adv Router is not-reachable
  LS age: 6
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 1.1.1.1
  Advertising Router: 1.1.1.1
  LS Seq Number: 80000003
  Checksum: 0x6889
  Length: 48
  Number of Links: 2

Link connected to: a Stub Network
  (Link ID) Network/subnet number: 1.1.1.1
  (Link Data) Network Mask: 255.255.255.255
  TOS 0 Metrics: 1

Link connected to: a Transit Network
  (Link ID) Designated Router address: 10.1.1.0
  (Link Data) Router Interface address: 10.1.1.1
  TOS 0 Metrics: 10

R1#sho ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	0	FULL/ -	00:00:36	10.1.1.1	Ethernet0/0

Why is the prefix 1.1.1.1/32 not present in the routing table of R1?

- A. There is a duplicate router ID.
- B. There is a subnet mask mismatch on Ethernet0/0.
- C. The router LSA has an invalid checksum.
- D. There is an OSPF network type mismatch that causes the advertising router to be unreachable.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

A common problem when using Open Shortest Path First (OSPF) is routes in the database don't appear in the routing table. In most cases OSPF finds a discrepancy in the database so it doesn't install the route in the routing table. Often, you can see the Adv Router is not-reachable message (which means that the router advertising the LSA is not reachable through OSPF) on top of the link-state advertisement (LSA) in the database when this problem occurs. Here is an example:

Adv Router is not-reachable

LS agE. 418

Options: (No TOS-capability, DC)

LS TypeE. Router Links

Link State ID. 172.16.32.2

Advertising Router: 172.16.32.2

LS Seq Number: 80000002

Checksum: 0xFA63


Length: 60

Number of Links: 3

There are several reasons for this problem, most of which deal with mis-configuration or a broken topology. When the configuration is corrected the OSPF database discrepancy goes away and the routes appear in the routing table.

Reason 1: Network Type Mismatch

Let's use the following network diagram as an example:



R4-4K	R1-7010
interface Loopback0 ip address 172.16.33.1 255.255.255.255	interface Loopback0 ip address 172.16.30.1 255.255.255.255
interface Serial2 ip address 172.16.32.1 255.255.255.0 ip ospf network broadcast	! interface Serial1/0 ip address 172.16.32.2 255.255.255.0 clockrate 64000
router ospf 20 network 172.16.0.0 0.0.255.255 area 0	router ospf 20 network 172.16.0.0 0.0.255.255 area 0

R4-4K(4)# show ip ospf interface serial 2
 Serial2 is up, line protocol is up
 Internet Address 172.16.32.1/24, Area 0
 Process ID 20, Router ID 172.16.33.1, Network Type BROADCAST, Cost: 64
 Transmit Delay is 1 sec, State DR, Priority 1
 Designated Router (ID) 172.16.33.1, Interface address 172.16.32.1
 Backup Designated router (ID) 172.16.32.2, Interface address 172.16.32.2
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in 00:00:08
 Neighbor Count is 1, Adjacent neighbor count is 1
 Adjacent with neighbor 172.16.32.2 (Backup Designated Router)
 Suppress hello for 0 neighbor(s)

R1-7010(5)# show ip ospf interface serial 1/0
 Serial1/0 is up, line protocol is up
 Internet Address 172.16.32.2/24, Area 0
 Process ID 20, Router ID 172.16.32.2, Network Type POINT_TO_POINT, Cost: 64
 Transmit Delay is 1 sec, State POINT_TO_POINT,
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in 00:00:02
 Neighbor Count is 1, Adjacent neighbor count is 1
 Adjacent with neighbor 172.16.33.1
 Suppress hello for 0 neighbor(s)

As you can see above, Router R4-4K is configured for broadcast, and Router R1-7010 is configured for point-to-point. This kind of network type mismatch makes the advertising router unreachable.

R4-4K(4)# show ip ospf database router 172.16.32.2

Adv Router is not-reachable

LS agE. 418
Options: (No TOS-capability, DC)
LS Type. Router Links
Link State ID. 172.16.32.2
Advertising Router: 172.16.32.2
LS Seq Number: 80000002
Checksum: 0xFA63
Length: 60
Number of Links: 3

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID. 172.16.33.1
(Link Data) Router Interface address: 172.16.32.2
Number of TOS metrics: 0
TOS 0 Metrics: 64

Link connected to: a Stub Network
(Link ID) Network/subnet number: 172.16.32.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 64

R1-7010(5)# show ip ospf database router 172.16.33.1

Adv Router is not-reachable
LS agE. 357
Options: (No TOS-capability, DC)
LS Type. Router Links
Link State ID. 172.16.33.1
Advertising Router: 172.16.33.1
LS Seq Number: 8000000A
Checksum: 0xD4AA
Length: 48
Number of Links: 2

Link connected to: a Transit Network
(Link ID) Designated Router address: 172.16.32.1
(Link Data) Router Interface address: 172.16.32.1
Number of TOS metrics: 0
TOS 0 Metrics: 64

You can see that for subnet 172.16.32.0/24, Router R1-7010 is generating a point-to-point link and Router R4-4K is generating a transit link. This creates a discrepancy in the link-state database, which means no routes are installed in the routing table.

```
R1-7010(5)# show ip route
172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C 172.16.32.0/24 is directly connected, Serial1/0
C 172.16.30.1/32 is directly connected, Loopback0
```

Solution

To solve this problem, configure both routers for the same network type. You can either change the network type of Router R1-7010 to broadcast, or change Router R4-4K's serial interface to point-to-point.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7112-26.html>

QUESTION 125

Which authentication method does OSPFv3 use to secure communication between neighbors?

- A. plaintext
- B. MD5 HMAC
- C. PKI
- D. IPSec

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

In order to ensure that OSPFv3 packets are not altered and re-sent to the device, causing the device to behave in a way not desired by its system administrators, OSPFv3 packets must be authenticated. OSPFv3 uses the IPsec secure socket API to add authentication to OSPFv3 packets. This API supports IPv6.

OSPFv3 requires the use of IPsec to enable authentication. Crypto images are required to use authentication, because only crypto images include the IPsec API needed for use with OSPFv3.

Reference. http://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

QUESTION 126

Which three statements are true about OSPFv3? (Choose three.)

- A. The only method to enable OSPFv3 on an interface is via the interface configuration mode.
- B. Multiple instances of OSPFv3 can be enabled on a single link.
- C. There are two methods to enable OSPFv3 on an interface, either via the interface configuration mode or via the router configuration mode.
- D. For OSPFv3 to function, IPv6 unicast routing must be enabled.
- E. For OSPFv3 to function, IPv6 must be enabled on the interface.

F. Only one instance of OSPFv3 can be enabled on a single link.

Correct Answer: BDE

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Here is a list of the differences between OSPFv2 and OSPFv3:

- They use different address families (OSPFv2 is for IPv4-only, OSPFv3 can be used for IPv6-only or both protocols (more on this following))
- OSPFv3 introduces new LSA types
- OSPFv3 has different packet format
- OSPFv3 uses different flooding scope bits (U/S2/S1)
- OSPFv3 adjacencies are formed over link-local IPv6 communications OSPFv3 runs per-link rather than per-subnet
- OSPFv3 supports multiple instances on a single link, Interfaces can have multiple IPv6 addresses
- OSPFv3 uses multicast addresses FF02::5 (all OSPF routers), FF02::6 (all OSPF DRs) OSPFv3 Neighbor Authentication done with IPsec (AH)
- OSPFv2 Router ID (RID) must be manually configured, still a 32-bit number
-

Following is a simple example of OSPFv3 configuration on a Cisco IOS 12.4T router.

```
ipv6 unicast-routing
ipv6 cef
!
interface GigabitEthernet 0/0
 description Area 0.0.0.0 backbone interface
 ipv6 address 2001:DB8:100:1::1/64
 ipv6 ospf network broadcast
 ipv6 ospf 100 area 0.0.0.0
```

Reference. <http://www.networkworld.com/article/2225270/cisco-subnet/ospfv3-for-ipv4-and-ipv6.html>

QUESTION 127

Which statement about OSPF multiaccess segments is true?

- A. The designated router is elected first.
- B. The designated and backup designated routers are elected at the same time.
- C. The router that sent the first hello message is elected first.
- D. The backup designated router is elected first.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

According to the RFC, the BDR is actually elected first, followed by the DR. The RFC explains why:

"The reason behind the election algorithm's complexity is the desire for an orderly transition from Backup Designated Router to Designated Router, when the current Designated Router fails. This orderly transition is ensured through the introduction of hysteresis: no new Backup Designated Router can be chosen until the old Backup accepts its new Designated Router responsibilities.

The above procedure may elect the same router to be both Designated Router and Backup Designated Router, although that router will never be the calculating router (Router X) itself."

Reference. <http://www.ietf.org/rfc/rfc2328.txt> Page 76

QUESTION 128

What are the minimal configuration steps that are required to configure EIGRP HMAC-SHA2 authentication?

- A. classic router mode, interface XX, authentication mode hmac-sha-256 <password>
- B. named router mode, address-family statement, authentication mode hmac-sha-256 <password>
- C. named router mode, address-family statement, af-interface default, authentication mode hmac-sha-256 <password>
- D. named router mode, address-family statement, authentication mode hmac-sha-256 <password>

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The example below shows how to configure EIGRP HMAC-SHA2 on Cisco router:

```
Device(config)# router eigrp name1
```

```
Device(config-router)# address-family ipv4 autonomous-system 45000
```

```
Device(config-router-af)# af-interface ethernet 0/0
```

```
Device(config-router-af-interface)# authentication mode hmac-sha-256 0 password1
```

```
Device(config-router-af-interface)# end
```


Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/15-sy/ire-15-sy-book/ire-sha-256.html

QUESTION 129

Refer to the exhibit.

R1:

```
interface Loopback0
 ip address 1.1.1.1 255.255.255.0
 ipv6 address 2001:12::1/128
 ipv6 eigrp 100
!
interface FastEthernet0/0
 ip address 10.1.12.1
 255.255.255.0
 duplex auto
 speed auto
 ipv6 address 2001:112::1/64
 ipv6 eigrp 100
!
interface FastEthernet0/1
 no ip address
 shutdown
 duplex auto
 speed auto
!
ip forward-protocol nd
!
!
no ip http server
no ip http secure-server
!
!
ipv6 router eigrp 100
 no shutdown
!
!
control-plane
!
```



R2:

```
interface Loopback0
 ip address 2.2.2.2 255.255.255.0
 ipv6 address 2001:12::2/128
 ipv6 eigrp 100
!
interface FastEthernet0/0
 ip address 10.1.12.2
 255.255.255.0
 duplex auto
 speed auto
 ipv6 address 2001:112::2/64
 ipv6 eigrp 100
!
interface FastEthernet0/1
 no ip address
 shutdown
 duplex auto
 speed auto
!
ip forward-protocol nd
!
!
no ip http server
no ip http secure-server
!
!
ipv6 router eigrp 100
 shutdown
!
!
control-plane
!
```

How many EIGRP routes will appear in the routing table of R2?

- A. 0
- B. 1
- C. 2

D. 3

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

EIGRPv6 on R2 was shut down so there is no EIGRP routes on the routing table of R2. If we turn on EIGRPv6 on R2 (with "no shutdown" command) then we would see the prefix of the loopback interface of R1 in the routing table of R2.

```
R2#sh ipv6 route
IPv6 Routing Table - 5 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
        U - Per-user Static route, M - MIPv6
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
        D - EIGRP, EX - EIGRP external
D 2001:12::1/128 [90/156160]
  via FE80::200:CFF:FE07:BB01, FastEthernet0/0
C 2001:12::2/128 [0/0]
  via ::, Loopback0
C 2001:112::/64 [0/0]
  via ::, FastEthernet0/0
L 2001:112::2/128 [0/0]
  via ::, FastEthernet0/0
L FF00::/8 [0/0]
  via ::, Null0
R2#
```

Note. EIGRPv6 requires the "ipv6 unicast-routing" global command to be turned on first or it will not work.

QUESTION 130

Which two configuration changes should be made on the OTP interface of an EIGRP OTP route reflector? (Choose two.)

- A. passive-interface
- B. no split-horizon
- C. no next-hop-self
- D. hello-interval 60, hold-time 180

Correct Answer: BC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The EIGRP Over the Top feature enables a single end-to-end Enhanced Interior Gateway Routing Protocol (EIGRP) routing domain that is transparent to the underlying public or private WAN transport that is used for connecting disparate EIGRP customer sites. When an enterprise extends its connectivity across multiple sites through a private or a public WAN connection, the service provider mandates that the enterprise use an additional routing protocol, typically the Border Gateway Protocol (BGP), over the WAN links to ensure end-to-end routing. The use of an additional protocol causes additional complexities for the enterprise, such as additional routing processes and sustained interaction between EIGRP and the routing protocol to ensure connectivity, for the enterprise. With the EIGRP Over the Top feature, routing is consolidated into a single protocol (EIGRP) across the WAN.

Perform this task to configure a customer edge (CE) device in a network to function as an EIGRP Route Reflector:

1. enable
2. configure terminal
3. router eigrp virtual-name
4. address-family ipv4 unicast autonomous-system as-number
5. af-interface interface-type interface-number
6. no next-hop-self
7. no split-horizon
8. exit
9. remote-neighbors source interface-type interface-number unicast-listen lisp-encap
10. network ip-address
11. end

Note. Use no next-hop-self to instruct EIGRP to use the received next hop and not the local outbound interface address as the next hop to be advertised to neighboring devices. If no next-hop-self is not configured, the data traffic will flow through the EIGRP Route Reflector.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/xr-3s/ire-xr-3s-book/ire-eigrp-over-the-top.html

QUESTION 131

Which statement about the function of poison reverse in EIGRP is true?

- A. It tells peers to remove paths that previously might have pointed to this router.
- B. It tells peers to remove paths to save memory and bandwidth.
- C. It provides reverse path information for multicast routing.
- D. It tells peers that a prefix is no longer reachable.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Poison Reverse in EIGRP states: "Once you learn of a route through an interface, advertise it as unreachable back through that same interface". For more information please read here. <http://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html#splithorizon>.

QUESTION 132

What is the preferred method to improve neighbor loss detection in EIGRP?

- A. EIGRP natively detects neighbor down immediately, and no additional feature or configuration is required.
- B. BFD should be used on interfaces that support it for rapid neighbor loss detection.
- C. Fast hellos (subsecond) are preferred for EIGRP, so that it learns rapidly through its own mechanisms.
- D. Fast hellos (one-second hellos) are preferred for EIGRP, so that it learns rapidly through its own mechanisms.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Bi-directional Forwarding Detection (BFD) provides rapid failure detection times between forwarding engines, while maintaining low overhead. It also provides a single, standardized method of link/device/protocol failure detection at any protocol layer and over any media.

Reference. "Bidirectional Forwarding Detection for EIGRP"

http://www.cisco.com/en/US/technologies/tk648/tk365/tk207/technologies_white_paper0900aecd8_0243fe7.html

QUESTION 133

How does EIGRP derive the metric for manual summary routes?

- A. It uses the best composite metric of any component route in the topology table.
- B. It uses the worst composite metric of any component route in the topology table.
- C. It uses the best metric vectors of all component routes in the topology table.
- D. It uses the worst metric vectors of all component routes in the topology table.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

For example if your router has a routing table like this:

```
D 192.168.8.0/24 [90/2632528] via 192.168.0.1, 00:00:12, Serial0/0
D 192.168.9.0/24 [90/2323456] via 192.168.0.1, 00:00:12, Serial0/0
D 192.168.10.0/24 [90/2195456] via 192.168.0.1, 00:00:12, Serial0/0
D 192.168.11.0/24 [90/2323456] via 192.168.0.1, 00:00:12, Serial0/0
```

Now suppose you want to manually summarize all the routes above, you can use this command (on the router that advertised these routes to our router):

```
Router(config-if)#ip summary-address eigrp 1 192.168.8.0 255.255.248.0
```

After that the routing table of your router will look like this:

```
D 192.168.8.0/21 [90/2195456] via 192.168.0.1, 00:01:42, Serial0/0
```

And we can see the manual summary route takes the smallest metric of the specific routes.

QUESTION 134

Refer to the exhibit.

```
R2#show ipv6 interface e0/0
Ethernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8BB:CCFF:FE00:200
  No Virtual link-local address(es):
  No global unicast address is configured
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::A
    FF02::1:FE00:200
```

Which part of the joined group addresses list indicates that the interface has joined the EIGRP multicast group address?

- A. FF02::1
- B. FF02::1:FE00:200
- C. FF02::A
- D. FF02::2

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

FF02::A is an IPv6 link-local scope multicast addresses. This address is for all devices on a wire that want to "talk" EIGRP with one another.

Focusing specifically on FF02::A and how routers join it, we can see and say three things:

- Local: FF02::A is local to the wire.
- Join: Each device "joins" FF02::A by just "deciding to listen" to the IPv6 link-local scope multicast address FF02::A. Then, by extension, it listens to the corresponding MAC address for that multicast IPv6 address (33:33:00:00:00:0A).
- Common interest: As we can see, these varying groups have something in common that they would all like to hear about. For FF02::A, the common interest -- the "connection" among the devices joining that group is that they all want to listen to or participate in EIGRP.

Reference. <http://www.networkcomputing.com/networking/understanding-ipv6-what-is-solicited-node-multicast/a/d-id/1315703>

QUESTION 135

EIGRP allows configuration of multiple MD5 keys for packet authentication to support easy rollover from an old key to a new key. Which two statements are true regarding the usage of multiple authentication keys? (Choose two.)

- A. Received packets are authenticated by the key with the smallest key ID.
- B. Sent packets are authenticated by all valid keys, which means that each packet is replicated as many times as the number of existing valid keys.
- C. Received packets are authenticated by any valid key that is chosen.
- D. Sent packets are authenticated by the key with the smallest key ID.

Correct Answer: CD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Suppose two routers are connected with each other via Fa0/0 interfaces and they are configured to authenticate via MD5. Below is a simple configuration on both routers so that they will work:

```
Router1(config)#key chain KeyChainR1
Router1(config-keychain)#key 1
Router1(config-keychain-key)#key-string FirstKey
Router1(config-keychain-key)#key 2
Router1(config-keychain-key)#key-string SecondKey
Router2(config)#key chain KeyChainR2
Router2(config-keychain)#key 1
```



```
Router2(config-keychain-key)#key-string FirstKey
Router2(config-keychain-key)#key 2
Router2(config-keychain-key)#key-string SecondKey
```

Apply these key chains to R1 & R2:

```
Router1(config)#interface fastEthernet 0/0
Router1(config-if)#ip authentication mode eigrp 1 md5
Router1(config-if)#ip authentication key-chain eigrp 1 KeyChainR1
Router2(config)#interface fastEthernet 0/0
Router2(config-if)#ip authentication mode eigrp 1 md5
Router2(config-if)#ip authentication key-chain eigrp 1 KeyChainR2
```

There are some rules to configure MD5 authentication with EIGRP:

- + The **key chain names** on two routers do not have to match (in this case the name "KeyChainR1 & "KeyChainR2 do not match)
 - + **The key number** and **key-string** on the two potential neighbors must match (for example "key 1 & "key-string FirstKey" must match on "key 1" & "key-string FirstKey" of neighboring router) Also some facts about MD5 authentication with EIGRP
 - + When sending EIGRP messages the lowest valid key number is used -> D is correct.
 - + When receiving EIGRP messages all currently configured valid keys are verified but the lowest valid one will be used -> Although answer C does not totally mention like that but it is the most suitable answer because A and B are totally wrong.
- Answer A is not correct because we need valid key to authenticate. As mentioned above, although answer C is not totally correct but it puts some light on why answer B is not correct: each packet is NOT "replicated as many times as the number of existing valid keys". All currently configured valid keys are verified but the lowest valid one will be used.

QUESTION 136

Refer to the exhibit.

```
R1
!
interface FastEthernet0/0
ip address 10.1.1.5 255.255.255.0
!
router ospf 1
network 10.1.1.5 0.0.0.0 area 0
passive-interface default
!

R2
!
interface FastEthernet0/1
ip address 10.1.1.6 255.255.255.0
!
router ospf 10
network 10.1.1.6 0.0.0.0 area 0
!
```

Which additional configuration is necessary for R1 and R2 to become OSPF neighbors?

- A. R1
!
router ospf 1
no passive-interface FastEthernet0/0
!
- B. R2
!
router ospf 10
no network 10.1.1.6 0.0.0.0 area 0
network 10.1.1.6 0.0.0.0 area 1
!
- C. R1
!
interface FastEthernet0/0
ip ospf mtu-ignore
!
R2
!
interface FastEthernet0/1

```
ip ospf mtu-ignore
!
D. R1
!
no router ospf 1
router ospf 10
network 10.1.1.5 0.0.0.0 area 0
```

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Because the passive interface default command is used, by default all interfaces are passive and no neighbors will form on these interfaces. We need to disable passive interface on the link to R2 by using the "no passive-interface FastEthernet0/0" on R1 under OSPF.

QUESTION 137

Refer to the exhibit.

```
Load for five secs: 12%/0%; one minute: 4%; five minutes: 5%
Time source is NTP, 11:19:50.533 US/Ariz Tue Oct 1 2013

(10.10.76.191, 239.93.200.8), 7w0d/00:02:55, flags: sTI
  Incoming interface: TenGigabitEthernet8/2, RPF nbr 70.169.73.188, RPF-MFD
  Outgoing interface list:
    GigabitEthernet1/5, Forward/Sparse, 2w5d/00:02:25, H
    GigabitEthernet1/2, Forward/Sparse, 5w3d/00:02:25, H
    GigabitEthernet1/1, Forward/Sparse, 25w6d/00:02:49, H

(10.10.76.191, 239.93.200.9), 7w0d/00:02:55, flags: sTI
  Incoming interface: TenGigabitEthernet8/2, RPF nbr 70.169.73.188, RPF-MFD
  Outgoing interface list:
    GigabitEthernet1/5, Forward/Sparse, 2w5d/00:02:25, H
```

Which two statements about the device that generated the output are true? (Choose two.)

- A. The SPT-bit is set.
- B. The sparse-mode flag is set.

- C. The RP-bit is set.
- D. The source-specific host report was received.

Correct Answer: AD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

In this example we can see that the s, T, and I flags are set. Here is a list of the flags and their meanings:

show ip mroute Field Descriptions

Field	Description
Flags:	Provides information about the entry.
D - Dense	Entry is operating in dense mode.
S - Sparse	Entry is operating in sparse mode.
B - Bidir Group	Indicates that a multicast group is operating in bidirectional mode.
s - SSM Group	Indicates that a multicast group is within the SSM range of IP addresses. This flag is reset if the SSM range changes.
C - Connected	A member of the multicast group is present on the directly connected interface.
L - Local	The router itself is a member of the multicast group.
P - Pruned	Route has been pruned. The Cisco IOS software keeps this information so that a downstream member can join the source.
R - RP-bit set	Indicates that the (S, G) entry is pointing toward the RP. This is typically prune state along the shared tree for a particular source.
F - Register flag	Indicates that the software is registering for a multicast source.
T - SPT-bit set	Indicates that packets have been received on the shortest path source tree.
J - Join SPT	For (*, G) entries, indicates that the rate of traffic flowing down the shared tree is exceeding the SPT-Threshold set for the group. (The default SPT-Threshold setting is 0 kbps.) When the J- Join shortest path tree (SPT) flag is set, the next (S, G) packet received down the shared tree triggers an (S, G) join in the direction of the source, thereby causing the router to join the source tree. For (S, G) entries, indicates that the entry was created because the SPT-Threshold for the group was exceeded. When the J- Join SPT flag is set for (S, G) entries, the router monitors the traffic rate on the source tree and attempts to switch back to the shared tree for this source if the traffic rate on the source tree falls below the SPT-Threshold of the group for more than 1 minute.
M - MSDP created entry	Indicates that a (*, G) entry was learned through a Multicast Source Discovery Protocol (MSDP) peer. This flag is only applicable for a rendezvous point (RP) running MSDP.
X - Proxy Join Timer Running	Indicates that the proxy join timer is running. This flag is only set for (S, G) entries of an RP or "turnaround" router. A "turnaround" router is located at the intersection of a shared path (*, G) tree and the shortest path from the source to the RP.
A - Advertised via MSDP	Indicates that an (S, G) entry was advertised through an MSDP peer. This flag is only applicable for an RP running MSDP.
U - URD	Indicates that a URD channel subscription report was received for the (S, G) entry.
I - Received Source Specific Host Report	Indicates that an (S, G) entry was created by an (S, G) report. This (S, G) report could have been created by IGMPv3, URD, or IGMP v3lite. This flag is only set on the designated router (DR).

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/12s_ssm.html

QUESTION 138

Refer to the exhibit.

```
Switch#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, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.192.1.1), 00:01:43/stopped, RP 10.210.150.1, flags: SJC
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Vlan150, Forward/Sparse-Dense, 00:01:43/00:02:55

(10.210.168.132, 239.192.1.1), 00:00:25/00:02:38, flags: T
  Incoming interface: Port-channel1, RPF nbr 10.85.20.20
  Outgoing interface list:
    Vlan150, Forward/Sparse-Dense, 00:00:25/00:02:34

(*, 224.0.1.40), 00:01:57/00:02:53, RP 10.210.150.1, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Port-channel1, Forward/Sparse-Dense, 00:01:09/00:03:18
    Vlan150, Forward/Sparse-Dense, 00:01:39/00:02:55
```


Which three statements about the output are true? (Choose three.)

- A. This switch is currently receiving a multicast data stream that is being forwarded out VLAN 150.
- B. A multicast receiver has requested to join one or more of the multicast groups.
- C. Group 224.0.1.40 is a reserved address, and it should not be used for multicast user data transfer.
- D. One or more multicast groups are operating in PIM dense mode.
- E. One or more of the multicast data streams will be forwarded out to neighbor 10.85.20.20.
- F. Group 239.192.1.1 is a reserved address, and it should not be used for multicast user data transfer.

Correct Answer: ABC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

- A. VLAN 150 shows up in the outgoing interface list so those specific multicast streams are being forwarded to this VLAN.
- B. A receiver has requested to receive the multicast stream associated with the multicast address of 239.192.1.1, that is why this stream appears in the mroute table.
- C. The 224.0.1.40 is a reserved multicast group for cisco's Rp discovery. All cisco routers are members of this grup by default and listen to this group for Cisco RP discovery messages advertised by mapping agent even if it is not configured

QUESTION 139

Which statement about the RPF interface in a BIDIR-PIM network is true?

- A. In a BIDIR-PIM network, the RPF interface is always the interface that is used to reach the PIM rendezvous point.
- B. In a BIDIR-PIM network, the RPF interface can be the interface that is used to reach the PIM rendezvous point or the interface that is used to reach the source.
- C. In a BIDIR-PIM network, the RPF interface is always the interface that is used to reach the source.
- D. There is no RPF interface concept in BIDIR-PIM networks.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

RPF stands for "Reverse Path Forwarding". The RPF Interface of a router with respect to an address is the interface that the MRIB indicates should be used to reach that address. In the case of a BIDIR-PIM multicast group, the RPF interface is determined by looking up the Rendezvous Point Address in the MRIB. The RPF information determines the interface of the router that would be used to send packets towards the Rendezvous Point Link for the group.

Reference. <https://tools.ietf.org/html/rfc5015>

QUESTION 140

Which technology is an application of MSDP, and provides load balancing and redundancy between the RPs?

- A. static RP
- B. PIM BSR
- C. auto RP
- D. anycast RP

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Using Anycast RP is an implementation strategy that provides load sharing and redundancy in Protocol Independent Multicast sparse mode (PIM-SM) networks. Anycast RP allows two or more rendezvous points (RPs) to share the load for source registration and the ability to act as hot backup routers for each other. Multicast Source Discovery Protocol (MSDP) is the key protocol that makes Anycast RP possible.

Reference.

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

QUESTION 141

Which two statements are true about IPv6 multicast? (Choose two.)

- A. Receivers interested in IPv6 multicast traffic use IGMPv6 to signal their interest in the IPv6 multicast group.
- B. The PIM router with the lowest IPv6 address becomes the DR for the LAN.
- C. An IPv6 multicast address is an IPv6 address that has a prefix of FF00::/8.
- D. The IPv6 all-routers multicast group is FF02:0:0:0:0:0:0:2.

Correct Answer: CD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Multicast addresses in IPv6 have the prefix ff00::/8.

Well-known IPv6 multicast addresses	
Address	Description
ff02::1	All nodes on the local network segment
ff02::2	All routers on the local network segment

Reference. http://en.wikipedia.org/wiki/Multicast_address

QUESTION 142

DRAG DROP

Drag and drop the IPv6 address on the left to the correct IPv6 address type on the right.

Select and Place:

Drag and drop the IPv6 address on the left to the correct IPv6 address type on the right.	
FF01::2	Link Local Unicast
FE80:2a5b::5	Global Unicast
FDF8:E5F3:83E4:FEAA::53	Multicast
2005:CA75:D095::5	Unique Local Unicast
F880:E6F4:B665::44	

Correct Answer:

Drag and drop the IPv6 address on the left to the correct IPv6 address type on the right.

	FE80:2a5b::5
	2005:CA75:D095::5
	FF01::2
	FDF8:E5F3:83E4:FEAA::53
F880:E6F4:B665::44	

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 143

DRAG DROP

Drag and drop the BGP attribute on the left to the correct category on the right.

Select and Place:

Drag and drop the BGP attribute on the left to the correct category on the right.

Originator ID

Community

Local-Pref

AS_path

Aggregator

Next-Hop

BGP Well-Known Mandatory Attribute

Target

Target

BGP Well-Known Discretionary Attribute

Target

BGP Optional Nontransitive Attribute

Target

Correct Answer:

Drag and drop the BGP attribute on the left to the correct category on the right.

	BGP Well-Known Mandatory Attribute
Community	AS_path
	Next-Hop
	BGP Well-Known Discretionary Attribute
Aggregator	Local-Pref
	BGP Optional Nontransitive Attribute
	Originator ID

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 144

DRAG DROP

Drag and drop the IPv6 multicast feature or protocol on the left to the correct address space on the right.

Select and Place:

Drag and drop the IPv6 multicast feature or protocol on the left to the correct address space on the right.

All nodes	FF02::D
All routers	FF02::6
EIGRP	FF02::2
PIM routers	FF02::A
RIP routers	FF02::1
OSPFv3 all DR routers	FF02::9

Correct Answer:

Drag and drop the IPv6 multicast feature or protocol on the left to the correct address space on the right.

	PIM routers
	OSPFv3 all DR routers
	All routers
	EIGRP
	All nodes
	RIP routers

Section: Layer 3 Technologies**Explanation****Explanation/Reference:****QUESTION 145****DRAG DROP**

Drag and drop the multicast protocol or feature on the left to the correct address space on the right.

Select and Place:

Drag and drop the multicast protocol or feature on the left to the correct address space on the right.	
Auto-RP announcement	224.0.0.13
PIMv2	232.0.0.0/8
GLBP	224.0.1.40
Auto-RP discovery	224.0.0.102
Source Specific Multicast (SSM)	224.0.1.39

Correct Answer:

Drag and drop the multicast protocol or feature on the left to the correct address space on the right.

	PIMv2
	Source Specific Multicast (SSM)
	Auto-RP discovery
	GLBP
	Auto-RP announcement

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 146

DRAG DROP

Drag and drop the router preference on the left to the correct routing sequence (from most preferred to least preferred) on the right.

Select and Place:

Drag and drop the router preference on the left to the correct routing sequence (from most preferred to least preferred) on the right.

EBGP route	1
Static route	2
Most specific prefix	3
Directly connected route	4

Correct Answer:

Drag and drop the router preference on the left to the correct routing sequence (from most preferred to least preferred) on the right.

	Most specific prefix
	Directly connected route
	Static route
	EBGP route

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 147

DRAG DROP

Drag and drop the OSPF network type on the left to the correct traffic type category on the right.

Select and Place:

Drag and drop the OSPF network type on the left to the correct traffic type category on the right.

Broadcast
Nonbroadcast
Point-to-Point
Loopback
Point-to-Multipoint
Point-to-Multipoint Nonbroadcast

Unicast
Multicast
Stub

Correct Answer:

Drag and drop the OSPF network type on the left to the correct traffic type category on the right.

Unicast

Nonbroadcast

Point-to-Multipoint Nonbroadcast

Multicast

Broadcast

Point-to-Point

Point-to-Multipoint

Stub

Loopback

Section: Layer 3 Technologies
Explanation

Explanation/Reference:

QUESTION 148
Refer to the exhibit.

```
Internet Protocol, Src: 192.168.0.2 (192.168.0.2), Dst: 224.0.0.10 (224.0.0.10)
Version: 4
Header length: 20 bytes
Differentiated Services Field: 0xc0 (DSCP 0x30: Class Selector 6; ECN: 0x00)
Total Length: 60
Identification: 0x0000 (0)
Flags: 0x00
Fragment offset: 0
Time to live: 2
Protocol: EIGRP (88)
Header checksum: 0x16f6 [correct]
Source: 192.168.0.2 (192.168.0.2)
Destination: 224.0.0.10 (224.0.0.10)
```

Which two pieces of information in this Wireshark capture indicate that you are viewing EIGRP traffic? (Choose two.)

- A. the header length
- B. the protocol number
- C. the destination address
- D. the Class Selector
- E. the source address
- F. the header checksum

Correct Answer: BC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

EIGRP uses protocol number 88, which shows as EIGRP in the capture. Also, we in the capture that the destination IP address is 224.0.0.10, which is the Enhanced Interior Gateway Routing Protocol (EIGRP) group address is used to send routing information to all EIGRP routers on a network segment.

QUESTION 149

When BGP route reflectors are used, which attribute ensures that a routing loop is not created?

- A. weight
- B. local preference
- C. multiexit discriminator

D. originator ID

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

As the iBGP learned routes are reflected, routing information may loop. The route reflector model has the following mechanisms to avoid routing loops:

- Originator ID is an optional, nontransitive BGP attribute. It is a 4-byte attributed created by a route reflector. The attribute carries the router ID of the originator of the route in the local autonomous system. Therefore, if a misconfiguration causes routing information to come back to the originator, the information is ignored.
- Cluster-list is an optional, nontransitive BGP attribute. It is a sequence of cluster IDs that the route has passed. When a route reflector reflects a route from its clients to nonclient peers, and vice versa, it appends the local cluster ID to the cluster-list. If the cluster-list is empty, a new cluster-list is created. Using this attribute, a route reflector can identify if routing information is looped back to the same cluster due to misconfiguration. If the local cluster ID is found in the cluster-list, the advertisement is ignored.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfbgp.html

QUESTION 150

Refer to the exhibit.

```
Router#sh ip osp data summ

      OSPF Router with ID (100.1.1.1) (Process ID 1)

      Summary Net Link States (Area 0)

LS age: 22
Options: (No TOS-capability, DC, Downward)
LS Type: Summary Links(Network)
Link State ID: 2.2.0.0 (summary Network Number)
Advertising Router: 2.3.4.101
LS Seq Number: 80000001
Checksum: 0x3316
Length: 28
Network Mask: /24
      MTID: 0      Metric: 1
```

Which statement is true about the downward bit?

- A. It forces the CE router to use a backup link instead of sending traffic via MPLS VPN.
- B. It informs the PE router that the LSA metric has been recently decreased to 1 and that partial SPF calculation cannot be delayed.
- C. It forces the CE router to install the LSA with the downward bit set into its routing table as a discard route.
- D. It informs the PE router that the LSA was already redistributed into BGP by another PE router and that the LSA must not be redistributed into BGP again.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

From RFC 4577, specifically section 4.2.5.1

When a type 3 LSA is sent from a PE router to a CE router, the DN bit [OSPF-DN] in the LSA Options field MUST be set. This is used to ensure that if any CE router sends this type 3 LSA to a PE router, the PE router will not redistribute it further.

When a PE router needs to distribute to a CE router a route that comes from a site outside the latter's OSPF domain, the PE router presents itself as an ASBR (Autonomous System Border Router), and distributes the route in a type 5 LSA. The DN bit [OSPF-DN] MUST be set in these LSAs to ensure that they will be ignored by any other PE routers that receive them.

QUESTION 151

Which regular expression will match prefixes that originated from AS200?

- A. ^\$
- B. ^200_
- C. _200\$
- D. ^200)
- E. _200_

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Example on how to deny all prefixes originated in Autonomous System 200

```
router bgp 100
  neighbor 10.1.1.1 remote-as 65535
  neighbor 10.1.1.1 route-map map1 in
!
```

```
route-map map1 permit 10
  match as-path 1
!
ip as-path access-list 5 deny _200$
ip as-path access-list 5 permit .*
```

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/termserv/configuration/guide/12_4t/tsv_12_4t_book/tsv_r eg_express.html

QUESTION 152

Which statement describes the difference between a stub area and a totally stub area?

- A. The ABR advertises a default route to a totally stub area and not to a stub area.
- B. Stub areas do not allow LSA types 4 and 5, while totally stub areas do not allow LSA types 3, 4, and 5.
- C. Totally stub areas allow limited external routes in the area via a special type 7 LSA, while stub areas do not
- D. Stub areas do not allow external LSAs, ASBR summary LSAs, or summary LSAs with the exception of a default route originated by the ABR via a summary LSA.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

- **Standard areas** can contain LSAs of type 1, 2, 3, 4, and 5, and may contain an ASBR. The backbone is considered a standard area.
- **Stub areas** can contain type 1, 2, and 3 LSAs. A default route is substituted for external routes.
- **Totally stubby areas** can only contain type 1 and 2 LSAs, and a single type 3 LSA. The type 3 LSA describes a default route, substituted for all external and inter-area routes.
- **Not-so-stubby areas** implement stub or totally stubby functionality yet contain an ASBR. Type 7 LSAs generated by the ASBR are converted to type 5 by ABRs to be flooded to the rest of the OSPF domain.

Reference. <http://packetlife.net/blog/2008/jun/24/ospf-area-types/>

QUESTION 153

Which two statements are true about IS-IS? (Choose two.)

- A. IS-IS DIS election is nondeterministic.
- B. IS-IS SPF calculation is performed in three phases.
- C. IS-IS works over the data link layer, which does not provide for fragmentation and reassembly.
- D. IS-IS can never be routed beyond the immediate next hop.

Correct Answer: CD
Section: Layer 3 Technologies
Explanation

Explanation/Reference:

Explanation:

IS-IS runs directly over the data link alongside IP. On Ethernet, IS-IS packets are always 802.3 frames, with LSAPs 0xFEFE while IP packets are either Ethernet II frames or SNAP frames identified with the protocol number 0x800. OSPF runs over IP as protocol number 89.

IS-IS runs directly over layer 2 and hence:

- cannot support virtual links unless some explicit tunneling is implemented
- packets are kept small so that they don't require hop-by-hop fragmentation
- uses ATM/SNAP encapsulation on ATM but there are hacks to make it use VcMux encapsulation
- some operating systems that support IP networking have been implemented to differentiate Layer 3 packets in kernel. Such OSs require a lot of kernel modifications to support IS-IS for IP routing.
- can never be routed beyond the immediate next hop and hence shielded from IP spoofing and similar Denial of Service attacks.

Reference. <https://tools.ietf.org/html/draft-bhatia-manral-diff-isis-ospf-00>

QUESTION 154

Which command do you use to connect a dense-mode domain to a sparse-mode multicast domain?

- A. none, because there is no such command
- B. ip pim spt-threshold infinity
- C. ip pim register dense-mode
- D. ip pim dense-mode proxy-register

Correct Answer: D
Section: Layer 3 Technologies
Explanation

Explanation/Reference:

Explanation:

For IP PIM multicast, Cisco recommends Sparse-Mode over Dense-Mode. In the midst of our network migration, we have a new network operating in Sparse-Mode with Anycast rendezvous point (RP) but our existing network is still operating in Dense-Mode. To bridge two different modes across both PIM domains, we should use the ip pim dense-mode proxy-register command on the interface leading toward the bordering dense mode region. This configuration will enable the border router to register traffic from the dense mode region (which has no concept of registration) with the RP in the sparse mode domain.

Reference. <http://networkerslog.blogspot.com/2010/12/bridging-dense-mode-pim-to-sparse-mode.html>

QUESTION 155

Which two statements about the function of a PIM designated router are true? (Choose two.)

- A. It forwards multicast traffic from the source into the PIM network.
- B. It registers directly connected sources to the PIM rendezvous point.
- C. It sends PIM Join/Prune messages for directly connected receivers.
- D. It sends IGMP queries.
- E. It sends PIM asserts on the interfaces of the outgoing interface list.

Correct Answer: BC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

In PIM ASM and SSM modes, the software chooses a designated router (DR) from the routers on each network segment. The DR is responsible for forwarding multicast data for specified groups and sources on that segment.

In ASM mode, the DR is responsible for unicasting PIM register packets to the RP. When a DR receives an IGMP membership report from a directly connected receiver, the shortest path is formed to the RP, which may or may not go through the DR. The result is a shared tree that connects all sources transmitting on the same multicast group to all receivers of that group.

In SSM mode, the DR triggers (*, G) or (S, G) PIM join messages toward the RP or the source. The path from the receiver to the source is determined hop by hop. The source must be known to the receiver or the DR.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/multicast/configuration/guide/n7k_multic_cli_5x/pim.html#wp1054047

QUESTION 156

Refer to the exhibit.

```
router bgp 1
neighbor 10.1.1.1 remote-as 2
neighbor 10.1.1.1 ttl-security hops 2
```

Which IP packets will be accepted from EBGP neighbor 10.1.1.1?

- A. IP packets with a TTL count in the header that is equal to or greater than 253
- B. IP packets with a TTL count in the header that is equal to 253
- C. IP packets with a TTL count in the header that is equal to or greater than 2
- D. IP packets with a TTL count in the header that is equal to 2

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:


Explanation:

neighbor ip-address	Configures the maximum number of hops that
ttl-security hops	separate two peers.
hop-count	• __ The <i>hop-count</i> argument is set to number of
Example.	hops that separate the local and remote peer. If the
Router(config-	expected TTL value in the IP packet header is 254,
router)# neighbor	then the number 1 should be configured for the
10.1.1.1 ttl-security	<i>hop-count</i> argument. The range of values is a
hops 2	number from 1 to 254.
	• __ When this feature is enabled, BGP will accept
	incoming IP packets with a TTL value that is equal to
	or greater than the expected TTL value. Packets that
	are not accepted are silently discarded.
	• __ The example configuration sets the expected
	incoming TTL value to at least 253, which is 255
	minus the TTL value of 2, and this is the minimum
	TTL value expected from the BGP peer. The local
	router will accept the peering session from the
	10.1.1.1 neighbor only if it is 1 or 2 hops away.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_2s/feature/guide/fs_btsh.html

QUESTION 157

Refer to the exhibit.



A tunnel is configured between R3 to R4 sourced with their loopback interfaces. The `ip pim sparse-dense mode` command is configured on the tunnel interfaces and multicast-routing is enabled on R3 and R4. The IP backbone is not configured for multicast routing.

The RPF check has failed toward the multicast source.

Which two conditions could have caused the failure? (Choose two.)

- A. The route back to the RP is through a different interface than tunnel 0.
- B. The backbone devices can only route unicast traffic.
- C. The route back to the RP is through the same tunnel interface.
- D. A static route that points the RP to GigabitEthernet1/0 is configured.

Correct Answer: AD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

- For a successful RPF verification of multicast traffic flowing over the shared tree (*,G) from RP, an `ip mroute rp-address nexthop` command needs to be configured for the RP address, that points to the tunnel interface.

A very similar scenario can be found at the reference link below:

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/ip-multicast/43584-mcast-over-gre.html>

QUESTION 158

Which option is the default number of routes over which EIGRP can load balance?

- A. 1
- B. 4
- C. 8
- D. 16

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

By default, EIGRP load-shares over four equal-cost paths. For load sharing to happen, the routes to load-share over must show up in the IP forwarding table or with the show ip route command. Only when a route shows up in the forwarding table with multiple paths to it will load sharing occur.

Reference. http://www.informit.com/library/content.aspx?b=CCIE_Practical_Studies_I&seqNum=126

QUESTION 159

When EIGRP is used as the IPv4 PE-CE protocol, which two requirements must be configured before the BGP IPv4 address family can be configured? (Choose two.)

- A. the route distinguisher
- B. the virtual routing and forwarding instance
- C. the loopback interface
- D. the router ID

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

A VRF must be created, and a route distinguisher and route target must be configured in order for the PE routers in the BGP network to carry EIGRP routes to the EIGRP CE site. The VRF must also be associated with an interface in order for the PE router to send routing updates to the CE router.

Prerequisites

Before this feature can be configured, MPLS and CEF must be configured in the BGP network, and multiprotocol BGP and EIGRP must be configured

on all PE routers that provide VPN services to CE routers.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/seipecec.html#wp1083316

QUESTION 160

Which three EIGRP packet types are valid? (Choose three.)

- A. open
- B. notification
- C. keep-alive
- D. hello
- E. query
- F. reply

Correct Answer: DEF

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

EIGRP uses the following packet types: hello and acknowledgment, update, and query and reply. Hello packets are multicast for neighbor discovery/recovery and do not require acknowledgment. An acknowledgment packet is a hello packet that has no data. Acknowledgment packets contain a nonzero acknowledgment number and always are sent by using a unicast address.

Update packets are used to convey reachability of destinations. When a new neighbor is discovered, unicast update packets are sent so that the neighbor can build up its topology table. In other cases, such as a link-cost change, updates are multicast. Updates always are transmitted reliably.

Query and reply packets are sent when a destination has no feasible successors. Query packets are always multicast. Reply packets are sent in response to query packets to instruct the originator not to recompute the route because feasible successors exist. Reply packets are unicast to the originator of the query. Both query and reply packets are transmitted reliably.

Reference. http://docwiki.cisco.com/wiki/Enhanced_Interior_Gateway_Routing_Protocol

QUESTION 161

Which term describes an EIGRP route that has feasible successors?

- A. active
- B. passive
- C. redistributed
- D. invalid

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

A topology table entry for a destination can have one of two states. A route is considered in the Passive state when a router is not performing a route recomputation. The route is in Active state when a router is undergoing a route recomputation. If there are always feasible successors, a route never has to go into Active state and avoids a route recomputation.

When there are no feasible successors, a route goes into Active state and a route recomputation occurs. A route recomputation commences with a router sending a query packet to all neighbors. Neighboring routers can either reply if they have feasible successors for the destination or optionally return a query indicating that they are performing a route recomputation. While in Active state, a router cannot change the next-hop neighbor it is using to forward packets. Once all replies are received for a given query, the destination can transition to Passive state and a new successor can be selected.

Reference: http://docwiki.cisco.com/wiki/Enhanced_Interior_Gateway_Routing_Protocol

QUESTION 162

Refer to the exhibit.

```
Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=1, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 1
  EIGRP NSF-aware route hold timer is 240s
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    10.1.24.0/24
    10.1.34.0/24
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.24.2         90           00:00:28
    10.1.34.3         90           00:00:28
  Distance: internal 90 external 170
```

If EIGRP is configured between two routers as shown in this output, which statement about their EIGRP relationship is true?

- A. The routers will establish an EIGRP relationship successfully.
- B. The routers are using different authentication key-strings.
- C. The reliability metric is enabled.
- D. The delay metric is disabled.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The 5 K values used in EIGRP are:

K1 = Bandwidth modifier

K2 = Load modifier

K3 = Delay modifier

K4 = Reliability modifier

K5 = Additional Reliability modifier

However, by default, only K1 and K3 are used (bandwidth and delay). In this output we see that K1, K3, and K4 (Reliability) are all set.

QUESTION 163

Which type of OSPF packet is an OSPF link state update packet?

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

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Link State Update packets are OSPF packet type 4. These packets implement the flooding of link state advertisements. Each Link State Update packet carries a collection of link state advertisements one hop further from its origin. Several link state advertisements may be included in a single packet.

Reference. <http://www.freesoft.org/CIE/RFC/1583/107.htm>

QUESTION 164

If two OSPF type 3 prefixes have the same metric, and are within the same process, which prefix(es) are installed into the routing table?

- A. The route whose originator has the lower router ID.
- B. Both routes are installed.
- C. The route whose originator has the higher router ID.
- D. The first route that is learned.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

OSPF allows multiple equal-cost paths to the same destination. Since all link-state information is flooded and used in the SPF calculation, multiple equal cost paths can be computed and used for routing, and each route will be installed in the routing table.

QUESTION 165

Which OSPF feature supports LSA rate limiting in milliseconds to provide faster convergence?

- A. LSA throttling
- B. incremental SPF
- C. fast hello
- D. SPF tuning

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The OSPF Link-State Advertisement (LSA) Throttling feature provides a dynamic mechanism to slow down link-state advertisement (LSA) updates in OSPF during times of network instability. It also allows faster Open Shortest Path First (OSPF) convergence by providing LSA rate limiting in milliseconds.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fsolsath.html

QUESTION 166

Which two options are BGP attributes that are updated when router sends an update to its eBGP peer? (Choose two.)

- A. weight

- B. local preference
- C. AS_path
- D. next-hop

Correct Answer: CD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:


AS_Path describes the inter-AS path taken to reach a destination. It gives a list of AS Numbers traversed when reaching to a destination. Every BGP speaker when advertising a route to a peer will include its own AS number in the NLRI. The subsequent BGP speakers who advertise this route will add their own AS number to the AS_Path, the subsequent AS numbers get prepended to the list. The end result is the AS_Path attribute is able to describe all the autonomous systems it has traversed, beginning with the most recent AS and ending with the originating AS.

NEXT_HOP Attribute specifies the next hop IP address to reach the destination advertised in the NLRI. NEXT_HOP is a well-known mandatory attribute that is included in every eBGP update.

Reference. <http://netcerts.net/bgp-path-attributes-and-the-decision-process/>

QUESTION 167

Refer to the exhibit.



If ISIS is configured utilizing default metrics, what is the cost for Router 4 to reach the 10.2.2.0/24 network?

- A. 1
- B. 20
- C. 30
- D. 63

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

By default, every link in an IS-IS network has a metric of 10.

QUESTION 168

Refer to the exhibit.

```
ip route vrf red 0.0.0.0 0.0.0.0 192.168.1.1 global
```

Which three statements about this configuration are true? (Choose three.)

- A. The default route appears in the global routing table.
- B. The static route appears in the VRF red routing table.
- C. The subnet 192.168.1.0 is unique to the VRF red routing table.
- D. The static route is added to the global routing table and leaked from the VRF red.
- E. The subnet 192.168.1.0 is unique to the global routing table.
- F. 192.168.1.1 is reachable using any of the addresses on the router where the static route is configured.

Correct Answer: ABE

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

This is an example of the route leaking feature. Here, this static route is created for the red VRF so it will be installed into the red VRF routing table, but the use of the global keyword will cause this default route to appear in the global routing table.

QUESTION 169

Refer to the exhibit.

```
ip route 10.0.0.0 255.255.255.0 192.168.1.2
interface loopback0
 ip address 10.0.0.1 255.255.255.0
router rip
 network 10.0.0.0
router eigrp 1
 network 10.0.0.0
router ospf 1
 network 10.0.0.0 0.0.0.255 area 0
```

Which route type is displayed when you enter the command show ip route supernets-only on a device with this configuration?

- A. Connected
- B. OSPF
- C. RIP
- D. EIGRP
- E. An empty route set

Correct Answer: E

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

This command shows supernets only; it does not show subnets. In this case, the routing table would contain the 10.0.0.0/24 subnet, but not the 10.0.0.0/8 supernet.

QUESTION 170

Which statement about passive interfaces is true?

- A. The interface with the OSPF passive interface configuration appears as a not-so-stubby network.
- B. The interface with the EIGRP passive interface configuration ignores routes after the exchange of hello packets.
- C. The interface with the IS-IS passive interface configuration sends the IP address of that interface in the link-state protocol data units.
- D. Passive interface can be configured on the interface for IS-IS.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

With IS-IS, passive interfaces are used to prevent unnecessary LSA packets out that interface, but the IP address of passive interfaces are still included in updates going out the other interfaces. This behavior is what enables the best practice of configuring loopback interfaces as passive, but still having the loopback be reachable.

QUESTION 171

Refer to the exhibit.

```
access-list 1 permit 10.3.5.0 0.0.3.255
router eigrp 1
  network 10.0.0.0
  no auto-summary
  distribute-list 1 out
```

Which two routes are included in the route update? (Choose two.)

- A. 10.3.0.0
- B. 10.3.2.0
- C. 10.3.4.0
- D. 10.3.6.0

Correct Answer: CD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

This access list will permit the 10.3.4.0, 10.3.5.0, 10.3.6.0, and 10.3.7.0 subnets.

QUESTION 172

Which two statements about the metric-style wide statement as it applies to route redistribution are true? (Choose two.)

- A. It is used in IS-IS.
- B. It is used in OSPF.
- C. It is used in EIGRP.
- D. It is used for accepting TLV.
- E. It is used in PIM for accepting mroutes.
- F. It is used for accepting external routes.

Correct Answer: AD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

To configure a router running IS-IS to generate and accept only new-style TLVs (TLV stands for type, length, and value object), use the metric-style wide

command.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/TE_1208S.html#wp49409

QUESTION 173

You are tasked with configuring a router on an OSPF domain to import routes from an EIGRP domain and summarize the routes to 192.168.64.0.

Which statement configures the summarized route and provides equal-path route redundancy?

- A. area 32 range 192.168.64.0 255.255.192.0 cost 100
- B. area 32 range 192.168.64.0 255.255.63.0 cost 100
- C. area 32 range 192.168.64.0 255.255.64.0 cost 100
- D. area 32 range 192.168.64.0 255.255.192.0 multi-path

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

area range

To consolidate and summarize routes at an area boundary, use the **area range** command in router configuration mode. To disable this function, use the **no** form of this command.

Syntax Description

<i>area-id</i>	Identifier of the area about which routes are to be summarized. It can be specified as either a decimal value or as an IP address.
<i>ip-address</i>	IP address.
<i>mask</i>	IP address mask
advertise	(Optional) Sets the address range status to advertise and generates a Type 3 summary link-state advertisement (LSA).
not-advertise	(Optional) Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed, and the component networks remain hidden from other networks.
cost cost	(Optional) Metric or cost for this summary route, which is used during OSPF SPF calculation to determine the shortest paths to the destination. The value can be 0 to 16777215.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_2/iproute/command/reference/fiprrp_r/1rfospf.html#w p1017596

QUESTION 174

Packets from a router with policy-based routing configured are failing to reach the next hop.

Which two additions can you make to the router configuration to enable the packets to flow correctly? (Choose two.)

- A. Enable ip proxy-arp on the exiting interface.
- B. Specify the next hop as an address.
- C. Specify the next hop as an interface.
- D. Add a match-any permit statement to the route map.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Here is an example:

```
Router(config)#route-map Engineers permit 20
Router(config-route-map)#match ip address 2
Router(config-route-map)#set interface Ethernet1
```

Here, instead of specifying a next-hop, it specifies that any packets matching this rule will be forwarded directly out the interface Ethernet1. This means that either the destination device must be on this segment, or there must be a router configured with Proxy ARP that can forward the packet to the ultimate destination.

QUESTION 175

Which two options are EIGRP route authentication encryption modes? (Choose two.)

- A. MD5
- B. HMAC-SHA-256bit
- C. ESP-AES
- D. HMAC-AES

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Packets exchanged between neighbors must be authenticated to ensure that a device accepts packets only from devices that have the same preshared authentication key. Enhanced Interior Gateway Routing Protocol (EIGRP) authentication is configurable on a per-interface basis; this means that packets exchanged between neighbors connected through an interface are authenticated. EIGRP supports message digest algorithm 5 (MD5) authentication to prevent the introduction of unauthorized information from unapproved sources. MD5 authentication is defined in RFC 1321. EIGRP also supports the Hashed Message Authentication Code-Secure Hash Algorithm-256 (HMAC-SHA-256) authentication method.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/xr-3s/ire-xr-3s-book/ire-sha-256.html

QUESTION 176

Which neighbor-discovery message type is used to verify connectivity to a neighbor when the link-layer address of the neighbor is known?

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

Correct Answer: A

Section: Layer 3 Technologies

Explanation

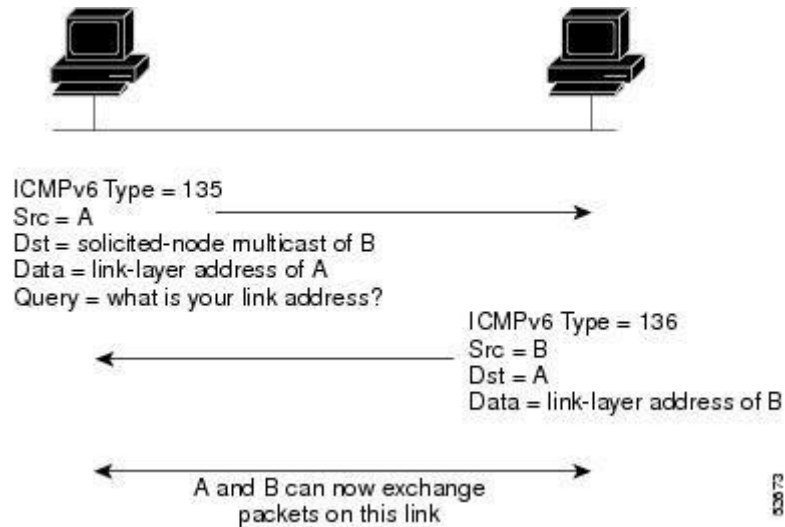
Explanation/Reference:

Explanation:

IPv6 Neighbor Solicitation Message

A value of 135 in the Type field of the ICMP packet header identifies a neighbor solicitation message. Neighbor solicitation messages are sent on the local link when a node wants to determine the link-layer address of another node on the same local link (see the figure below). When a node wants to determine the link-layer address of another node, the source address in a neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The destination address in the neighbor solicitation message is the solicited-node multicast address that corresponds to the IPv6 address of the destination node. The neighbor solicitation message also includes the link-layer address of the source node.

Figure 1. IPv6 Neighbor Discovery: Neighbor Solicitation Message



Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_basic/configuration/xr-3s/ipv6b-xr-3s-book/ipv6-neighb-disc-xr.html

QUESTION 177

DRAG DROP

Drag and drop the IPv6 prefix on the left to the correct address type on the right.

Select and Place:

Drag and drop the IPv6 prefix on the left to the correct address type on the right.

FF00::/8	Unique Local Unicast
FEC0::/10	Global Unicast
2000::/3	Link Local Unicast
FE80::/10	Multicast
FC00::/7	
FE00::/9	

Correct Answer:

Drag and drop the IPv6 prefix on the left to the correct address type on the right.

	FC00::/7
FEC0::/10	2000::/3
	FE80::/10
	FF00::/8
FE00::/9	

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 178

DRAG DROP

Drag and drop the BGP attribute on the left to the correct category on the right.

Select and Place:

Drag and drop the BGP attribute on the left to the correct category on the right.	
Community	BGP Well-Known Mandatory Attribute
Atomic-Aggregate	Target
Aggregator	BGP Well-Known Discretionary Attribute
Cluster List	Target
Next-Hop	BGP Optional Nontransitive Attribute
MED	Target
	Target

Correct Answer:

Drag and drop the BGP attribute on the left to the correct category on the right.

Community
Aggregator

BGP Well-Known Mandatory Attribute

Next-Hop

BGP Well-Known Discretionary Attribute

Atomic-Aggregate

BGP Optional Nontransitive Attribute

Cluster List

MED

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 179

DRAG DROP

Drag and drop the protocol on the left to the corresponding administrative distance on the right.

Select and Place:

Drag and drop the protocol on the left to the corresponding administrative distance on the right.

ODR	0
connected	1
external EIGRP	160
static	115
IS-IS	200
iBGP	170

Correct Answer:

Drag and drop the protocol on the left to the corresponding administrative distance on the right.

	connected
	static
	ODR
	IS-IS
	iBGP
	external EIGRP

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 180

DRAG DROP

Drag and drop the BGP state on the left to the action that defines it on the right.

Select and Place:

Drag and drop the BGP state on the left to the action that defines it on the right.	
OpenConfirm	The BGP routing process detects that a peer is trying to establish a TCP session with a local BGP speaker.
Idle	The BGP routing process tries to establish a TCP session with a peer device.
Active	The TCP connection is established.
Connect	The BGP routing process waits to receive an initial keepalive message from the peer.
Established	The initial BGP state.
OpenSent	The router exchanges update messages with the peer.

Correct Answer:

Drag and drop the BGP state on the left to the action that defines it on the right.

	Connect
	Active
	OpenSent
	OpenConfirm
	Idle
	Established

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 181

DRAG DROP

Drag and drop the method for refreshing BGP prefixes on the left to the corresponding description on the right.

Select and Place:

Drag and drop the method for refreshing BGP prefixes on the left to the corresponding description on the right.

hard reset	requests a complete refresh of the Adj-RIB-Out
soft reset	tears down the peering session and deletes prefixes from the peer
dynamic inbound soft reset	uses extra prefix information stored locally
Enhanced Route Refresh	finds route inconsistencies and synchronizes with the peer

Correct Answer:

Drag and drop the method for refreshing BGP prefixes on the left to the corresponding description on the right.

	dynamic inbound soft reset
	hard reset
	soft reset
	Enhanced Route Refresh

Section: Layer 3 Technologies
Explanation

Explanation/Reference:

QUESTION 182

DRAG DROP

Drag and drop the IS-IS component on the left to the function that it performs on the right.

Select and Place:

Drag and drop the IS-IS component on the left to the function that it performs on the right.

attached bits	instructs other devices to route around the sending device until its LSDB is fully converged
overload bit	discovers neighboring IS-IS systems
TLV	carries additional data within an IS-IS packet
IIH	synchronizes the LSDB within an IS-IS domain
PNSP	indicates to a Level 1 device that the sending device has reachability to other areas
CNSP	requests retransmission of the latest version of an LSP

Correct Answer:

Drag and drop the IS-IS component on the left to the function that it performs on the right.

	overload bit
	IIH
	TLV
	CNSP
	attached bits
	PNSP

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 183

DRAG DROP

Drag and drop the RIP configuration command on the left to the function it performs on the right.

Select and Place:

Drag and drop the RIP configuration command on the left to the function it performs on the right.

ip rip triggered	configures the router to verify the IP address of routers that send updates
output-delay	configures the router to send information only when the routing database is updated
validate-update-source	configures the router to modify routing metrics
offset-list	configures the router to throttle RIP updates

Correct Answer:

Drag and drop the RIP configuration command on the left to the function it performs on the right.


	validate-update-source
	ip rip triggered
	offset-list
	output-delay

Section: Layer 3 Technologies
Explanation

Explanation/Reference:

QUESTION 184

Refer to the exhibit.



Which action must you take to enable full reachability from router C to router D?

- A. Build an OSPF virtual link.
- B. Build an OSPF sham link.
- C. Configure mutual redistribution between OSPF and EIGRP on routers A and B.
- D. Add a static route on router D.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

For full connectivity, we need to configure mutual redistribution to advertise the EIGRP routes into OSPF and to advertise the OSPF routes into the EIGRP network. This needs to be done at the two border routers that connect to both the EIGRP and OSPF domains.

QUESTION 185

Refer to the exhibit.

```
Router#show version
```

```
Router processor (revision 0x00) with 524288K bytes of memory.
```

```
Router#show memory statistics
```

	Head	Total (b)	Used (b)	Free (b)	Lowest (b)	Largest (b)
Processor	38A6400	405117952	360086164	1031788	37130412	34036896

```
Router#show process memory
```

PID	TTY	Allocated	Freed	Holding	Getbufs	Retbufs	Process
0	0	73373216	1706280	69497168	0	0	*Init*
154	0	1103256760	1247933568	311905892	204360	0	BGP Router
327	0	212528944	322521272	44071084	0	0	IP RIB Update

```
Router#show ip bgp summary
```

```
BGP router identifier 1.1.1.1, local AS number 65000
BGP table version is 310248959, main routing table version 310248959
246316 network entries using 29557920 bytes of memory
1586197 path entries using 76137456 bytes of memory
256960/41528 BGP path/bestpath attribute entries using 27751680 bytes of memory
440 BGP rrinfo entries using 10560 bytes of memory
115467 BGP AS-PATH entries using 3047538 bytes of memory
5952 BGP community entries using 479704 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
230723 BGP filter-list cache entries using 2768676 bytes of memory
BGP using 139753534 total bytes of memory
Dampening enabled. 8 history paths, 0 dampened paths
631350 received paths for inbound soft reconfiguration
BGP activity 9798913/9552597 prefixes, 220384574/218798377 paths, scan interval 60 secs
Neighbor      V    AS  MsgRcvd  MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
1.1.1.2        4  65001  39985912  1384531  310248959    0    0  9w1d    277030
1.1.1.3        4  65001  12269759   529250  310248959    0    0  26w0d    276929
1.1.1.4        4  65001  42728751 20209410  310248959    0    0  32w2d    200372
1.1.1.5        4  65001  46624114 20179383  310248959    0    0  1y14w    200372
```

Why is the router out of memory?

- A. The router is experiencing a BGP memory leak software defect.
- B. The BGP peers have been up for too long.
- C. The amount of BGP update traffic in the network is too high.
- D. The router has insufficient memory due to the size of the BGP database.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Here we can see that this router is running out of memory due to the large size of the BGP routing database. In this case, this router is receiving over 200,000 routes from each of the 4 peers.

QUESTION 186

Refer to the exhibit.

```
R1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.2.2	0	2WAY/DROTHER	00:00:35	10.25.123.2	Ethernet0/0
192.168.3.3	0	2WAY/DROTHER	00:00:38	10.25.123.3	Ethernet0/0

```
R1#
```

Why is the OSPF state in 2WAY/DROTHER?

- A. This is the expected output when the interface Ethernet0/0 of R1 is configured with OSPF Priority 0.
- B. There is a duplicate router ID.
- C. There is an MTU mismatch.
- D. There is an OSPF timer (hello/dead) mismatch.
- E. This is the expected output when R1 is the DR.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Sometimes it is desirable for a router to be configured so that it is not eligible to become the DR or BDR. You can do this by setting the OSPF priority to zero with the `ip ospf priority priority# interface` subcommand. If two OSPF neighbors both have their OSPF interface priority set to zero, they establish two-way adjacency instead of full adjacency.

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

QUESTION 187

In a nonbackbone OSPF area, all traffic that is destined to the Internet is routed by using a default route that is originated by the ABR. Which change in the configuration of the OSPF area type causes traffic from that area that is destined to the Internet to be dropped?

- A. The OSPF area changes from NSSA to totally stubby area.
- B. The OSPF area changes from NSSA to regular area.
- C. The OSPF area changes from stub area to totally stubby area.
- D. The OSPF area changes from stub area to NSSA.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

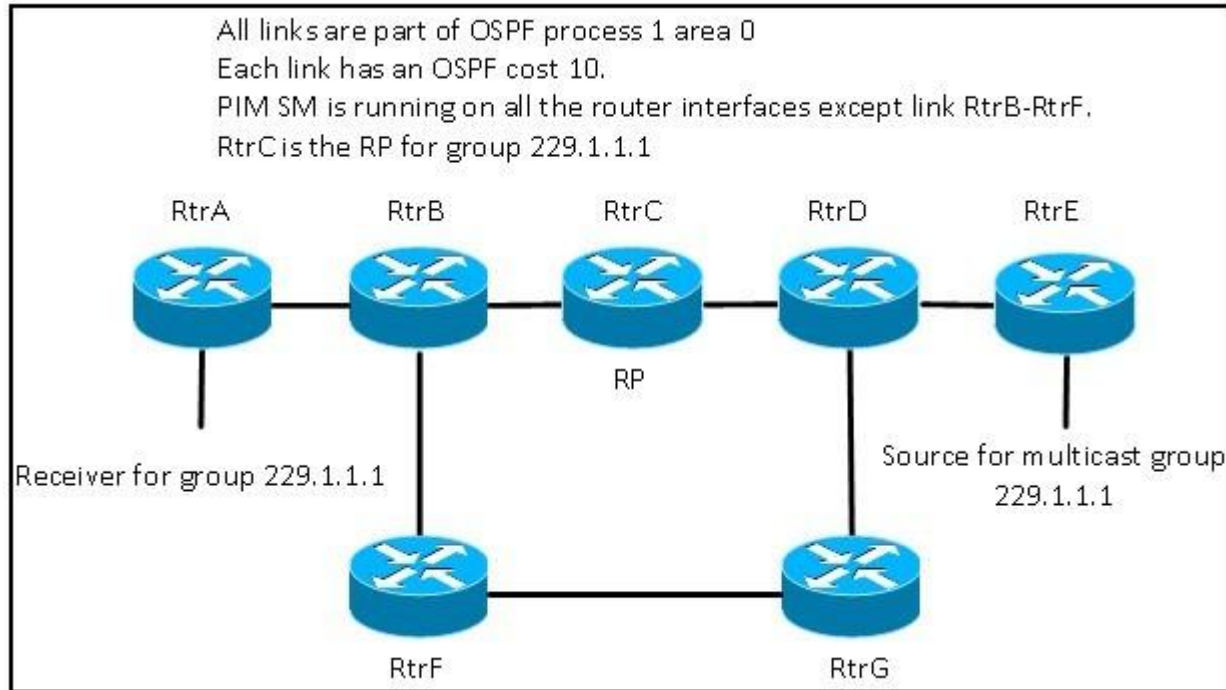
Explanation:

The ABR for the NSSA generates the default route, but not by default. To force the ABR to generate the default route, use the area `<area id> nssa default-information originate` command. The ABR generates a Type 7 LSA with the link-state ID 0.0.0.0 and is advertised inside the NSSA. This default route will be propagated inside the NSSA as Type 7 LSA

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13692-21.html#nssas>

QUESTION 188

Refer to the exhibit.



When the link between RtrB and RtrC goes down, multicast receivers stop receiving traffic from the source for multicast group 229.1.1.1. Which solution will resolve this?

- A. adding a static mroute on RtrB and RtrF
- B. adding a static unicast route on RtrB and RtrF
- C. creating a GRE tunnel between RtrB and RtrD
- D. enabling PIM sparse mode on both ends of the link between RtrB and RtrF

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

For multicast traffic to flow, PIM must be enabled on all routers in the path of the multicast stream.

QUESTION 189

Which measure does ISIS use to avoid sending traffic with a wrong MTU configuration?

- A. ISIS does not protect from MTU mismatch.
- B. MTU value is communicated in ISIS Sequence Number PDUs (SNP), and ISIS adjacency is not established if an MTU mismatch is detected.
- C. ISIS uses path MTU discovery as specified in RFC 1063.
- D. ISIS uses padding of hello packets to full MTU.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Intermediate System-to-Intermediate System (IS-IS) hellos are padded to the full maximum transmission unit (MTU) size. The benefit of padding IS-IS Hellos (IIHs) to the full MTU is that it allows for early detection of errors due to transmission problems with large frames or due to mismatched MTUs on adjacent interfaces.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/integrated-intermediate-system-to-intermediate-system-is-is/47201-isis-mtu.html>

QUESTION 190

Which regular expression will match prefixes from the AS 200 that is directly connected to our AS?

- A. ^\$
- B. ^200)
- C. _200\$
- D. _200_
- E. ^200_

Correct Answer: E

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Table 2 Commonly Used Regular Expressions

Expression	Meaning
.	Anything
^\$	Locally originated routes
^100_	Learned from autonomous system 100
_100\$	Originated in autonomous system 100
100	Any instance of autonomous system 100
^[0-9]+\$	Directly connected autonomous system paths

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/termserv/configuration/guide/12_4t/tsv_12_4t_book/tsv_reg_express.html

QUESTION 191

Refer to the exhibit.

```
!
interface Loopback10
  no ip address
  ipv6 address 6010:AB8::/64 eui-64
!
interface Loopback20
  no ip address
  ipv6 address 6020:AB8::/64 eui-64
!
interface Ethernet0/0
  no ip address
  ipv6 enable
  ipv6 eigrp 50
!
ipv6 router eigrp 50
!
```

Assuming that the peer is configured correctly and the interface is up, how many neighbors will be seen in the EIGRPv6 neighbor table on this IPv6-only router?

- A. one neighbor, which will use a local router-id of 6010. AB8. . /64
- B. one neighbor, which will use a local router-id of 6020. AB8. . /64
- C. none, because EIGRPv6 only supports authenticated peers
- D. none, because of the mismatch of timers
- E. none, because there is no EIGRP router ID configured

Correct Answer: E

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Configuring EIGRP for IPv6 has some restrictions; they are listed below:

- The interfaces can be directly configured with EIGRP for IPv6, without the use of a global IPv6 address. There is no network statement in EIGRP for IPv6.
- The router ID needs to be configured for an EIGRPv6 protocol instance before it can run.
- EIGRP for IPv6 has a shutdown feature. Ensure that the routing process is in "no shut" mode to start running the protocol.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/113267-eigrp-ipv6-00.html>

QUESTION 192

What does a nonzero forwarding address indicate in a type-5 LSA?

- A. It indicates that this link-state ID is eligible for ECMP.
- B. It indicates that this router should have an OSPF neighbor relationship with the forwarding address before using this link-state ID.
- C. It indicates that the receiving router must check that the next hop is reachable in its routing table before using this link-state ID.
- D. It indicates that traffic can be directly routed to this next hop in shared segment scenarios where the external route source is directly connected.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The value of the forwarding address specified by the autonomous system boundary router (ASBR) can be either 0.0.0.0 or non-zero. The 0.0.0.0 address indicates that the originating router (the ASBR) is the next hop. The forwarding address is determined by these conditions:

- The forwarding address is set to 0.0.0.0 if the ASBR redistributes routes and OSPF *is not enabled* on the next hop interface for those routes.
- These conditions set the forwarding address field to a non-zero address:
 - OSPF is enabled on the ASBR's next hop interface AND
 - ASBR's next hop interface is non-passive under OSPF AND

ASBR's next hop interface is not point-to-point AND
 ASBR's next hop interface is not point-to-multipoint AND
 ASBR's next hop interface address falls under the network range specified in the router ospf command.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13682-10.html>

QUESTION 193

Which type of EIGRP routes are summarized by the auto-summary command?

- A. internal routes that are learned from a peer that is outside the range of local network statements
- B. external routes that are learned from a peer that is inside the range of local network statements
- C. locally created routes that are outside the range of local network statements
- D. external routes that are learned from a peer that is outside the range of local network statements

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Auto-Summarization of External Routes

EIGRP will not auto-summarize external routes unless there is a component of the same major network that is an internal route. To illustrate, let us look at Figure 15.




Figure 15

Router Three is injecting external routes to 192.1.2.0/26 and 192.1.2.64/26 into EIGRP using the **redistribute connected** command, as shown in the configurations below.

Router Three

```
interface Ethernet0
ip address 192.1.2.1 255.255.255.192
!
```

```
interface Ethernet1
ip address 192.1.2.65 255.255.255.192
!
```

```
interface Ethernet2
ip address 10.1.2.1 255.255.255.0
!router eigrp 2000
redistribute connected
network 10.0.0.0
default-metric 10000 1 255 1 1500
```

With this configuration on Router Three, the routing table on Router One shows:

```
one# show ip route
```

```
....
 10.0.0.0/8 is subnetted, 2 subnets
D   10.1.2.0 [90/11023872] via 10.1.50.2, 00:02:03, Serial0
C   10.1.50.0 is directly connected, Serial0
 192.1.2.0/26 is subnetted, 1 subnets
D EX 192.1.2.0 [170/11049472] via 10.1.50.2, 00:00:53, Serial0
D EX 192.1.2.64 [170/11049472] via 10.1.50.2, 00:00:53, Serial0
```

Although auto-summary normally causes Router Three to summarize the 192.1.2.0/26 and 192.1.2.64/26 routes into one major net destination (192.1.2.0/24), it does not do this because both routes are external. However, if you reconfigure the link between Routers Two and Three to 192.1.2.128/26, and add network statements for this network on Routers Two and Three, the 192.1.2.0/24 auto-summary is then generated on Router Two.

Router Three

```
interface Ethernet0
ip address 192.1.2.1 255.255.255.192
!
interface Ethernet1
ip address 192.1.2.65 255.255.255.192
!
interface Serial0
ip address 192.1.2.130 255.255.255.192
!
```

```
router eigrp 2000
network 192.1.2.0
```

Now Router Two generates the summary for 192.1.2.0/24:

```
two# show ip route
```

```
....
D   192.1.2.0/24 is a summary, 00:06:48, Null0
....
```

And Router One shows only the summary route.

```
one# show ip route
```

```
....
 10.0.0.0/8 is subnetted, 1 subnets
C   10.1.1.0 is directly connected, Serial0
```

D 192.1.2.0/24 [90/11023872] via 10.1.50.2, 00:00:36, Serial0

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

QUESTION 194

Refer to the exhibit.

```
router eigrp foo
!
address-family ipv4 unicast autonomous-system 1
!
af-interface default
hello-interval 10
hold-time 30
exit-af-interface
!
topology base
exit-af-topology
network 10.0.0.0
exit-address-family
```

How can the EIGRP hello and hold time for Gig0/0 be changed to 5 and 15?

- A. No action is required, since Gig0/0 is not listed with a nondefault hello and hold time.
- B. Add the commands `ip hello-interval eigrp 1 5` and `ip hold-time eigrp 1 15` under interface Gig0/0.
- C. Add the commands `hello-interval 5` and `hold-time 15` under "af-interface Gig0/0" under the address family.
- D. Add the commands `default hello-interval` and `default hold-time` under the af-interface Gig0/0 statement under the address family.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

To configure the hello interval for an interface, use the `hello-interval` command in interface configuration mode

To configure the hold time for an interface, use the `hold-time` command in interface configuration mode.

Reference. http://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-1/routing/command/reference/b_routing_cr41crs/b_routing_cr41crs_chapter_010.html#wp2323069

QUESTION 195

What is the range of addresses that is used for IPv4-mapped IPv6 addresses?

- A. 2001. db9. . /32
- B. 2001. db8. . /32
- C. 2002. . /16
- D. . . ffff. /16
- E. . . ffff. 0. 0/96

Correct Answer: E

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

IPv4-Mapped Addresses

::FFFF:0:0/96 are the IPv4-mapped addresses [RFC4291]. Addresses within this block should not appear on the public Internet.

Reference. <https://tools.ietf.org/html/rfc5156>

QUESTION 196

Which statement about the overload bit in IS-IS is true?

- A. The IS-IS adjacencies on the links for which the overload bit is set are brought down.
- B. Routers running SPF ignore LSPs with the overload bit set and hence avoid blackholing traffic.
- C. A router setting the overload bit becomes unreachable to all other routers in the IS-IS area.
- D. The overload bit in IS-IS is used only for external prefixes.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The OL bit is used to prevent unintentional blackholing of packets in BGP transit networks. Due to the nature of these protocols, IS-IS and OSPF converge must faster than BGP. Thus there is a possibility that while the IGP has converged, IBGP is still learning the routes. In that case if other IBGP routers start sending traffic towards this IBGP router that has not yet completely converged it will start dropping traffic. This is because it isn't yet aware of the complete BGP routes. OL bit comes handy in such situations. When a new IBGP neighbor is added or a router restarts, the IS-IS OL bit is set. Since directly connected (including loopbacks) addresses on an "overloaded" router are considered by other routers, IBGP can be brought up and can begin exchanging routes. Other routers will not use this router for transit traffic and will route the packets out through an alternate path. Once BGP has

converged, the OL bit is cleared and this router can begin forwarding transit traffic.
Reference: <https://routingfreak.wordpress.com/category/ospf-vs-is-is/>

QUESTION 197

Refer to the exhibit.

```
R2#show ip mroute 225.1.1.1
(*, 225.1.1.1), 01:32:54/00:03:06, RP 10.100.1.2, flags: SJC
  Incoming interface: Ethernet1/0, RPF nbr 10.1.3.2
  Outgoing interface list:
    Ethernet3/0, Forward/Sparse, 01:32:54/00:03:06

(10.1.4.7, 225.1.1.1), 01:32:54/00:01:05, flags: JT
  Incoming interface: Ethernet1/0, RPF nbr 10.1.3.2
  Outgoing interface list:
    Ethernet3/0, Forward/Sparse, 00:37:38/00:02:26, A
```

Which statement is true?

- A. R2 is directly connected to the receiver for this group and is the winner of an assert mechanism.
- B. R2 is directly connected to the receiver for this group, and it forwards the traffic onto Ethernet3/0, but it is forwarding duplicate traffic onto Ethernet3/0.
- C. R2 has the A flag (Accept flag) set on Ethernet 3/0. This is fine, since the group is in BIDIR PIM mode.
- D. R2 is directly connected to the receiver for this group and is the loser of an assert mechanism.
- E. The A flag is set until the SPT threshold is reached for this multicast group.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

show ip mroute Field Descriptions

Field	Description
RPF neighbor or RPF nbr	IP address of the upstream router to the source. Tunneling indicates that this router is sending data to the RP encapsulated in register packets. The hexadecimal number in parentheses indicates to which RP it is registering. Each bit indicates a different RP if multiple RPs per group are used. If an asterisk (*) appears after the IP address in this field, the RPF neighbor has been learned through an assert.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_2/ipmulti/command/reference/fiprmc_r/rfmult3.html

QUESTION 198

Which three statements about IS-IS are true? (Choose three.)

- A. IS-IS is not encapsulated in IP.
- B. IS-IS is directly encapsulated in the data link layer.
- C. 0xFEFE is used in the Layer 2 header to identify the Layer 3 protocol.
- D. IS-IS uses protocol ID 93.
- E. IS-IS can be used to route the IPX protocol.
- F. IS-IS is an IETF standard.

Correct Answer: ABC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

IS-IS is an Interior Gateway Protocol (IGP) for routing OSI. IS-IS packets are not encapsulated in CLNS or IP but are encapsulated directly in the data-link layer. The IS-IS protocol family is OSI, and values such as 0xFE and 0xFEFE are used by the data-link protocol to identify the Layer 3 protocol as OSI.

Reference.

http://www.cisco.com/en/US/products/ps6599/products_white_paper09186a00800a3e6f.shtml

QUESTION 199

Refer to the exhibit.

```
PE1#show ip rpf 10.100.1.4
RPF information for ? (10.100.1.4)
  RPF interface: Ethernet1/0
  RPF neighbor: ? (10.1.1.4)
  RPF route/mask: 10.100.1.4/32
  RPF type: multicast (isis)
  Doing distance-preferred lookups across tables
  RPF topology: ipv4 multicast base
```

Which statement is true?

- A. The command `ip multicast rpf multitopology` is missing from the configuration.
- B. Multitopology routing for multicast has been enabled for IS-IS.
- C. This output is invalid.
- D. The command `mpls traffic-eng multicast-intact` is configured on this router.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The following is sample output from the `show ip rpf` command in a Multi-Topology Routing (MTR) routing environment. In Cisco IOS releases that support MTR, the "RPF topology" field was introduced to indicate which RIB topology is being used for the RPF lookup. For the "RPF topology" field in this example, the first topology listed (`ipv4 multicast base`) indicates where the nexthop of the RPF lookup is being conducted and the second topology listed (`ipv4 unicast data`) indicates where the route originated from.

Router# **show ip rpf 10.30.30.32**

RPF information for ? (10.30.30.32)

RPF interface: Ethernet1/0

RPF neighbor: ? (10.1.1.32)

RPF route/mask: 10.30.30.32/32

RPF type: unicast (ospf 100)

Doing distance-preferred lookups across tables

RPF topology: ipv4 multicast base, originated from ipv4 unicast data

The table below describes the fields shown in the displays.

Table 15 show ip rpf Field Descriptions

Field	Description
RPF information for	Hostname and source address for which RPF information is displayed.
RPF interface	For the given source, the interface from which the router expects to receive packets.
RPF neighbor	For the given source, the neighbor from which the router expects to receive packets.
RPF route/mask	Route number and mask that matched against this source.
RPF type	Routing table from which this route was obtained, either unicast, MBGP, DVMRP, or static mroutes.
RPF recursion count	The number of times the route is recursively resolved.
Doing distance-preferred	Whether RPF was determined based on distance or length of mask.
Using Group Based VRF Select, RPF VRF.	The RPF lookup was based on the group address and the VRF where the RPF lookup is being performed.
Metric preference	The preference value used for selecting the unicast routing metric to the RP announced by the designated forwarder (DF).
Metric	Unicast routing metric to the RP announced by the DF.
RPF topology	RIB topology being used for the RPF lookup, and, if originated from a different RIB topology, which RIB topology the route originated from.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/command/imc-xe-3se-5700-cr-book/imc-xe-3se-3850-cr-book_chapter_010.html

QUESTION 200

As a best practice, when a router is configured as an EIGRP Stub, which routes should be received from its distribution neighbor?

- A. the default route
- B. static routes
- C. internal routes only
- D. internal and external routes

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Stub routing is commonly used in a hub and spoke network topology. In a hub and spoke network, one or more end (stub) networks are connected to a remote router (the spoke) that is connected to one or more distribution routers (the hub). The remote router is adjacent only to one or more distribution routers. The only route for IP traffic to follow into the remote router is through a distribution router. This type of configuration is commonly used in WAN topologies where the distribution router is directly connected to a WAN. The distribution router can be connected to many more remote routers. Often, the distribution router will be connected to 100 or more remote routers. In a hub and spoke topology, the remote router must forward all nonlocal traffic to a distribution router, so it becomes unnecessary for the remote router to hold a complete routing table. Generally, the distribution router need not send anything more than a default route to the remote router.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/eigrpstb.html

QUESTION 201

Which BGP feature allows BGP routing tables to be refreshed without impacting established BGP sessions?

- A. BGP synchronization
- B. soft reconfiguration
- C. confederations
- D. hard reset

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Clearing a BGP session using a hard reset invalidates the cache and results in a negative impact on the operation of networks as the information in the cache becomes unavailable.

Soft reset is recommended because it allows routing tables to be reconfigured and activated without clearing the BGP session. Soft reset is done on a per-neighbor basis.

Reference.

http://www.cisco.com/en/US/products/ps6599/products_data_sheet09186a0080087b3a.html

QUESTION 202

Which two options describe two functions of a neighbor solicitation message? (Choose two.)

- A. It requests the link-layer address of the target.
- B. It provides its own link-layer address to the target.
- C. It requests the site-local address of the target.
- D. It provides its own site-local address to the target.

- E. It requests the admin-local address of the target.
- F. It provides its own admin-local address to the target.

Correct Answer: AB

Section: Layer 3 Technologies

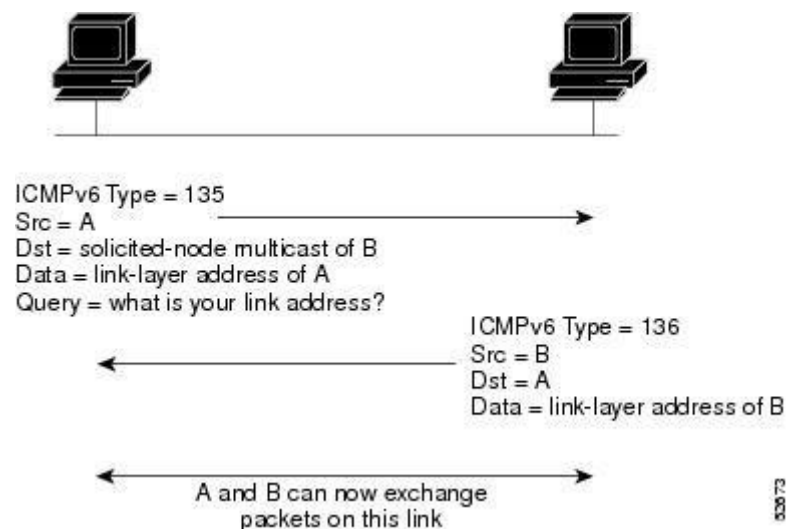
Explanation

Explanation/Reference:

Explanation:

Neighbor solicitation messages are sent on the local link when a node wants to determine the link-layer address of another node on the same local link (see the figure below). When a node wants to determine the link-layer address of another node, the source address in a neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The destination address in the neighbor solicitation message is the solicited-node multicast address that corresponds to the IPv6 address of the destination node. The neighbor solicitation message also includes the link-layer address of the source node.

Figure 1. IPv6 Neighbor Discovery: Neighbor Solicitation Message



After receiving the neighbor solicitation message, the destination node replies by sending a neighbor advertisement message, which has a value of 136 in the Type field of the ICMP packet header, on the local link. The source address in the neighbor advertisement message is the IPv6 address of the node (more specifically, the IPv6 address of the node interface) sending the neighbor advertisement message. The destination address in the neighbor advertisement message is the IPv6 address of the node that sent the neighbor solicitation message. The data portion of the neighbor advertisement message includes the link-layer address of the node sending the neighbor advertisement message.

After the source node receives the neighbor advertisement, the source node and destination node can communicate.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_basic/configuration/xr-3s/ip6b-xe-3s-book/ip6-neighb-disc-xe.html

QUESTION 203

Which three options are three of the default EIGRP administrative distances? (Choose three.)

- A. Internal, 90
- B. External, 170
- C. Summary, 5
- D. Outside Local, 100
- E. Inside Local, 180
- F. Inside Global, 1

Correct Answer: ABC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The following table lists the default administrative distances for various routing protocols used on Cisco routers.

Routing Protocol	Administrative distance
Directly connected interface	0
Static route out an interface	1
Static route to next-hop address	1
DMNR - Dynamic Mobile Network Routing	3
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
Routing Information Protocol (RIP)	120
Exterior Gateway Protocol (EGP)	140
On Demand Routing (ODR)	160
External EIGRP	170
Internal BGP	200
Floating Static Route (ex. DHCP-learned)	254
Unknown	255

Reference. http://en.wikipedia.org/wiki/Administrative_distance

QUESTION 204

Refer to the exhibit.

```

O E2    172.17.108.128/25
        [110/20] via 10.169.73.12, 3d07h, TenGigabitEthernet8/0/0
O E2    10.167.111.216/29
        [110/20] via 10.169.73.12, 3d07h, TenGigabitEthernet8/0/0
O IA    10.68.2.0/31
        [110/489] via 10.169.73.12, 3d07h, TenGigabitEthernet8/0/0
O IA    10.68.2.2/31
        [110/488] via 10.169.73.12, 3d07h, TenGigabitEthernet8/0/0
B       10.1.50.0/24 [200/0] via 172.16.189.9, 3d07h
B       10.1.51.0/24 [200/0] via 172.16.189.9, 3d07h

```


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

- A. The BGP routes are internal.
- B. The OSPF routes with the E2 flag retain the same metric as they leave the router.
- C. The OSPF routes with the IA flag have their administrative distances incremented as they leave the router.
- D. The BGP routes are external.
- E. The OSPF routes with the E2 flag have their metrics incremented as they leave the router.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

IBGP routes have an Administrative distance of 200, while EBGP have an AD of 20. Here we see that the BGP routes have an AD value of 200.

With OSPF, external routes fall under two categories, external type 1 and external type 2. The difference between the two is in the way the cost (metric) of the route is being calculated. The cost of a type 2 route is always the external cost, irrespective of the interior cost to reach that route. A type 1 cost is the addition of the external cost and the internal cost used to reach that route. The metric for E2 routes do not change when advertising to other routers.

QUESTION 205

Refer to the exhibit.

```
interface GigabitEthernet0/1
 ip address 192.168.1.5 255.255.255.0
 prefix-list FILTER seq 5 permit 172.16.0.0/16
 prefix-list FILTER seq 10 permit 0.0.0.0/0
 router eigrp 65000
  no auto-summary
  network 192.168.1.5 0.0.0.0
  distribute-list prefix FILTER out
```

Which two statements about this configuration are true? (Choose two.)

- A. It allows 172.16.0.0/16 to be distributed into EIGRP.
- B. It allows a default route to be distributed into EIGRP.
- C. It allows 172.16.0.0/16 and larger subnets to be distributed into EIGRP.
- D. It prevents 172.16.0.0/16 from being distributed into EIGRP.

- E. It prevents a default route from being distributed into EIGRP.
- F. It creates summary routes and injects them into EIGRP.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

In this example, the prefix list is configured to only allow the two specific routes of 172.16.0.0/16 and the default route. Any other routes will be filtered.

QUESTION 206

Refer to the exhibit.

```
R2
interface Loopback2
 ip address 172.16.2.2 255.255.255.0
interface Loopback3
 ip address 172.16.3.3 255.255.255.0
interface Loopback4
 ip address 172.16.5.4 255.255.255.0
interface GigabitEthernet1/0
 ip address 10.0.78.8 255.255.255.0
 ip router isis
router isis
 net 49.0001.0031.0031.00
 redistribute connected route-map LOOPBACKS
 ip access-list standard LOOPBACKS
 permit 172.16.0.0 0.0.3.255
 route-map LOOPBACKS permit 10
 match ip address LOOPBACKS
```

R1 is able to reach only some of the subnets that R2 is advertising. Which two configuration changes can you make to ensure that R1 can reach all routes from R2? (Choose two.)

- A. Add an additional permit statement to the LOOPBACKS route map.
- B. Modify the LOOPBACKS access list to include all loopback subnets.
- C. Add an additional statement in the LOOPBACKS route map to match both Level 1 and Level 2 circuits.
- D. Add an additional statement in the LOOPBACKS route map to match the R1 CLNS address.

- E. Configure the interfaces between R1 and R2 with a Level 1 IS-IS circuit.
- F. Configure the interfaces between R1 and R2 with a Level 2 IS-IS circuit.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation


Explanation/Reference:

Explanation:

In this example, the access list is using a 0.0.3.255 wildcard mask, so only the loopback IP's of 172.16.0.0 172.16.3.255 will be included. We need to add another statement to allow loopback 4 to be advertised, or modify the wildcard mask to include them all.

QUESTION 207

Refer to the exhibit.



R1, R2, and R3 have full network connectivity to each other, but R2 prefers the path through R3 to reach network 172.17.1.0/24. Which two actions can you take so that R2 prefers the path through R1 to reach 172.17.1.0/24? (Choose two.)

- A. Set the reference bandwidth to 10000 on R1, R2, and R3.
- B. Configure the cost on the link between R1 and R3 to be greater than 100 Mbps.
- C. Set the reference bandwidth on R2 only.
- D. Configure a manual bandwidth statement with a value of 1 Gbps on the link between R1 and R3.
- E. Modify the cost on the link between R1 and R2 to be greater than 10 Gbps.
- F. Configure a manual bandwidth statement with a value of 100 Mbps on the link between R1 and R2.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

By default, the reference bandwidth used in Cisco routers is 100Mbps, so FastEthernet and above will have a cost of 1, so a gigabit interface and 10GE interface will be equal with a fastethernet. This is not ideal. If we change the reference bandwidth to 100000 then the faster links will be used. Changing the reference bandwidth needs to be done on all routers in the OSPF network.

Increasing the cost on the R1-R3 link will also cause the traffic to take the more direct route.

QUESTION 208

What are two advantages to using Asynchronous mode instead of Demand mode for BFD? (Choose two.)

- A. Asynchronous mode requires half as many packets as Demand mode for failure detection.
- B. Asynchronous mode can be used in place of the echo function.
- C. Asynchronous mode supports a larger number of BFD sessions.
- D. Asynchronous mode requires one fourth as many packets as Demand mode for failure detection.
- E. Asynchronous mode's round-trip jitter is less than that of Demand mode.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Pure Asynchronous mode is advantageous in that it requires half as many packets to achieve a particular Detection Time as does the Echo function. It is also used when the Echo function cannot be supported for some reason.

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

QUESTION 209

Which action does route poisoning take that serves as a loop-prevention method?

- A. It immediately sends routing updates with an unreachable metric to all devices.
- B. It immediately sends routing updates with a metric of 255 to all devices.
- C. It prohibits a router from advertising back onto the interface from which it was learned.
- D. It advertises a route with an unreachable metric back onto the interface from which it was learned.
- E. It poisons the route by tagging it uniquely within the network.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

With route poisoning, when a router detects that one of its connected routes has failed, the router will poison the route by assigning an infinite metric to it and advertising it to neighbors.

QUESTION 210

Which two statements about the ipv6 ospf authentication command are true? (Choose two.)

- A. The command is required if you implement the IPsec AH header.
- B. The command configures an SPI.
- C. The command is required if you implement the IPsec TLV.
- D. The command can be used in conjunction with the SPI authentication algorithm.
- E. The command must be configured under the OSPFv3 process.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

OSPFv3 requires the use of IPsec to enable authentication. Crypto images are required to use authentication, because only crypto images include the IPsec API needed for use with OSPFv3.

In OSPFv3, authentication fields have been removed from OSPFv3 packet headers. When OSPFv3 runs on IPv6, OSPFv3 requires the IPv6 authentication header (AH) or IPv6 ESP header to ensure integrity, authentication, and confidentiality of routing exchanges. IPv6 AH and ESP extension headers can be used to provide authentication and confidentiality to OSPFv3.

To use the IPsec AH, you must enable the ipv6 ospf authentication command. To use the IPsec ESP header, you must enable the ipv6 ospf encryption command. The ESP header may be applied alone or in combination with the AH, and when ESP is used, both encryption and authentication

are provided. Security services can be provided between a pair of communicating hosts, between a pair of communicating security gateways, or between a security gateway and a host.

To configure IPsec, you configure a security policy, which is a combination of the security policy index (SPI) and the key (the key is used to create and validate the hash value). IPsec for OSPFv3 can be configured on an interface or on an OSPFv3 area. For higher security, you should configure a different policy on each interface configured with IPsec. If you configure IPsec for an OSPFv3 area, the policy is applied to all of the interfaces in that area, except for the interfaces that have IPsec configured directly. Once IPsec is configured for OSPFv3, IPsec is invisible to you.

Reference. http://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

QUESTION 211

Which two statements about SoO checking in EIGRP OTP deployments are true? (Choose two).

- A. During the import process, the SoO value in BGP is checked against the SoO value of the site map.
- B. During the reception of an EIGRP update, the SoO value in the EIGRP update is checked against the SoO value of the site map on the ingress interface.
- C. At the ingress of the PE/CE link, the SoO in the EIGRP update is checked against the SoO within the PE/CE routing protocol.
- D. At the egress of the PE/CE link, the SoO is checked against the SoO within the PE/CE routing protocol.
- E. The SoO is checked at the ingress of the backdoor link.
- F. The SoO is checked at the egress of the backdoor link.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

- - SoO checking:
 - During the import process the SoO value in BGP update is checked against the SoO value of the site-map attached to VRF interface. The update is propagated to CE only if there is no match (this check is done regardless of protocol used on PE/CE link).
 - At reception of EIGRP update, the SoO value in the EIGRP update is checked against the SoO value of site-map attached to the incoming interface. This update is accepted only if there is no match (this check can optionally be done on backdoor router).

Reference. http://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/ip-routing/whitepaper_C11-730404.html

QUESTION 212

Which two OSPF LSA types are flooded within the originating area? (Choose two.)

- A. type 1, Router LSA
- B. type 2, Network LSA
- C. type 3, Network Summary LSA
- D. type 4, ASBR Summary LSA

- E. type 6, Group Membership LSA
- F. type 9, Opaque LSA

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

OSPF relies on several types of *Link State Advertisements (LSAs)* to communicate link state information between neighbors. A brief review of the most applicable LSA types:

- **Type 1** - Represents a router
- **Type 2** - Represents the pseudonode (designated router) for a multiaccess link
- **Type 3** - A network link summary (internal route)
- **Type 4** - Represents an ASBR
- **Type 5** - A route external to the OSPF domain
- **Type 7** - Used in stub areas in place of a type 5 LSA

LSA types 1 and 2 are found in all areas, and are never flooded outside of an area. They are only flooded within the area that they originated from.

Reference. <http://packetlife.net/blog/2008/jun/24/ospf-area-types/>

QUESTION 213

Which statement about the OSPF Loop-Free Alternate feature is true?

- A. It is supported on routers that are configured with virtual links.
- B. It is supported in VRF OSPF instances.
- C. It is supported when a traffic engineering tunnel interface is protected.
- D. It is supported when traffic can be redirected to a primary neighbor.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Restrictions for OSPF IPv4 Remote Loop-Free Alternate IP Fast Reroute


- The OSPF IPv4 Remote Loop-Free Alternate IP Fast Reroute feature is not supported on devices that are virtual links headends.
- The feature is supported only in global VPN routing and forwarding (VRF) OSPF instances.
- The only supported tunneling method is MPLS.
- You cannot configure a traffic engineering (TE) tunnel interface as a protected interface. Use the MPLS Traffic Engineering--Fast Reroute Link and Node Protection feature to protect these tunnels. For more information, see the "MPLS Traffic Engineering--Fast Reroute Link and Node Protection" section in the *Multiprotocol Label Switching Configuration Guide*.

- You can configure a TE tunnel interface in a repair path, but OSPF will not verify the tunnel's placement; you must ensure that it is not crossing the physical interface that it is intended to protect.
- Not all routes can have repair paths. Multipath primary routes might have repair paths for all, some, or no primary paths, depending on the network topology, the connectivity of the computing router, and the attributes required of repair paths.
- Devices that can be selected as tunnel termination points must have a /32 address advertised in the area in which remote LFA is enabled. This address will be used as a tunnel termination IP. If the device does not advertise a /32 address, it may not be used for remote LFA tunnel termination.
- All devices in the network that can be selected as tunnel termination points must be configured to accept targeted LDP sessions using the `mpls ldp discovery targeted-hello accept` command.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xr-3s/iro-3s-book/iro-ipfrr-lfa.html

QUESTION 214

Refer to the exhibit.



ASN 64523 has a multi-homed BGP setup to ISP A and ISP B. Which BGP attribute can you set to allow traffic that originates in ASN 64523 to exit the ASN through ISP B?

- A. origin
- B. next-hop
- C. weight
- D. multi-exit discriminator

Correct Answer: D
Section: Layer 3 Technologies
Explanation

Explanation/Reference:

Explanation:

MED is an optional nontransitive attribute. MED is a hint to external neighbors about the preferred path into an autonomous system (AS) that has multiple entry points. The MED is also known as the external metric of a route. A lower MED value is preferred over a higher value.

Example at reference link below:

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13759-37.html>

QUESTION 215

When deploying redundant route reflectors in BGP, which attribute can you configure on the route reflector to allow routes to be identified as belonging to the same group?

- A. ROUTER_ID
- B. CLUSTER_ID
- C. ORIGINATOR_ID
- D. PEER_GROUP

Correct Answer: B
Section: Layer 3 Technologies
Explanation

Explanation/Reference:

Explanation:

Together, a route reflector and its clients form a cluster. When a single route reflector is deployed in a cluster, the cluster is identified by the router ID of the route reflector.

The bgp cluster-id command is used to assign a cluster ID to a route reflector when the cluster has one or more route reflectors. Multiple route reflectors are deployed in a cluster to increase redundancy and avoid a single point of failure. When multiple route reflectors are configured in a cluster, the same cluster ID is assigned to all route reflectors. This allows all route reflectors in the cluster to recognize updates from peers in the same cluster and reduces the number of updates that need to be stored in BGP routing tables.

Reference. <http://ieoc.com/forums/t/5326.aspx>

QUESTION 216

Which two options are mandatory components of a multiprotocol BGP VPN-IPv4 address? (Choose two.)

- A. a route distinguisher
- B. an IPv4 address
- C. a route target

- D. an MPLS label
- E. a system ID
- F. an area ID

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

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

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-mt/mp-l3-vpns-15-mt-book/mp-bgp-mpls-vpn.html

QUESTION 217

Which BGP feature enables you to install a backup path in the forwarding table?

- A. soft reconfiguration
- B. prefix independent convergence
- C. route refresh
- D. synchronization

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

To install a backup path into the forwarding table and provide prefix independent convergence (PIC) in case of a PE-CE link failure, use the additional-paths install backup command in an appropriate address family configuration mode. To prevent installing the backup path, use the no form of this command. To disable prefix independent convergence, use the disable keyword.

Reference. http://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-2/routing/command/reference/b_routing_cr42crs/b_routing_cr42crs_chapter_01.html

QUESTION 218

Refer to the exhibit.

```
R1
ip vrf VPN
  rd 1:1
  import-map INBOUND
  route-target both 1:1
interface GigabitEthernet0/0
  ip vrf forwarding VPN
  ip address 192.168.0.1 255.255.255.0
access-list 10 5 permit 192.168.0.0 255.255.0.0
access-list 10 10 permit 10.2.2.2 255.255.255.255
route-map INBOUND 10
  match ip address 10
router ospf 1 vrf VPN
  network 192.168.0.0 0.0.0.255 area 0
```

```
R2
ip vrf VPN
  rd 1:1
  route-target both 1:1
interface Loopback1
  ip vrf forwarding VPN
  ip address 10.1.1.1 255.255.255.0
interface Loopback1
  ip vrf forwarding VPN
  ip address 10.2.2.2 255.255.255.0
interface Loopback2
  ip vrf forwarding VPN
  ip address 10.3.3.3 255.255.255.0
interface GigabitEthernet0/0
  ip vrf forwarding VPN
  ip address 192.168.0.2 255.255.255.0
router ospf 2 vrf VPN
  network 192.168.0.0 0.0.0.255 area 0
  network 10.1.1.1 0.0.0.0 area 0
  network 10.2.2.2 0.0.0.0 area 0
  network 10.3.3.0 0.0.0.15 area 0
```

R1 and R2 have a working VRF-Lite configuration, but R1 is receiving a route only to 10.2.2.2 from R2. Which two changes can you make so that R1 receives all routes from R2? (Choose two.)

A. Create an additional permit statement in the access list that is referenced by the import-map on R1.

- B. Disable VRF filtering on R1.
- C. Set the R1 and R2 OSPF process IDs to match.
- D. Change the wildcard mask for the network 10.3.3.0 to 0.0.0.0.
- E. Create a matching export map in the VRF for R2.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

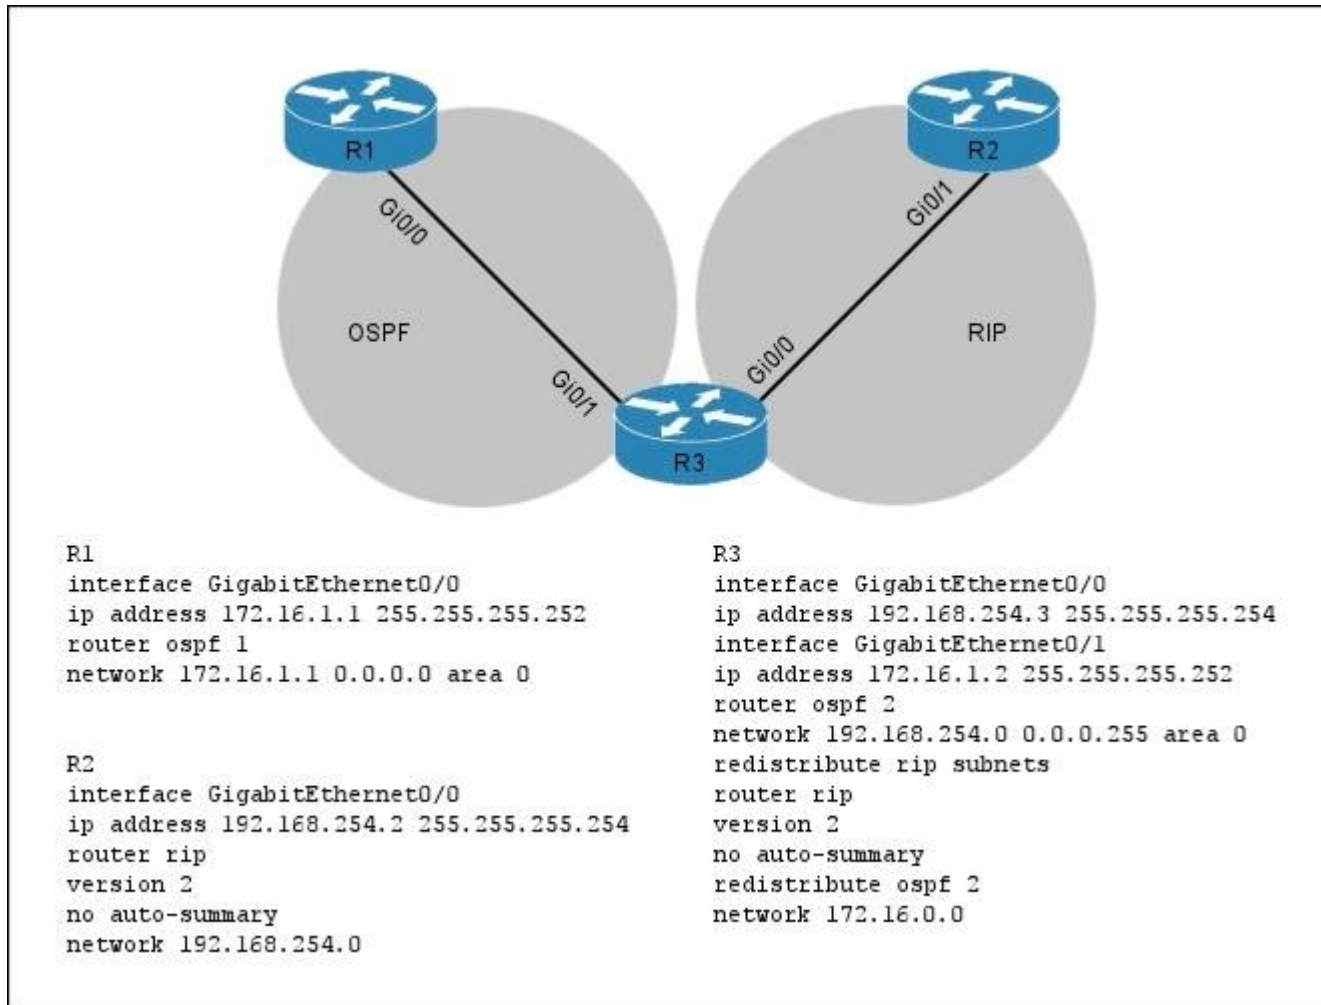
Explanation/Reference:

Explanation:

The access list in this example is only allowing the 192.168.0.0/16 and 10.2.2.2/32 routes to be advertised, so an additional permit statement is needed to allow the other routes. Alternatively, we could simply disable VRF filtering, then all routes would be advertised.

QUESTION 219

Refer to the exhibit.



R2 is unable to access the 172.16.1.0/30 network between R1 and R3. Which option is a possible reason for the failure?

- A. The seed metric for redistributing into RIP on R3 is missing.
- B. The OSPF processes on R2 and R3 are different.
- C. Auto-summary is misconfigured under the RIP process of R3.
- D. The subnet mask on the link between R2 and R3 is smaller than /30.
- E. The wildcard mask on R3 is misconfigured.

Correct Answer: A
Section: Layer 3 Technologies
Explanation

Explanation/Reference:

Explanation:

The problem is that RIP requires a seed metric to be specified when redistributing routes into that protocol. A seed metric is a "starter metric" that gives the RIP process a metric it can work with. The OSPF metric of cost is incomprehensible to RIP, since RIP's sole metric is hop count.

Reference. [http://www.thebryantadvantage.com/CCNP%20Certification%20BSCI%20Exam %20Tutorial%20Route%20Redistribution%20Seed%20Metric.htm](http://www.thebryantadvantage.com/CCNP%20Certification%20BSCI%20Exam%20Tutorial%20Route%20Redistribution%20Seed%20Metric.htm)

QUESTION 220

DRAG DROP

Drag and drop the BGP attribute on the left to the correct category on the right.

Select and Place:

Drag and drop the BGP attribute on the left to the correct category on the right.	
Local-Pref	BGP Well-Known Mandatory Attribute
Community	Target
Atomic-Aggregate	BGP Optional Nontransitive Attribute
AS_path	Target
Cluster List	Target
Originator ID	BGP Optional Transitive Attribute
	Target

Correct Answer:

Drag and drop the BGP attribute on the left to the correct category on the right.

Local-Pref
Atomic-Aggregate

BGP Well-Known Mandatory Attribute

AS_path

BGP Optional Nontransitive Attribute

Originator ID

Cluster List

BGP Optional Transitive Attribute

Community

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 221

DRAG DROP

Drag and drop the RIP configuration command on the left to the function it performs on the right.

Select and Place:

Drag and drop the RIP configuration command on the left to the function it performs on the right.

ip rip triggered	controls the advertisement of routes on an interface
default-information originate	divides traffic among routes with the lowest cost
ip split-horizon	configures the router to send information only when the routing database is updated
traffic-share min	configures the router to source the network with RIP

Correct Answer:

Drag and drop the RIP configuration command on the left to the function it performs on the right.

	ip split-horizon
	traffic-share min
	ip rip triggered
	default-information originate

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 222**DRAG DROP**

Drag and drop each EIGRP element on the left to the corresponding definition on the right.

Select and Place:

Drag and drop each EIGRP element on the left to the corresponding definition on the right.	
Feasibility Condition	the metric for a route advertised by EIGRP
Feasible Distance	the lowest sum of the EIGRP metric and the metric used to reach the next hop
Feasible Successor	a route that could become the best path
Neighbor Table	the route currently in use as the best path
Reported Distance	a list of EIGRP devices that have a direct physical connection
Successor	the requirement that the RD of a new route is lower than the FD of the current route

Correct Answer:

Drag and drop each EIGRP element on the left to the corresponding definition on the right.

Reported Distance

Feasible Distance

Feasible Successor

Successor

Neighbor Table

Feasibility Condition

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 223

DRAG DROP

Drag and drop each BGP attribute on the left to the matching description on the right.

Select and Place:

Drag and drop each BGP attribute on the left to the matching description on the right.

AS_PATH	sets the value used to reach the advertising router
community	an attribute whose value can affect the preferred path for eBGP peers
LOCAL_PREF	an attribute whose value is shared within iBGP
MED	supports values of IGP, EGP, and INCOMPLETE
NEXT_HOP	a Cisco proprietary attribute that is local to the individual router
origin	allows the administrator to customize path selection by grouping routes
weight	a list that shows the path through which a route has passed

Correct Answer:

Drag and drop each BGP attribute on the left to the matching description on the right.

	NEXT_HOP
	MED
	LOCAL_PREF
	origin
	weight
	community
	AS_PATH

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 224

Which option describes the effect of the OSPF default-information originate always command?

- A. It creates a stub area.
- B. It configures the device to advertise a default route regardless of whether it exists in the routing table.
- C. It configures the device to automatically redistribute a default route.

D. It adds a static default route to the device configuration.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

default-information originate

To generate a default external route into an Open Shortest Path First (OSPF) routing domain, use the default-information originate command in router configuration mode. To disable this feature, use the no form of this command.

default-information originate [always] metric *metric-value* [metric-type *type-value*] [route-map *map-name*]

Syntax Description

always	(Optional) Always advertises the default route regardless of whether the software has a default route.
--------	--

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/command/ipv6-cr-book/ipv6-d2.html>

QUESTION 225

Which two options are reasons to manipulate the delay metric instead of the bandwidth metric for EIGRP routing? (Choose two.)

- A. Because the delay metric provides better handling for bursty traffic
- B. Because manipulating the bandwidth metric can also affect QoS
- C. Because manipulating the bandwidth affects only a particular path
- D. Because changes to the delay metric are propagated to all neighbors on a segment

Correct Answer: BD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Using the bandwidth to influence EIGRP paths is discouraged for two reasons:

- Changing the bandwidth can have impact beyond affecting the EIGRP metrics. For example, quality of service (QoS) also looks at the bandwidth on an interface.
- EIGRP throttles to use 50 percent of the configured bandwidth. Lowering the bandwidth can cause problems like staving EIGRP neighbors from getting hello packets because of the throttling back.

Because changes to the delay metric are propagated to all downstream routers, changing the interface delay parameter is the preferred method of

influencing path selection

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/13673-14.html>

QUESTION 226

What is the maximum number of secondary IP addresses that can be configured on a router interface?

- A. 1
- B. 2
- C. 4
- D. 1024
- E. 65535
- F. no limit to the number of addresses

Correct Answer: F

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

From "IP Routing Frequently Asked Questions"

Q. What are the maximum number of secondary IP addresses that can be configured on a router interface?

A. There are no limits on configuring secondary IP addresses on a router interface.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/28745-44.html#q21>

QUESTION 227

Which address is a MAC address that is mapped from an IPv6 address (RFC 2464)?

- A. 3333.FF17.FC0F
- B. FFFE.FF17.FC0F
- C. FF34.3333.FF17
- D. FF7E.FF17.FC0F

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

An IPv6 packet with a multicast destination address DST, consisting of the sixteen octets DST through DST, is transmitted to the Ethernet multicast address whose first two octets are the value 3333 hexadecimal and whose last four octets are the last four octets of DST.

Reference. <https://tools.ietf.org/html/rfc2464>

QUESTION 228

Which multicast protocol uses source trees and RPF?

- A. DVMRP
- B. PIM sparse mode
- C. CBT
- D. mOSPF

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

DVMRP builds a parent-child database using a constrained multicast model to build a forwarding tree rooted at the source of the multicast packets. Multicast packets are initially flooded down this source tree. If redundant paths are on the source tree, packets are not forwarded along those paths. Forwarding occurs until prune messages are received on those parent-child links, which further constrains the broadcast of multicast packets.

Reference. DVMRP and dense-mode PIM use only source trees and use RPF as previously described.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3550/software/release/12-1_19_ea1/configuration/guide/3550scg/swmcast.html

QUESTION 229

What is the function of the command `ip pim autorp listener`?

- A. It allows a border PIM sparse mode router to accept autorp information from another autonomous system.
- B. It allows the mapping agents to accept autorp information from the PIM rendezvous point.
- C. It allows the routers to flood the autorp information in a sparse-mode-only network.
- D. It allows a BSR to accept autorp information and translate it into BSR messages.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

To cause IP multicast traffic for the two Auto-RP groups 224.0.1.39 and 224.0.1.40 to be Protocol Independent Multicast (PIM) dense mode flooded

across interfaces operating in PIM sparse mode, use the `ip pim autorp listener` command in global configuration mode. To disable this feature, use the no form of this command.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipmulti/command/imc-cr-book/imc_i3.html#wp3085748429

QUESTION 230

Refer to the exhibit.

```
FHR#show ipv6 mroute FF7E::1234
(2001:db8::7, FF7E::1234), 00:02:27/00:01:02, flags: SFT
Incoming interface: Ethernet1/0
RPF nbr: FE80::A8BB:CCFF:FE00:701, Registering
Immediate Outgoing interface list:
Tunnel2, Forward, 00:01:38/never
```

Which statement is true about why the first-hop PIM IPv6 router is stuck in registering?

- A. The scope of the IPv6 multicast address is link-local.
- B. The outgoing interface for the IPv6 multicast group should not be a tunnel interface.
- C. The R-bit is set in the IPv6 address, but this is not an embedded RP multicast IPv6 address.
- D. The S flag should not be set on a first-hop PIM router.
- E. A multicast IPv6 address does not start with FF.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

-R bit: RP bit: R = 1 indicates the address embeds the address of the Rendezvous Point (RP). The embedded RP address needs to begin with the prefix FF70::/12, But here we see that the address is FF7E::1234.

QUESTION 231

Refer to the exhibit.

```
!  
ip access-list extended REDIRECT  
permit tcp any any eq 25  
!  
route-map REDIRECT 10  
match ip address REDIRECT-SNMP  
set interface GigabitEthernet1/0  
!  
interface loopback0  
ip address 172.21.254.254 255.255.252.0  
!  
ip local policy route-map REDIRECT-SNMP  
!
```

Which option is the result of this configuration?

- A. All SNMP traffic coming into the router is redirected to interface GigabitEthernet1/0.
- B. All SNMP traffic generated from the router is redirected to interface GigabitEthernet1/0.
- C. All SMTP traffic generated from the router is redirected to interface GigabitEthernet1/0.
- D. All POP3 traffic coming into the router is redirected to interface GigabitEthernet1/0.
- E. All SMTP traffic coming into the router is redirected to interface GigabitEthernet1/0.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

This is an example of policy based routing, where traffic sourced from this router that matches the access list (all traffic with port 25 which is SMTP) will be forced out the Gig 0/1 interface.

QUESTION 232

Which three statements about EIGRP and BFD are true? (Choose three.)

- A. BFD is independent of the routing protocol, so it can be used as a generic failure detection mechanism for EIGRP.
- B. Some parts of BFD can be distributed to the data plane, so it can be less CPU-intensive than reduced timers, which exist wholly at the control plane.
- C. Reduced EIGRP timers have an absolute minimum detection timer of 1-2 seconds; BFD can provide sub-second failure detection.
- D. BFD is tied to specific routing protocols and can be used for generic fault detection for the OSPF, EIGRP, and BGP routing protocols.

- E. BFD is dependent on the EIGRP routing protocol, so it can be used as a specific failure detection mechanism.
- F. BFD resides on the control plane, so it is less CPU-intensive than if it resided on the data plane.

Correct Answer: ABC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

There are several advantages to implementing BFD over reduced timer mechanisms for routing protocols:

- Although reducing the EIGRP, IS-IS, and OSPF timers can result in minimum detection timer of one to two seconds, BFD can provide failure detection in less than one second.
- Because BFD is not tied to any particular routing protocol, it can be used as a generic and consistent failure detection mechanism for EIGRP, IS-IS, and OSPF.
- Because some parts of BFD can be distributed to the data plane, it can be less CPU-intensive than the reduced EIGRP, IS-IS, and OSPF timers, which exist wholly at the control plane.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html

QUESTION 233

You are implementing new addressing with EIGRP routing and must use secondary addresses, which are missing from the routing table. Which action is the most efficient solution to the problem?

- A. Disable split-horizon on the interfaces with secondary addresses.
- B. Disable split-horizon inside the EIGRP process on the router with the secondary interface addresses.
- C. Add additional router interfaces and move the secondary addresses to the new interfaces.
- D. Use a different routing protocol and redistribute the routes between EIGRP and the new protocol.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Normally, routers that are connected to broadcast-type IP networks and that use distance-vector routing protocols employ the split horizon mechanism to reduce the possibility of routing loops. Split horizon blocks information about routes from being advertised by a router out of any interface from which that information originated. This behavior usually optimizes communications among multiple routers, particularly when links are broken. However, with nonbroadcast networks, situations can arise for which this behavior is less than ideal. For these situations, you might want to disable split horizon with EIGRP and RIP.

If an interface is configured with secondary IP addresses and split horizon is enabled, updates might not be sourced by every secondary address. One routing update is sourced per network number unless split horizon is disabled.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfrip.html

QUESTION 234

Refer to the exhibit.

```
R1(config-if)#ipv6 ospf authentication ipsec spi 256 md5 0 o-routes
```

Which two options are possible states for the interface configured with the given OSPFv3 authentication? (Choose two.)

- A. GOING UP
- B. DOWN
- C. UNCONFIGURED
- D. GOING DOWN

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

To configure IPsec, you configure a security policy, which is a combination of the security policy index (SPI) and the key (the key is used to create and validate the hash value). IPsec for OSPFv3 can be configured on an interface or on an OSPFv3 area. For higher security, you should configure a different policy on each interface configured with IPsec. If you configure IPsec for an OSPFv3 area, the policy is applied to all of the interfaces in that area, except for the interfaces that have IPsec configured directly. Once IPsec is configured for OSPFv3, IPsec is invisible to you.

The secure socket API is used by applications to secure traffic. The API needs to allow the application to open, listen, and close secure sockets. The binding between the application and the secure socket layer also allows the secure socket layer to inform the application of changes to the socket, such as connection open and close events. The secure socket API is able to identify the socket; that is, it can identify the local and remote addresses, masks, ports, and protocol that carry the traffic requiring security.

Each interface has a secure socket state, which can be one of the following:

- NULL: Do not create a secure socket for the interface if authentication is configured for the area.
- DOWN: IPsec has been configured for the interface (or the area that contains the interface), but OSPFv3 either has not requested IPsec to create a secure socket for this interface, or there is an error condition.
- GOING UP: OSPFv3 has requested a secure socket from IPsec and is waiting for a CRYPTO_SS_SOCKET_UP message from IPsec.
- UP: OSPFv3 has received a CRYPTO_SS_SOCKET_UP message from IPsec.
- CLOSING: The secure socket for the interface has been closed. A new socket may be opened for the interface, in which case the current secure socket makes the transition to the DOWN state. Otherwise, the interface will become UNCONFIGURED.
- UNCONFIGURED. Authentication is not configured on the interface.

OSPFv3 will not send or accept packets while in the DOWN state.

ReferenE. http://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

QUESTION 235

Refer to the exhibit.

```
ip route 10.0.0.0 255.255.255.0 192.168.192.9
ip default-network 172.16.199.9
```

The device with this configuration is unable to reach network 172.31.31.0/24. The next hop router has been verified to have full connectivity to the network. Which two actions can you take to establish connectivity to the network? (Choose two.)

- A. Create a static route to 172.16.199.0 using the address of the next hop router.
- B. Create a default route to the link address of the next hop router.
- C. Create a static route to the loopback address of the next hop router.
- D. Create a default route to 172.16.199.9.
- E. Modify the existing static route so that the next hop is 0.0.0.0.
- F. Replace the ip default-network command with the ip default-gateway command.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Unlike the ip default-gateway command, you can use ip default-network when ip routing is enabled on the Cisco router. When you configure ip default-network the router considers routes to that network for installation as the gateway of last resort on the router.

For every network configured with ip default-network, if a router has a route to that network, that route is flagged as a candidate default route. However, in this case if the router does not a route to the drfault network of 172.16.199.9, then you would need to ensure that this route exisits by creating a static route to 172.16.199.0 using the address of the next hop router, or simply create a default route using the address of the next hop router.

QUESTION 236

Which algorithm heavily influenced the algorithm used by path-vector protocols?

- A. Bellman-Ford
- B. SPF
- C. DUAL
- D. Spanning-Tree
- E. Adaptive

F. Deflection

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:


Explanation:

A path vector protocol is a computer network routing protocol which maintains the path information that gets updated dynamically. Updates which have looped through the network and returned to the same node are easily detected and discarded. This algorithm is sometimes used in BellmanFord routing algorithms to avoid "Count to Infinity" problems.

Reference. http://en.wikipedia.org/wiki/Path_vector_protocol

QUESTION 237

Refer to the exhibit.



All of the routers on this network are running RIP. If you edit the R3 RIP process configuration to reduce the number of hops from R3 to R1, which statement about the configuration change is true?

- A. Configuring no passive-interface for GigabitEthernet0/0 in the R3 RIP process reduces the number of hops to R1 by 2.
- B. Configuring no passive-interface for GigabitEthernet0/0 in the R3 RIP process reduces the number of hops to R1 by 1.
- C. Configuring no passive-interface for GigabitEthernet0/1 in the R3 RIP process reduces the number of hops to R1 by 3.
- D. Configuring no passive-interface for GigabitEthernet0/1 in the R3 RIP process reduces the number of hops to R1 by 1.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

By changing the link from R3 to R2 to not be passive, traffic can then take the direct route from R3-R2-R1 instead of the longer path of R3-R6-R5-R4-R1, resulting in two less hops.

QUESTION 238

Where should the passive-interface command be used?

- A. Under the routing process for interfaces that need to be routed, but prevented from peering
- B. under the routing process for interfaces that need to be routed and allowed to peer
- C. under the interface configuration for interfaces that need to be routed, but prevented from peering
- D. under the interface configuration for interfaces that need to be routed and allowed to peer
- E. under the VTY configuration within global configuration mode

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Passive-interface is a feature you enable on a per interface basis which allows a particular interface to participate in a routing process but prevents that interface from forming neighbor relationships by not sending hello packets and discarding received hello packets.

QUESTION 239

Refer to the exhibit.

```
ip prefix-list EIGRP-ROUTES seq 5 permit 10.10.10.0/24 le 32
ip prefix-list OUTBOUND seq 5 permit 192.168.168.1/32
router eigrp 65535
network 192.168.168.0 0.0.255.255
network 172.31.10.0 0.0.0.255
distribute-list prefix EIGRP-ROUTES gateway OUTBOUND in
```

Which statement about the device routing table is true?

- A. Only networks 10.10.10.0/24 and smaller from host 192.168.168.1 are in the routing table.
- B. Only networks 10.10.10.0/24 and larger from host 192.168.168.1 are in the routing table.
- C. Only network 10.10.10.0/24 from host 192.168.168.1 is in the routing table.
- D. Networks 10.10.10.0/24 and smaller from any host are in the routing table.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

When you add the keywords "GE" and "LE" to the prefix-list, the "len" value changes its meaning. When using GE and LE, the len value specifies how many bits of the prefix you are checking, starting with the most significant bit.

```
ip prefix-list LIST permit 1.2.3.0/24 le 32
```

This means:


Check the first 24 bits of the prefix 1.2.3.0

The subnet mask must be less than or equal to 32

Reference. <http://blog.ine.com/2007/12/26/how-do-prefix-lists-work/>

QUESTION 240

Refer to the exhibit.



R1 is performing mutual redistribution, but OSPF routes from R3 are unable to reach R2. Which three options are possible reasons for this behavior? (Choose three.)

- A. R1 requires a seed metric to redistribute RIP.
- B. The RIP version supports only classful subnet masks.
- C. R1 is filtering OSPF routes when redistributing into RIP.
- D. R3 and R1 have the same router ID.
- E. R1 and R3 have an MTU mismatch.
- F. R2 is configured to offset OSPF routes with a metric of 16.

Correct Answer: ACF

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

A. RIP requires a seed metric to be specified when redistributing routes into that protocol. A seed metric is a "starter metric" that gives the RIP process a metric it can work with. The OSPF metric of cost is incomprehensible to RIP, since RIP's sole metric is hop count. We've got to give RIP a metric it understands when redistributing routes into that protocol, so let's go back to R1 and do so.

C. Filtering routes is another explanation, if the routes to R2 are being filtered from being advertised to R1.

F. If the metric is offset to 16, then the routes will have reached the maximum hop count when redistributed to RIP. The max hop count for RIP is 16.

QUESTION 241

Refer to the exhibit.

```
R1#sh ip eigrp 1 topology all
IP-EIGRP Topology Table for AS(1)/ID(10.1.1.2)

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

P 10.8.1.0/24, 1 successors, FD is 156160, serno 7
  via 10.1.1.1 (156160/128256), FastEthernet1/0
P 10.1.1.0/24, 1 successors, FD is 28160, serno 1
  via Connected, FastEthernet1/0
P 10.6.1.0/24, 1 successors, FD is 156160, serno 8
  via 10.1.1.1 (156160/128256), FastEthernet1/0
```

If the downstream router has a summary route configured, which two actions must you take on the local router to create the summary route that summarizes all routes from the downstream router? (Choose two.)

- A. Configure the summary address on the interface.
- B. Use 10.0.0.0 255.248.0.0 as the summary route.
- C. Configure the summary address in the EIGRP process.
- D. Use 10.0.0.0 255.252.0.0 as the summary route.
- E. Configure a route map to permit the route.
- F. Configure a distribute list in.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Route summarization works in conjunction with the ip summary-address eigrp interface configuration command, in which additional summarization can be performed. To correctly summarize all the networks shown, the correct route to use is 10.0.0.0 255.248.0.0

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfeigrp.html

QUESTION 242

Which three statements about RIP timers are true? (Choose three.)

- A. The default update timer is 30 seconds.
- B. The default invalid timer is 180 seconds.
- C. The default holddown timer is 180 seconds.
- D. The default flush timer is 60 seconds.
- E. The default scan timer is 60 seconds.
- F. The default hello timer is 5 seconds.

Correct Answer: ABC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The routing information protocol uses the following timers as part of its operation:

- Update Timer
- Invalid Timer
- Flush Timer
- Holddown Timer

Update Timer

The update timer controls the interval between two gratuitous Response Message. By default the value is 30 seconds. The response message is broadcast to all its RIP enabled interface.

Invalid Timer

The invalid timer specifies how long a routing entry can be in the routing table without being updated. This is also called as expiration Timer. By default, the value is 180 seconds. After the timer expires the hop count of the routing entry will be set to 16, marking the destination as unreachable.

Flush Timer

The flush timer controls the time between the route is invalidated or marked as unreachable and removal of entry from the routing table. By default the value is 240 seconds. This is 60 seconds longer than Invalid timer. So for 60 seconds the router will be advertising about this unreachable route to all its neighbors. This timer must be set to a higher value than the invalid timer.

Hold-down Timer

The hold-down timer is started per route entry, when the hop count is changing from lower value to higher value. This allows the route to get stabilized. During this time no update can be done to that routing entry. This is not part of the RFC 1058. This is Cisco's implementation. The default value of this timer is 180 seconds.

Reference. http://en.wikipedia.org/wiki/Routing_Information_Protocol#Timers

QUESTION 243

Which timer expiration can lead to an EIGRP route becoming stuck in active?

- A. hello
- B. active

- C. query
- D. hold

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

As noted above, when a route goes into the active state, the router queries its neighbors to find a path to the pertinent network. At this point, the router starts a three minute active timer by which time it must receive replies from all queried neighbors. If a neighbor has feasible successors for the route, it will recalculate its own local distance to the network and report this back. However, if a neighbor does not have a feasible successor, it also goes into active state. In some cases, multiple routers along multiple query paths will go into active state as routers continue to query for the desired route. In most cases, this process will yield responses from all queried routers and the sought after route will transition back into the passive state within the three minute SIA query timer. In the case that none of the queried routers can provide a feasible successor, the route is cleared.

In some cases, a response is not received between two neighbor routers because of link failures, congestion or some other adverse condition in either the network or on the queried router, and the three minute active timer expires on the router originating the query. When this happens, the querying router that did not receive a response logs a "DUAL-3-SIA" or "stuck-in-active" error for the route and then drops and restarts its adjacency with the non-responding router

Reference. <http://www.packetdesign.com/resources/technical-briefs/diagnosing-eigrp-stuck-active>

QUESTION 244

Which three values can be used to tag external EIGRP routes? (Choose three.)

- A. The router ID of the router that redistributed the route
- B. The administrative distance of the external protocol
- C. The protocol ID of the external protocol
- D. The cost to reach the router that redistributed the route
- E. The metric from the external protocol
- F. The router ID of the router from which the external protocol route was learned

Correct Answer: ACE

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

EIGRP has the notion of internal and external routes. Internal routes are ones that have been originated within an EIGRP autonomous system (AS).

Therefore, a directly attached network that is configured to run EIGRP is considered an internal route and is propagated with this information throughout the EIGRP AS. External routes are ones that have been learned by another routing protocol or reside in the routing table as static routes. These routes

are tagged individually with the identity of their origination.

External routes are tagged with the following information:

- The router ID of the EIGRP router that redistributed the route.
- The AS number where the destination resides.
- A configurable administrator tag.
- Protocol ID of the external protocol.
- The metric from the external protocol.
- Bit flags for default routing.

Reference. http://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/13669-1.html#route_tagging

QUESTION 245

Which data plane protocol does EIGRP Over the Top use?

- A. MPLS
- B. GRE
- C. LISP
- D. IP-in-IP

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The EIGRP Over the Top solution can be used to ensure connectivity between disparate Enhanced Interior Gateway Routing Protocol (EIGRP) sites. This feature uses EIGRP on the control plane and Locator ID Separation Protocol (LISP) encapsulation on the data plane to route traffic across the underlying WAN architecture. EIGRP is used to distribute routes between customer edge (CE) devices within the network, and the traffic forwarded across the WAN architecture is LISP encapsulated. Therefore, to connect disparate EIGRP sites, you must configure the neighbor command with LISP encapsulation on every CE in the network.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/xr-3s/ire-xr-3s-book/ire-eigrp-over-the-top.html

QUESTION 246

Which statement about the feasible distance in EIGRP is true?

- A. It is the maximum metric that should feasibly be considered for installation in the RIB.
- B. It is the minimum metric to reach the destination as stored in the topology table.
- C. It is the metric that is supplied by the best next hop toward the destination.
- D. It is the maximum metric possible based on the maximum hop count that is allowed.

Correct Answer: B
Section: Layer 3 Technologies
Explanation

Explanation/Reference:

Explanation:

An EIGRP router advertises each destination it can reach as a route with an attached metric. This metric is called the route's reported distance (the term advertised distance has also been used in older documentation). A successor route for any given destination is chosen as having the lowest computed feasible distance; that is, the lowest sum of reported distance plus the cost to get to the advertising router.

By default, an EIGRP router will store only the route with the best (lowest) feasible distance in the routing table (or, multiple routes with equivalent feasible distances).

Reference. <http://packetlife.net/blog/2010/aug/9/eigrp-feasible-successor-routes/>

QUESTION 247

Which statement about the EIGRP RTO is true?

- A. It is six times the SRTT.
- B. It is the time that it normally takes for an update to be received by a peer.
- C. It is the time that it normally takes to receive a reply to a query.
- D. It is the average time that it takes for a reliable packet to be acknowledged.

Correct Answer: A
Section: Layer 3 Technologies
Explanation

Explanation/Reference:

Explanation:

The RTO is typically six times the SRTT, the value may vary from a minimum of 200 microseconds (ms) to a maximum of 5 seconds (s).

Reference. EIGRP for IP: Basic Operation and Configuration, Alvaro Retana, Russ White, Don Slice - 2000

QUESTION 248

Which option describes the purpose of the leak-map keyword in the command `eigrp stub connected leak-map EigrpLeak`?

- A. It allows the specified static routes to be advertised.
- B. It allows exceptions to the route summarization that is configured.
- C. It allows specified EIGRP-learned routes to be advertised.
- D. It restricts specified connected routes from being advertised.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Example. eigrp stub leak-map Command

In the following example, the eigrp stub command is issued with the leak-map name keyword- argument pair to configure the device to reference a leak map that identifies routes to be advertised that would have been suppressed otherwise.

Device(config)# **router eigrp 1**


Device(config-router)# **network 10.0.0.0**

Device(config-router)# **eigrp stub leak-map map1**

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/15-mt/ire-15-mt-book/ire-eigrp-stub-rtg.html#GUID-FB899CA9-E9DE-48D8-8048-C971179E4E24

QUESTION 249

Refer to the exhibit.



If OSPF is implemented on the network, which additional configuration is needed to allow traffic from host 10.4.1.15/24 to host 10.1.2.20/24?

- A. A virtual link between router 2 and router 4
- B. A virtual link between router 3 and router 4
- C. A virtual link between router 2 and router 3
- D. The current design allows traffic between the two hosts.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:


Explanation:

This specific traffic from 10.4.1.0/24 to 10.1.2.0/24 would work because this traffic crosses only over the single OSPF area of 0.0.0.1.

However, traffic from hosts on R4 to R1 would indeed need a virtual link, since area 0.0.0.2 is not connected to the backbone area of 0.0.0.0.

QUESTION 250

Refer to the exhibit.



Which OSPFv3 routes will be visible in the routing table of R2?

A. 2001:12::1/128

- B. 2001:12::1/128, 2001:112::1/128
- C. 2001:12::2/128
- D. No OSPFv3 routes will be visible.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The command "ipv6 unicast-routing" needs to be configured on both routers before any IPv6 routes will be seen.

QUESTION 251

Refer to the exhibit.

R1 is configured as shown. R1 is able to establish a neighbor adjacency only with R2. Which addition must you make to the R1 configuration to allow it to establish an adjacency with R3?

- A. interface gigabitethernet 0/1
ip address 10.1.0.1 255.255.255.0
ip ospf network point-to-point
- B. interface gigabitethernet 0/1
ip address 10.1.0.1 255.255.255.0
ip ospf 1 area 0
- C. router ospf 1
network 10.1.0.0 0.0.0.255 area 1
- D. router ospf 1
area 0 stub

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

To enable interfaces and networks with OSPF, the networks need to be specified in the network statement. In the configuration shown, only 10.0.0.0/24 has been enabled, we are missing the network connecting to R3 (10.1.0.0/24).

QUESTION 252

Which option describes how a router responds if LSA throttling is configured and it receives the identical LSA before the interval is set?

- A. The LSA is added to the OSPF database and a notification is sent to the sending router to slow down its LSA packet updates.

- B. The LSA is added to the OSPF database.
- C. The LSA is ignored.
- D. The LSA is ignored and a notification is sent to the sending router to slow down its LSA packet updates.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

How OSPF LSA Throttling Works

The timers throttle lsa all command controls the generation (sending) of LSAs. The first LSA is always generated immediately upon an OSPF topology change, and the next LSA generated is controlled by the minimum start interval. The subsequent LSAs generated for the same LSA are rate-limited until the maximum interval is reached. The "same LSA" is defined as an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID.

The timers lsa arrival command controls the minimum interval for accepting the same LSA. If an instance of the same LSA arrives sooner than the interval that is set, the LSA is dropped. It is recommended that the arrival interval be less than or equal to the hold-time interval of the timers throttle lsa all command.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fsolsath.html

QUESTION 253

Which two options are valid for the number of bytes in a BGP AS number? (Choose two.)

- A. 2 bytes
- B. 4 bytes
- C. 6 bytes
- D. 8 bytes
- E. 16 bytes

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

During the early time of BGP development and standardization, it was assumed that availability of a 16 bit binary number to identify the Autonomous System (AS) within BGP would have been more than sufficient. The 16 bit AS number, also known as the 2-byte AS number, provides a pool of 65536 unique Autonomous System numbers. The IANA manages the available BGP Autonomous System Numbers (ASN) pool, with the assignments being carried out by the Regional Registries.


The current consumption rate of the publicly available AS numbers suggests that the entire public 2-byte ASN pool will be fully depleted. A solution to

this depletion is the expansion of the existing 2-byte AS number to a 4-byte AS number, which provides a theoretical 4,294,967,296 unique AS numbers. ARIN has made the following policy changes in conjunction with the adoption of the solution. The Cisco IOS BGP "4-byte ASN" feature allows BGP to carry a Autonomous System Number (ASN) encoded as a 4-byte entity. The addition of this feature allows an operator to use an expanded 4-byte AS number granted by IANA.

QUESTION 254

Refer to the exhibit.

```
hostname R3
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0
!
router bgp 65001
 network 3.3.3.0 mask 255.255.255.0
 neighbor 10.1.23.2 remote-as 65001
 neighbor 10.1.23.2 route-reflector-client
 neighbor 10.1.34.4 remote-as 65001
 neighbor 10.1.34.4 route-reflector-client
 no auto-summary
```



```
R2#sh ip bgp
BGP table version is 1, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
* 1.1.1.0/24	10.1.12.1	0	0	65105	i
* 2.2.2.0/24	0.0.0.0	0		32768	i

Why is the loopback 0 interface of R4 missing in the routing table of R2?

- A. R2 is configured as a route reflector client.
- B. There is no peering between R2 and R3.
- C. The next hop is not reachable from R2.
- D. The route originated within the same AS.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

In the "show ip bgp" output we see that there is no peering session between R2 and R3. Since R3 is the route reflector here, R3 would reflect routes advertised from R4 to R2, but the peer needs to be established first.

QUESTION 255

Which statement about the BGP scope of the cost community is true?

- A. It is shared with IBGP neighbors only.
- B. It is shared with IBGP neighbors and route reflectors.
- C. It is shared with EBGP neighbors only.
- D. It is shared with IBGP and EBGP neighbors.
- E. It is shared with IBGP and confederation peers.

Correct Answer: E

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The BGP Cost Community feature introduces the cost extended community attribute. The cost community is a non-transitive extended community attribute that is passed to internal BGP (iBGP) and confederation peers but not to external BGP (eBGP) peers. The cost community feature allows you to customize the local route preference and influence the best path selection process by assigning cost values to specific routes.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/s_bgpcc.html

QUESTION 256

Which statement is true about conditional advertisements?

- A. Conditional advertisements create routes when a predefined condition is met.
- B. Conditional advertisements create routes when a predefined condition is not met.

- C. Conditional advertisements delete routes when a predefined condition is met.
- D. Conditional advertisements create routes and withhold them until a predefined condition is met.
- E. Conditional advertisements do not create routes, they only withhold them until a predefined condition is met.

Correct Answer: E

Section: Layer 3 Technologies

Explanation

Explanation/Reference:


Explanation:

The Border Gateway Protocol (BGP) conditional advertisement feature provides additional control of route advertisement, depending on the existence of other prefixes in the BGP table.

Normally, routes are propagated regardless of the existence of a different path. The BGP conditional advertisement feature uses the non-exist-map and the advertise-map keywords of the neighbor advertise-map command in order to track routes by the route prefix. If a route prefix is not present in output of the non-exist-map command, then the route specified by the advertise-map command is announced. This feature is useful for multihomed networks, in which some prefixes are advertised to one of the providers only if information from the other provider is not present (this indicates a failure in the peering session or partial reachability). Reference. <http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/16137-cond-adv.html>

QUESTION 257

Refer to the exhibit.



How can Router X in AS70000 peer with Router Y in AS65000, in case Router Y supports only 2-byte ASNs?

- A. Router X should be configured with a remove-private-as command, because this will establish the peering session with a random private 2-byte ASN.
- B. It is not possible. Router Y must be upgraded to an image that supports 4-byte ASN.
- C. Router Y should be configured with a 4-byte AS using the local-as command.
- D. Router X should be configured with a 2-byte AS using the local-as command.

Correct Answer: D

Section: Layer 3 Technologies**Explanation****Explanation/Reference:**

Explanation:

Since router Y does not support 4-byte ASN,s it will not understand any AS numbers larger than 65535, so router X should use the local-as command on the peering statement to router Y to so that it sends in a 2-byte ASN to router Y.

QUESTION 258

Which statement about BGP and diverse path advertisement is true?

- A. The BGP best-path selection must be disabled.
- B. The BGP best-path selection algorithm has been changed to always ignore the IGP metric.
- C. The BGP best-path selection algorithm has been changed so that two BGP paths can be flagged as best in the BGP table.
- D. The BGP best-path selection algorithm has not been changed.
- E. The BGP best-path selection is disabled for BGP routes for which the feature is enabled.

Correct Answer: D

Section: Layer 3 Technologies**Explanation****Explanation/Reference:**

Explanation:

The BGP Diverse Path Using a Diverse-Path Route Reflector feature allows BGP to distribute an alternative path other than the best path between BGP speakers when route reflectors are deployed. This additional path is added to the best-path, and the best path algorithm still remains unchanged.

QUESTION 259

For which two conditions is Cisco Express Forwarding recursion disabled by default when the BGP Prefix Independent Convergence functionality is enabled? (Choose two.)

- A. next hops learned with a /24 mask
- B. next hops learned with any mask shorter than /32
- C. next hops learned with a /32 mask
- D. next hops that are directly connected

Correct Answer: CD

Section: Layer 3 Technologies**Explanation****Explanation/Reference:**

Explanation:

Recursion is the ability to find the next longest matching path when the primary path goes down. When the BGP PIC feature is not installed, and if the next hop to a prefix fails, Cisco Express Forwarding finds the next path to reach the prefix by recursing through the FIB to find the next longest matching path to the prefix. This is useful if the next hop is multiple hops away and there is more than one way of reaching the next hop.

However, with the BGP PIC feature, you may want to disable Cisco Express Forwarding recursion for the following reasons:

- Recursion slows down convergence when Cisco Express Forwarding searches all the FIB entries.
- BGP PIC Edge already precomputes an alternate path, thus eliminating the need for Cisco Express Forwarding recursion.

When the BGP PIC functionality is enabled, Cisco Express Forwarding recursion is disabled by default for two conditions:

- For next hops learned with a /32 network mask (host routes) For next hops that are directly connected
- For all other cases, Cisco Express Forwarding recursion is enabled.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xr-3s/irg-xr-3s-book/irg-bgp-mp-pic.html

QUESTION 260

How many bytes comprise the system ID within an IS-IS NET?

- A. 4 bytes
- B. 6 bytes
- C. 8 bytes
- D. 16 bytes
- E. 20 bytes

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Routers are identified with NETs of 8 to 20 bytes. ISO/IEC 10589 distinguishes only three fields in the NSAP address format: a variable-length area address beginning with a single octet, a system ID, and a 1-byte n-selector. Cisco implements a fixed length of 6 bytes for the system ID, which is like the OSPF router ID.

Reference.

http://www.cisco.com/en/US/products/ps6599/products_white_paper09186a00800a3e6f.shtml

QUESTION 261

Which two statements about IS-IS are true? (Choose two.)

- A. The default hello interval is 10 seconds and the default hold timer is 30 seconds.
- B. The hello interval can be changed on a per-interface basis with the command `isis hello- multiplier`.
- C. Both routers need to have the same hello intervals and hold timers in order to form IS-IS neighbors.
- D. Both IS-IS routers need to have the same capabilities in the hello packet in order to form neighbors.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

isis hello-interval

To specify the length of time between hello packets that the Cisco IOS software sends, use the **isis hello-interval** command in interface configuration mode.

By default, a value three times the hello interval seconds is advertised as the hold time in the hello packets sent. (Change the multiplier of 3 by specifying the isis hello-multiplier command.) With smaller hello intervals, topological changes are detected faster, but there is more routing traffic. The default is 10 seconds.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_2/iproute/command/reference/fiprrp_r/1rfisis.html

QUESTION 262

Which bit should be set in the link-state PDU of an IS-IS L1/L2 router to indicate that it is a potential exit point of the area?

- A. the ABR bit
- B. the ATT bit
- C. the down bit
- D. the P bit

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Default routing is achieved in two distinct ways with Integrated IS-IS:

- *Attached bit*--Set by a Level 1/Level 2 router in its own Level 1 LSP and used to indicate to all Level 1 routers (within the area) that this router is a potential exit point of the area. Level 1-only routers will default to the nearest attached Level 2 router.
- *Default information originate*--Can be configured in Level 1 as well as Level 2. The default route (0.0.0.0/0) is inserted in the router LSP (Level 1 or Level 2, according to the configuration command) and the LSP is flooded according to the router type (Level 1 or Level 2). A Level 2 router doesn't need to have a default route to originate a default route.

Reference.

http://www.cisco.com/en/US/products/ps6599/products_white_paper09186a00800a3e6f.shtml

QUESTION 263

Which two options are benefits of EIGRP OTP? (Choose two.)

- A. It allows EIGRP routers to peer across a service provider without the service provider involvement.
- B. It allows the customer EIGRP domain to remain contiguous.
- C. It requires only minimal support from the service provider.
- D. It allows EIGRP neighbors to be discovered dynamically.
- E. It fully supports multicast traffic.
- F. It allows the administrator to use different autonomous system numbers per EIGRP domain.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

EIGRP Over the Top (OTP) allows EIGRP routers to peer across a service provider infrastructure without the SP's involvement. In fact with OTP, the provider won't see customer routes at all. EIGRP OTP acts as a provider-independent overlay that transports customer data between the customer's routers.

To the customer, the EIGRP domain is contiguous. A customer's EIGRP router sits at the edge of the provider cloud, and peers with another EIGRP router a different location across the cloud. Learned routes feature a next hop of the customer router -- not the provider. Good news for service providers is that customers can deploy EIGRP OTP with their involvement

Reference. <http://ethancbanks.com/2013/08/01/an-overview-of-eigrp-over-the-top-otp/>

QUESTION 264

DRAG DROP

Drag and drop the OSPFv3 LSA type on the left to the functionality it provides on the right.

Select and Place:

Router LSA (Type 1)	advertises an internal network or set of networks to routers in other areas
Network LSA (Type 2)	associates a group of prefixes for transit networks or stub networks
Interarea-prefix LSA for ABRs (Type 3)	indicates whether the router is part of a virtual link
Interarea-router LSA for ASBRs (Type 4)	collects link-state information and cost information for the
Autonomous system external LSA (Type 5)	provides the link-local address of a router to other routers on
Link LSA (Type 8)	redistributes external routes
Intra-Area-Prefix LSAs (Type 9)	enables routers to determine the best path to an external network

Correct Answer:

	Interarea-prefix LSA for ABRs (Type 3)
	Intra-Area-Prefix LSAs (Type 9)
	Router LSA (Type 1)
	Network LSA (Type 2)
	Link LSA (Type 8)
	Autonomous system external LSA (Type 5)
	Interarea-router LSA for ASBRs (Type 4)

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 265

DRAG DROP

Drag and drop the OSPF network type on the left to the correct category of timers on the right.

Select and Place:

Point-to-Point
Loopback
Point-to-Multipoint Nonbroadcast
Broadcast
Point-to-Multipoint
Nonbroadcast

Hello 10, Dead 40, Wait 40
1
2

None
1

Hello 30, Dead 120, Wait 120
1
2
3

Correct Answer:

Drag and drop the BGP attribute on the left to the correct category on the right.

Hello 10, Dead 40, Wait 40

Point-to-Point

Broadcast

None

Loopback

Hello 30, Dead 120, Wait 120

Point-to-Multipoint Nonbroadcast

Point-to-Multipoint

Nonbroadcast

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 266

DRAG DROP

Drag and drop the BGP attribute on the left to the correct category on the right.

Select and Place:

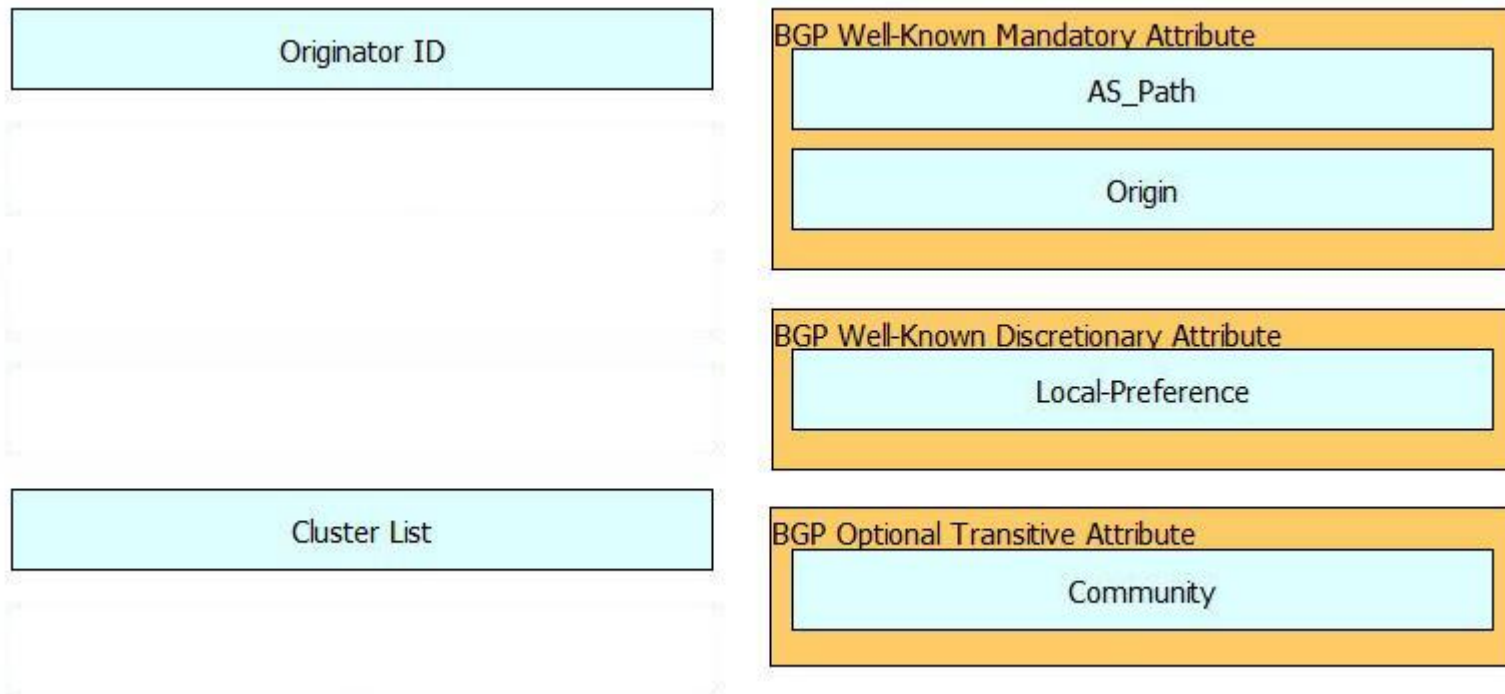
Originator ID
Community
Local-Preference
AS_Path
Cluster List
Origin

BGP Well-Known Mandatory Attribute
1
2

BGP Well-Known Discretionary Attribute
1

BGP Optional Transitive Attribute
1

Correct Answer:




Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

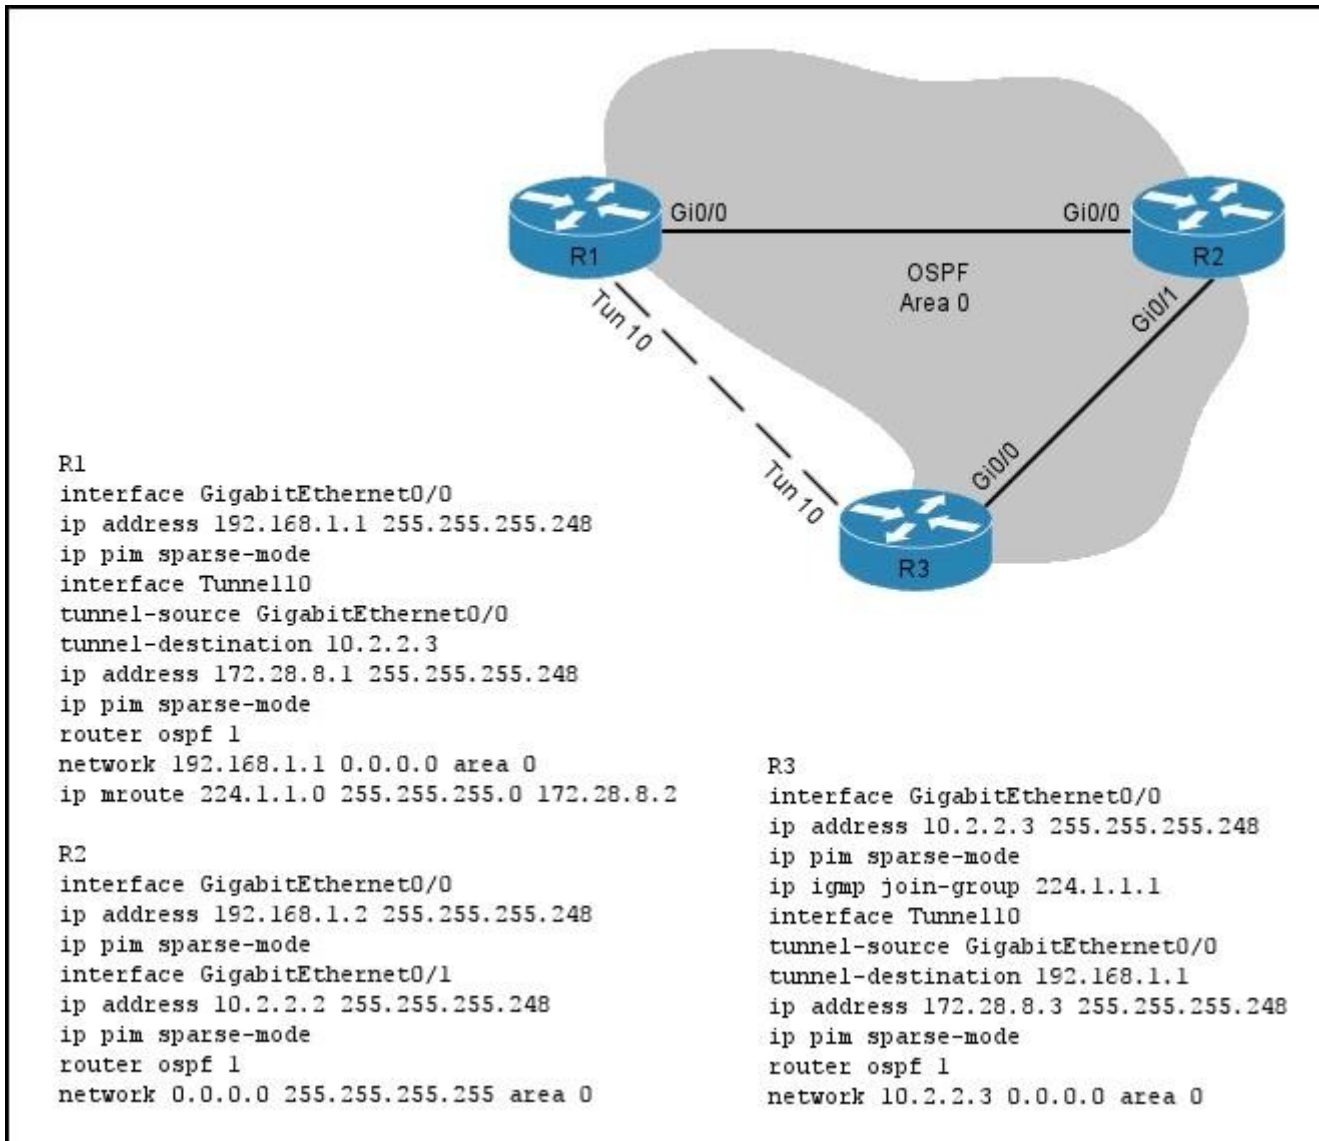


Reference. <http://www.deepsh.it/networking/BGP/bgp-attributes.png>

Reference. <http://www.deepsh.it/networking/BGP/bgp-attributes.png>

QUESTION 267

Refer to the exhibit.



R3 is failing to join the multicast group 224.1.1.1 that is sourcing from R1. Which two actions can you take to allow multicast traffic to flow correctly? (Choose two.)

A. Remove the static multicast route on R1.

- B. Configure OSPF on R1 and R3 to include the tunnel interfaces.
- C. Add an additional static multicast route on R2 for multicast group 224.1.1.1 toward R3.
- D. Replace the static multicast route on R1 to send traffic toward R2.
- E. Remove the static unicast route on R1.
- F. Add an additional static unicast route on R2 toward the loopback interface of R3.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Since the tunnel interfaces are not part of OSPF, the best path to the multicast source of R1 from R3 would be over the Gi0/0 path via OSPF. However, the static mroute is configured to use the tunnel, so this causes an RPF failure used in Sparse Mode. Best fix is to add the tunnel interfaces into OSPF and remove the static mroute so that the RPF check no longer fails.

QUESTION 268

Which two modes of operation does BFD support? (Choose two.)

- A. synchronous mode
- B. asynchronous mode
- C. demand mode
- D. echo mode
- E. aggressive mode
- F. passive mode

Correct Answer: BC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

A session may operate in one of two modes: asynchronous mode and demand mode. In asynchronous mode, both endpoints periodically send Hello packets to each other. If a number of those packets are not received, the session is considered down.

In demand mode, no Hello packets are exchanged after the session is established; it is assumed that the endpoints have another way to verify connectivity to each other, perhaps on the underlying physical layer. However, either host may still send Hello packets if needed.

Reference: http://en.wikipedia.org/wiki/Bidirectional_Forwarding_Detection

QUESTION 269

Which two loop-prevention mechanisms are implemented in BGP? (Choose two.)

- A. A route with its own AS in the AS_PATH is dropped automatically if the route reenters its own AS.
- B. A route with its own cluster ID in the CLUSTER_LIST is dropped automatically when the route reenters its own AS.
- C. The command `bgp allowas-in` enables a route with its own AS_PATH to be dropped when it reenters its own AS.
- D. The command `bgp bestpath as-path ignore` enables the strict checking of AS_PATH so that they drop routes with their own AS in the AS_PATH.
- E. The command `bgp bestpath med missing-as-worst` assigns the smallest possible MED, which directly prevents a loop.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

When dealing with the possibility of routing updates making their way back into an AS, BGP relies on the information in the AS_path for loop detection. An update that tries to make its way back into the AS it was originated from will be dropped by the border router.

With the introduction of route reflectors, there is a potential for having routing loops within an AS. A routing update that leaves a cluster might find its way back inside the cluster. Loops inside the AS cannot be detected by the traditional AS_path approach because the routing updates have not left the AS yet. BGP offers two extra measures for loop avoidance inside an AS when route reflectors are configured.

Using an Originator ID

The *originator* ID is a 4-byte, optional, nontransitive BGP attribute (type code 9) that is created by the route reflector. This attribute carries the router ID of the originator of the route in the local AS. If, because of poor configuration, the update comes back to the originator, the originator ignores it.

Using a Cluster List

The *cluster* list is an optional, nontransitive BGP attribute (type code 10). Each cluster is represented with a cluster ID.

A cluster list is a sequence of cluster IDs that an update has traversed. When a route reflector sends a route from its clients to nonclients outside the cluster, it appends the local cluster ID to the cluster list. If the route reflector receives an update whose cluster list contains the local cluster ID, the update is ignored. This is basically the same concept as the AS_path list applied between the clusters inside the AS.

Reference: http://borg.uu3.net/cisco/inter_arch/page11.html

QUESTION 270

Refer to the exhibit.

```
key chain kcl
  key 1
  key-string ripauth
interface Serial0
  ip address 10.1.1.1 255.255.255.252
  ip rip authentication key-chain kcl
router rip
  version 2
  network 10.0.0.0
```

RIPv2 authentication is failing on a device with this configuration. Which two actions can you take to enable it? (Choose two.)

- A. Set the RIP authentication mode to text.
- B. Set the RIP authentication mode to MD5.
- C. Configure the password encryption for the key.
- D. Set the password encryption to AES.

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

See the reference link below for information on configuring RIPv2 authentication, including both text and MD5 modes.

Reference: [http://www.cisco.com/c/en/us/support/docs/ip/routing-information-protocol-rip/13719- 50.html#configuringplain](http://www.cisco.com/c/en/us/support/docs/ip/routing-information-protocol-rip/13719-50.html#configuringplain)

QUESTION 271

Which three routing protocols utilize TLVs? (Choose three.)

- A. BGP
- B. IS-IS
- C. ODR
- D. OSPF
- E. EIGRP
- F. RIP

Correct Answer: ABE

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

IS-IS, originally designed for Open System Interconnection (OSI) routing, uses TLV parameters to carry information in Link State Packets (LSPs). The TLVs make IS-IS extendable. IS-IS can therefore carry different kinds of information in the LSPs.

Several routing protocols use TLVs to carry a variety of attributes. Cisco Discovery Protocol (CDP), Label Discovery Protocol (LDP), and Border Gateway Protocol (BGP) are examples of protocols that use TLVs. BGP uses TLVs to carry attributes such as Network Layer Reachability Information (NLRI), Multiple Exit Discriminator (MED), and local preference.

The IP header of the EIGRP packet specifies IP protocol number 88 within it, and the maximum length of the packet will be the IP MTU of the interface on which it is transmitted, most of the time 1500 octets. Following the IP header is the various Type/Length/Value (TLV) triplets. These TLVs will not only carry the route entries but also provide fields for the management of the DUAL process, multicast sequencing, and IOS software versions from the router.

References: <http://www.cisco.com/c/en/us/support/docs/ip/integrated-intermediate-system-to-intermediate-system-is-is/5739-tlvs-5739.html>
<http://ericleahy.com/index.php/eigrp-packets-neighborships/>

QUESTION 272

Which two statements about the command `distance bgp 90 60 120` are true? (Choose two.)

- A. Implementing the command is a Cisco best practice.
- B. The external distance it sets is preferred over the internal distance.
- C. The internal distance it sets is preferred over the external distance.
- D. The local distance it sets may conflict with the EIGRP administrative distance.
- E. The internal distance it sets may conflict with the EIGRP administrative distance.
- F. The local distance it sets may conflict with the RIP administrative distance.

Correct Answer: CF

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

To allow the use of external, internal, and local administrative distances that could be a better route than other external, internal, or local routes to a node, use the **distance bgp** command in address family or router configuration mode. To return to the default values, use the no form of this command.

distance bgp *external-distance internal-distance local-distance*

no distance bgp

Syntax Description

external-distance	Administrative distance for BGP external routes. External routes are routes for which the best path is learned from a neighbor external to the autonomous system. Accept table values are from 1 to 255. The default is 20. Routes with a distance of 255 are not installed in the routing table.
internal-distance	Administrative distance for BGP internal routes. Internal routes are those routes that are learned from another BGP entity within the same autonomous system. Accept table values are from 1 to 255. The default is 200. Routes with a distance of 255 are not installed in the routing table.
local-distance	Administrative distance for BGP local routes. Local routes are those networks listed with a network router configuration command, often as back doors, for that router or for networks that are being redistributed from another process. Accept table values are from 1 to 255. The default is 200. Routes with a distance of 255 are not installed in the routing table.

Defaults

external-distance: 20

internal-distance: 200

local-distance: 200

In this case, the internal distance is 60 and the external is 90, and the local distance is 120 (same as RIP).

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/12_2/iproute/command/reference/fiprrp_r/1rfbgp1.html#wp1113874

QUESTION 273

Refer to the exhibit.

```
vrf definition v1
  rd 1:1
  address-family ipv4
  exit-address-family
  address-family ipv6
  exit-address-family

vrf definition v2
  rd 2:2
  address-family ipv6
  exit-address-family

interface FastEthernet0/0
  no ip address
interface FastEthernet0/0.100
  encapsulation dot1Q 100
  vrf forwarding v1
  ip address 192.168.1.1 255.255.255.0
  ipv6 enable
  ospfv3 1 ipv6 area 0
  ospfv3 1 ipv4 area 0
interface FastEthernet0/0.200
  encapsulation dot1Q 200
  vrf forwarding v2
  ipv6 enable
  ospfv3 1 ipv6 area 0
interface FastEthernet0/1
  vrf forwarding v1
  ip address 10.1.1.1 255.255.255.0
  ipv6 enable
  ospfv3 1 ipv6 area 1
  ospfv3 1 ipv4 area 0
  no keepalive
interface FastEthernet0/2
  vrf forwarding v2
  no ip address
  ipv6 address 2001:DB8:1::1
  ipv6 enable
  ospfv3 1 ipv6 area 1

router ospfv3 1
  address-family ipv6 unicast vrf v2
  router-id 192.168.2.1
  exit-address-family

  address-family ipv4 unicast vrf v1
  router-id 192.168.1.4
  exit-address-family

  address-family ipv6 unicast vrf v1
  router-id 192.168.1.1
  exit-address-family
```

Route exchange is failing on a PE edge device configured with this VRF-Lite. Which action can you take to correct the problem?

- A. Configure the vrf-lite capability under the OSPF address families.
- B. Correct the route descriptors.

- C. Correct the OSPF router-ids.
- D. Configure the control plane with a larger memory allocation to allow the device to appear in the routing table.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Open Shortest Path First version 3 (OSPFv3) operates in nondefault VPN routing and forwarding (VRF) instances for both IPv6 and IPv4 address families and, transports the routes across a Border Gateway Protocol (BGP) or a Multiprotocol Label Switching (MPLS) backbone. On the provider edge (PE) device, customer routes are installed together by OSPFv3 and BGP in a common VRF or address family and each protocol is configured to redistribute the routes of the other. BGP combines the prefixes redistributed into it with a route-distinguisher value defined for the VRF and advertises them to other MPLS-BGP speakers in the same autonomous system using the VPNv4 or VPNv6 address family as appropriate.

The OSPFv3 route selection algorithm prefers intra-area routes across the back-door link over inter- area routes through the MPLS backbone. Sham-links are a type of virtual link across the MPLS backbone that connect OSPFv3 instances on different PEs. OSPFv3 instances tunnel protocol packets through the backbone and form adjacencies. Because OSPFv3 considers the sham-link as an intra-area connection, sham-link serves as a valid alternative to an intra-area back-door link. Domain IDs are used to determine whether the routes are internal or external. They describe the administrative domain of the OSPFv3 instance from which the route originates. Every PE has a 48- bit primary domain ID (which may be NULL) and zero or more secondary domain IDs.

How to Configure VRF-Lite/PE-CE

Configuring a VRF in an IPv6 Address Family for OSPFv3

SUMMARY STEPS

1. enable
2. configure terminal
3. vrf definition *vrf-name*
4. rd *route-distinguisher*
5. exit
6. router ospfv3 [*process-id*]
7. address-family ipv6 [unicast] [vrf *vrf-name*]
8. end

Reference: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/15-sy/iro-15-sy-book/iro-vrf-lite-pe-ce.html

QUESTION 274

Refer to the exhibit.


```
router ospf 1
network 192.168.10.0 0.0.0.255 area 0
network 172.22.19.0 0.0.0.255 area 15
area 15 range 192.168.0.0 255.255.0.0 not-advertise
!
```

Which option is the result of this configuration?

- A. Devices in OSPF area 15 can reach the summary route 192.168.0.0/16 and its more specific subnets.
- B. Devices in OSPF area 15 can reach only the more specific routes of 192.168.0.0/16.
- C. Devices in OSPF area 0 can reach the summary route 192.168.0.0/16 and its more specific subnets.
- D. Devices in OSPF area 0 can reach only the summary route of 192.168.0.0/16.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

area range

To consolidate and summarize routes at an area boundary, use the **area range** command in router configuration mode. To disable this function, use the **no** form of this command. **Area** *area-id* **range** *ip-address mask* [**advertise** | **not-advertise**] [**cost cost**] **no area** *area-id* **range** *ip-address mask* [**advertise** | **not-advertise**] [**cost cost**]

Syntax Description

<i>area-id</i>	Identifier of the area about which routes are to be summarized. It can be specified as either a decimal value or as an IP address.
<i>ip-address</i>	IP address.
<i>mask</i>	IP address mask.
advertise	(Optional) Sets the address range status to advertise and generates a Type 3 summary link-state advertisement (LSA).
not-advertise	(Optional) Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed, and the component networks remain hidden from other networks.

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/12_2/iproute/command/reference/fiprrp_r/1rfospf.html

QUESTION 275

Which two technologies are supported by EIGRP? (Choose two.)

- A. clear-text authentication
- B. MD5 authentication
- C. stub routing
- D. multiple areas

Correct Answer: BC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The IP Enhanced IGRP Route Authentication feature provides MD5 authentication of routing updates from the EIGRP routing protocol. The MD5 keyed digest in each EIGRP packet prevents the introduction of unauthorized or false routing messages from unapproved sources. The EIGRP stub routing feature improves network stability, reduces resource utilization, and simplifies the stub device configuration.

Stub routing is commonly used in hub-and-spoke network topologies. In a hub-and-spoke network, one or more end (stub) networks are connected to a remote device (the spoke) that is connected to one or more distribution devices (the hub). The remote device is adjacent to one or more distribution devices. The only route for IP traffic to reach the remote device is through a distribution device.

References:

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/15-s/ire-15-s-book/ire-rte-auth.html

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/15-s/ire-15-s-book/ire-eigrp-stub-rtg.html

QUESTION 276

How does having an EIGRP feasible successor speed up convergence?

- A. EIGRP sends queries only if there is a feasible successor, which decreases the number of routers that are involved in convergence.
- B. EIGRP sends queries only if there is not a feasible successor, which causes less control traffic to compete with data.
- C. EIGRP immediately installs the loop-free alternative path in the RIB.
- D. EIGRP preinstalls the feasible successor in the RIB in all cases, which causes traffic to switch more quickly.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Feasible Successor

- A next-hop router that serves as backup to the current successor.
- The condition is that the said router's AD (or RD) is less than the FD of the current successor route.
- Once the feasible successor is selected, they are placed in the topology table. If a change in topology occurs which requires a new route, DUAL looks for the feasible successor and uses it as new route immediately, resulting in fast convergence.

Reference: <http://routemyworld.com/2008/07/page/2/>

QUESTION 277

Which two options are ways in which an OSPFv3 router handles hello packets with a clear address- family bit? (Choose two.)

- A. IPv4 unicast packets are discarded.
- B. IPv6 unicast packets are discarded.
- C. IPv4 unicast packets are forwarded.
- D. IPv6 unicast packets are forwarded.

Correct Answer: AD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

A typical distance vector protocol saves the following information when computing the best path to a destination: the distance (total metric or distance, such as hop count) and the vector (the next hop). For instance, all the routers in the network in Figure 1 are running Routing Information Protocol (RIP). Router Two chooses the path to Network A by examining the hop count through each available path.

Since the path through Router Three is three hops, and the path through Router One is two hops, Router Two chooses the path through One and discards the information it learned through Three. If the path between Router One and Network A goes down, Router Two loses all connectivity with this destination until it times out the route of its routing table (three update periods, or 90 seconds), and Router Three re-advertises the route (which occurs every 30 seconds in RIP). Not including any hold-down time, it will take between 90 and 120 seconds for Router Two to switch the path from Router One to Router Three.

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

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

- Correct Answer: BF**

Section: Layer 3 Technologies**Explanation****Explanation/Reference:**

Explanation:

External routes fall under two categories, external type 1 and external type 2. The difference between the two is in the way the cost (metric) of the route is being calculated. The cost of a type 2 route is always the external cost, irrespective of the interior cost to reach that route. A type 1 cost is the addition of the external cost and the internal cost used to reach that route. A type 1 route is always preferred over a type 2 route for the same destination.

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

QUESTION 279

A company is multihomed to several Internet providers using EBGp. Which two measures guarantee that the network of the company does not become a transit AS for Internet traffic? (Choose two.)


- A. Prepend three times the AS number of the company to the AS path list.
- B. Add the community NO_EXPORT when sending updates to EBGp neighbors.
- C. Write AS-path access-list which permits one AS long paths only and use it to filter updates sent to EBGp neighbors.
- D. Add the community NO_EXPORT when receiving updates from EBGp neighbors.

Correct Answer: CD

Section: Layer 3 Technologies**Explanation****Explanation/Reference:**

Explanation:

By default BGP will advertise all prefixes to EBGp (External BGP) neighbors. This means that if you are multi-homed (connected to two or more ISPs) that you might become a transit AS. Let me show you an example:



R1 is connected to ISP1 and ISP2 and each router is in a different AS (Autonomous System). Since R1 is multi-homed it's possible that the ISPs will use R1 to reach each other. In order to prevent this we'll have to ensure that R1 only advertises prefixes from its own autonomous system. As far as I know there are 4 methods how you can prevent becoming a transit AS:

- **Filter-list with AS PATH access-list.**
- **No-Export Community.**
- Prefix-list Filtering
- Distribute-list Filtering

Reference: <http://networklessons.com/bgp/bgp-prevent-transit-as/>

QUESTION 280

Which BGP feature allows a router to maintain its current BGP configuration while it advertises a different AS number to new connections?

- A. local-AS
- B. next-hop-self
- C. allow-AS in
- D. soft reset

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The local-AS feature allows a router to appear to be a member of a second autonomous system (AS), in addition to its real AS. This feature can only be used for true eBGP peers. The local-AS feature is useful if ISP-A purchases ISP-B, but ISP-B's customers do not want to modify any peering arrangements or configurations. The local-AS feature allows routers in ISP-B to become members of ISP-A's AS. At the same time, these routers appear to their customers to retain their ISP-B AS number.

Reference: <http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13761-39.html>

QUESTION 281

Which problem can result when private AS numbers are included in advertisements that are sent to the global Internet BGP table?

- A. The prefixes sent with private AS numbers are always discarded on the Internet.
- B. The prefixes sent with private AS numbers are always tagged as invalid on the Internet.
- C. The prefixes sent with private AS numbers lack uniqueness, which can lead to a loss of connectivity.
- D. The prefixes sent with private AS numbers are sometimes tagged as invalid on the Internet.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Private AS numbers are not meant to be used for global Internet BGP routing, as they are assigned locally and can be used by any organization. They are meant to enable BGP within a enterprise or VPN, but since these numbers can be used by any organization they are not unique and could cause connectivity loss if leaked to the Internet.

QUESTION 282

Which two statements about the BGP community attribute are true? (Choose two.)

- A. Routers send the community attribute to all BGP neighbors automatically.
- B. A router can change a received community attribute before advertising it to peers.
- C. It is a well-known, discretionary BGP attribute.
- D. It is an optional transitive BGP attribute.
- E. A prefix can support only one community attribute.

Correct Answer: BD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

A community is a group of prefixes that share some common property and can be configured with the BGP community attribute. The BGP Community attribute is an optional transitive attribute of variable length. The attribute consists of a set of four octet values that specify a community. The community attribute values are encoded with an Autonomous System (AS) number in the first two octets, with the remaining two octets defined by the AS. A prefix can have more than one community attribute. A BGP speaker that sees multiple community attributes in a prefix can act based on one, some or all the attributes. A router has the option to add or modify a community attribute before the router passes the attribute on to other peers.

Reference:

<http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/28784-bgp-community.html>

QUESTION 283

Refer to the exhibit.

```
ip as-path access-list 1 permit ^64496_[0-9]*$
```

Which AS paths are matched by this access list?

- A. the origin AS 64496 only
- B. the origin AS 64496 and any ASs after AS 64496
- C. the directly attached AS 64496 and any ASs directly attached to AS 64496
- D. the directly attached AS 64496 and any longer AS paths

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

If you want AS 1 to get networks originated from AS 4 and all directly attached ASs of AS 4, apply the following inbound filter on Router 1.

```
ip as-path access-list 1 permit ^4_[0-9]*$ router bgp 1 neighbor 4.4.4.4 remote-as 4 neighbor 4.4.4.4 route-map foo in route-map foo permit 10 match as-path 1
```

In the **ip as-path access-list** command, the carat (^) starts the input string and designates "AS". The underscore (_) means there is a null string in the string that follows "AS 4". The [0-9]* specifies that any connected AS with a valid AS number can pass the filter. The advantage of using the [0-9]* syntax is that it gives you the flexibility to add any number of ASs without modifying this command string.

Reference: <http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13754-26.html>

QUESTION 284

Which two features improve BGP convergence? (Choose two.)

- A. next-hop address tracking
- B. additional paths
- C. advertise map
- D. communities
- E. soft reconfiguration

Correct Answer: AB

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The BGP Support for Next-Hop Address Tracking feature is enabled by default when a supporting Cisco software image is installed. BGP next-hop address tracking is event driven. BGP prefixes are automatically tracked as peering sessions are established. Next-hop changes are rapidly reported to the BGP routing process as they are updated in the RIB. This optimization improves overall BGP convergence by reducing the response time to next-hop changes for routes installed in the RIB. When a best path calculation is run in between BGP scanner cycles, only next-hop changes are tracked and processed.

BGP routers and route reflectors (RRs) propagate only their best path over their sessions. The advertisement of a prefix replaces the previous announcement of that prefix (this behavior is known as an implicit withdraw). The implicit withdraw can achieve better scaling, but at the cost of path diversity.

Path hiding can prevent efficient use of BGP multipath, prevent hitless planned maintenance, and can lead to MED oscillations and suboptimal hot-potato routing. Upon nexthop failures, path hiding also inhibits fast and local recovery because the network has to wait for BGP control plane convergence to restore traffic. The BGP Additional Paths feature provides a generic way of offering path diversity; the Best External or Best Internal features offer path diversity only in limited scenarios.

The BGP Additional Paths feature provides a way for multiple paths for the same prefix to be advertised without the new paths implicitly replacing the previous paths. Thus, path diversity is achieved instead of path hiding.

References: http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_bgp/configuration/15-1sg/irg-nexthop-track.html

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/irg-xe-3s-book/bgp_additional_paths.html

QUESTION 285

Which three statements about the route preference of IS-IS are true? (Choose three.)

- A. An L1 path is preferred over an L2 path.
- B. An L2 path is preferred over an L1 path.
- C. Within each level, a path that supports optional metrics is preferred over a path that supports only the default metric.
- D. Within each level of metric support, the path with the lowest metric is preferred.
- E. The Cisco IS-IS implementation usually performs equal cost path load balancing on up to eight paths.

F. Both L1 and L2 routes will be installed in the routing table at the same time.

Correct Answer: ACD

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Given multiple possible routes to a particular destination, an L1 path is preferred over an L2 path. Within each level, a path that supports the optional metrics is preferred over a path that supports only the default metric. (Again, Cisco supports only the default metric, so the second order of preference is not relevant to Cisco routers.) Within each level of metric support, the path with the lowest metric is preferred. If multiple equal-cost, equal-level paths are found by the Decision process, they are all entered into the route table. The Cisco IS-IS implementation usually performs equal-cost load balancing on up to six paths.

Reference: <http://www.realccielab.org/operation-of-integrated-is-is.html>

QUESTION 286

DRAG DROP

Drag and drop the IPv6 multicast feature on the left to its corresponding function on the right.

Select and Place:

PIMv2	communicates multicast group membership source awareness
MLDv1	uses only shared tree forwarding
MLDv2	communicates multicast group membership states from the
PIM-SSM	provides intradomain multicast forwarding for all underlying unicast routing protocols
PIM Bi-dir	aids IPv6 multicast deployment
IPv6 multicast over IPv4 tunnels	defined for interdomain use to support broadcast applications

Correct Answer:

- MLDv2
- PIM Bi-dir
- MLDv1
- PIMv2
- IPv6 multicast over IPv4 tunnels
- PIM-SSM

Section: Layer 3 Technologies**Explanation****Explanation/Reference:****QUESTION 287****DRAG DROP**

Drag and drop the EIGRP term on the left to the corresponding definition on the right.

Select and Place:

adjacency	the neighbor with the route that has the lowest metric
split horizon	a neighbor whose advertised distance is lower than the
successor	a feature that prevents routing loops
feasible successor	the logical association between two neighbors over which routing information is exchanged

Correct Answer:

	successor
	feasible successor
	split horizon
	adjacency

Section: Layer 3 Technologies
Explanation

Explanation/Reference:

QUESTION 288

DRAG DROP

Drag and drop the EIGRP query condition on the left to the corresponding action taken by the router on the right.

Select and Place:

The EIGRP table is missing an entry for the route	A feasible successor is installed in the routing table, and a reply
The EIGRP table lists the querying router as the successor for	The router replies with the successor information
The querying router is the successor, and no feasible successor exists	The router replies to the query with an unreachable message
the EIGRP table has a successor	The router send a query on all interfaces except the interface that had the successor route

Correct Answer:

	The EIGRP table lists the querying router as the successor for
	the EIGRP table has a successor
	The EIGRP table is missing an entry for the route
	The querying router is the successor, and no feasible successor exists

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 289**DRAG DROP**

Drag and drop the path-selection criteria on the left into the correct route-selection order on the right, that a router will use when having multiple routes toward the same destination.

Select and Place:

external Type 2 routes	first
external Type 1 routes	second
intra-area routes	third
inter-area routes	fourth

Correct Answer:

	intra-area routes
	inter-area routes
	external Type 1 routes
	external Type 2 routes

Section: Layer 3 Technologies**Explanation****Explanation/Reference:****QUESTION 290****DRAG DROP**

Drag and drop the multiprotocol BGP feature on the left to the corresponding description on the right.

Select and Place:

Multiprotocol Reachable NLRI	an optional, nontransitive attribute used to withdraw a route
Multiprotocol Unreachable NLRI	a value that indicates whether multiprotocol extensions are supported
AFI	an optional, nontransitive attribute used to advertise a
SAFI	a value that identifies a network protocol
Capability code	a value that identifies a subtype of network protocol

Correct Answer:

	Multiprotocol Unreachable NLRI
	Capability code
	Multiprotocol Reachable NLRI
	AFI
	SAFI

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

QUESTION 291

Refer to the exhibit.

```
R6#debug nhrp
NHRP protocol debugging is on
*Apr 14 02:05:29.416: NHRP: Attempting to send packet through interface Tunnel0 via DEST dst 10.250.20.1
*Apr 14 02:05:29.416: NHRP: Encapsulation succeeded. Sending NHRP Control Packet NBMA Address: 192.168.1.1
*Apr 14 02:05:29.416: NHRP: Send Registration Request via Tunnel0 vrf 0, packet size: 105
*Apr 14 02:05:29.416:      src: 10.250.20.6, dst: 10.250.20.1
*Apr 14 02:05:29.416: NHRP: 133 bytes out Tunnel0
*Apr 14 02:05:29.416: NHRP: Resetting retransmit due to hold-timer for 10.250.20.1
*Apr 14 02:05:30.306: NHRP: Setting retrans delay to 2 for nhs dst 10.250.20.1

R6#sh ip nhrp brief
  Target                Via                NBMA                Mode    Intfc    Claimed
R6#
```

NHRP registration is failing; what might be the problem?

- A. invalid IP addressing
- B. fragmentation
- C. incorrect NHRP mapping
- D. incorrect NHRP authentication

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Configuring an authentication string ensures that only routers configured with the same string can communicate using NHRP. Therefore, if the authentication scheme is to be used, the same string must be configured in all devices configured for NHRP on a fabric

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_4/ip_addr/configuration/guide/hadnhrp.html#wp1055

QUESTION 292

In GETVPN, which key is used to secure the control plane?

- A. Traffic Encryption Key (TEK)
- B. content encryption key (CEK)
- C. message encryption key (MEK)
- D. Key Encryption Key (KEK).

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

GDOI introduces two different encryption keys. One key secures the GET VPN control plane; the other key secures the data traffic. The key used to secure the control plane is commonly called the Key Encryption Key (KEK), and the key used to encrypt data traffic is known as Traffic Encryption Key (TEK).

Reference. Group Encrypted Transport VPN (Get VPN) Design and Implementation Guide PDF

QUESTION 293

Which statement is true comparing L2TPv3 to EoMPLS?

- A. L2TPv3 requires OSPF routing, whereas EoMPLS does not.
- B. EoMPLS requires BGP routing, whereas L2TPv3 does not.
- C. L2TPv3 carries L2 frames inside MPLS tagged packets, whereas EoMPLS carries L2 frames inside IPv4 packets.
- D. L2TPv3 carries L2 frames inside IPv4 packets, whereas EoMPLS carries L2 frames inside MPLS packets.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Ethernet-over-MPLS (EoMPLS) provides a tunneling mechanism for Ethernet traffic through an MPLS-enabled L3 core and encapsulates Ethernet protocol data units (PDUs) inside MPLS packets (using label stacking) to forward them across the MPLS network. Another technology that more or less achieves the result of AToM is L2TPV3. In the case of L2TPV3 Layer 2 frames are encapsulated into an IP packet instead of a labelled MPLS packet.

Reference. http://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r4-3/lxvpn/configuration/guide/lesc43xbook/lesc43p2ps.html

QUESTION 294

Which statement is true about VPLS?

- A. MPLS is not required for VPLS to work.
- B. VPLS carries packets as Layer 3 multicast.
- C. VPLS has been introduced to address some shortcomings of OTV.
- D. VPLS requires an MPLS network.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:


Explanation:

VPLS uses MPLS labels so an MPLS network is required. VPLS MPLS packets have a two-label stack. The outer label is used for normal MPLS forwarding in the service provider's network. If BGP is used to establish the VPLS, the inner label is allocated by a PE as part of a label block. If LDP is used, the inner label is a virtual circuit ID assigned by LDP when it first established a mesh between the participating PEs. Every PE keeps track of assigned inner label, and associates these with the VPLS instance.

Reference. http://en.wikipedia.org/wiki/Virtual_Private_LAN_Service

QUESTION 295

Refer to the exhibit.



Service provider SP 1 is running the MPLS-VPN service. The MPLS core network has MP-BGP configured with RR-1 as route reflector. What will be the effect on traffic between PE1 and PE2 if router P1 goes down?

- A. No effect, because all traffic between PE1 and PE2 will be rerouted through P2.
- B. No effect, because P1 was not the only P router in the forwarding path of traffic.
- C. No effect, because RR-1 will find an alternative path for MP-BGP sessions to PE-1 and PE-2.
- D. All traffic will be lost because RR-1 will lose the MP-BGP sessions to PE-1 and PE-2.

Correct Answer: D

Section: VPN Technologies**Explanation****Explanation/Reference:**

Explanation:

If the connection to the route reflector goes down, then routes from PE-1 will not get advertised to PE2, and vice versa. Route reflectors are critical in an MPLS VPN such as the one shown, which is why it is a best practice to have multiple route reflectors in this kind of network.

QUESTION 296

According to RFC 4577, OSPF for BGP/MPLS IP VPNs, when must the down bit be set?

- A. when an OSPF route is distributed from the PE to the CE, for Type 3 LSAs
- B. when an OSPF route is distributed from the PE to the CE, for Type 5 LSAs
- C. when an OSPF route is distributed from the PE to the CE, for Type 3 and Type 5 LSAs
- D. when an OSPF route is distributed from the PE to the CE, for all types of LSAs

Correct Answer: C

Section: VPN Technologies**Explanation****Explanation/Reference:**

Explanation:

If an OSPF route is advertised from a PE router into an OSPF area, the Down bit (DN) is set. Another PE router in the same area does not redistribute this route into iBGP of the MPLS VPN network if down is set.

RFC 4577 says:

"When a type 3 LSA is sent from a PE router to a CE router, the DN bit in the LSA Options field MUST be set. This is used to ensure that if any CE router sends this type 3 LSA to a PE router, the PE router will not redistribute it further. When a PE router needs to distribute to a CE router a route that comes from a site outside the latter's OSPF domain, the PE router presents itself as an ASBR (Autonomous System Border Router), and distributes the route in a type 5 LSA. The DN bit [OSPF-DN] MUST be set in these LSAs to ensure that they will be ignored by any other PE routers that receive them."

For more information about Down bit according to RFC 4577 please read more here.

<http://tools.ietf.org/html/rfc4577#section-4.2.5.1>.

QUESTION 297

Refer to the exhibit.

```
IPSEC(ipsec_process_proposal): proxy identities not supported
```

What is a possible reason for the IPSEC tunnel not establishing?

- A. The peer is unreachable.
- B. The transform sets do not match.
- C. The proxy IDs are invalid.
- D. The access lists do not match.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Proxy Identities Not Supported

This message appears in debugs if the access list for IPsec traffic does not match.

1d00h: IPSec(validate_transform_proposal): proxy identities not supported

1d00h: ISAKMP: IPSec policy invalidated proposal

1d00h: ISAKMP (0:2): SA not acceptable!

The access lists on each peer needs to mirror each other (all entries need to be reversible). This example illustrates this point.

Peer A

```
access-list 150 permit ip 172.21.113.0 0.0.0.255 172.21.114.0 0.0.0.255
```

```
access-list 150 permit ip host 15.15.15.1 host 172.21.114.123
```

Peer B

```
access-list 150 permit ip 172.21.114.0 0.0.0.255 172.21.113.0 0.0.0.255
```

```
access-list 150 permit ip host 172.21.114.123 host 15.15.15.1
```

Reference. <http://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/5409-ipsec-debug-00.html#proxy>

QUESTION 298

What is a key advantage of Cisco GET VPN over DMVPN?

- A. Cisco GET VPN provides zero-touch deployment of IPSEC VPNs.
- B. Cisco GET VPN supports certificate authentication for tunnel establishment.
- C. Cisco GET VPN has a better anti-replay mechanism.
- D. Cisco GET VPN does not require a secondary overlay routing infrastructure.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

DMVPN requires overlaying a secondary routing infrastructure through the tunnels, which results in suboptimal routing while the dynamic tunnels are

built. The overlay routing topology also reduces the inherent scalability of the underlying IP VPN network topology.

Traditional point-to-point IPsec tunneling solutions suffer from multicast replication issues because multicast replication must be performed before tunnel encapsulation and encryption at the IPsec CE (customer edge) router closest to the multicast source. Multicast replication cannot be performed in the provider network because encapsulated multicasts appear to the core network as unicast data.

Cisco's Group Encrypted Transport VPN (GET VPN) introduces the concept of a trusted group to eliminate point-to-point tunnels and their associated overlay routing. All group members (GMs) share a common security association (SA), also known as a group SA. This enables GMs to decrypt traffic that was encrypted by any other GM. (Note that IPsec CE acts as a GM.) In GET VPN networks, there is no need to negotiate point-to-point IPsec tunnels between the members of a group, because GET VPN is "tunnel-less."

Reference. Group Encrypted Transport VPN (Get VPN) Design and Implementation Guide PDF

QUESTION 299

Refer to the exhibit.

```
interface Tunnel0
 ip address 172.16.1.2 255.255.255.0
 ip nhrp map 172.16.1.1 192.168.1.1
 ip nhrp network-id 1
 ip nhrp nhs 172.16.1.1
 tunnel source 192.168.2.2
 ip mtu 1416
```

What is wrong with the configuration of the tunnel interface of this DMVPN Phase II spoke router?

- A. The interface MTU is too high.
- B. The tunnel destination is missing.
- C. The NHRP NHS IP address is wrong.
- D. The tunnel mode is wrong.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

By default, tunnel interfaces use GRE as the tunnel mode, but a DMVPN router needs to be configured for GRE multipoint by using the "tunnel mode gre multipoint" interface command.

QUESTION 300

Which two statements are true about VPLS? (Choose two.)

- A. It can work over any transport that can forward IP packets.
- B. It provides integrated mechanisms to maintain First Hop Resiliency Protocols such as HSRP, VRRP, or GLBP.
- C. It includes automatic detection of multihoming.
- D. It relies on flooding to propagate MAC address reachability information.
- E. It can carry a single VLAN per VPLS instance.

Correct Answer: DE

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

VPLS relies on flooding to propagate MAC address reachability information. Therefore, flooding cannot be prevented.

VPLS can carry a single VLAN per VPLS instance. To multiplex multiple VLANs on a single instance, VPLS uses IEEE QinQ.

Reference. http://www.cisco.com/c/en/us/products/collateral/switches/nexus-7000-series-switches/white_paper_c11-574984.html

QUESTION 301

Refer to the exhibit.

```
!  
ip vrf Cust123  
  rd 200:3000  
  export map Cust123mgmt  
  route-target export 200:3000  
!  
route-map Cust123mgmt permit 10  
  set extcommunity rt 200:9999  
!
```

What will be the extended community value of this route?

- A. RT:200:3000 RT:200:9999
- B. RT:200:9999 RT:200:3000
- C. RT:200:3000
- D. RT:200:9999

Correct Answer: D

Section: VPN Technologies**Explanation****Explanation/Reference:**

Explanation:

Here the route map is being used to manually set the extended community RT to 200:9999

QUESTION 302

Refer to the exhibit.

```
CE1#trace
Protocol [ip]: ipv6
Target IPv6 address: 2001:db8:100:1::7
Source address: 2001:db8:100:1::5
Insert source routing header? [no]:
Numeric display? [no]:
Timeout in seconds [3]:
Probe count [3]:
Minimum Time to Live [1]:
Maximum Time to Live [30]:
Priority [0]:
Port Number [0]:
Type escape sequence to abort.
Tracing the route to 2001:10:100:1::7

 1 2001:db8:1:5::1 1 msec 1 msec 1 msec
 2 ::FFFF:10.1.2.4 [MPLS: Labels 17/23 Exp 0] 2 msec 2 msec 2 msec
 3 2001:db8:1:7::2 [AS 1] [MPLS: Label 23 Exp 0] 2 msec 1 msec 1 msec
 4 2001:db8:1:7::7 [AS 1] 2 msec 1 msec 2 msec
```

Which statement is true?

- A. There is an MPLS network that is running 6PE, and the ingress PE router has no mpls ip propagate-ttl.
- B. There is an MPLS network that is running 6VPE, and the ingress PE router has no mpls ip propagate-ttl.
- C. There is an MPLS network that is running 6PE or 6VPE, and the ingress PE router has mpls ip propagate-ttl.
- D. There is an MPLS network that is running 6PE, and the ingress PE router has mpls ip propagate- ttl.

E. There is an MPLS network that is running 6VPE, and the ingress PE router has mpls ip propagate-ttl.

Correct Answer: C

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

The second hop shows an IPv6 address over MPLS, so we know that there is an MPLS network running 6PE or 6VPE. And because the second and third hops show up in the traceroute. Then TTL is being propagated because if the "no ip propagate-ttl" command was used these devices would be hidden in the traceroute.

QUESTION 303

Refer to the exhibit.

```
vrf definition one
 rd 1:1
  route-target export 100:1
  route-target import 100:1
 !
 address-family ipv4
  route-target import 100:2
 exit-address-family
 !
 address-family ipv6
  route-target export 100:3
  route-target import 100:3
 exit-address-family
```

Which statement is true about a VPNv4 prefix that is present in the routing table of vrf one and is advertised from this router?

- A. The prefix is advertised only with route target 100:1.
- B. The prefix is advertised with route targets 100:1 and 100:2.
- C. The prefix is advertised only with route target 100:3.
- D. The prefix is not advertised.
- E. The prefix is advertised with route targets 100:1, 100:2, and 100:3.

Correct Answer: A

Section: VPN Technologies**Explanation****Explanation/Reference:**

Explanation:

The route target used for prefix advertisements to other routers is defined on the route-target export command, which shows 100:1 in this case for VPNv4 routes.

QUESTION 304

Which is the way to enable the control word in an L2 VPN dynamic pseudowire connection on router R1?

- A. R1(config)# pseudowire-class cw-enable
R1(config-pw-class)# encapsulation mpls
R1(config-pw-class)# set control-word
- B. R1(config)# pseudowire-class cw-enable
R1(config-pw-class)# encapsulation mpls
R1(config-pw-class)# enable control-word
- C. R1(config)# pseudowire-class cw-enable
R1(config-pw-class)# encapsulation mpls
R1(config-pw-class)# default control-word
- D. R1(config)# pseudowire-class cw-enable
R1(config-pw-class)# encapsulation mpls
R1(config-pw-class)# control-word

Correct Answer: D

Section: VPN Technologies**Explanation****Explanation/Reference:**

Explanation:

The following example shows how to enable the control word in an AToM dynamic pseudowire connection:

```
Device(config)# pseudowire-class cw-enable
Device(config-pw-class)# encapsulation mpls
Device(config-pw-class)# control-word
Device(config-pw-class)# exit
```

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mps/command/mp-cr-book/mp-a1.html>

QUESTION 305

Where is multicast traffic sent, when it is originated from a spoke site in a DMVPN phase 2 cloud?

- A. spoke-spoke

- B. nowhere, because multicast does not work over DMVPN
- C. spoke-spoke and spoke-hub
- D. spoke-hub

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Spokes map multicasts to the static NBMA IP address of the hub, but hub maps multicast packets to the "dynamic" mappings that is, the hub replicates multicast packets to all spokes registered via NHRP, so multicast traffic is sent to the hub from a spoke instead of to the other spokes directly.

QUESTION 306

Refer to the exhibit.

```
Router-A# show ip nhrp
10.0.2.1/32 via 10.0.2.1, Tunnel0 created 00:00:21, expire 00:05:38
  Type: dynamic, Flags: authoritative unique registered used
NBMA address: 144.254.21.2
  (Claimed NBMA address: 172.16.2.1)

Router-B# show ip nhrp
10.0.1.1/32 via 10.0.1.1, Tunnel0 created 00:00:13, expire 00:05:48
  Type: dynamic, Flags: authoritative unique registered used
NBMA address: 72.34.1.2
```

A spoke site that is connected to Router-A cannot reach a spoke site that is connected to Router-B, but both spoke sites can reach the hub. What is the likely cause of this issue?

- A. There is a router doing PAT at site B.
- B. There is a router doing PAT at site A.
- C. NHRP is learning the IP address of the remote spoke site as a /32 address rather than a /24 address.
- D. There is a routing issue, as NHRP registration is working.

Correct Answer: B

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

If one spoke is behind one NAT device and another different spoke is behind another NAT device, and Peer Address Translation (PAT) is the type of NAT used on both NAT devices, then a session initiated between the two spokes cannot be established.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/ios_xe/sec_secure_connectivity/configuration/guide/convert/sec_dmvpn_xe_3s_book/sec_dmvpn_dt_spokes_b_nat_xe.html

QUESTION 307

Which three statements are functions that are performed by IKE phase 1? (Choose three.)

- A. It builds a secure tunnel to negotiate IKE phase 1 parameters.
- B. It establishes IPsec security associations.
- C. It authenticates the identities of the IPsec peers.
- D. It protects the IKE exchange by negotiating a matching IKE SA policy.
- E. It protects the identities of IPsec peers.
- F. It negotiates IPsec SA parameters.

Correct Answer: CDE

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

The basic purpose of IKE phase 1 is to authenticate the IPSec peers and to set up a secure channel between the peers to enable IKE exchanges. IKE phase 1 performs the following functions:

- Authenticates and protects the identities of the IPSec peers
- Negotiates a matching IKE SA policy between peers to protect the IKE exchange
- Performs an authenticated Diffie-Hellman exchange with the end result of having matching shared secret keys
- Sets up a secure tunnel to negotiate IKE phase 2 parameters

Reference. <http://www.ciscopress.com/articles/article.asp?p=25474&seqNum=7>

QUESTION 308

The session status for an IPsec tunnel with IPv6-in-IPv4 is down with the error message IKE message from 10.10.1.1 failed its sanity check or is malformed.

Which statement describes a possible cause of this error?

- A. There is a verification failure on the IPsec packet.

- B. The SA has expired or has been cleared.
- C. The pre-shared keys on the peers are mismatched.
- D. There is a failure due to a transform set mismatch.
- E. An incorrect packet was sent by an IPsec peer.

Correct Answer: C

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

IKE Message from X.X.X.X Failed its Sanity Check or is Malformed

This **debug** error appears if the pre-shared keys on the peers do not match. In order to fix this issue, check the pre-shared keys on both sides. 1d00H:%CRPTO-4-IKMP_BAD_MESSAGE. IKE message from 150.150.150.1 failed its sanity check or is malformed.

Reference. <http://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/5409-ipsec-debug-00.html#ike>

QUESTION 309

Which three statements describe the characteristics of a VPLS architecture? (Choose three.)

- A. It forwards Ethernet frames.
- B. It maps MAC address destinations to IP next hops.
- C. It supports MAC address aging.
- D. It replicates broadcast and multicast frames to multiple ports.
- E. It conveys MAC address reachability information in a separate control protocol.
- F. It can suppress the flooding of traffic.

Correct Answer: ACD

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

As a VPLS forwards Ethernet frames at Layer 2, the operation of VPLS is exactly the same as that found within IEEE 802.1 bridges in that VPLS will self learn source MAC address to port associations, and frames are forwarded based upon the destination MAC address. Like other 802.1 bridges, MAC address aging is supported.

Reference.

http://www.cisco.com/en/US/products/hw/routers/ps368/products_white_paper09186a00801f6084.shtml

QUESTION 310

A GRE tunnel is down with the error message %TUN-5-RECURDOWN: Tunnel0 temporarily disabled due to recursive routing error.

Which two options describe possible causes of the error? (Choose two.)

- A. Incorrect destination IP addresses are configured on the tunnel.
- B. There is link flapping on the tunnel.
- C. There is instability in the network due to route flapping.
- D. The tunnel mode and tunnel IP address are misconfigured.
- E. The tunnel destination is being routed out of the tunnel interface.

Correct Answer: CE

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

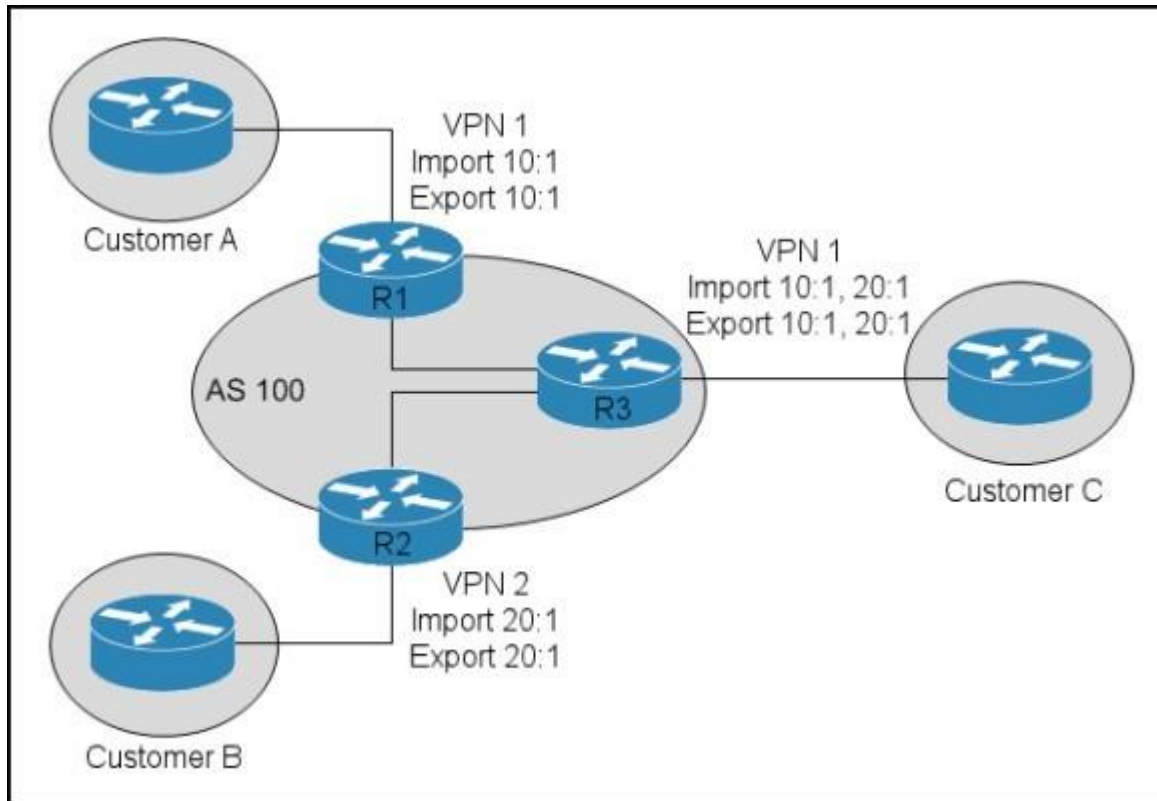
The %TUN-5-RECURDOWN: Tunnel0 temporarily disabled due to recursive routing error message means that the generic routing encapsulation (GRE) tunnel router has discovered a recursive routing problem. This condition is usually due to one of these causes:

- A misconfiguration that causes the router to try to route to the tunnel destination address using the tunnel interface itself (recursive routing)
- A temporary instability caused by route flapping elsewhere in the network

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

QUESTION 311

Refer to the exhibit.



Which two statements about the VPN solution are true? (Choose two.)

- A. Customer A and customer B will exchange routes with each other.
- B. R3 will advertise routes received from R1 to R2.
- C. Customer C will communicate with customer A and B.
- D. Communication between sites in VPN1 and VPN2 will be blocked.
- E. R1 and R2 will receive VPN routes advertised by R3.

Correct Answer: CE

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

+ VPN1 exports 10:1 while VPN3 imports 10:1 so VPN3 can learn routes of VPN1. + VNP1 imports 10:1 while VNP3 export 10:1 so VNP1 can learn routes of VPN3.

-> Customer A can communicate with Customer C

+ VPN2 exports 20:1 while VPN3 imports 20:1 so VPN3 can learn routes of VPN2. + VPN2 imports 20:1 while VPN3 exports 20:1 so VPN2 can learn routes of VPN3.

-> Customer B can communicate with Customer C

Therefore answer C is correct.

Also answer E is correct because R1 & R2 import R3 routes.

Answer A is not correct because Customer A & Customer B do not import routes which are exported by other router. Customer A & B can only see Customer C.

Answer B is not correct because a router never exports what it has learned through importation. It only exports its own routes.

Answer D is correct because two VPN1 and VPN2 cannot see each other. Maybe in this question there are three correct answers.

QUESTION 312

Which mechanism does Cisco recommend for CE router interfaces that face the service provider for an EVPL circuit with multiple EVCs and multiple traffic classes?

- A. HCBWFQ
- B. LLQ
- C. tail drop
- D. WRED

Correct Answer: A

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

In a simple handoff, packets may be discarded in the service provider network, either because of congestion on a link without an appropriate QoS policy or because of a policer QoS configuration on the service provider network that serves to rate limit traffic accessing the WAN core. To address these issues, QoS on the CE device is applied at a per-port level. A QoS service policy is configured on the outside Ethernet interface, and this parent policy includes a shaper that then references a second or subordinate (child) policy that enables queueing within the shaped rate. This is called a hierarchical CBWFQ (HCBWFQ) configuration.

Reference.

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/Ethernet_Access_for_NG_MAN_WAN_V3-1_external.html

QUESTION 313

Which Carrier Ethernet service supports the multiplexing of multiple point-to-point EVCs across as a single UNI?

- A. EPL
- B. EVPL
- C. EMS
- D. ERMS

Correct Answer: B

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Ethernet Relay Service (ERS or EVPL)

An Ethernet Virtual Circuit (EVC) is used to logically connect endpoints, but multiple EVCs could exist per single UNI. Each EVC is distinguished by 802.1q VLAN tag identification. The ERS network acts as if the Ethernet frames have crossed a switched network, and certain control traffic is not carried between ends of the EVC. ERS is analogous to Frame Relay where the CE-VLAN tag plays the role of a Data-Link Connection Identifier (DLCI). The MEF term for this service is EVPL.

Reference. http://www.cisco.com/c/en/us/td/docs/net_mgmt/ip_solution_center/5-1/carrier_ethernet/user/guide/l2vpn51book/concepts.html

QUESTION 314

What is the purpose of Route Target Constraint?

- A. to avoid using route reflectors in MPLS VPN networks
- B. to avoid using multiple route distinguishers per VPN in MPLS VPN networks
- C. to be able to implement VPLS with BGP signaling
- D. to avoid sending unnecessary BGP VPNv4 or VPNv6 updates to the PE router
- E. to avoid BGP having to perform route refreshes

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Some service providers have a very large number of routing updates being sent from RRs to PEs, using considerable resources. A PE does not need routing updates for VRFs that are not on the PE; therefore, the PE determines that many routing updates it receives are "unwanted." The PE can filter out the unwanted updates using Route Target Constraint.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/ios_xe/iproute_bgp/configuration/guide/2_xe/irg_xe_book/irg_rt_filter_xe.html

QUESTION 315

Refer to the exhibit.

```
P#show mpls forwarding-table
```

Local Label	Outgoing Label	Prefix or Tunnel Id	Bytes Switched	Outgoing interface	Next Hop
19	21	[mdt 1000:2000 0]	\		
			33516	Et2/0	10.1.2.2
19	19	[mdt 1000:2000 0]	\		
			912	Et1/0	10.1.1.1
20	24	[mdt 1000:2000 0]	\		
			1932	Et3/0	10.1.3.3
21	21	[mdt 1000:2000 0]	\		
			1932	Et2/0	10.1.2.2
23	24	[mdt 1000:2000 0]	\		
			33940	Et3/0	10.1.3.3
19	19	[mdt 1000:2000 0]	\		
			912	Et1/0	10.1.1.1

Which statement is true?

A. This is an MPLS TE point-to-multipoint LSP in an MPLS network.

- B. This is an MPLS TE multipoint-to-point LSP in an MPLS network.
- C. This is a point-to-multipoint LSP in an MPLS network.
- D. This is a multipoint-to-multipoint LSP in an MPLS network.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Same example of this provided on slide 24 at the reference link below:

Reference. "mVPN Deployment Models" Cisco Live Presentation <http://d2zmdbbm9feqrf.cloudfront.net/2014/eur/pdf/BRKIPM-2011.pdf>, slide 24

QUESTION 316

Refer to the exhibit.

```
R1#show mpls l2transport vc 100 detail
Local interface: Fa2/6 up, line protocol up, Ethernet up
  Destination address: 2.2.2.3, VC ID: 100, VC status: up
    Preferred path: Tunnel1, active
    Default path: ready
    Tunnel label: 12307, next hop point2point
    Output interface: Tu1, imposed label stack {12307 20}
  Create time: 00:00:11, last status change time: 00:00:11
  Signaling protocol: LDP, peer 2.2.2.3:0 up
    MPLS VC labels: local 21, remote 20
    Group ID: local 0, remote 2
    MTU: local 1500, remote 1500
    Remote interface description:
  Sequencing: receive disabled, send disabled
  VC statistics:
    packet totals: receive 1, send 6
    byte totals:   receive 368, send 0
    packet drops:  receive 0, send 0
```

Which statement is true?

- A. R1 routes this pseudowire over MPLS TE tunnel 1 with transport label 20.
- B. The default route 0.0.0.0/0 is available in the IPv4 routing table.
- C. R1 is using an MPLS TE tunnel for this pseudowire, because the IP path is not available.
- D. R1 has preferred-path configured for the pseudowire.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Verifying the Configuration: Example

In the following example, the **show mpls l2transport vc** command shows the following information (in bold) about the VCs:

- VC 101 has been assigned a preferred path called Tunnel1. The default path is disabled because the preferred path specified that the default path should not be used if the preferred path fails.
- VC 150 has been assigned an IP address of a loopback address on PE2. The default path can be used if the preferred path fails.

Router# **show mpls l2transport vc detail**

Local interfaceE. Gi0/0/0.1 up, line protocol up, Eth VLAN 222 up

Destination address: 10.16.16.16, VC ID. 101, VC status: up

Preferred path: Tunnel1, active

Default path: disabled

Tunnel label: 3, next hop point2point

Output interfaceE. Tu1, imposed label stack {17 16}

Create timeE. 00:27:31, last status change timeE. 00:27:31

Signaling protocol: LDP, peer 10.16.16.16:0 up

MPLS VC labels: local 25, remote 16

Group ID. local 0, remote 6

MTU: local 1500, remote 1500

Remote interface description:

Sequencing: receive disabled, send disabled

VC statistics:

packet totals: receive 10, send 10

byte totals: receive 1260, send 1300

packet drops: receive 0, send 0

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_2sr/12_2sra/feature/guide/srtunsel.html#wp1057815

QUESTION 317

For which kind of MPLS deployment is the next-hop-self all keyword used on a BGP neighbor command?

- A. 6VPE

- B. MPLS Carrier's carrier
- C. inter-AS MPLS VPN option D
- D. inter-AS MPLS VPN option C
- E. Unified MPLS

Correct Answer: E

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Since the core and aggregation parts of the network are integrated and end-to-end LSPs are provided, the Unified MPLS solution is also referred to as "Seamless MPLS."

New technologies or protocols are not used here, only MPLS, Label Distribution Protocol (LDP), IGP, and BGP. Since you do not want to distribute the loopback prefixes of the PE routers from one part of the network into another part, you need to carry the prefixes in BGP. The Internal Border Gateway Protocol (iBGP) is used in one network, so the next hop address of the prefixes is the loopback prefixes of the PE routers, which is not known by the IGP in the other parts of the network. This means that the next hop address cannot be used to recurse to an IGP prefix. The trick is to make the ABR routers Route Reflectors (RR) and set the next hop to self, even for the reflected iBGP prefixes. In order for this to work, a new knob is needed.

Only the RRs need newer software to support this architecture. Since the RRs advertise the BGP prefixes with the next hop set to themselves, they assign a local MPLS label to the BGP prefixes. This means that in the data plane, the packets forwarded on these end-to-end LSPs have an extra MPLS label in the label stack. The RRs are in the forwarding path.

There are two possible scenarios:

- The ABR does not set the next hop to self for the prefixes advertised (reflected by BGP) by the ABR into the aggregation part of the network. Because of this, the ABR needs to redistribute the loopback prefixes of the ABRs from the core IGP into the aggregation IGP. If this is done, there is still scalability. Only the ABR loopback prefixes (from the core) need to be advertised into the aggregation part, not the loopback prefixes from the PE routers from the remote aggregation parts.
- The ABR sets the next hop to self for the prefixes advertised (reflected by BGP) by the ABR into the aggregation part. Because of this, the ABR does not need to redistribute the loopback prefixes of the ABRs from the core IGP into the aggregation IGP. In both scenarios, the ABR sets the next hop to self for the prefixes advertised (reflected by BGP) by the ABR from the aggregation part of the network into the core part. If this is not done, the ABR needs to redistribute the loopback prefixes of the PEs from the aggregation IGP into the core IGP.

If this is done, there is no scalability.

In order to set the next hop to self for reflected iBGP routes, you must configure the **neighbor x.x.x.x next-hop-self all** command.

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

QUESTION 318

What is a reason for 6PE to use two MPLS labels in the data plane instead of one?

- A. 6PE allows penultimate hop popping and has a requirement that all P routers do not have to be IPv6 aware.
- B. 6PE does not allow penultimate hop popping.
- C. It allows MPLS traffic engineering to work in a 6PE network.
- D. It allows 6PE to work in an MPLS network where 6VPE is also deployed.

Correct Answer: A

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Q. Why does 6PE use two MPLS labels in the data plane?

A. 6PE uses two labels:

- The top label is the transport label, which is assigned hop-by-hop by the Label Distribution Protocol (LDP) or by MPLS traffic engineering (TE).
- The bottom label is the label assigned by the Border Gateway Protocol (BGP) and advertised by the internal BGP (iBGP) between the Provider Edge (PE) routers.

When the 6PE was released, a main requirement was that none of the MPLS core routers (the P routers) had to be IPv6-aware. That requirement drove the need for two labels in the data plane.

There are two reasons why the 6PE needs both labels.

PHP Functionality

If only the transport label were used, and if penultimate hop popping (PHP) were used, the penultimate hop router (the P router) would need to understand IPv6.

With PHP, this penultimate hop router would need to remove the MPLS label and forward the packet as an IPv6 packet. This P router would need to know that the packet is IPv6 because the P router would need to use the correct Layer 2 encapsulation type for IPv6. (The encapsulation type is different for IPv6 and IPv4; for example, for Ethernet, the encapsulation type is 0x86DD for IPv6, while it is 0x0800 for IPv4.) If the penultimate hop router is not IPv6-capable, it would likely put the Layer 2 encapsulation type for IPv4 for the IPv6 packet. The egress PE router would then believe that the packet was IPv4.

There is time-to-live (TTL) processing in both the IPv4 and IPv6 headers. In IPv6, the field is called Hop Limit. The IPv4 and IPv6 fields are at different locations in the headers. Also, the Header Checksum in the IPv4 header would also need to be changed; there is no Header Checksum field in IPv6. If the penultimate hop router is not IPv6-capable, it would cause the IPv6 packet to be malformed since the router expects to find the TTL field and Header Checksum field in the header.

Because of these differences, the penultimate hop router would need to know it is an IPv6 packet. How would this router know that the packet is an IPv6 packet, since it did not assign a label to the IPv6 Forwarding Equivalence Class (FEC), and there is no encapsulation field in the MPLS header? It could scan for the first nibble after the label stack and determine that the packet is IPv6 if the value is 6. However, that implies that the penultimate hop router needs to be IPv6-capable. This scenario could work if the explicit null label is used (hence no PHP). However, the decision was to require PHP.

Load Balancing

Typical load balancing on a P router follows this process. The P router goes to the end of the label stack and determines if it is an IPv4 packet by looking at the first nibble after the label stack.

- If the nibble has a value of 4, the MPLS payload is an IPv4 packet, and the P router load balances by hashing the source and destination IPv4 addresses.
- If the P router is IPv6-capable and the value of the nibble is 6, the P router load balances by hashing the source and destination IPv6 addresses.
- If the P router is not IPv6-capable and the value of the nibble is not 4 (it could be 6 if the packet is an IPv6 packet), the P router determines it is not an IPv4 packet and makes the load balancing decision based on the bottom label.

In the 6PE scenario, imagine there are two egress PE routers advertising one IPv6 prefix in BGP towards the ingress PE router. This IPv6 prefix would be advertised with two different labels in BGP. Hence, in the data plane, the bottom label would be either of the two labels. This would allow a P router to load balance on the bottom label on a per-flow basis.

If 6PE used only the transport label to transport the 6PE packets through the MPLS core, the P routers would not be able to load balance these packets

on a per-flow basis unless the P routers were IPv6-capable. If the P routers were IPv6-capable, they could use the source and destination IPv6 addresses in order to make a load balancing decision.

Reference. <http://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/116061-qa-6pe-00.html>

QUESTION 319

Refer to the exhibit.

```
R1
!
ip vrf R2
rd 1:1
!
interface FastEthernet0/0
ip address 192.168.0.1 255.255.255.252
!
router eigrp 100
no auto-summary
address-family ipv4 vrf R2
network 192.168.0.0 0.0.0.255
!

R2
!
interface FastEthernet0/0
ip address 192.168.0.2 255.255.255.252
!
router eigrp 100
no auto-summary
network 192.168.0.2 0.0.0.1
!
```

Which two corrective actions could you take if EIGRP routes from R2 fail to reach R1? (Choose two.)

- A. Configure R2 to use a VRF to send routes to R1.
- B. Configure the autonomous system in the EIGRP configuration of R1.
- C. Correct the network statement on R2.
- D. Add the interface on R1 that is connected to R2 into a VRF.

Correct Answer: BD

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

In this question we are running VRF Lite on R1. VRF Lite is also known as "VRF without running MPLS". This is an example of how to configure VRF Lite with EIGRP:

```
ip vrf FIRST
rd 1:1
!
ip vrf SECOND
rd 1:2
!
router eigrp 1
no auto-summary
!
address-family ipv4 vrf FIRST
network 10.1.1.1 0.0.0.0
no auto-summary
autonomous-system 200
exit-address-family
!
address-family ipv4 vrf SECOND
network 10.1.2.1 0.0.0.0
no auto-summary
autonomous-system 100
exit-address-family
!
interface FastEthernet0/0
ip vrf forwarding FIRST
ip address 10.1.1.1
255.255.255.0
!
interface FastEthernet0/1
ip vrf forwarding SECOND
ip address 10.1.2.1
255.255.255.0
```

The above example creates two VRFs (named "FIRST" and "SECOND"). VRF "FIRST" runs on EIGRP AS 200 while VRF "SECOND" runs on EIGRP AS 100. After that we have to add interfaces to the appropriate VRFs. From this example, back to our question we can see that R1 is missing the "autonomous-system ..." command under "address-family ipv4 vrf R2". And R1 needs an interface configured under that VRF.

Note. R2 does not run VRF at all! Usually R2 resides on customer side.

QUESTION 320

A service provider is deploying L2VPN LAN services in its MPLS cloud. Which statement is true regarding LDP signaling and autodiscovery?

- A. LDP signaling requires that each PE is identified, and that an LDP session is active with its P neighbor for autodiscovery to take place.
- B. LDP signaling requires that each P is identified, and that a targeted LDP session is active for autodiscovery to take place.
- C. LDP signaling requires that each PE is identified, and that a targeted LDP session with a BGP route reflector is active for autodiscovery to take place.
- D. LDP signaling requires that each PE is identified, and that a targeted LDP session is active for autodiscovery to take place.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

LDP signaling requires that each PE is identified and a targeted LDP session is active for autodiscovery to take place. Although the configuration can be automated using NMS/OSS the overall scalability of the solution is poor as a PE must be associated with all other PEs for LDP discovery to work, which can lead to a large number of targeted LDP sessions (n^2), which may be largely unused as not all VPLS will be associated with every PE. The security attributes of LDP are reasonably good, although additional configuration is required to prevent unauthorized sessions being set up. Although LDP can signal additional attributes, it requires additional configuration either from an NMS/OSS or static configuration.

Reference.

http://www.cisco.com/en/US/products/hw/routers/ps368/products_white_paper09186a00801f6084.shtml

QUESTION 321

Which attribute is not part of the BGP extended community when a PE creates a VPN-IPv4 route while running OSPF between PE-CE?

- A. OSPF domain identifier
- B. OSPF route type
- C. OSPF router ID
- D. MED
- E. OSPF network type

Correct Answer: E

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

By process of elimination, from RFC 4577:

For every address prefix that was installed in the VRF by one of its associated OSPF instances, the PE must create a VPN-IPv4 route in BGP. Each

such route will have some of the following Extended Communities attributes:

- The OSPF Domain Identifier Extended Communities attribute. If the OSPF instance that installed the route has a non-NULL primary Domain Identifier, this MUST be present; if that OSPF instance has only a NULL Domain Identifier, it MAY be omitted.
- OSPF Route Type Extended Communities Attribute. This attribute MUST be present. It is encoded with a two-byte type field, and its type is 0306.
- OSPF Router ID Extended Communities Attribute. This OPTIONAL attribute specifies the OSPF Router ID of the system that is identified in the BGP Next Hop attribute. More precisely, it specifies the OSPF Router Id of the PE in the OSPF instance that installed the route into the VRF from which this route was exported.
- MED (Multi_EXIT_DISC attribute). By default, this SHOULD be set to the value of the OSPF distance associated with the route, plus 1.

Reference. <https://tools.ietf.org/html/rfc4577>

QUESTION 322

What is a disadvantage of using aggressive mode instead of main mode for ISAKMP/IPsec establishment?

- A. It does not use Diffie-Hellman for secret exchange.
- B. It does not support dead peer detection.
- C. It does not support NAT traversal.
- D. It does not hide the identity of the peer.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

IKE phase 1's purpose is to establish a secure authenticated communication channel by using the DiffieHellman key exchange algorithm to generate a shared secret key to encrypt further IKE communications. This negotiation results in one single bi-directional ISAKMP Security Association (SA). The authentication can be performed using either pre-shared key (shared secret), signatures, or public key encryption. Phase 1 operates in either Main Mode or Aggressive Mode. Main Mode protects the identity of the peers; Aggressive Mode does not.

Reference. http://en.wikipedia.org/wiki/Internet_Key_Exchange

QUESTION 323

Which two statements are true about an EVPL? (Choose two.)

- A. It has a high degree of transparency.
- B. It does not allow for service multiplexing.
- C. The EVPL service is also referred to as E-line.
- D. It is a point-to-point Ethernet connection between a pair of UNIs.

Correct Answer: CD

Section: VPN Technologies**Explanation****Explanation/Reference:****Explanation:**

Following the MEF approach, the services that comprise the Metro Ethernet (ME) solution can be classified into the following two general categories:

- · Point-to-point (PtP)--A single point-to-point Ethernet circuit provisioned between two User Network Interfaces (UNIs).
- · Multipoint-to-multipoint (MPtMP)--A single multipoint-to-multipoint Ethernet circuit provisioned between two or more UNIs. When there are only two UNIs in the circuit, more UNIs can be added to the same Ethernet virtual connection if required, which distinguishes this from the point-to-point type.

In the MEF terminology, this maps to the following Ethernet service types:

- · Ethernet Line Service Type (E-Line)--Point-to-point Ethernet service
- · Ethernet LAN Service Type (E-LAN)--Multipoint-to-multipoint Ethernet service

Reference.

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/HA_Clusters/HA_Clusters/HA_ME3_6.pdf

QUESTION 324

Which two statements are true about OTV? (Choose two.)

- A. It relies on flooding to propagate MAC address reachability information.
- B. It uses a full mesh of point-to-multipoint tunnels to prevent head-end replication of multicast traffic.
- C. It can work over any transport that can forward IP packets.
- D. It supports automatic detection of multihoming.

Correct Answer: CD**Section: VPN Technologies****Explanation****Explanation/Reference:****Explanation:**

The overlay nature of OTV allows it to work over any transport as long as this transport can forward IP packets. Any optimizations performed for IP in the transport will benefit the OTV encapsulated traffic.

As part of the OTV control protocol, automatic detection of multihoming is included. This feature enables the multihoming of sites without requiring additional configuration or protocols

Reference. http://www.cisco.com/c/en/us/products/collateral/switches/nexus-7000-series-switches/white_paper_c11-574984.html

QUESTION 325

Which technology facilitates neighbor IP address resolution in DMVPN?

- A. CEF
- B. mGRE

- C. a dynamic routing protocol
- D. NHRP

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

NHRP Used with a DMVPN

NHRP is used to facilitate building a VPN and provides address resolution in DMVPN. In this context, a VPN consists of a virtual Layer 3 network that is built on top of an actual Layer 3 network. The topology you use over the VPN is largely independent of the underlying network, and the protocols you run over it are completely independent of it. The VPN network (DMVPN) is based on GRE IP logical tunnels that can be protected by adding in IPsec to encrypt the GRE IP tunnels.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_4/ip_addr/configuration/guide/hadnhrp.html#wp1057

QUESTION 326

Which two are features of DMVPN? (Choose two.)

- A. It does not support spoke routers behind dynamic NAT.
- B. It requires IPsec encryption.
- C. It only supports remote peers with statically assigned addresses.
- D. It supports multicast traffic.
- E. It offers configuration reduction.

Correct Answer: DE

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

DMVPN Hub-and-spoke deployment model: In this traditional topology, remote sites (spokes) are aggregated into a headend VPN device at the corporate headquarters (hub). Traffic from any remote site to other remote sites would need to pass through the headend device. Cisco DMVPN supports dynamic routing, QoS, and IP Multicast while significantly reducing the configuration effort.

Reference. http://www.cisco.com/c/en/us/products/collateral/security/dynamic-multipoint-vpn-dmvpn/data_sheet_c78-468520.html

QUESTION 327

Refer to the exhibit.

```
interface tunnel 1
 tunnel source ethernet 0
 tunnel mode ipv6ip isatap
 ipv6 address 2001:DB8::/64 eui-64
```

What is wrong with the configuration of this tunnel interface?

- A. ISATAP tunnels cannot use the EUI-64 address format.
- B. No tunnel destination has been specified.
- C. The tunnel source of an ISATAP tunnel must always point to a loopback interface.
- D. Router advertisements are disabled on this tunnel interface.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Sending of IPv6 router advertisements is disabled by default on tunnel interfaces. This command reenables the sending of IPv6 router advertisements to allow client autoconfiguration:

Example:

Router(config-if)# no ipv6 nd ra suppress

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/configuration/xe-3s/ir-xe-3s-book/ip6-isatap-xe.html>

QUESTION 328

Which two statements are true about a 6to4 tunnel connecting two IPv6 islands over the IPv4 Internet? (Choose two.)

- A. It embeds the IPv6 packet into the IPv4 payload with the protocol type set to 51.
- B. It works by appending the private IPv4 address (converted into hexadecimal format) to the 2002::/16 prefix.
- C. It embeds the IPv6 packet into the IPv4 payload with the protocol type set to 41.
- D. It works by appending the public IPv4 address (converted into hexadecimal format) to the 2002::/16 prefix.

Correct Answer: CD

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

6to4 embeds an IPv6 packet in the payload portion of an IPv4 packet with protocol type 41. To send an IPv6 packet over an IPv4 network to a 6to4 destination address, an IPv4 header with protocol type 41 is prepended to the IPv6 packet. The IPv4 destination address for the prepended packet header is derived from the IPv6 destination address of the inner packet (which is in the format of a 6to4 address), by extracting the 32 bits immediately following the IPv6 destination address's 2002::/16 prefix. The IPv4 source address in the prepended packet header is the IPv4 address of the host or router which is sending the packet over IPv4. The resulting IPv4 packet is then routed to its IPv4 destination address just like any other IPv4 packet.

Reference. <http://en.wikipedia.org/wiki/6to4>

QUESTION 329

Refer to the exhibit.

```
interface GigabitEthernet0/0
ip address 10.10.10.1 255.255.255.0
duplex auto
speed auto
media-type rj45
!
interface Tunnel0
ip address 192.168.1.1 255.255.255.252
tunnel source GigabitEthernet0/0
tunnel destination 192.168.1.240
```

What will be the IP MTU of tunnel 0?

- A. 1500
- B. 1524
- C. 1476
- D. 1452
- E. 1548

Correct Answer: C

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

In the case of the GRE tunnel interface, the IP maximum transmission unit (MTU) is 24 bytes less than the IP MTU of the real outgoing interface. For an Ethernet outgoing interface that means the IP MTU on the tunnel interface would be 1500 minus 24, or 1476 bytes.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/generic-routing-encapsulation-gre/13725-56.html>

QUESTION 330

On an MPLS L3VPN, which two tasks are performed by the PE router? (Choose two.)

- A. It exchanges VPNv4 routes with other PE routers.
- B. It typically exchanges iBGP routing updates with the CE device.
- C. It distributes labels and forwards labeled packets.
- D. It exchanges VPNv4 routes with CE devices.
- E. It forwards labeled packets between CE devices.

Correct Answer: AC

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

MPLS VPN functionality is enabled at the edge of an MPLS network. The PE router performs these tasks:

- Exchanges routing updates with the CE router
- Translates the CE routing information into VPN version 4 (VPNv4) routes
- Exchanges VPNv4 routes with other PE routers through the Multiprotocol Border Gateway Protocol (MP-BGP)

A PE router binds a label to each customer prefix learned from a CE router and includes the label in the network reachability information for the prefix that it advertises to other PE routers. When a PE router forwards a packet received from a CE router across the provider network, it labels the packet with the label learned from the destination PE router. When the destination PE router receives the labeled packet, it pops the label and uses it to direct the packet to the correct CE router. Label forwarding across the provider backbone is based on either dynamic label switching or traffic engineered paths. A customer data packet carries two levels of labels when traversing the backbone

Reference. http://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r4-2/lxvpn/configuration/guide/vcasr9kv342/vcasr9k42v3.html

QUESTION 331

Refer to the exhibit.

```
R1#show ip nhrp detail
10.1.0.2/32 via 10.1.0.2, Tunnel0 created 00:06:35, expire 00:00:29
  Type: dynamic, Flags: authoritative unique registered used
  NBMA address: 192.168.2.2
10.1.0.3/32 via 10.1.0.3, Tunnel0 created 00:05:28, expire 00:00:52
  Type: dynamic, Flags: authoritative unique registered used
  NBMA address: 192.168.3.3
```

Which statement describes what the authoritative flag indicates?

- A. Authentication was used for the mapping.
- B. R1 learned about the NHRP mapping from a registration request.
- C. Duplicate mapping in the NHRP cache is prevented.
- D. The registration request had the same flag set.

Correct Answer: B

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Show NHRP: Examples

The following is sample output from the show ip nhrp command:

Router# show ip nhrp

10.0.0.2 255.255.255.255, tunnel 100 created 0:00:43 expire 1:59:16

- TypE. dynamic Flags: authoritative
- NBMA address: 10.1111.1111.1111.1111.1111.1111.1111.1111.11

10.0.0.1 255.255.255.255, Tunnel0 created 0:10:03 expire 1:49:56

- TypE. static Flags: authoritative

The fields in the sample display are as follows:

Flags:

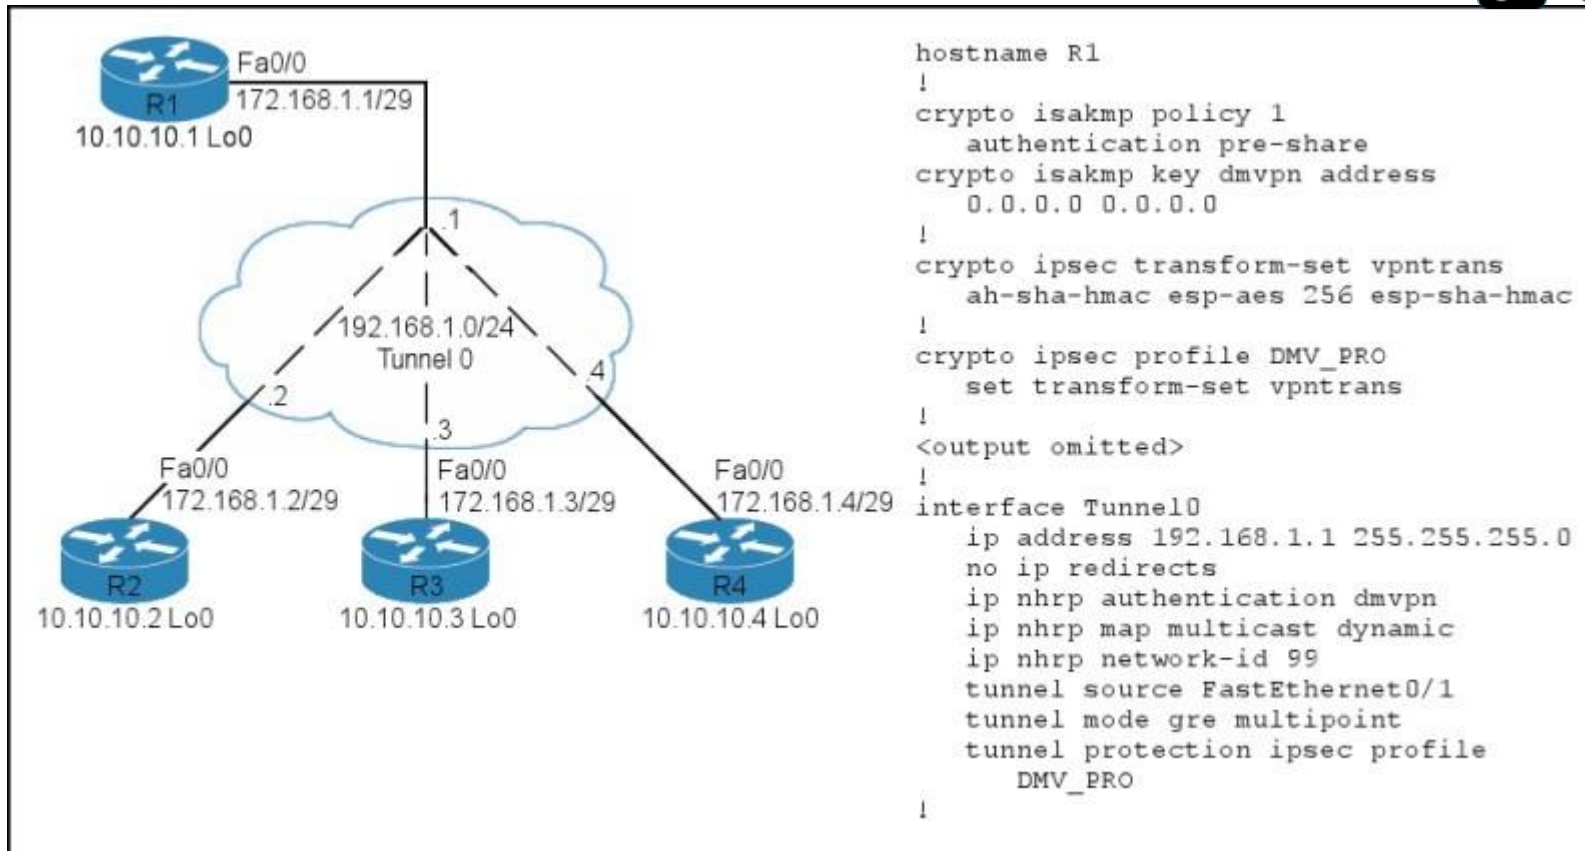
- authoritative--Indicates that the NHRP information was obtained from the Next Hop Server or router that maintains the NBMA-to-IP address mapping for a particular destination.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_4/ip_addr/configuration/guide/hadnhrp.html

QUESTION 332

Refer to the exhibit.



Which two statements about this configuration are true? (Choose two.)

- A. Spoke devices will be dynamically added to the NHRP mappings.
- B. The next-hop server address must be configured to 172.168.1.1 on all spokes.
- C. The next-hop server address must be configured to 192.168.1.1 on all spokes.
- D. R1 will create a static mapping for each spoke.

Correct Answer: AC

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

NHRP is a client/server model protocol which is defined by RFC2332. The hub is considered to be the Next Hop Server (NHS) and the spokes are considered to be the Next Hop Client (NHC). The hub must be configured as the next-hop server.

NHRP provides a mapping between the inside and outside address of a tunnel endpoint. These mappings can be static or dynamic. In a dynamic scenario, a next-hop server (NHS) is used to maintain a list of possible tunnel endpoints. Each endpoint using the NHS registers its own public and private mapping with the NHS. The local mapping of the NHS must always be static. It is important to note that the branch points to the inside or protected address of the NHS server. This scenario is an example of dynamic mappings.

Reference.

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/DMVPDG/DMVPN_2_Phase2.html

QUESTION 333

Which two tunneling techniques determine the IPv4 destination address on a per-packet basis? (Choose two.)

- A. 6to4 tunneling
- B. ISATAP tunneling
- C. manual tunneling
- D. GRE tunneling

Correct Answer: AB

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Tunnel Configuration Parameters by Tunneling Type				
Tunneling Type	Tunnel Configuration Parameter			
	Tunnel Mode	Tunnel Source	Tunnel Destination	Interface Prefix or Address
Manual	ipv6ip	An IPv4 address, or a reference to an interface on which IPv4 is configured.	An IPv4 address.	An IPv6 address.
GRE/IPv4	gre ip		An IPv4 address.	An IPv6 address.
IPv4-compatible	ipv6ip auto-tunnel		Not required. These are all point-to-multipoint tunneling types. The IPv4 destination address is calculated, on a per-packet basis, from the IPv6 destination.	Not required. The interface address is generated as <code>::tunnel-source/96</code> .
6to4	ipv6ip 6to4			An IPv6 address. The prefix must embed the tunnel source IPv4 address
ISATAP	ipv6ip isatap			An IPv6 prefix in modified eui-64 format. The IPv6 address is generated from the prefix and the tunnel source IPv4 address.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/ipv6/configuration/guide/12_4t/ipv6_12_4t_book/ip6-tunnel.html

QUESTION 334

Which two services are used to transport Layer 2 frames across a packet-switched network? (Choose two.)

- A. Frame Relay
- B. ATM
- C. ATOM
- D. L2TPv3

Correct Answer: CD

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Both AToM and L2TPv3 have the common objective of transmitting packet switched traffic of L2 frames (Frame Relay, ATM, and Ethernet) across a packet-switched network.

Reference. Layer 2 VPN Architectures - Google Books Result Wei Luo, Carlos Pignataro, Anthony Chan
<https://books.google.com/books?isbn=0132796864>

QUESTION 335

Which two statements about the C-bit and PW type are true? (Choose two.)

- A. The C-bit is 1 byte and the PW type is 15 bytes.
- B. The PW type indicates the type of pseudowire.
- C. The C-bit is 3 bits and the PW type is 10 bits.
- D. The C-bit set to 1 indicates a control word is present.
- E. The PW type indicates the encryption type.

Correct Answer: BD

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

The control word carries generic and Layer 2 payload-specific information. If the C-bit is set to 1, the advertising PE expects the control word to be present in every pseudowire packet on the pseudowire that is being signaled. If the C-bit is set to 0, no control word is expected to be present. Pseudowire Type--PW Type is a 15-bit field that represents the type of pseudowire.

Reference. <http://www.ciscopress.com/articles/article.asp?p=386788&seqNum=2>

QUESTION 336

Which statement describes the function of rekey messages?

- A. They prevent unencrypted traffic from passing through a group member before registration.
- B. They refresh IPsec SAs when the key is about to expire.
- C. They trigger a rekey from the server when configuring the rekey ACL.
- D. They authenticate traffic passing through a particular group member.

Correct Answer: B

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Rekey messages are used to refresh IPsec SAs. When the IPsec SAs or the rekey SAs are about to expire, one single rekey message for a particular group is generated on the key server. No new IKE sessions are created for the rekey message distribution. The rekey messages are distributed by the key server over an existing IKE SA. Rekeying can use multicast or unicast messages.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_getvpn/configuration/x3/sec-get-vpn-xe-3s-book/sec-get-vpn.html

QUESTION 337

Which three statements about GET VPN are true? (Choose three.)

- A. It encrypts WAN traffic to increase data security and provide transport authentication.
- B. It provides direct communication between sites, which reduces latency and jitter.
- C. It can secure IP multicast, unicast, and broadcast group traffic.
- D. It uses a centralized key server for membership control.
- E. It enables the router to configure tunnels.
- F. It maintains full-mesh connectivity for IP networks.

Correct Answer: ABD

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

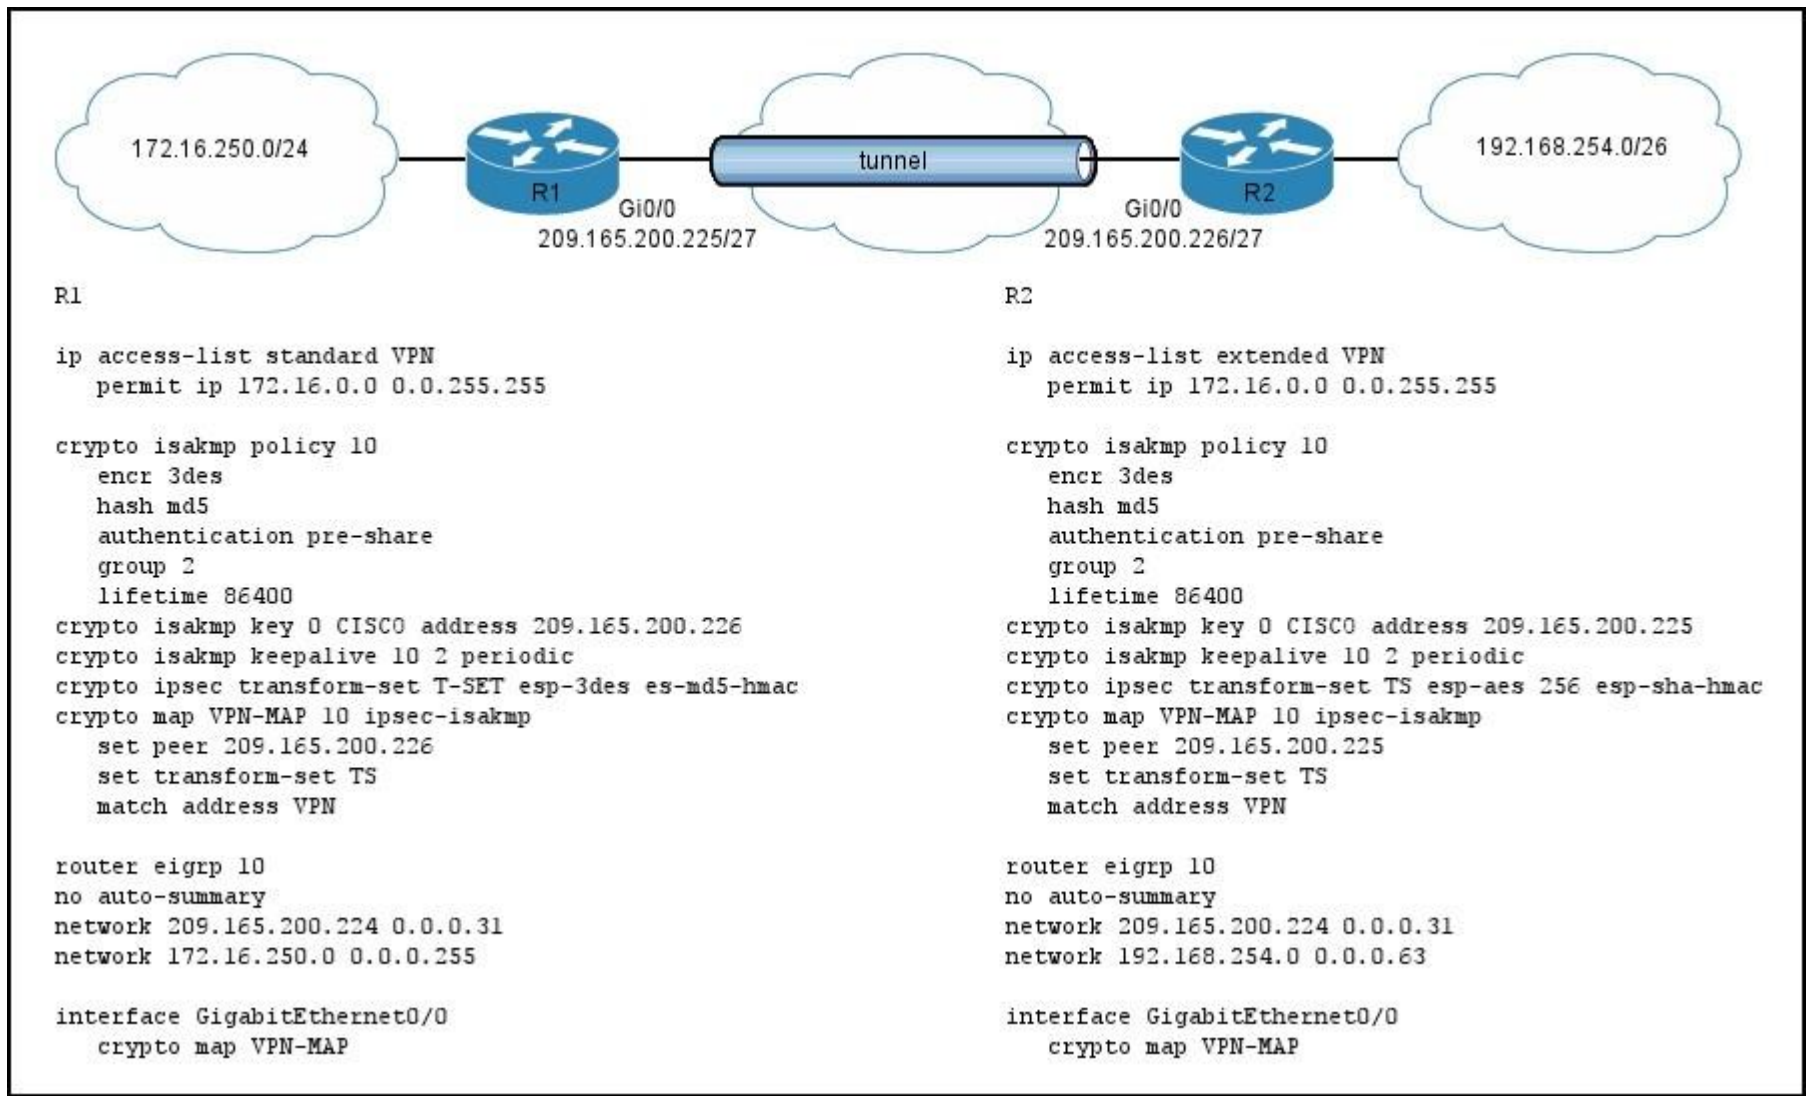
Cisco GET VPN Features and Benefits

Feature	Description and Benefit
Key Services	<p>Key Servers are responsible for ensuring that keys are granted to authenticated and authorized devices only. They maintain the freshness of the key material, pushing re-key messages as well as security policies on a regular basis. The chief characteristics include:</p> <ul style="list-style-type: none">• Key Servers can be located centrally, granting easy control over membership.• Key Servers are not in the "line of fire" - encrypted application traffic flows directly between VPN end points without a bottleneck or an additional point of failure.• Supports both local and global policies, applicable to all members in a group - such as "Permit any any", a policy to encrypt all traffic.• Supports IP Multicast to distribute and manage keys, for improved efficiency; Unicast is also supported where IP Multicast is not possible.
Scalability and Throughput	<ul style="list-style-type: none">• The full mesh nature of the solution allows devices to communicate directly with each other, without requiring transport through a central hub; this minimizes extra encrypts and decrypts at the hub router; it also helps minimize latency and jitter.• Efficient handling of IP Multicast traffic by using the core network for replication can boost effective throughput further
Security	<p>Provides data security and transport authentication, helping to meet security compliance and internal regulation by encrypting all WAN traffic</p>

Reference. http://www.cisco.com/c/en/us/products/collateral/security/group-encrypted-transport-vpn/product_data_sheet0900aecd80582067.html

QUESTION 338

Refer to the exhibit.



If the traffic flowing from network 192.168.254.0 to 172.16.250.0 is unencrypted, which two actions must you take to enable encryption? (Choose two).

- A. Configure the transform-set on R2 to match the configuration on R1.
- B. Configure the crypto map on R2 to include the correct subnet.
- C. Configure the ISAKMP policy names to match on R1 and R2.
- D. Configure the crypto map names to match on R1 and R2.
- E. Configure the Diffie-Hellman keys used in the ISAKMP policies to be different on R1 and R2.

Correct Answer: AB

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

A transform set combines an encryption method and an authentication method. During the IPsec security association negotiation with ISAKMP, the peers agree to use a particular transform set to protect a particular data flow. The transform set must be the same for both peers. Also, the crypto map on R2 points to the address name of VPN, which includes 172.16.0.0/16, but it should be the local subnet of 192.168.0.0/16

QUESTION 339

DRAG DROP

Drag and drop the NHRP flag on the left to the corresponding meaning on the right.

Select and Place:

Drag and drop the NHRP flag on the left to the corresponding meaning on the right.	
authoritative	NHRP information was learned from a forwarded NHRP packet.
implicit	The NHRP mapping entry is active and process-switched.
negative	NHRP information was obtained from the next hop server that maintains the NBMA mapping.
used	The requested NBMA mapping failed.

Correct Answer:

Drag and drop the NHRP flag on the left to the corresponding meaning on the right.

implicit

used

authoritative

negative

Section: VPN Technologies**Explanation****Explanation/Reference:****QUESTION 340**

Refer to the exhibit.

```
interface tunnel0
 tunnel mode ipv6ip 6to4
 tunnel source 125.203.89.1
 ipv6 address ?
```

Which statement is true about a valid IPv6 address that can be configured on interface tunnel0?

- A. There is not enough information to calculate the IPv6 address.
- B. 6to4 tunneling allows you to use any IPv6 address.
- C. 2001::7DCB::5901::/128 is a valid IPv6 address.
- D. 2002: 7DCB. 5901. ::/128 is a valid IPv6 address.

Correct Answer: D
Section: VPN Technologies
Explanation

Explanation/Reference:

Explanation:

Most IPv6 networks use autoconfiguration, which requires the last 64 bits for the host. The first 64 bits are the IPv6 prefix. The first 16 bits of the prefix are always 2002:, the next 32 bits are the IPv4 address, and the last 16 bits of the prefix are available for addressing multiple IPv6 subnets behind the same 6to4 router. Since the IPv6 hosts using autoconfiguration already have determined the unique 64 bit host portion of their address, they must simply wait for a Router Advertisement indicating the first 64 bits of prefix to have a complete IPv6 address. A 6to4 router will know to send an encapsulated packet directly over IPv4 if the first 16 bits are 2002, using the next 32 as the destination, or otherwise send the packet to a well-known relay server, which has access to native IPv6.

Reference. <http://en.wikipedia.org/wiki/6to4>

QUESTION 341

Which technology is not necessary to set up a basic MPLS domain?

- A. IP addressing
- B. an IGP
- C. LDP or TDP
- D. CEF
- E. a VRF

Correct Answer: E
Section: VPN Technologies
Explanation

Explanation/Reference:

Explanation:

The simplest form of VRF implementation is VRF Lite. In this implementation, each router within the network participates in the virtual routing environment in a peer-based fashion. While simple to deploy and appropriate for small to medium enterprises and shared data centres, VRF Lite does not scale to the size required by global enterprises or large carriers, as there is the need to implement each VRF instance on every router, including intermediate routers. VRFs were initially introduced in combination with MPLS, but VRF proved to be so useful that it eventually evolved to live independent of MPLS. This is the historical explanation of the term VRF Lite. usage of VRFs without MPLS.

Reference. http://en.wikipedia.org/wiki/Virtual_routing_and_forwarding

QUESTION 342

What is the main component of Unified MPLS?

- A. Multiple IGPs in the network are used, where the loopback IP addresses of the PE routers are aggregated on the area border routers.

- B. Confederations are used to provide scalability.
- C. The loopback prefixes from one IGP area are redistributed into BGP without changing the next hop.
- D. The ABR is a BGP route reflector and sets next-hop to self for all reflected routes.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Since the core and aggregation parts of the network are integrated and end-to-end LSPs are provided, the Unified MPLS solution is also referred to as "Seamless MPLS."

New technologies or protocols are not used here, only MPLS, Label Distribution Protocol (LDP), IGP, and BGP. Since you do not want to distribute the loopback prefixes of the PE routers from one part of the network into another part, you need to carry the prefixes in BGP. The Internal Border Gateway Protocol (iBGP) is used in one network, so the next hop address of the prefixes is the loopback prefixes of the PE routers, which is not known by the IGP in the other parts of the network. This means that the next hop address cannot be used to recurse to an IGP prefix. The trick is to make the ABR routers Route Reflectors (RR) and set the next hop to self, even for the reflected iBGP prefixes. In order for this to work, a new knob is needed.

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

QUESTION 343

For which feature is the address family "rtfilter" used?

- A. Enhanced Route Refresh
- B. MPLS VPN filtering
- C. Route Target Constraint
- D. Unified MPLS

Correct Answer: C

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

With Multiprotocol Label Switching (MPLS) VPN, the internal Border Gateway Protocol (iBGP) peer or Route Reflector (RR) sends all VPN4 and/or VPN6 prefixes to the PE routers. The PE router drops the VPN4/6 prefixes for which there is no importing VPN routing and forwarding (VRF). This is a behavior where the RR sends VPN4/6 prefixes to the PE router, which it does not need. This is a waste of processing power on the RR and the PE and a waste of bandwidth.

With Route Target Constraint (RTC), the RR sends only wanted VPN4/6 prefixes to the PE. 'Wanted' means that the PE has VRF importing the specific prefixes. RFC 4684 specifies Route Target Constraint (RTC). The support is through a new address family rtfilter for both VPNv4 and VPNv6.

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

QUESTION 344

Refer to the exhibit.

```
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 label entry,
      '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 70.169.72.33, return code 3
! size 100, reply addr 70.169.72.33, return code 3
! size 100, reply addr 70.169.72.33, return code 3
! size 100, reply addr 70.169.72.33, return code 3
! size 100, reply addr 70.169.72.33, return code 3

Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
```

What does the return code 3 represent in this output?

- A. The mapping of the replying router for the FEC is different.
- B. The packet is label-switched at stack depth.
- C. The return code is reserved.
- D. The upstream index is unknown.
- E. The replying router was the proper egress for the FEC.

Correct Answer: E

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Return Codes

The Return Code is set to zero by the sender. The receiver can set it to one of the values listed below. The notation <RSC> refers to the Return Subcode. This field is filled in with the stack-depth for those codes that specify that. For all other codes, the Return Subcode MUST be set to zero.

Value Meaning

- 0 No return code
- 1 Malformed echo request received
- 2 One or more of the TLVs was not understood
- 3 Replying router is an egress for the FEC at stack-depth <RSC>
- 4 Replying router has no mapping for the FEC at stack-depth <RSC>

Reference. <https://www.ietf.org/rfc/rfc4379.txt>

QUESTION 345

Which two values comprise the VPN ID for an MPLS VPN? (Choose two.)

- A. an OUI
- B. a VPN index
- C. a route distinguisher
- D. a 16-bit AS number
- E. a 32-bit IP address

Correct Answer: AB

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:


Each MPLS VPN ID defined by RFC 2685 consists of the following elements:

- An Organizational Unique Identifier (OUI), a three-octet hex number: The IEEE Registration Authority assigns OUIs to any company that manufactures components under the ISO/IEC 8802 standard. The OUI is used to generate universal LAN MAC addresses and protocol identifiers for use in local and metropolitan area network applications. For example, an OUI for Cisco Systems is 00-03-6B (hex).
- A Virtual Private Network (VPN) index: a four-octet hex number, which identifies the VPN within the company.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-mt/mp-l3-vpns-15-mt-book/mp-assgn-id-vpn.html

QUESTION 346

Refer to the exhibit.



Which LISP component do routers in the public IP network use to forward traffic between the two networks?

- A. EID
- B. RLOC
- C. map server
- D. map resolver

Correct Answer: B

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Locator ID Separation Protocol (LISP) is a network architecture and protocol that implements the use of two namespaces instead of a single IP address:

- Endpoint identifiers (EIDs)--assigned to end hosts.
- Routing locators (RLOCs)--assigned to devices (primarily routers) that make up the global routing system. The public networks use the RLOC to forward traffic between networks.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_lisp/configuration/15-mt/irl-15-mt-book/irl-overview.html

QUESTION 347

Refer to the exhibit.


```
NHRP: Send Registration Request via Tunnel1 vrf 0, packet size: 108
src: 172.30.10.66, dst: 172.30.10.1
(F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
  shtl: 4(NSAP), sstl: 0(NSAP)
  pktsz: 108 extoff: 52
(M) flags: "unique nat ", reqid: 113922
  src NBMA: 10.100.100.193
  src protocol: 172.30.10.66, dst protocol: 172.30.10.1
(C-1) code: no error(0)
  prefix: 32, mtu: 17912, hd_time: 600
  addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: Receive Registration Reply via
  addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
```

Which device role could have generated this debug output?

- A. an NHS only
- B. an NHC only
- C. an NHS or an NHC
- D. a DMVPN hub router

Correct Answer: B

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

NHRP works off a server/client relationship, where the NHRP clients (let's call them next hop clients/NHCs) register with their next hop server (NHS), it's the responsibility of the NHS to track all of its NHCs this is done with registration request and reply packets. Here we see a registration request, which can only be sent by an NHC.

QUESTION 348

Which statement about the NHRP network ID is true?

- A. It is sent from the spoke to the hub to identify the spoke as a member of the same NHRP domain.
- B. It is sent from the hub to the spoke to identify the hub as a member of the same NHRP domain.
- C. It is sent between spokes to identify the spokes as members of the same NHRP domain.
- D. It is a locally significant ID used to define the NHRP domain for an interface.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

The NHRP network ID is used to define the NHRP domain for an NHRP interface and differentiate between multiple NHRP domains or networks, when two or more NHRP domains (GRE tunnel interfaces) are available on the same NHRP node (router). The NHRP network ID is used to help keep two NHRP networks (clouds) separate from each other when both are configured on the same router.

The NHRP network ID is a local only parameter. It is significant only to the local router and it is not transmitted in NHRP packets to other NHRP nodes. For this reason the actual value of the NHRP network ID configured on a router need not match the same NHRP network ID on another router where both of these routers are in the same NHRP domain. As NHRP packets arrive on a GRE interface, they are assigned to the local NHRP domain in the NHRP network ID that is configured on that interface.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_4/ip_addr/configuration/guide/hadnhrp.html

QUESTION 349

You are configuring a DMVPN spoke to use IPsec over a physical interface that is located within a VRF. For which three configuration sections must you specify the VRF name? (Choose three.)

- A. the ISAKMP profile
- B. the crypto keyring
- C. the IPsec profile
- D. the IPsec transform set
- E. the tunnel interface
- F. the physical interface

Correct Answer: BEF

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

ip vrf forwarding <i>vrf-name</i> Example: Router(config-if)# ip vrf forwarding green	Associates a virtual private network (VPN) routing and forwarding (VRF) instance with an interface or subinterface. • <i>__vrf-name</i> is the name assigned to a VRF.
Router(config-if)# tunnel vrf <i>vrf-name</i> Example: Router(config-if)# tunnel vrf financial	Associates a VPN routing and forwarding (VRF) instance with a specific tunnel destination. <i>vrf-name</i> is the name assigned to a VRF.
Router(config)# crypto keyring <i>keyring-name</i> [<i>vrf fvr</i> <i>fvr-name</i>]	Defines a crypto keyring to be used during IKE authentication and enters keyring configuration mode. • <i>__keyring-name</i> —Name of the crypto keyring. • <i>__fvr-name</i> —(Optional) Front door virtual routing and forwarding (FVRF) name to which the keyring will be referenced. <i>fvr-name</i> must match the FVRF name that was defined during virtual routing and forwarding (VRF) configuration

QUESTION 350

Which IPv6 prefix is used for 6to4 tunnel addresses?

- A. 2001. . /23
- B. 2002. . /16
- C. 3ffe. . /16
- D. 5f00. . /8
- E. 2001. . /32

Correct Answer: B

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

6to4 works by taking advantage of a reserved IPv6 prefix, 2002::/16. A 6to4 tunnel interface automatically converts the 32 bits in its IPv6 address following this prefix to a global unicast IPv4 address for transport across an IPv4 network such as the public Internet.

Reference. <http://packetlife.net/blog/2010/mar/15/6to4-ipv6-tunneling/>

QUESTION 351

When you configure the `ip pmtu` command under an L2TPv3 pseudowire class, which two things can happen when a packet exceeds the L2TP path MTU? (Choose two.)

- A. The router drops the packet.
- B. The router always fragments the packet after L2TP/IP encapsulation.
- C. The router drops the packet and sends an ICMP unreachable message back to the sender only if the DF bit is set to 1.
- D. The router always fragments the packet before L2TP/IP encapsulation.
- E. The router fragments the packet after L2TP/IP encapsulation only if the DF bit is set to 0.
- F. The router fragments the packet before L2TP/IP encapsulation only if the DF bit is set to 0.

Correct Answer: CF

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

If you enable the `ip pmtu` command in the pseudowire class, the L2TPv3 control channel participates in the path MTU discovery. When you enable this feature, the following processing is performed:

- ICMP unreachable messages sent back to the L2TPv3 router are deciphered and the tunnel MTU is updated accordingly. In order to receive ICMP unreachable messages for fragmentation errors, the DF bit in the tunnel header is set according to the DF bit value received from the CE, or statically if the `ip dfbit set` option is enabled. The tunnel MTU is periodically reset to the default value based on a periodic timer.
- ICMP unreachable messages are sent back to the clients on the CE side. ICMP unreachable messages are sent to the CE whenever IP packets arrive on the CE-PE interface and have a packet size greater than the tunnel MTU. A Layer 2 header calculation is performed before the ICMP unreachable message is sent to the CE.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/l2tpv325.html

QUESTION 352

Which two parameters does the Tunnel Mode Auto Selection feature select automatically? (Choose two.)

- A. the tunneling protocol

- B. the transport protocol
- C. the ISAKMP profile
- D. the transform-set
- E. the tunnel peer

Correct Answer: AB

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

The Tunnel Mode Auto Selection feature eases the configuration and spares you about knowing the responder's details. This feature automatically applies the tunneling protocol (GRE or IPsec) and transport protocol (IPv4 or IPv6) on the virtual template as soon as the IKE profile creates the virtual access interface. This feature is useful on dual stack hubs aggregating multivendor remote access, such as Cisco AnyConnect VPN Client, Microsoft Windows7 Client, and so on.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_vpniips/configuration/xen-3s/sec-sec-for-vpns-w-ipsec-xe-3s-book/sec-ipsec-virt-tunnl.html

QUESTION 353

By default, how does a GET VPN group member router handle traffic when it is unable to register to a key server?

- A. All traffic is queued until registration is successful or the queue is full.
- B. All traffic is forwarded through the router unencrypted.
- C. All traffic is forwarded through the router encrypted.
- D. All traffic through the router is dropped.

Correct Answer: B

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

In the basic GETVPN configuration, the traffic passing through group members will be sent in clear until it registers with the Key Server. This is because the crypto ACL is configured on the KS and GM will get that information only after the registration is successful. This means for a short period of time the traffic can go out unencrypted after a GM is booted up or the existing GETVPN session is cleared manually. This mode is called "fail open" and it is the default behavior. This behavior can be turned off by configuring "Fail Close" mode on the GMs.

Reference. http://www.cisco.com/c/en/us/products/collateral/security/group-encrypted-transport-vpn/deployment_guide_c07_554713.html

QUESTION 354

DRAG DROP

Drag and drop each GET VPN feature on the left to the corresponding function it performs on the right.

Select and Place:

Drag and drop each GET VPN feature on the left to the corresponding function it performs on the right.	
GDOI	uses pseudotime to prevent replay
KEK	encrypts the rekey message
SAR	encrypts data between group members
TEK	handles communication between group members and a group controller or key server

Correct Answer:

Drag and drop each GET VPN feature on the left to the corresponding function it performs on the right.	
	SAR
	KEK
	TEK
	GDOI

Section: VPN Technologies
Explanation

Explanation/Reference:

QUESTION 355

MPLS LDP IGP synchronization is configured on a link. The OSPF adjacency on that link is UP but MPLS LDP synchronization is not achieved. Which statement about this scenario is true?

- A. The router excludes the link from its OSPF LSA type 1.
- B. The router flushes its own router LSA.
- C. The router advertises the link in its router LSA with max-metric.
- D. The router advertises an LSA type 2 for this link, with the metric set to max-metric.
- E. The router advertises the link and OSPF adjacency as it would when the synchronization is achieved.

Correct Answer: C

Section: VPN Technologies
Explanation

Explanation/Reference:

Explanation:

To enable LDP-IGP Synchronization on each interface that belongs to an OSPF or IS-IS process, enter the `mpls ldp sync` command. If you do not want some of the interfaces to have LDP-IGP Synchronization enabled, issue the `no mpls ldp igp sync` command on those interfaces. If the LDP peer is reachable, the IGP waits indefinitely (by default) for synchronization to be achieved. To limit the length of time the IGP session must wait, enter the `mpls ldp igp sync holddown` command. If the LDP peer is not reachable, the IGP establishes the adjacency to enable the LDP session to be established. When an IGP adjacency is established on a link but LDP-IGP Synchronization is not yet achieved or is lost, the IGP advertises the max-metric on that link.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fsldpsyn.html

QUESTION 356

What is the new designation for the MPLS EXP (experimental) bits?

- A. QoS bits
- B. traffic class bits
- C. flow bits
- D. precedence bits

Correct Answer: B

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

To avoid misunderstanding about how this field may be used, it has become increasingly necessary to rename this field. This document changes the name of the EXP field to the "Traffic Class field" ("TC field"). In doing so, it also updates documents that define the current use of the EXP field.

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

QUESTION 357

Which two options are signaling protocols that are used in MPLS? (Choose two.)

- A. LDP
- B. RSVP
- C. BFD
- D. LISP
- E. CLNS
- F. CDP

Correct Answer: AB

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

* *Signaling* is the means by which LSRs all along the path know that they are a part of a given LSP. It is a signaling function by which the LSR knows that the internal transit path for the LSP depicted goes from Interface 2 to Interface 4.

* *Label distribution* is the means by which an LSR tells an upstream LSR what label value to use for a particular LSP.

There are four protocols that can perform the label distribution function:

- * Label Distribution Protocol (LDP)
- * Resource Reservation Protocol with Traffic Engineering Extensions (RSVP-TE)
- * Constraint-Based Routed LDP (CR-LDP)
- * Multiprotocol BGP

LDP and RSVP-TE are the two most commonly used label distribution protocols

Reference. <http://www.networkworld.com/article/2237487/cisco-subnet/understanding-mpls-label-distribution.html>

QUESTION 358

Which option is an incorrect design consideration when deploying OSPF areas?

- A. area 1 - area 0 - MPLS VPN backbone - area 0 - area 2

- B. area 1 - MPLS VPN backbone - area 2
- C. area 1 - MPLS VPN backbone - area 1
- D. area 2 - area 0 - MPLS VPN backbone - area 1
- E. area 0 - area 2 - MPLS VPN superbackbone - area 1

Correct Answer: E

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

In the case of MPLS-VPN Backbone as The OSPF superbackbone behaves exactly like Area 0 in regular OSPF, so we cannot have two different area 0's that are not directly connected to each other. When area 0 connects to the superbackbone, it simply becomes an extension of area 0.

QUESTION 359

Refer to the exhibit.

```
ip vrf 10051
  rd 10.1.1.1:10051
  route-target export 64512:100010051
  route-target import 64512:100010051

ip access-list standard mgmt1-10051
  permit 192.168.1.0 0.0.0.255

route-map 10051-export permit 10
  match ip address mgmt1-10051
  set extcommunity rt 64512:3002300

route-map 10051-export permit 20
  match ip address mgmt1-10051
  set extcommunity rt 64512:2002250 64512:3002300 additive
```

Which statement about the route target for 192.168.1.0/24 is true?

- A. Its route target is 64512:100010051.
- B. Its route targets are 64512:100010051, 64512:2002250, and 64512:3002300.
- C. Its route target is 64512:3002300.
- D. Its route targets are 64512:100010051 and 64512:3002300.

E. Its route targets are 64512:2002250 and 64512:3002300.

Correct Answer: C

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Here we are using route maps to change the route target for the 192.168.1.0/24 network from the default route target of 64512:100010051 to 64512:3002300.

QUESTION 360

Which three options are best practices for implementing a DMVPN? (Choose three.)

- A. Use IPsec in tunnel mode.
- B. Implement Dead Peer Detection to detect communication loss.
- C. Configure AES for encryption of transported data.
- D. Configure SHA-1 for encryption of transported data.
- E. Deploy IPsec hardware acceleration to minimize router memory overhead.
- F. Configure QoS services only on the head-end router.

Correct Answer: ABC

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Best Practices Summary for Hub-and-Spoke Deployment Model

This section describes the best practices for a dual DMVPN cloud topology with the hub-and-spoke deployment, supporting IP multicast (IPmc) traffic including routing protocols.

The following are general best practices:

- · Use IPsec in transport mode
- · Configure Triple DES (3DES) or AES for encryption of transported data (exports of encryption algorithms to certain countries may be prohibited by law).
- · Implement Dead Peer Detection (DPD) on the spokes to detect loss of communication between peers.
- · Deploy hardware-acceleration of IPsec to minimize router CPU overhead, to support traffic with low latency and jitter requirements, and for the highest performance for cost.
- · Keep IPsec packet fragmentation to a minimum on the customer network by setting MTU size or using Path MTU Discovery (PMTUD).
- · Use Digital Certificates/Public Key Infrastructure (PKI) for scalable tunnel authentication.
- · Configure a routing protocol (for example, EIGRP, BGP or OSPF) with route summarization for dynamic routing.
- · Set up QoS service policies as appropriate on headend and branch router interfaces to help alleviate interface congestion issues and to attempt to keep higher priority traffic from being dropped during times of congestion.

Reference.

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/DMVPDG/DMVPN_1.html

QUESTION 361

Which three components comprise the structure of a pseudowire FEC element? (Choose three.)

- A. pseudowire ID
- B. pseudowire type
- C. control word
- D. Layer 3 PDU
- E. header checksum
- F. type of service

Correct Answer: ABC

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

The Pseudowire ID FEC element has the following components:

- Pseudowire ID FEC -- The first octet has a value of 128 that identifies it as a Pseudowire ID FEC element.
- Control Word Bit (C-Bit) -- The C-bit indicates whether the advertising PE expects the control word to be present for pseudowire packets. A control word is an optional 4-byte field located between the MPLS label stack and the Layer 2 payload in the pseudowire packet. The control word carries generic and Layer 2 payload-specific information. If the C-bit is set to 1, the advertising PE expects the control word to be present in every pseudowire packet on the pseudowire that is being signaled. If the C-bit is set to 0, no control word is expected to be present.
- Pseudowire Type -- PW Type is a 15-bit field that represents the type of pseudowire. Examples of pseudowire types are shown in Table 6-1.
- Pseudowire Information Length -- Pseudowire Information Length is the length of the Pseudowire ID field and the interface parameters in octets. When the length is set to 0, this FEC element stands for all pseudowires using the specified Group ID. The Pseudowire ID and Interface Parameters fields are not present.
- Group ID -- The Group ID field is a 32-bit arbitrary value that is assigned to a group of pseudowires.
- Pseudowire ID -- The Pseudowire ID, also known as VC ID, is a non-zero, 32-bit identifier that distinguishes one pseudowire from another. To connect two attachment circuits through a pseudowire, you need to associate each one with the same Pseudowire ID.
- Interface Parameters -- The variable-length Interface Parameters field provides attachment circuit-specific information, such as interface MTU, maximum number of concatenated ATM cells, interface description, and so on.

Reference. <http://www.ciscopress.com/articles/article.asp?p=386788&seqNum=2>

QUESTION 362

Which IPv6 tunneling type establishes a permanent link between IPv6 domains over IPv4?

- A. IPv4-compatible tunneling

- B. ISATAP tunneling
- C. 6to4 tunneling
- D. manual tunneling

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

A manually configured tunnel is equivalent to a permanent link between two IPv6 domains over an IPv4 backbone. The primary use is for stable connections that require regular secure communication between two edge routers or between an end system and an edge router, or for connection to remote IPv6 networks.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/ipv6/configuration/guide/12_4t/ipv6_12_4t_book/ip6-tunnel.html

QUESTION 363

In which two modes do IPv6-in-IPv4 tunnels operate? (Choose two.)

- A. tunnel mode
- B. transport mode
- C. 6to4 mode
- D. 4to6 mode
- E. ISATAP mode

Correct Answer: CE

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

There are 5 tunneling solution in IPv6:

*1. Using the "Tunnel mode ipv6ip", in this case the tunnel source and destination are configured with IPv4 addressing and the tunnel interface is configured with IPv6. This will use protocol 41.

This is used for IPv6/IPv4.

*

R1(config)#int tunnel 1

R1(config-if)#ipv6 address 12:1:12::1/64

```
R1(config-if)#tunnel source 10.1.12.1
R1(config-if)#tunnel destination 10.1.12.2
R1(config-if)#*tunnel mode ipv6ip*
```

*2. Using the "Tunnel mode gre ipv6, in this case the tunnel source and destination are all configured with IPv6 addressing. This is used for IPv6/IPv6.
*

```
BB1(config)#int tunnel 1
BB1(config-if)#ipv6 address 121:1:121::111/64
BB1(config-if)#tunnel source 10:1:111::111
BB1(config-if)#tunnel destination 10:1:112::112
BB1(config-if)#*tunnel mode gre ipv6*
```

*3. In this case, the third type, the tunnel mode is NOT used at all, note that the tunnel interface is configured with IPv6 and the tunnel source and destination is configured with IPv4 but no mention of tunnel mode. This configuration will use protocol 47. This is used for IPv6/IPv4.
*

```
R1(config)#int tunnel 13
R1(config-if)#ipv6 address 13:1:13::1/64
R1(config-if)#tunnel source 10.1.13.1
R1(config-if)#tunnel destination 10.1.13.3
```

*4. Note in this case a special addressing is assigned to the tunnel interface which is a concatenation of a reserved IPv6 address of 2002 followed by the translated IPv4 address of a given interface on the router. In this configuration ONLY the tunnel source address is used and since the tunnel is automatic, the destination address is NOT configured. The tunnel mode is set to "Tunnel mode ipv6ip 6to4. Note the IPv4 address of 10.1.1.1 is translated to 0A.01.01.01 and once concatenated, it will be "2002:0A01:0101: or 2002:A01:101. This is used for IPv6/IPv4.
*

```
R1(config)#interface Tunnel14
R1(config-if)#ipv6 address 2002:A01:101::128
R1(config-if)#tunnel source 10.1.1.1
R1(config-if)#*tunnel mode ipv6ip 6to4*
```

*5. ISATAP, ISATAP works like 6to4 tunnels, with one major difference, it uses a special IPv6 address which is formed as follows: *

*In this tunnel mode, the network portion can be any IPv6 address, whereas in 6to4 it had to start with 2002. *

Note when the IPv6 address is assigned to the tunnel interface, the "eui-64 is used, in this case the host portion of the IPv6 address starts with "0000.5EFE" and then the rest of the host portion is the translated IPv4 address of the tunnel's source IPv4 address. This translation is performed automatically unlike 6to4. This is used for IPv6/IPv4.

```
R4(config)#int tunnel 46
R4(config-if)#ipv6 address 46:1:46::/64 eui-64
R4(config-if)#tunnel source 10.44.44.44
R4(config-if)#*tunnel mode ipv6ip ISATAP*
```

QUESTION 364

Which VPN technology requires the use of an external key server?

- A. GETVPN
- B. GDOI
- C. SSL
- D. DMVPN
- E. IPsec
- F. L2TPv3

Correct Answer: A

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

A GETVPN deployment has primarily three components, Key Server (KS), Group Member (GM), and Group Domain of Interpretation (GDOI) protocol. GMs do encrypt/decrypt the traffic and KS distribute the encryption key to all the group members. The KS decides on one single data encryption key for a given life time. Since all GMs use the same key, any GM can decrypt the traffic encrypted by any other GM. GDOI protocol is used between the GM and KS for group key and group SA management. Minimum one KS is required for a GETVPN deployment.

Reference. http://www.cisco.com/c/en/us/products/collateral/security/group-encrypted-transport-vpn/deployment_guide_c07_554713.html

QUESTION 365

Which three roles does a key server perform when used with GETVPN? (Choose three.)

- A. It authenticates group members.
- B. It manages security policies.
- C. It creates group keys.
- D. It distributes multicast replication policies.
- E. It distributes multicast replication keys.
- F. It configures and routes the GDOI protocol.

Correct Answer: ABC

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Key server is responsible for maintaining security policies, authenticating the Group Members and providing the session key for encrypting traffic. KS

authenticates the individual GMs at the time of registration. Only after successful registration the GMs can participate in group SA. Reference. http://www.cisco.com/c/en/us/products/collateral/security/group-encrypted-transport-vpn/deployment_guide_c07_554713.html

QUESTION 366

What is the most secure way to store ISAKMP/IPSec preshared keys in Cisco IOS?

- A. Use the service password-encryption command.
- B. Encrypt the ISAKMP preshared key in secure type 5 format.
- C. Encrypt the ISAKMP preshared key in secure type 7 format.
- D. Encrypt the ISAKMP preshared key in secure type 6 format.

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Using the Encrypted Preshared Key feature, you can securely store plain text passwords in type 6 format in NVRAM using a command-line interface (CLI). Type 6 passwords are encrypted. Although the encrypted passwords can be seen or retrieved, it is difficult to decrypt them to find out the actual password. This is currently the most secure way to store keys.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_ikevpn/configuration/xen-3s/asr1000/sec-ike-for-ipsec-vpns-xe-3s-asr1000-book/sec-encrypt-preshare.html

QUESTION 367

DRAG DROP

Drag and drop the DMVPN command on the left to the corresponding function on the right.

Select and Place:

ip nhrp map group	configures mapping from an ip adress to an NBMA mapping
ip nhrp group	associates an NHRP group to a QoS policy
ip nhrp map	allows broadcast packets to be sent over a tunnel
ip nhrp map multicast	configured an NHRP group
ip nhrp nhs	designates the IP to use for communication to the next hop server
ip nhrp responder	specifies the next hop server

Correct Answer:

- ip nhrp map
- ip nhrp map group
- ip nhrp map multicast
- ip nhrp group
- ip nhrp nhs
- ip nhrp responder

Section: VPN Technologies

Explanation

Explanation/Reference:

QUESTION 368

DRAG DROP

Drag and drop the OTV component on the left to the function it performs on the right.

Select and Place:

edge device	elected by the OTV to provide loop-free multihoming
join interface	connects VLANs to be extended
internal interface	receives local OTV hello messages
overlay interface	provides an uplink to the overlay network
site VLAN	encapsulates layer 2 frames within an IP header
authoritative edge device	connects a site to an overlay network

Correct Answer:

- authoritative edge device
- internal interface
- site VLAN
- join interface
- overlay interface
- edge device

Section: VPN Technologies

Explanation

Explanation/Reference:

QUESTION 369

Which two events occur when a packet is decapsulated in a GRE tunnel? (Choose two.)

- A. The destination IPv4 address in the IPv4 payload is used to forward the packet.
- B. The TTL of the payload packet is decremented.
- C. The source IPv4 address in the IPv4 payload is used to forward the packet.
- D. The TTL of the payload packet is incremented.
- E. The version field in the GRE header is incremented.
- F. The GRE keepalive mechanism is reset.

Correct Answer: AB

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

After the GRE encapsulated packet reaches the remote tunnel endpoint router, the GRE packet is decapsulated. The destination address lookup of the outer IP header (this is the same as the tunnel destination address) will find a local address (receive) entry on the ingress line card.

The first step in GRE decapsulation is to qualify the tunnel endpoint, before admitting the GRE packet into the router, based on the combination of tunnel source (the same as source IP address of outer IP header) and tunnel destination (the same as destination IP address of outer IP header). If the received packet fails tunnel admittance qualification check, the packet is dropped by the decapsulation router. On successful tunnel admittance check, the decapsulation strips the outer IP and GRE header off the packet, then starts processing the inner payload packet as a regular packet.

When a tunnel endpoint decapsulates a GRE packet, which has an IPv4/IPv6 packet as the payload, the destination address in the IPv4/IPv6 payload packet header is used to forward the packet, and the TTL of the payload packet is decremented.

Reference. http://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r5-3/addr-serv/configuration/guide/b-ipaddr-cg53asr9k/b-ipaddr-cg53asr9k_chapter_01001.html

QUESTION 370

Refer to the exhibit.

```
crypto isakmp policy 1
  authentication pre-share
crypto isakmp key dmvpn address 0.0.0.0 0.0.0.0
crypto ipsec transform-set vpntrans ah-sha-hmac esp-aes 256 esp-sha-hmac
crypto ipsec profile DMVPN-PROF
  set transform-set vpntrans

policy-map SHAPE
  class class-default
    shape average 200000

interface Loopback0
  ip address 10.1.1.1 255.255.255.0

interface Tunnel0
  ip address 192.168.1.1 255.255.255.0
  no ip next-hop-self eigrp 1
  ip nhrp authentication dmvpn
  ip nhrp map multicast dynamic
  ip nhrp network-id 99
  ip tcp adjust-mss 1360
  no ip split-horizon eigrp 1
  tunnel source Serial1/0
  tunnel mode gre multipoint
  tunnel protection ipsec profile DMVPN-PROF

interface Serial1/0
  ip address 172.16.1.1 255.255.255.248
  encapsulation frame-relay
  frame-relay inverse-arp

router eigrp 1
  network 10.0.0.0
  network 192.168.1.0
```

The spokes of the DMVPN with the given configuration are having QoS issues.

Which two actions can you take to resolve the problem? (Choose two.)

A. Configure qos pre-classify on the tunnel interface.

- B. Configure an NHRP group on the tunnel interface and associate it to a QoS policy.
- C. Modify the configuration of the IPsec policy to accept QoS policies.
- D. Manually configure a QoS policy on the serial interface.
- E. Configure the bandwidth statement on the tunnel interface.
- F. Configure the bandwidth statement on the serial interface.

Correct Answer: AB

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

It is possible to classify based on information that is encrypted, which is needed in this example. You can use an access-list, configured to match the private subnet behind the remote spoke. The qos pre-classify command is used on the tunnel interface, and is required because the traffic is classified by a parameter that is encrypted as the traffic leaves the physical outbound interface. L4 information from the IP data packet can also classify traffic destined to the same private subnet.

The "**nhrp map group** *group-name* **service-policy output** *parent-policy-name* " command adds the NHRP group to the QoS policy map on the hub.

QUESTION 371

Which two statements about 6VPE are true? (Choose two.)

- A. It allows a service provider to use an existing MPLS network to provide VPN services to IPv6 customers.
- B. It uses MP-BGP as the carrier protocol to transport IPv6 connectivity.
- C. It provides IPv6 connectivity to MPLS-VPN customers when IPv6 overlay tunneling is also configured.
- D. It allows a service provider to use an existing MPLS network to provide global addressing to their IPv6 customers.
- E. It requires the configuration of a GRE tunnel tagged with a VLAN ID.
- F. It allows a service provider to use an existing L2TPv3 network to provide VPN services to IPv6 customers.

Correct Answer: AB

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

The IPv6 MPLS VPN service model is similar to that of IPv4 MPLS VPNs. Service providers who have already deployed MPLS IPv4 VPN services over an IPv4 backbone can deploy IPv6 MPLS VPN services over the same IPv4 backbone by upgrading the PE router IOS version and dual-stack configuration, without any change on the core routers. IPv4 services can be provided in parallel with IPv6 services. IPv6 VPN service is exactly the same as MPLS VPN for IPv4. 6VPE offers the same architectural features as MPLS VPN for IPv4. It offers IPv6 VPN and uses the same components, such as:

- Multiprotocol BGP (MP-BGP) VPN address family

- . Route distinguishers
- . VPN Routing and Forwarding (VRF) instances
- . Site of Origin (SOO)
- . Extended community
- . MP-BGP

Reference: http://www.cisco.com/c/en/us/td/docs/net_mgmt/ip_solution_center/5-2/mpls_vpn/user/guide/mpls52book/ipv6.html

QUESTION 372

Refer to the exhibit.

```
R1#show mpls l2transport vc 1611 detail

Local interface: Gi4/0/2 up, line protocol up, Eth VLAN 1611 up
  Destination address: 172.16.12.70, VC ID: 1611, VC status: down
    Output interface: none, imposed label stack {}
    Preferred path: not configured
    Default path: no route
    No adjacency
  Create time: 4w2d, last status change time: 4w2d
  Signaling protocol: LDP, peer 172.16.12.70:0 up
    Targeted Hello: 172.16.192.80(LDP Id) -> 172.16.12.70
    Status TLV support (local/remote)      : enabled/unknown (no remote binding)
      Label/status state machine           : local ready, LruRnd
      Last local dataplane status rcvd: no fault
      Last local SSS circuit status rcvd: no fault
      Last local SSS circuit status sent: not sent
      Last local LDP TLV status sent: no fault
      Last remote LDP TLV status rcvd: unknown (no remote binding)
    MPLS VC labels: local 4006, remote unassigned
    Group ID: local 0, remote unknown
    MTU: local 1500, remote unknown
    Remote interface description:
  Sequencing: receive disabled, send disabled
  VC statistics:
    packet totals: receive 0, send 0
    byte totals:   receive 0, send 0
    packet drops:  receive 0, seq error 0, send 0
```

Which three statements about the R1 configuration are true? (Choose three.)

- A. The virtual circuit identifier is 1611 and the virtual circuit is down.
- B. The local label for the circuit is 4006.
- C. The targeted LDP session to the remote peer is up.
- D. The local label for the circuit is 1611.
- E. The virtual circuit identifier is 4006 and the virtual circuit is down.
- F. The circuit is using MPLS VC type 4.

Correct Answer: ABC

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

The number after the vc is the identifier, which is 1611 in this case. Here, the VC status is shown as down.

As shown, the MPLS VC labels: local 4006, remote unassigned shows the local label used is 4006. The targeted LDP session is up as verified by the "Signalling protocol: LDP, peer 172.16.12.70 up" statement in the output.

QUESTION 373

Which statement about OTV is true?

- A. The overlay interface becomes active only when configuration is complete and it is manually enabled.
- B. OTV data groups can operate only in PIM sparse-mode.
- C. The overlay interface becomes active immediately when it is configured.
- D. The interface facing the OTV groups must be configured with the highest MTU possible.

Correct Answer: A

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

OTV has the following configuration guidelines and limitations:

- If the same device serves as the default gateway in a VLAN interface and the OTV edge device for the VLANs being extended, configure OTV on a device (VDC or switch) that is separate from the VLAN interfaces (SVIs).
- When possible, we recommend that you use a separate nondefault VDC for OTV to allow for better manageability and maintenance.
- An overlay interface will only be in an up state if the overlay interface configuration is complete and enabled (no shutdown). The join interface has to be in an up state.

Reference: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/OTV/config_guide/b_Cisco_Nexus_7000_Series_NX-OS_OTV_Configuration_Guide/basic-otv.html

QUESTION 374**DRAG DROP**

Drag and drop the MPLS term on the left to the function it performs on the right.

Select and Place:

label	instructs the router to keep the label when forwarding
implicit-null	groups IP packets so that they are given the same forwarding treatment
explicit-null	identifies the group to which an IP packet belongs
penultimate hop popping	instructs the penultimate router to pop the label before
FEC	identifies a layer 2 MPLS connection from one device to
virtual circuit	pops an MPLS label off one hop before its final destination

Correct Answer:

	explicit-null
	FEC
	label
	implicit-null
	virtual circuit
	penultimate hop popping

Section: VPN Technologies**Explanation****Explanation/Reference:****QUESTION 375****DRAG DROP**

Drag and drop the NHRP flag on the left to the corresponding meaning on the right.

Select and Place:

authoritative

The NHRP mapping entry was created from an NHRP registration request

implicit

NHRP information was obtained from the next hop server that maintains the NBMA-to-IP mapping

unique

NHRP information was learned from a forwarded NHRP packet

registered

The NHRP mapping entry is protected from being overwritten by a mapping address that has the same IP address and a different NBMA address.

Correct Answer:

registered

authoritative

implicit

unique

Section: VPN Technologies
Explanation

Explanation/Reference:

QUESTION 376

Refer to the exhibit.

```
DOT1X-SP-5-SECURITY_VIOLATION: Security violation on interface GigabitEthernet4/8,  
New MAC address 0080.ad00.c2e4 is seen on the interface in Single host mode  
%PM-SP-4-ERR_DISABLE: security-violation error detected on Gi4/8, putting Gi4/8 in  
err-disable state
```

Which action will solve the error state of this interface when connecting a host behind a Cisco IP phone?

- A. Configure dot1x-port control auto on this interface
- B. Enable errdisable recovery for security violation errors
- C. Enable port security on this interface
- D. Configure multidomain authentication on this interface

Correct Answer: D

Section: Infrastructure Security
Explanation

Explanation/Reference:

Explanation:

In single-host mode, a security violation is triggered when more than one device are detected on the data vlan. In multidomain authentication mode, a security violation is triggered when more than one device are detected on the data or voice VLAN. Here we see that single host mode is being used, not multidomain mode.

Reference. <http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/50sg/configuration/guide/Wrapper-46SG/dot1x.html#wp1309041>

QUESTION 377

What is the goal of Unicast Reverse Path Forwarding?

- A. to verify the reachability of the destination address in forwarded packets
- B. to help control network congestion
- C. to verify the reachability of the destination address in multicast packets

D. to verify the reachability of the source address in forwarded packets

Correct Answer: D

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Network administrators can use Unicast Reverse Path Forwarding (Unicast RPF) to help limit the malicious traffic on an enterprise network. This security feature works by enabling a router to verify the reachability of the source address in packets being forwarded. This capability can limit the appearance of spoofed addresses on a network. If the source IP address is not valid, the packet is discarded.

Reference. <http://www.cisco.com/web/about/security/intelligence/unicast-rpf.html>

QUESTION 378

Which three features are considered part of the IPv6 first-hop security suite? (Choose three.)

- A. DNS guard
- B. destination guard
- C. DHCP guard
- D. ICMP guard
- E. RA guard
- F. DoS guard

Correct Answer: BCE

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Cisco IOS has (at least) these IPv6 first-hop security features:

IPv6 RA Guard rejects fake RA messages coming from host (non-router) ports (not sure whether it handles all possible IPv6 header fragmentation attacks). Interestingly, it can also validate the contents of RA messages (configuration flags, list of prefixes) received through router-facing ports, potentially giving you a safeguard against an attack of fat fingers.

DHCPv6 Guard blocks DHCPv6 messages coming from unauthorized DHCPv6 servers and relays. Like IPv6 RA Guard it also validates the DHCPv6 replies coming from authorized DHCPv6 servers, potentially providing protection against DHCPv6 server misconfiguration.

IPv6 Snooping and device tracking builds a IPv6 First-Hop Security Binding Table (nicer name for ND table) by monitoring DHCPv6 and ND messages as well as regular IPv6 traffic. The binding table can be used to stop ND spoofing (in IPv4 world we'd call this feature DHCP Snooping and Dynamic ARP Inspection).

IPv6 Source Guard uses the IPv6 First-Hop Security Binding Table to drop traffic from unknown sources or bogus IPv6 addresses not in the binding table. The switch also tries to recover from lost address information, querying DHCPv6 server or using IPv6 neighbor discovery to verify the source IPv6

address after dropping the offending packet(s).

IPv6 Prefix Guard denies illegal off-subnet traffic. It uses information gleaned from RA messages and IA_PD option of DHCPv6 replies (delegated prefixes) to build the table of valid prefixes.

IPv6 Destination Guard drops IPv6 traffic sent to directly connected destination addresses not in IPv6 First-Hop Security Binding Table, effectively stopping ND exhaustion attacks.

Reference. <http://blog.ipspace.net/2013/07/first-hop-ipv6-security-features-in.html>

QUESTION 379

Refer to the exhibit.

```
interface GigabitEthernet0
 ip vrf forwarding Mgmt-intf
 ip address 1.1.1.1 255.255.255.0

ip access-list extended telnet-acl
 permit tcp any 1.1.1.1 0.0.0.0 eq 23 log

line vty 0 4
 access-class telnet-acl in
 transport input telnet
```

Why is the router not accessible via Telnet on the GigabitEthernet0 management interface?

- A. The wrong port is being used in the telnet-acl access list.
- B. The subnet mask is incorrect in the telnet-acl access list.
- C. The log keyword needs to be removed from the telnet-acl access list.
- D. The access class needs to have the vrf-also keyword added.

Correct Answer: D

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

The correct command should be "access-class telnet-acl in vrf-also". If you do not specify the vrf-also keyword, incoming Telnet connections from interfaces that are part of a VRF are rejected.

QUESTION 380

Which two features does the show ipv6 snooping features command show information about? (Choose two.)

- A. RA guard
- B. DHCP guard
- C. ND inspection
- D. source guard

Correct Answer: AC

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

The show ipv6 snooping features command displays the first-hop features that are configured on the router.

Examples

The following example shows that both IPv6 NDP inspection and IPv6 RA guard are configured on the router:

```
Router# show ipv6 snooping features
```

```
Feature name priority state
```

```
RA guard 100 READY
```

```
NDP inspection 20 READY
```

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/command/ipv6-cr-book/ipv6-s5.html>

QUESTION 381

Refer to the exhibit.

```
R1#sh policy-map control-plane
Service-policy input: CoPP-POLICY

Class-map: CoPP-CLASS (match-all)
  8 packets, 480 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: access-group name R9-T0-R2
  police:
    rate 10 pps, burst 0 packets
    conformed 0 packets; actions:
      drop
    exceeded 8 packets; actions:
      drop
    conformed 0 pps, exceed 0 pps

Class-map: class-default (match-any)
  929 packets, 86395 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: any

R1#sh access-lists
Extended IP access list R9-T0-R2
  10 permit tcp host 10.1.1.9 host 10.10.10.1 eq telnet (4 matches)
  20 deny tcp any any eq telnet (9 matches)
```

Which two statements about how the configuration processes Telnet traffic are true? (Choose two.)

- A. Telnet traffic from 10.1.1.9 to 10.10.10.1 is dropped.
- B. All Telnet traffic is dropped.
- C. Telnet traffic from 10.10.10.1 to 10.1.1.9 is permitted.
- D. Telnet traffic from 10.1.1.9 to 10.10.10.1 is permitted.
- E. Telnet traffic is permitted to all IP addresses.

Correct Answer: AC

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

The ACL applied to the COPP policy matches only telnet traffic from 10.1.1.9 to 10.10.10.1, all other telnet traffic is not matched and therefore not used in the COPP policy, which means this traffic will be handled normally (accepted). For telnet traffic from 10.1.1.9 to 10.10.10.1, the COPP policy has defined this traffic as an exceed, and dropped.

QUESTION 382

Which two statements about port ACLs are true? (Choose two.)

- A. Port ACLs are supported on physical interfaces and are configured on a Layer 2 interface on a switch.
- B. Port ACLs support both outbound and inbound traffic filtering.
- C. When it is applied to trunk ports, the port ACL filters only native VLAN traffic.
- D. When it is applied to a port with voice VLAN, the port ACL filters both voice and data VLAN traffic.

Correct Answer: AD

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

PACLs filter incoming traffic on Layer 2 interfaces, using Layer 3 information, Layer 4 header information, or non-IP Layer 2 information. The port ACL (PACL) feature provides the ability to perform access control on specific Layer 2 ports. A Layer 2 port is a physical LAN or trunk port that belongs to a VLAN. Port ACLs perform access control on all traffic entering the specified Layer 2 port, including voice and data VLANs that may be configured on the port. Port ACLs are applied only on the ingress traffic.

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

QUESTION 383

Which two statements about private VLANs are true? (Choose two.)

- A. Only one isolated VLAN can be mapped to a primary VLAN.
- B. Only one community VLAN can be mapped to a primary VLAN.
- C. Multiple isolated VLANs can be mapped to a primary VLAN.
- D. Multiple community VLANs can be mapped to a primary VLAN.

Correct Answer: AD

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

An isolated VLAN is a secondary VLAN that carries unidirectional traffic upstream from the hosts toward the promiscuous ports. You can configure only

one isolated VLAN in a PVLAN domain. An isolated VLAN can have several isolated ports. The traffic from each isolated port also remains completely separate. Only one isolated VLAN can be mapped under a given primary VLAN. A community VLAN is a secondary VLAN that carries upstream traffic from the community ports to the promiscuous port and to other host ports in the same community. You can configure multiple community VLANs in a PVLAN domain. The ports within one community can communicate, but these ports cannot communicate with ports in any other community or isolated VLAN in the private VLAN.

Reference.

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus6000/sw/layer2/6x/b_6k_Layer2_C onfig_6x/b_6k_Layer2_Config_602N12_chapter_011.html

QUESTION 384

Refer to the exhibit.

```
aaa new-model
aaa authentication login default local
username cisco privilege 15 password cisco

User Access Verification

Username: cisco
Password:

Router>en
% Error in authentication.

Router>
```

While configuring AAA with a local database, users can log in via Telnet, but receive the message "error in authentication" when they try to go into enable mode. Which action can solve this problem?

- A. Configure authorization to allow the enable command.
- B. Use aaa authentication login default enable to allow authentication when using the enable command.
- C. Verify whether an enable password has been configured.
- D. Use aaa authentication enable default enable to allow authentication when using the enable command.

Correct Answer: C

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

If a different enable password is configured, it will override the privilege level 15 of that user and force the existing password to be used for enable access.

QUESTION 385

DRAG DROP

Select and Place:

Encrypts the entire session	RADIUS
Uses less memory and CPU on a router	
Combines authentication and authorization	
Can limit router commands based on user groups	TACACS+

Correct Answer:

	RADIUS
	Uses less memory and CPU on a router
	Combines authentication and authorization
	TACACS+
	Encrypts the entire session
	Can limit router commands based on user groups

Section: Infrastructure Security
Explanation

Explanation/Reference:

QUESTION 386

DRAG DROP

Drag and drop the events on the left to display the correct sequence on the right when CoPP is enabled.

Select and Place:

Drag and drop the events on the left to display the correct sequence on the right when CoPP is enabled.

The packet gets forwarded to the switch CPU.	1
A packet enters the switch that is configured with CoPP on the ingress port.	2
The switch makes a routing or a switching decision, which determines whether or not the packet is destined for the control plane.	3
The port performs any applicable input port and QoS services.	4
Packets that are destined for the control plane are processed by CoPP and are dropped or delivered to the control plane according to each traffic class policy. Packets that have other destinations are forwarded normally.	5

Correct Answer:

Drag and drop the events on the left to display the correct sequence on the right when CoPP is enabled.

	A packet enters the switch that is configured with CoPP on the ingress port.
	The port performs any applicable input port and QoS services.
	The packet gets forwarded to the switch CPU.
	The switch makes a routing or a switching decision, which determines whether or not the packet is destined for the control plane.
	Packets that are destined for the control plane are processed by CoPP and are dropped or delivered to the control plane according to each traffic class policy. Packets that are destined for other destinations are forwarded normally.

Section: Infrastructure Security
Explanation

Explanation/Reference:

QUESTION 387

Which three condition types can be monitored by crypto conditional debug? (Choose three.)

- A. Peer hostname
- B. SSL
- C. ISAKMP
- D. Flow ID

- E. IPsec
- F. Connection ID

Correct Answer: ADF

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Supported Condition Types

The new crypto conditional debug CLIs--**debug crypto condition**, **debug crypto condition unmatched**, and **show crypto debug-condition**--allow you to specify conditions (filter values) in which to generate and display debug messages related only to the specified conditions. The table below lists the supported condition types.

Table 1 Supported Condition Types for Crypto Debug CLI

Condition Type (Keyword)	Description
<u>1</u> connid	An integer between 1-32766. Relevant debug messages will be shown if the current IPSec operation uses this value as the connection ID to interface with the crypto engine.
flowid 1	An integer between 1-32766. Relevant debug messages will be shown if the current IPSec operation uses this value as the flow-ID to interface with the crypto engine.
FVRF	The name string of a virtual private network (VPN) routing and forwarding (VRF) instance. Relevant debug messages will be shown if the current IPSec operation uses this VRF instance as its front-door VRF (FVRF).
IVRF	The name string of a VRF instance. Relevant debug messages will be shown if the current IPSec operation uses this VRF instance as its inside VRF (IVRF).
peer group	A Unity group-name string. Relevant debug messages will be shown if the peer is using this group name as its identity.
peer hostname	A fully qualified domain name (FQDN) string. Relevant debug messages will be shown if the peer is using this string as its identity; for example, if the peer is enabling IKE Xauth with this FQDN string.
peeripaddress	A single IP address. Relevant debug messages will be shown if the current IPSec operation is related to the IP address of this peer.
peer subnet	A subnet and a subnet mask that specify a range of peer IP addresses. Relevant debug messages will be shown if the IP address of the current IPSec peer falls into the specified subnet range.
peer username	A username string. Relevant debug messages will be shown if the peer is using this username as its identity; for example, if the peer is enabling IKE Extended Authentication (Xauth) with this username.
SPI 1	A 32-bit unsigned integer. Relevant debug messages will be shown if the current IPSec operation uses this value as the SPI.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_vpnips/configuration/xr-3s/sec-sec-for-vpns-w-ipsec-xr-3s-book/sec-crypto-debug-sup.html

QUESTION 388

What is the ip dhcp snooping information option command used for?

- A. It displays information about the DHCP snooping table.
- B. It sends a syslog and an SNMP trap for a DHCP snooping violation.
- C. It enables the DHCP snooping host tracking feature.
- D. It enables DHCP option 82 data insertion.

Correct Answer: D

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

To enable DHCP option-82 data insertion, perform this task:

	Command	Purpose
Step 1	Router(config)# ip dhcp snooping information option	Enables DHCP option- 82 data insertion.
Step 2	Router(config)# ip dhcp snooping information option replace Or: Router(config-if)# ip dhcp snooping information option replace	(Optional) Replaces the DHCP relay information option received in snooped packets with the switch's option-82 data.
Step 3	Router(config)# do show ip dhcp snooping include 82	Verifies

Reference. <http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SX/configuration/guide/book/snoodhcp.html>

QUESTION 389

Which two statements are true about unicast RPF? (Choose two.)

- A. Unicast RPF requires CEF to be enabled.
- B. Unicast RPF strict mode works better with multihomed networks.
- C. Unicast RPF strict mode supports symmetric paths.
- D. Unicast RPF strict mode supports asymmetric paths.
- E. CEF is optional with Unicast RPF, but when CEF is enabled it provides better performance.

Correct Answer: AC

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Unicast RPF requires Cisco express forwarding (CEF) to function properly on the router.

Strict Versus Loose Checking Mode

The Unicast RPF in Strict Mode feature filters ingress IPv4 traffic in strict checking mode and forwards packets only if the following conditions are satisfied.

- · An IPv4 packet must be received at an interface with the best return path (route) to the packet source (a process called symmetric routing). There must be a route in the Forwarding Information Base (FIB) that matches the route to the receiving interface. Adding a route in the FIB can be done via static route, network statement, or dynamic routing.
- · IPv4 source addresses at the receiving interface must match the routing entry for the interface.

References:

http://www.cisco.com/c/en/us/td/docs/ios/12_2/security/configuration/guide/fsecur_c/scfprpf.html http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/srpf_gsr.html

QUESTION 390

Under Cisco IOS Software, which two features are supported in RADIUS Change of Authorization requests? (Choose two.)

- A. session identification
- B. session reauthentication
- C. session termination
- D. host termination

Correct Answer: AC

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

CoA requests, as described in RFC 5176, are used in a pushed model to allow for session identification, host reauthentication, and session termination. The model comprises one request (CoA-Request) and two possible response codes.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_aaa/configuration/15-sy/sec_usr_aaa-15-sy-book/sec-rad-coa.html

QUESTION 391

DRAG DROP

Select and Place:

Uses UDP	TACACS+
Separates authentication, authorization, and accountability	
Is proprietary to Cisco	
Encrypts only the password	RADIUS

Correct Answer:

	TACACS+
	Is proprietary to Cisco
	Separates authentication, authorization, and accountability
	RADIUS
	Uses UDP
	Encrypts only the password

Section: Infrastructure Security**Explanation****Explanation/Reference:****QUESTION 392**

Which technology can be used to secure the core of an STP domain?

- A. UplinkFast
- B. BPDU guard
- C. BPDU filter
- D. root guard

Correct Answer: D

Section: Infrastructure Security**Explanation****Explanation/Reference:****Explanation:**

Since STP does not implement any authentication or encryption to protect the exchange of BPDUs, it is vulnerable to unauthorized participation and attacks. Cisco IOS offers the STP Root Guard feature to enforce the placement of the root bridge and secure the core of the STP domain. STP root guard forces a port to become a designated port so that no switch on the other end of the link can become a root switch. If a port configured for root

guard receives a superior BPDU, the port it is received on is blocked. In this way, STP root guard blocks other devices from trying to become the root bridge.

STP root guard should be enabled on all ports that will never connect to a root bridge, for example, all end user ports. This ensures that a root bridge will never be negotiated on those ports.

Reference.

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Security/Baseline_Security/securebasebook/sec_chap7.html

QUESTION 393

Which two protocols are not protected in an edge router by using control plane policing? (Choose two.)

- A. SMTP
- B. RPC
- C. SSH
- D. Telnet

Correct Answer: AB

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

A CoPP policy can limit a number of different packet types that are forwarded to the control plane.

Traffic destined for the switch CPU includes:

- . Address Resolution Protocol (ARP)
- . First-hop redundancy protocol packets
- . Layer 2 control packets
- . **Management packets (telnet, Secure Shell [SSH] Protocol, Simple Network Management Protocol [SNMP])** <--- C and D are not correct.
- . Multicast control packets
- . Routing protocol packets
- . Packets with IP options
- . Packets with time to live (TTL) set to 1
- . Packets that require ACL logging
- . Packets that require an initial lookup (first packet in a flow: FIB miss) · Packets that have don't support hardware switching/routing

Reference. http://www.cisco.com/c/en/us/products/collateral/switches/catalyst-6500-series-switches/white_paper_c11_553261.html

QUESTION 394

Which two statements are true about AAA? (Choose two.)

- A. AAA can use RADIUS, TACACS+, or Windows AD to authenticate users.
- B. If RADIUS is the only method configured in AAA, and the server becomes unreachable, the user will be able to log in to the router using a local

username and password.

- C. If the local keyword is not included and the AAA server does not respond, then authorization will never be possible and the connection will fail.
- D. AAA can be used to authenticate the enable password with a AAA server.

Correct Answer: CD

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

AAA can be used to authenticate user login and the enable passwords.

Example 1: Same Exec Authentication Methods for All Users

Once authenticated with:

aaa authentication login default group radius local

All users who want to log in to the access server have to be authorized using Radius (first method) or local database (second method).

We configure:

aaa authorization exec default group radius local

Note. On the AAA server, Service-Type=1 (login) must be selected.

Note. With this example, if the local keyword is not included and the AAA server does not respond, then authorization will never be possible and the connection will fail.

Reference. <http://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/10384-security.html>

QUESTION 395

Which three types of traffic are allowed by IEEE 802.1X access control prior to getting authenticated? (Choose three.)

- A. EAPOL
- B. VTP
- C. STP
- D. ARP
- E. CDP
- F. HTTP

Correct Answer: ACE

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Until the client is authenticated, IEEE 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL), Cisco Discovery Protocol (CDP), and Spanning Tree Protocol (STP) traffic through the port to which the client is connected. After authentication, normal traffic passes through the

port.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750/software/release/15-0_2_se/configuration/guide/scg3750/sw8021x.pdf

QUESTION 396

Which two statements about MAC ACLs are true? (Choose two.)

- A. They support only inbound filtering.
- B. They support both inbound and outbound filtering.
- C. They are configured with the command `mac access-list standard`.
- D. They can filter non-IP traffic on a VLAN and on a physical interface.

Correct Answer: AD

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

MAC ACL, also known as Ethernet ACL, can filter non-IP traffic on a VLAN and on a physical Layer 2 interface by using MAC addresses in a named MAC extended ACL. The steps to configure a MAC ACL are similar to those of extended named ACLs. MAC ACL supports only inbound traffic filtering.

Reference. <http://www.ciscopress.com/articles/article.asp?p=1181682&seqNum=4>

QUESTION 397

Refer to the exhibit.

```
Router#sh policy-map control-plane
Control Plane

Service-policy output: control-plane-out

Class-map: icmp-class (match-all)
 197314985 packets, 11510114428 bytes
 5 minute offered rate 1000 bps, drop rate 0000 bps
 Match: access-group name killicmpv2
 police:
   cir 1000000 bps, bc 31250 bytes
   conformed 197138885 packets, 11499818077 bytes; actions:
     transmit
   exceeded 176100 packets, 10296351 bytes; actions:
     drop
   conformed 1000 bps, exceed 0000 bps

Class-map: class-default (match-any)
 1126224901 packets, 158790413979 bytes
 5 minute offered rate 41000 bps, drop rate 0000 bps
 Match: any
```

What happens to packets when traffic in the icmp-class class exceeds the policed amount?

- A. Packets are discarded and a message is logged.
- B. Packets are discarded and a trap is sent to any servers that are configured to receive traps.
- C. Packets are discarded silently.
- D. Packets are discarded and an inform is sent to any servers that are configured to receive informs.

Correct Answer: C

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

service-policy {input output policy-map-name Example: Router(config-cp)# service-policy input control-plane-policy	Attaches a QoS service policy to the control plane. Note the following points: <ul style="list-style-type: none">• input --Applies the specified service policy to packets received on the control plane.• output --Applies the specified service policy to packets transmitted from the control plane and enables the router to silently discard packets.• policy-map-name --Name of a service policy map (created using the policy-map command) to be attached. The name can be a maximum of 40 alphanumeric characters.
---	--

In this case, the service policy is set to output, which drops the traffic silently per above.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_plcshp/configuration/15-mt/qos-plcshp-15-mt-book/qos-plcshp-ctrl-pln-plc.html

QUESTION 398

Which option is the Cisco recommended method to secure access to the console port?

- A. Configure the activation-character command.
- B. Configure a very short timeout (less than 100 milliseconds) for the port.
- C. Set the privilege level to a value less than 15.
- D. Configure an ACL.

Correct Answer: A

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

The activation-character command defines a session activation character. Entering this character at a vacant terminal begins a terminal session. The default activation character is the Return key.

To secure the console port, you should change this character to a different one as most people simply hit the enter key when trying to access the console.

QUESTION 399

DRAG DROP

Drag and drop each SNMP security model and level on the left to the corresponding mode of authentication on the right.

Select and Place:

Drag and drop each SNMP security model and level on the left to the corresponding mode of authentication on the right.	
SNMPv2c-noAuthNoPriv	provides HMAC-MD5 or HMAC-SHA authentication with DES 56-bit encryption
SNMPv3-authNoPriv	authenticates with a user name match
SNMPv3-authPriv	provides HMAC-MD5 or HMAC-SHA authentication without encryption
SNMPv3-noAuthNoPriv	authenticates with a community string match

Correct Answer:

Drag and drop each SNMP security model and level on the left to the corresponding mode of authentication on the right.

	SNMPv3-authPriv
	SNMPv3-noAuthNoPriv
	SNMPv3-authNoPriv
	SNMPv2c-noAuthNoPriv

Section: Infrastructure Security

Explanation

Explanation/Reference:

QUESTION 400

Which two Cisco IOS AAA features are available with the local database? (Choose two.)

- A. command authorization
- B. network access authorization
- C. network accounting
- D. network access authentication

Correct Answer: AD

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Configuring the Local Database

This section describes how to manage users in the local database. You can use the local database for CLI access authentication, privileged mode authentication, command authorization, network access authentication, and VPN authentication and authorization. You cannot use the local database for

network access authorization. The local database does not support accounting.

Reference.

http://www.cisco.com/c/en/us/td/docs/security/asa/asa80/configuration/guide/conf_gd/aaa.html

QUESTION 401

Which two features are used for inspection when IPv6 address glean is enabled? (Choose two.)

- A. DHCP messages
- B. ND messages
- C. ICMPv6 messages
- D. UDP messages
- E. TCP messages

Correct Answer: AB

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

IPv6 address glean is the foundation for many other IPv6 features that depend on an accurate binding table. It inspects ND and DHCP messages on a link to glean addresses, and then populates the binding table with these addresses. This feature also enforces address ownership and limits the number of addresses any given node is allowed to claim.

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/15-s/ipv6f-15-s-book/ip6-snooping.html

QUESTION 402

Which two statements about the protected ports feature and the private VLAN feature are true? (Choose two.)

- A. The protected ports feature is limited to the local switch.
- B. The protected ports feature can isolate traffic between two "protected" ports on different switches.
- C. The private VLAN feature is limited to the local switch.
- D. The private VLAN feature prevents interhost communication within a VLAN across one or more switches.

Correct Answer: AD

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Protected Ports (PVLAN Edge)

In some network environments, there is a requirement for no traffic to be seen or forwarded between host(s) on the same LAN segment, thereby

preventing interhost communications. The PVLAN edge feature provisions this isolation by creating a firewall-like barrier, thereby blocking any unicast, broadcast, or multicast traffic among the protected ports on the switch. Note that the significance of the protected port feature is limited to the local switch, and there is no provision in the PVLAN edge feature to isolate traffic between two "protected" ports located on different switches. For this purpose, the PVLAN feature can be used.

Reference. <http://www.ciscopress.com/articles/article.asp?p=1181682&seqNum=2>

QUESTION 403

DRAG DROP

Drag and drop the TACACS+ configuration command on the left to the correct function it performs on the right.

Select and Place:

tacacs-server host 192.168.1.250	globally configures a pre-shared TACACS+ key
tacacs-server host 10.1.1.93 key CISCO	configures a device to send only a portion of the username
tacacs-server key CISCO	configures the device to send TACACS+ requests to a
tacacs-server directed-request	maintains a single open connection between the device and
tacacs-server packet 12000	configures the device to securely send TACACS+ requests to a TACACS+ server
tacacs-server host 172.16.16.25 single-connection	configured the maximum TACACS+ packet size

Correct Answer:

tacacs-server key CISCO

tacacs-server directed-request

tacacs-server host 192.168.1.250

tacacs-server host 172.16.16.25 single-connection

tacacs-server host 10.1.1.93 key CISCO

tacacs-server packet 12000

Section: Infrastructure Security
Explanation

Explanation/Reference:

QUESTION 404
Refer to the exhibit.

```
username admin privilege 15 password SECUREPASSWORD
aaa new-model
aaa authentication login default group tacacs+ local-case
aaa authentication enable default group tacacs+ enable
aaa authorization console
aaa authorization exec default group tacacs+ if-authenticated
aaa authorization commands 1 default group tacacs+ local if-authenticated
aaa authorization commands 4 default group tacacs+ if-authenticated
aaa authorization commands 15 default group tacacs+ local if-authenticated
aaa accounting exec default start-stop group tacacs+
aaa accounting commands 1 default start-stop group tacacs+
aaa accounting commands 15 default start-stop group tacacs+
aaa session-id common
```

Which two configuration changes enable the user admin to log in to the device? (Choose two.)

- A. Configure the login authentication to be case-insensitive.
- B. Configure the user admin with a password and appropriate privileges.
- C. Configure the login authentication to be case-sensitive.
- D. Modify the configuration to use a named group.
- E. Configure additional login authentication under the terminal lines.

Correct Answer: AB

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Usernames and passwords are case-sensitive. Users attempting to log in with an incorrectly cased username or password will be rejected. If users are unable to log into the router with their specific passwords, reconfigure the username and password on the router.

Reference: <http://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/ios-software-releases-110/45843-configpasswords.html>

QUESTION 405

Which two advantages does CoPP have over receive path ACLs? (Choose two.)

- A. Only CoPP applies to IP packets and non-IP packets.
- B. Only CoPP applies to receive destination IP packets.
- C. A single instance of CoPP can be applied to all packets to the router, while rACLs require multiple instances.

D. Only CoPP can rate-limit packets.

Correct Answer: AD

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Control Plane Policing CoPP is the Cisco IOS-wide route processor protection mechanism. As illustrated in Figure 2, and similar to rACLs, CoPP is deployed once to the punt path of the router. However, unlike rACLs that only apply to receive destination IP packets, CoPP applies to all packets that punt to the route processor for handling. CoPP therefore covers not only receive destination IP packets, it also exceptions IP packets and non-IP packets. In addition, CoPP is implemented using the Modular QoS CLI (MQC) framework for policy construction. In this way, in addition to simply permit and deny functions, specific packets may be permitted but rate-limited. This behavior substantially improves the ability to define an effective CoPP policy. (Note: that "Control Plane Policing" is something of a misnomer because CoPP generally protects the punt path to the route processor and not solely the control plane.)

Reference: http://www.cisco.com/web/about/security/intelligence/coppwp_gs.html

QUESTION 406

Which command drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses to drop below the maximum value, and also causes the Security Violation counter to increment?

- A. switchport port-security violation protect
- B. switchport port-security violation drop
- C. switchport port-security violation shutdown
- D. switchport port-security violation restrict

Correct Answer: D

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

When configuring port security violation modes, note the following information:

- **protect**--Drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses to drop below the maximum value.
- **restrict**--Drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses to drop below the maximum value and causes the SecurityViolation counter to increment.
- **shutdown**--Puts the interface into the error-disabled state immediately and sends an SNMP trap notification.

Reference: http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SX/configuration/guide/book/port_sec.html

QUESTION 407

Refer to the exhibit.

```
interface GigabitEthernet0/0/0
  ip address 192.168.1.1 255.255.255.0
  !
  ip ssh version 2
  !
  ip access-list extended protect-ssh
    permit ip any any eq 22
  !
  line vty 0 4
    access-class protect-ssh in
    transport input ssh
```

Which configuration is missing that would enable SSH access on a router that is running Cisco IOS XE Software?

- A. int Gig0/0/0
management-interface
- B. class-map ssh-class
match access-group protect-ssh
policy-map control-plane-in
class ssh-class
police 80000 conform transmit exceed drop
control-plane
service-policy input control-plane-in
- C. control-plane host
management-interface GigabitEthernet0/0/0 allow ssh
- D. interface Gig0/0/0
ip access-group protect-ssh in

Correct Answer: C

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The feature Management Plane Protection (MPP) allows an administrator to restrict on which interfaces management traffic can be received by a device. This allows the administrator additional control over a device and how the device is accessed.

This example shows how to enable the MPP in order to only allow SSH and HTTPS on the GigabitEthernet0/1 interface:
!

```
control-plane host
management-interface GigabitEthernet 0/1 allow ssh https
!
```

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

QUESTION 408

Which three modes are valid PfR monitoring modes of operation? (Choose three.)

- A. route monitor mode (based on BGP route changes)
- B. RMON mode (based on RMONv1 and RMONv2 data)
- C. passive mode (based on NetFlow data)
- D. active mode (based on Cisco IP SLA probes)
- E. fast mode (based on Cisco IP SLA probes)
- F. passive mode (based on Cisco IP SLA probes)

Correct Answer: CDE

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Modes are:

Mode monitor passive

Passive monitoring is the act of PfR gathering information on user packets assembled into flows by Netflow. Passive monitoring is typically only recommended in Internet edge deployments because active probing is ineffective because of security policies that block probing. PfR, when enabled, automatically enables Netflow on the managed interfaces on the Border Routers. By aggregating this information on the Border Routers and periodically reporting the collected data to the Master Controller, the network prefixes and applications in use can automatically be learned.

Mode monitor active

Active monitoring is the act of generating Cisco IOS IP Service Level Agreements (SLAs) probes to generate test traffic for the purpose of obtaining information regarding the characteristics of the WAN links. PfR can either implicitly generate active probes when passive monitoring has identified destination hosts, or the network manager can explicitly configure probes in the PfR configuration. When jitter probes are used (common use case), Target Discovery is used to learn the respond address and to automatically generate the probes.

Mode monitor Fast

This mode generates active probes through all exists continuously at the configured probe frequency. This differs from either active or both modes in that these modes only generate probes through alternate paths (exits) in the event the current path is out-of-policy.

Reference. http://docwiki.cisco.com/wiki/PfR:Technology_Overview#Mode_monitor_passive

QUESTION 409

Refer to the exhibit.

```
MC#sh pfr master border detail
```

Border	Status	UP/DOWN	AuthFail	Version
10.1.1.1	ACTIVE	UP	00:52:21	0 3.0
Et0/0	INTERNAL	UP		
Et0/1	EXTERNAL	UP		

External Interface	Capacity (kbps)	Max BW (kbps)	BW Used (kbps)	Load (%)	Status	Exit Id
Et0/1	Tx 500	450	192	39	UP	2
	Rx	500	49	9		

```
MC#sh pfr master border detail
```

Border	Status	UP/DOWN	AuthFail	Version
10.1.1.2	ACTIVE	UP	00:52:21	0 3.0
Et0/0	INTERNAL	UP		
Et0/1	EXTERNAL	UP		

External Interface	Capacity (kbps)	Max BW (kbps)	BW Used (kbps)	Load (%)	Status	Exit Id
Et0/1	Tx 500	450	175	33	UP	1
	Rx	500	0	0		

Which statement is true?

- A. The Cisco PfR state is UP; however, the external interface Et0/1 of border router 10.1.1.1 has exceeded the maximum available bandwidth threshold.
- B. The Cisco PfR state is UP; however, an issue is preventing the border router from establishing a TCP session to the master controller.
- C. The Cisco PfR state is UP and is able to monitor traffic flows; however, MD5 authentication has not been successful between the master controller and the border routers.
- D. The Cisco PfR State is UP; however, the receive capacity was not configured for inbound traffic.

E. The Cisco PIR state is UP, and the link utilization out-of-policy threshold is set to 90 percent for traffic exiting the external links.

Correct Answer: E

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

All three interfaces show as UP, and the capacity is set to 500 kbps, with the max threshold set to 450 kbps (90% of 500kbps).

QUESTION 410

In the DiffServ model, which class represents the highest priority with the highest drop probability?

- A. AF11
- B. AF13
- C. AF41
- D. AF43

Correct Answer: D

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

AF43-- Assured forwarding, high drop probability, Class 4 DSCP, and Flash-override precedence.

Table of AF Classes and Drop Priority

Drop Precedence	Class 1	Class 2	Class 3	Class 4
Low drop	AF11	AF21	AF31	AF41
	DSCP 10	DSCP 18	DSCP 26	DSCP 34
	001010	010010	011010	100010
Medium drop	AF12	AF22	AF32	AF42
	DSCP 12	DSCP 20	DSCP 28	DSCP 36
	001100	010100	011100	100100
High drop	AF13	AF23	AF33	AF43
	DSCP 14	DSCP 22	DSCP 30	DSCP 38
	001110	010110	011110	100110

Reference. [https://www.informit.com/library/content.aspx?](https://www.informit.com/library/content.aspx?b=CCIE_Practical_Studies_II&seqNum=56)

[b=CCIE_Practical_Studies_II&seqNum=56](https://www.informit.com/library/content.aspx?b=CCIE_Practical_Studies_II&seqNum=56)

QUESTION 411

Refer to the exhibit.

```
Entry number: 1
Owner:
Tag:
Type of operation to perform: echo
Target address/Source address: 172.16.129.9/0.0.0.0
Type of Service parameter: 0x0
Request size (ARR data portion): 28
Operation timeout (milliseconds): 5000
Verify data: No
Vrf Name:
Schedule:
  Operation frequency (seconds): 10
  Next Scheduled Start Time: Pending trigger
  Group Scheduled : FALSE
  Randomly Scheduled : FALSE
  Life (seconds): 3600
  Entry Ageout (seconds): never
  Recurring (Starting Everyday): FALSE
  Status of entry (SNMP RowStatus): notInService
Threshold (milliseconds): 5000
Distribution Statistics:
  Number of statistic hours kept: 2
  Number of statistic distribution buckets kept: 1
  Statistic distribution interval (milliseconds): 20
History Statistics:
  Number of history Lives kept: 0
  Number of history Buckets kept: 15
  History Filter Type: None
Enhanced History:
```

Which statement about this IP SLA is true?

- A. The SLA must also have a schedule configured before it will start.
- B. The TTL of the SLA packets is 10.
- C. The SLA has a timeout of 3.6 seconds.
- D. The SLA has a lifetime of 5 seconds.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

When you configure an IP SLAs operation, you must schedule the operation to begin capturing statistics and collecting error information. You can schedule an operation to start immediately or to start at a certain month, day, and hour. You can use the pending option to set the operation to start at a later time. The pending option is an internal state of the operation that is visible through SNMP. The pending state is also used when an operation is a reaction (threshold) operation waiting to be triggered. You can schedule a single IP SLAs operation or a group of operations at one time. We can see in this output that the IP SLA is still in a pending trigger state.

Reference. <http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/44sg/configuration/guide/Wrapper-44SG/swipsla.html>

QUESTION 412

Which three actions are required when configuring NAT-PT? (Choose three.)

- A. Enable NAT-PT globally.
- B. Specify an IPv4-to-IPv6 translation.
- C. Specify an IPv6-to-IPv4 translation.
- D. Specify a ::/96 prefix that will map to an IPv4 address.
- E. Specify a ::/48 prefix that will map to a MAC address.
- F. Specify a ::/32 prefix that will map to an IPv6 address.

Correct Answer: BCD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The detailed steps on configuring NAT-PT is found at the reference link below:

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/ipv6/configuration/guide/12_4t/ipv6_12_4t_book/ip6-nat_trnsln.html

QUESTION 413

Refer to the exhibit.

```
Cos-dscp map:
cos:    0  1  2  3  4  5  6  7
-----
dscp:    0  8 16 46 36 38 42 32
```

Which statement about this COS-DSCP mapping is true?

- A. COS 3 is mapped to the expedited forwarding DSCP.
- B. COS 16 is mapped to DSCP 2.
- C. The default COS is mapped to DSCP 32.
- D. This mapping is the default COS-DSCP mapping on Cisco switches.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Here we see that COS 3 is mapped to DSCP 46, which is the Expedited forwarding class:

The Expedited Forwarding (EF) model is used to provide resources to latency (delay) sensitive real- time, interactive traffic. The EF model uses one marking -- DSCP 46.

QUESTION 414

Which three statements about implementing a NAT application layer gateway in a network are true? (Choose three.)

- A. It allows client applications to use dynamic ports to communicate with a server regardless of whether NAT is being used.
- B. It maintains granular security over application-specific data.
- C. It allows synchronization between multiple streams of data between two hosts.
- D. Application layer gateway is used only in VoIP/SIP deployments.
- E. Client applications require additional configuration to use an application layer gateway.
- F. An application layer gateway inspects only the first 64 bytes of a packet before forwarding it through the network.

Correct Answer: ABC

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

An application-level gateway (ALG), also known as an application-layer gateway, is an application that translates the IP address information inside the payload of an application packet. An ALG is used to interpret the application-layer protocol and perform firewall and Network Address Translation (NAT) actions. These actions can be one or more of the following depending on your configuration of the firewall and NAT:

- Allow client applications to use dynamic TCP or UDP ports to communicate with the server application.
- Recognize application-specific commands and offer granular security control over them. Synchronize multiple streams or sessions of data between two hosts that are exchanging data.
- Translate the network-layer address information that is available in the application payload

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_nat/configuration/xr-3s/asr1000/nat-xr-3s-asr1k-book/fw-msrpc-supp.html

QUESTION 415

Refer to the exhibit.

```
RouterA(config)#ip options drop
```

At which location will the benefit of this configuration be observed?

- A. on Router A and its upstream routers
- B. on Router A and its downstream routers
- C. on Router A only
- D. on Router A and all of its ARP neighbors

Correct Answer: B

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The following example shows how to configure the router (and downstream routers) to drop all options packets that enter the network:

```
Router(config)# ip options drop
```

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/sel_drop.html

QUESTION 416

Which statement about shaped round robin queuing is true?

- A. Queues with higher configured weights are serviced first.
- B. The device waits a period of time, set by the configured weight, before servicing the next queue.
- C. The device services a single queue completely before moving on to the next queue.

D. Shaped mode is available on both the ingress and egress queues.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

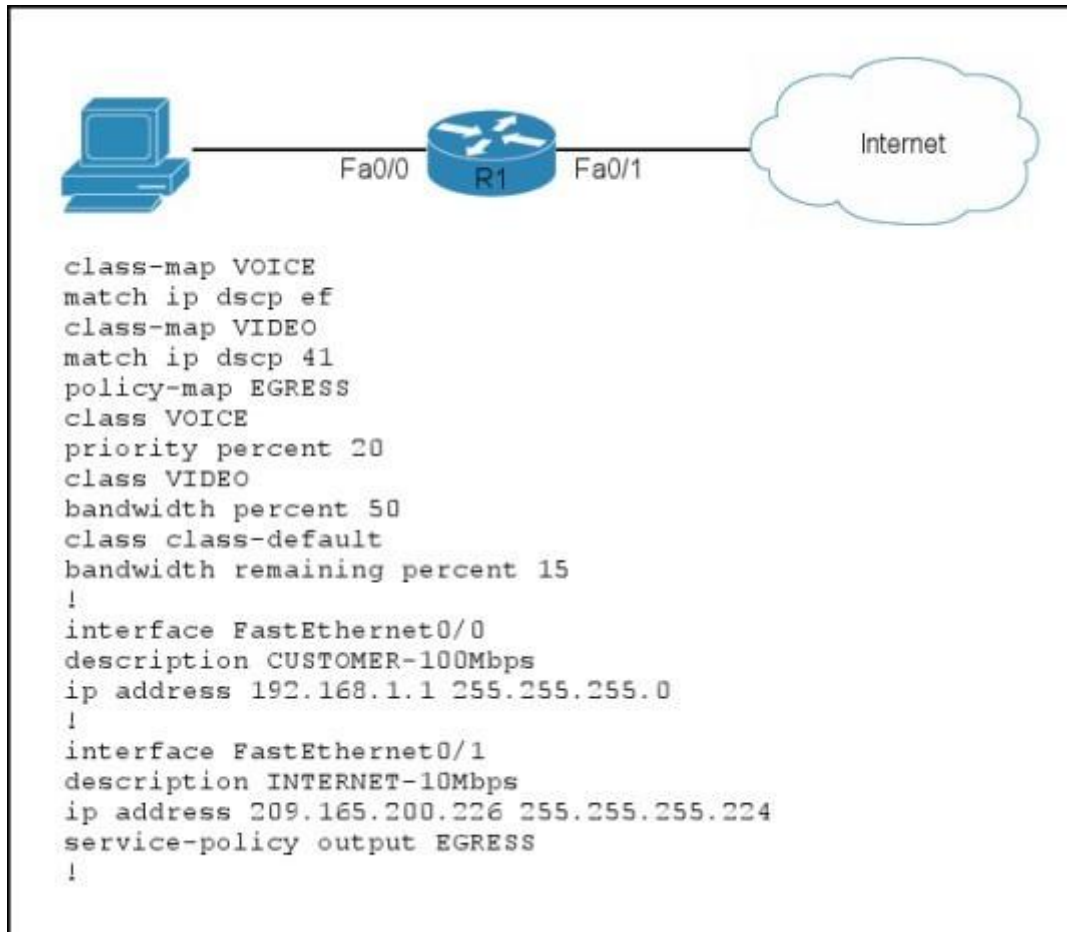
SRR is scheduling service for specifying the rate at which packets are dequeued. With SRR there are two modes, shaped and shared. Shaped mode is only available on the egress queues

SRR differs from typical WRR. With WRR queues are serviced based on the weight. Q1 is serviced for weight 1 period of time, Q2 is served for weight 2 period of time, and so forth. The servicing mechanism works by moving from queue to queue and services them for the weighted amount of time. With SRR weights are still followed; however, SRR services Q1, moves to Q2, then Q3 and Q4 in a different way. It does not wait at and service each queue for a weighted amount of time before moving on to the next queue. Instead, SRR makes several rapid passes at the queues; in each pass, each queue might or might not be serviced. For each given pass, the more highly weighted queues are more likely to be serviced than the lower priority queues.

Reference. http://www.cisco.com/c/en/us/products/collateral/switches/catalyst-3560-e-series-switches/prod_qas0900aecd805bacc7.html

QUESTION 417

Refer to the exhibit.



You discover that only 1.5 Mb/s of web traffic can pass during times of congestion on the given network.

Which two options are possible reasons for this limitation? (Choose two.)

- A. The web traffic class has too little bandwidth reservation.
- B. Video traffic is using too much bandwidth.
- C. The service-policy is on the wrong interface.
- D. The service-policy is going in the wrong direction.
- E. The NAT policy is adding too much overhead.

Correct Answer: AB

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

In this example, the web traffic will fall into the default class, which is only 15 percent of the 10Mbps Internet connection (1.5Mbps). Meanwhile, video traffic is allowed 50% of the 10 Mbps.

QUESTION 418

Refer to the exhibit.

```
snmp-server community public RO 2
snmp-server trap-source Loopback0
snmp-server chassis-id HONGKONG
snmp-server enable traps snmp linkdown linkup coldstart
snmp-server enable traps ospf state-change
snmp-server enable traps bgp state-changes
snmp-server enable traps pim neighbor-change
snmp-server enable traps cpu threshold
snmp-server enable traps mpls ldp
snmp-server host 192.168.252.254 version 2c public
```

Which statement about this device configuration is true?

- A. The NMS needs a specific route configured to enable it to reach the Loopback0 interface of the device.
- B. The ifindex of the device could be different when the device is reloaded.
- C. The device will allow anyone to poll it via the public community.
- D. The device configuration requires the AuthNoPriv security level.

Correct Answer: B

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

One of the most commonly used identifiers in SNMP-based network management applications is the Interface Index (ifIndex) value. IfIndex is a unique identifying number associated with a physical or logical interface. For most software, the ifIndex is the name of the interface. Although relevant RFCs do not require that the correspondence between particular ifIndex values and their interfaces be maintained across reboots, applications such as device

inventory, billing, and fault detection depend on this correspondence.

Consider a situation where a simple monitoring software (like MRTG) is polling the interface statistics of the router specific serial interface going to the internet.

As an example, you could have these conditions prior to re-initialization:

physical port	ifIndex
ethernet port	1
tokenring port	2
serial port	3

Therefore, the management application is polling the ifIndex 3, which corresponds to the serial port.

After the router re-initialization (reboot, reload and so on) the conditions change to something similar to this:

physical port	ifIndex
ethernet port	3
tokenring port	1
serial port	2

The management application continues polling the ifIndex 3, which corresponds now to the ethernet port. Therefore, if the management application is not warned by a trap, for example, that the router has been rebooted, the statistics polled could be completely wrong.

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/simple-network-management-protocol-snmpp/28420-ifIndex-Persistence.html>

QUESTION 419

Which three steps are necessary to enable SSH? (Choose three.)

- A. generating an RSA or DSA cryptographic key
- B. configuring the version of SSH
- C. configuring a domain name
- D. configuring VTY lines for use with SSH
- E. configuring the port for SSH to listen for connections
- F. generating an AES or SHA cryptographic key

Correct Answer: ACD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Here are the steps:

1. Configure a hostname for the router using these commands.

```
yourname#configure terminal
```

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

```
yourname (config)#hostname LabRouter
```

```
LabRouter(config)#
```

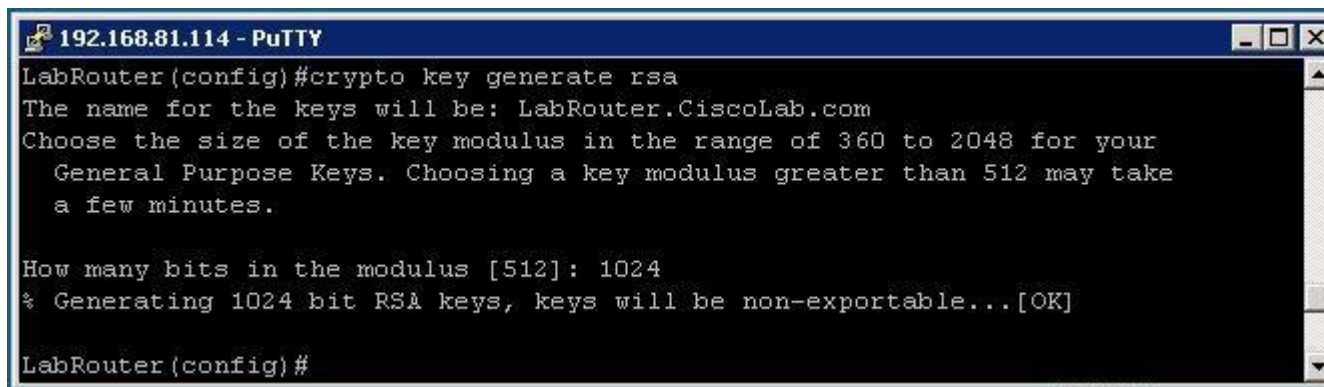
2. Configure a domain name with the **ip domain-name** command followed by whatever you would like your domain name to be. I used CiscoLab.com.

```
LabRouter(config)#ip domain-name CiscoLab.com
```

3. We generate a certificate that will be used to encrypt the SSH packets using the **crypto key generate rsa** command.

Take note of the message that is displayed right after we enter this command. "*The name for the keys will be: LabRouter.CiscoLab.com*" -- it combines the hostname of the router along with the domain name we configured to get the name of the encryption key generated; this is why it was important for us to, first of all, configure a hostname then a domain name before we generated the keys.

Notice also that it asks us to choose a size of modulus for the key we're about to generate. The higher the modulus, the stronger the encryption of the key. For our example, we'll use a modulus of 1024.



```
192.168.81.114 - PuTTY
LabRouter(config)#crypto key generate rsa
The name for the keys will be: LabRouter.CiscoLab.com
Choose the size of the key modulus in the range of 360 to 2048 for your
  General Purpose Keys. Choosing a key modulus greater than 512 may take
    a few minutes.

How many bits in the modulus [512]: 1024
% Generating 1024 bit RSA keys, keys will be non-exportable...[OK]

LabRouter (config) #
```

4. Now that we've generated the key, our next step would be to configure our vty lines for SSH access and specify which database we are going to use to provide authentication to the device. The local database on the router will do just fine for this example.

```
LabRouter(config)#line vty 0 4
```

```
LabRouter(config-line)#login local
```

```
LabRouter(config-line)#transport input ssh
```

5. You will need to create an account on the local router's database to be used for authenticating to the device. This can be accomplished with these commands.

```
LabRouter(config)#username XXXX privilege 15 secret XXXX
```

Reference. <http://blog.pluralsight.com/configure-secure-shell-ssh-on-cisco-router>

QUESTION 420

Refer to the exhibit.

```
event manager applet LARGECONFIG
  event cli pattern "show running-config" sync yes
  action 1.0 puts "Warning! This device has a VERY LARGE configuration
    and may take some time to process"
  action 1.1 puts newline "Do you wish to continue [Y/N]"
  action 1.2 gets response
  action 1.3 string toupper "$response"
  action 1.4 string match "$_string_result" "Y"
  action 2.0 if $_string_result eq 1
  action 2.1 cli command "enable"
  action 2.2 cli command "show running-config"
  action 2.3 puts $_cli_result
  action 2.4 cli command "exit"
  action 2.9 end
```

Which two statements about the EEM applet configuration are true? (Choose two.)

- A. The EEM applet runs before the CLI command is executed.
- B. The EEM applet runs after the CLI command is executed.
- C. The EEM applet requires a case-insensitive response.
- D. The running configuration is displayed only if the letter Y is entered at the CLI.

Correct Answer: AD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

sync	<p>Indicates whether the policy should be executed synchronously before the CLI command executes.</p> <ul style="list-style-type: none">• If the yes keyword is specified, the policy will run synchronously with the CLI command.• If the no keyword is specified, the policy will run asynchronously with the CLI command.
------	---

Here we see that the sync knob was enabled so A is correct. However, C is not correct as the nocase argument was not used, so the applet is configured to display the config only if a capital Y is issued.

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/eem/command/eem-cr-book/eem-cr-a2.html>

QUESTION 421

Which variable in an EEM applet is set when you use the sync yes option?

- A. \$_cli_result
- B. \$_result
- C. \$_string_result
- D. \$_exit_status

Correct Answer: D

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The CLI event detector screens CLI commands for a regular expression match. When a match is found, an event is published. The match logic is performed on the fully expanded CLI command after the command is successfully parsed and before it is executed. The CLI event detector supports three publish modes:

- Synchronous publishing of CLI events--The CLI command is not executed until the EEM policy exits, and the EEM policy can control whether the command is executed. The read/write variable, `_exit_status`, allows you to set the exit status at policy exit for policies triggered from synchronous events. If `_exit_status` is 0, the command is skipped, if `_exit_status` is 1, the command is run.

- Asynchronous publishing of CLI events--The CLI event is published, and then the CLI command is executed.
- Asynchronous publishing of CLI events with command skipping--The CLI event is published, but the CLI command is not executed.

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/eem/command/eem-cr-book/eem-cr-e1.html>

QUESTION 422

Which two options are advantages of NetFlow version 9 over NetFlow version 5? (Choose two.)

- A. NetFlow version 9 adds support for IPv6 headers.
- B. NetFlow version 9 adds support for MPLS labels.
- C. NetFlow version 9 adds support for the Type of Service field.
- D. NetFlow version 9 adds support for ICMP types and codes.

Correct Answer: AB

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

NetFlow version 9 includes support for all of these fields that version 5 supports and can optionally include additional information such as Multiprotocol Label Switching (MPLS) labels and IPv6 addresses and ports.

QUESTION 423

Refer to the exhibit.

```
Flow export v5 is enabled for main cache
Export source and destination details :
VRF ID : Default
Destination(1) 10.5.206.250 (9995)
Version 5 flow records
Cache for prefix aggregation:
Flow export is disabled
53 flows exported in 18 udp datagrams
0 flows failed due to lack of export packet
0 export packets were sent up to process level
0 export packets were dropped due to no fib
0 export packets were dropped due to adjacency issues
0 export packets were dropped due to fragmentation failures
0 export packets were dropped due to encapsulation fixup failures
0 export packets were dropped enqueueing for the RP
0 export packets were dropped due to IPC rate limiting
0 export packets were dropped due to Card not being able to export
```

Which two statements about the output are true? (Choose two.)

- A. It indicates that prefix aggregation cache export is enabled on the device.
- B. It was obtained with the show ip cache flow command.

- C. It indicates that the device is using NetFlow version 5.
- D. It indicates that the flows are being sent to a destination using an RFC1918 address.

Correct Answer: CD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

C. The fourth line shows that Version 5 is being used.

D. The third line shows that the destination server is 10.5.206.250, which of course is a private, RFC 1918 address.

QUESTION 424

In the DiffServ model, which class represents the lowest priority with the lowest drop probability?

- A. AF11
- B. AF13
- C. AF41
- D. AF43

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Assured Forwarding (AF) Behavior Group				
	Class 1	Class 2	Class 3	Class 4
Low Drop	AF11 (DSCP 10)	AF21 (DSCP 18)	AF31 (DSCP 26)	AF41 (DSCP 34)
Med Drop	AF12 (DSCP 12)	AF22 (DSCP 20)	AF32 (DSCP 28)	AF42 (DSCP 36)
High Drop	AF13 (DSCP 14)	AF23 (DSCP 22)	AF33 (DSCP 30)	AF43 (DSCP 38)

Reference. http://en.wikipedia.org/wiki/Differentiated_services

QUESTION 425

Which three factors does Cisco PfR use to calculate the best exit path? (Choose three.)

- A. quality of service

- B. packet size
- C. delay
- D. loss
- E. reachability
- F. administrative distance

Correct Answer: CDE

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Cisco Pfr selects an egress or ingress WAN path based on parameters that affect application performance, including reachability, delay, cost, jitter, and Mean Opinion Score (MOS).

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

QUESTION 426

What is a reason to use DHCPv6 on a network that uses SLAAC?

- A. To get a record of the IPs that are used by the clients
- B. To push DNS and other information to the clients
- C. No reason, because there is no need for DHCPv6 when using SLAAC
- D. Because DHCPv6 can be used only in stateful mode with SLAAC to record the IPs of the clients
- E. Because DHCPv6 can be used only in stateless mode with SLAAC to record the IPs of the clients
- F. Because DHCPv6 is required to use first-hop security features on the switches

Correct Answer: B

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

SLAAC is by far the easiest way to configure IPv6 addresses, simply because you don't have to configure any IPv6 address. With SLAAC, a host uses the IPv6 Neighbor Discovery Protocol (NDP) to determine its IP address and default routers. Using SLAAC, a host requests and listens for Router Advertisements (RA) messages, and then taking the prefix that is advertised to form a unique address that can be used on the network. For this to work, the prefix that is advertised must advertise a prefix length of 64 bits (i.e., /64). But the most significant of Stateless Address Autoconfiguration (SLAAC) is it provided no mechanism for configuring DNS resolver information. Therefore SLAAC can be used along with DHCPv6 (Stateless) to push DNS and other information to the clients.

QUESTION 427

What can PfR passive monitoring mode measure for TCP flows?

- A. only delay
- B. delay and packet loss
- C. delay and reachability
- D. delay, packet loss, and throughput
- E. delay, packet loss, throughput, and reachability

Correct Answer: E

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Passive monitoring metrics include the following:

- · Delay: Cisco PfR measures the average delay of TCP flows for a given prefix or traffic class. Delay is the measurement of the round-trip response time (RTT) between the transmission of a TCP synchronization message and receipt of the TCP acknowledgement.
- · Packet loss: Cisco PfR measures packet loss by tracking TCP sequence numbers for each TCP flow; it tracks the highest TCP sequence number. If it receives a subsequent packet with a lower sequence number, PfR increments the packet-loss counter. Packet loss is measured in packets per million.
- · Reachability: Cisco PfR measures reachability by tracking TCP synchronization messages that have been sent repeatedly without receiving a TCP acknowledgement.
- · Throughput: Cisco PfR measures TCP throughput by measuring the total number of bytes and packets for each interesting traffic class or prefix for a given interval of time.

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

QUESTION 428

Refer to the exhibit.

```
police cir percent 10 conform-action proceed exceed-action set-mpls-experimental-topmost 6
```

A PE router is configured with a policy map that contains the policer shown. The policy map is configured in the inbound direction of an interface facing a CE router. If the PE router receives 12Mb/s of traffic with the CoS value set to 7 on a 100-Mb/s interface from the CE router, what value of MPLS EXP is set when this traffic goes through the policer shown?

- A. 0
- B. 6

- C. 7
- D. 8

Correct Answer: B

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Here, the policer is set where the conforming traffic is set to 10 percent of the 100 Mbps interface, so anything more than 10 Mbps will be placed into the exceeding traffic class, the traffic EXP value will be changed from 7 to 6 per the configuration.

QUESTION 429

DRAG DROP

What is the correct order of the VSS initialization process? Drag the actions on the left to the correct initialization step on the right.

Select and Place:

What is the correct order of the VSS initialization process? Drag the actions on the left to the correct initialization step on the right.	
bring up VSL links	initialization step 1
run VSLP	initialization step 2
preparse config	initialization step 3
run RRP	initialization step 4
continue system bootstrap	initialization step 5
interchassis SSO	initialization step 6

Correct Answer:

What is the correct order of the VSS initialization process? Drag the actions on the left to the correct initialization step on the right.

	preparse config
	bring up VSL links
	run VSLP
	run RRP
	interchassis SSO
	continue system bootup

Section: Infrastructure Services

Explanation

Explanation/Reference:

QUESTION 430

DRAG DROP

Drag and drop the QoS requirement on the left to the correct QoS technology on the right.

Select and Place:

Drag and drop the QoS requirement on the left to the correct QoS technology on the right.

Guarantees an amount of bandwidth	Police
Is an application classification	CBWFQ
Prioritizes real-time voice traffic	Shaping
Buffers bursting traffic	LLQ
Limits an amount of bandwidth	NBAR

Correct Answer:

Drag and drop the QoS requirement on the left to the correct QoS technology on the right.

	Limits an amount of bandwidth
	Guarantees an amount of bandwidth
	Buffers bursting traffic
	Prioritizes real-time voice traffic
	Is an application classification

Section: Infrastructure Services
Explanation

Explanation/Reference:

QUESTION 431

In a PfR environment, which two statements best describe the difference between active mode monitoring and fast mode monitoring? (Choose two.)

- A. Active mode monitoring can monitor and measure actual traffic via NetFlow data collection.
- B. Fast mode monitoring can measure bursty traffic better than active mode.
- C. Active mode monitoring uses IP SLA probes for the purpose of obtaining performance characteristics of the current WAN exit link.
- D. Fast mode monitoring uses IP SLA probes via all valid exits continuously to quickly determine an alternate exit link.

Correct Answer: CD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Active Monitoring

PfR uses Cisco IOS IP Service Level Agreements (SLAs) to enable active monitoring. IP SLAs support is enabled by default. IP SLAs support allows PfR to be configured to send active probes to target IP addresses to measure the jitter and delay, determining if a prefix is out-of-policy and if the best exit is selected. The border router collects these performance statistics from the active probe and transmits this information to the master controller.

Fast Failover Monitoring

Fast failover monitoring enables passive and active monitoring and sets the active probes to continuously monitor all the exits (probe-all). Fast failover monitoring can be used with all types of active probes: Internet Control Message Protocol (ICMP) echo, jitter, TCP connection, and UDP echo.

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/pfr/command/pfr-cr-book/pfr-s1.html>

QUESTION 432

Refer to the exhibit.

```
MC#sh pfr master traffic-class
```

OER Prefix Statistics:

Pas - Passive, Act - Active, S - Short term, L - Long term, Dly - Delay (ms),
 P - Percentage below threshold, Jit - Jitter (ms),
 MOS - Mean Opinion Score
 Los - Packet Loss (packets-per-million), Un - Unreachable (flows-per-million),
 E - Egress, I - Ingress, Bw - Bandwidth (kbps), N - Not applicable
 U - unknown, * - uncontrolled, + - control more specific, @ - active probe all
 # - Prefix monitor mode is Special, & - Blackholed Prefix
 % - Force Next-Hop, ^ - Prefix is denied

DstPrefix	Flags	Appl_ID	Dscp	Prot	SrcPort	DstPort	SrcPrefix	Protocol
	PasSDly	PasLDly	PasSUn	PasLUn	PasSJos	PasLJos	EBw	IBw
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSJos	ActLJos
10.1.0.0/24			N	N	N	N	N	
			INPOLICY*		@83	10.4.5.4	Et0/1	U
	52	52	0	0	0	0	67	7
	51	51	0	0	N	N	N	N

Which two statements are true regarding prefix 10.1.0.0/24? (Choose two.)

- A. The prefix is in policy, and Cisco PfR rerouted the traffic via 10.4.5.3 Et0/1 because of an OOP event.
- B. Cisco PfR is monitoring the prefix via passive NetFlow mode only.
- C. Cisco PfR is monitoring the prefix via active, fast, or active throughput IP SLA probe mode only.
- D. The prefix is in policy, and Cisco PfR did not reroute the traffic via 10.4.5.3 Et0/1 because the traffic was previously in policy.
- E. Cisco PfR is monitoring the prefix via mode monitor, which provides both NetFlow and IP SLA measurements.

Correct Answer: DE

Section: Infrastructure Services

Explanation

Explanation/Reference:

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/pfr/command/pfr-cr-book/pfr-s1.html#wp2707728086>

QUESTION 433

In the DiffServ model, which class represents the lowest priority with the highest drop probability?

- A. AF11
- B. AF13
- C. AF41
- D. AF43

Correct Answer: B

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Assured Forwarding (AF) Behavior Group				
	Class 1	Class 2	Class 3	Class 4
Low Drop	AF11 (DSCP 10)	AF21 (DSCP 18)	AF31 (DSCP 26)	AF41 (DSCP 34)
Med Drop	AF12 (DSCP 12)	AF22 (DSCP 20)	AF32 (DSCP 28)	AF42 (DSCP 36)
High Drop	AF13 (DSCP 14)	AF23 (DSCP 22)	AF33 (DSCP 30)	AF43 (DSCP 38)

Reference. http://en.wikipedia.org/wiki/Differentiated_services

QUESTION 434

Which two hashing algorithms can be used when configuring SNMPv3? (Choose two.)

- A. MD5
- B. SHA-1
- C. Blowfish
- D. DES
- E. AES
- F. SSL

Correct Answer: AB

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Note that SNMPv3 does not send passwords in clear-text and uses hash-based authentication with either MD5 or SHA1 functions (HMAC authentication the packet content is hashed along with authentication key to produce the authentication string).

Reference. <http://blog.ine.com/2008/07/19/snmpv3-tutorial/>

QUESTION 435

Which two statements about the default router settings for SSH connections are true? (Choose two.)

- A. The default timeout value for the SSH negotiation phase is 120 seconds.
- B. Data is exchanged in clear text by default unless AAA authentication is enabled on the console.
- C. The default number of authentication retries is 3.
- D. SSH is enabled by default when you configure the username command.

Correct Answer: AC

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

<code>ip ssh {timeout seconds authentication- retries number}</code>	<p>Configures the SSH control parameters:</p> <ul style="list-style-type: none">• Specify the time-out value in seconds; the default is 120 seconds. The range is 0 to 120 seconds. This parameter applies to the SSH negotiation phase. After the connection is established, the Switch uses the default time-out values of the CLI-based sessions. By default, up to five simultaneous, encrypted SSH connections for multiple CLI-based sessions over the network are available (session 0 to session 4). After the execution shell starts, the CLI-based session time-out value returns to the default of 10 minutes.• Specify the number of times that a client can re-authenticate to the server. The default is 3; the range is 0 to 5.
--	---

Reference.

http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3850/software/release/3se/security/configuration_guide/b_sec_3se_3850_cg/b_sec_3se_3850_cg_chapter_01000.html

QUESTION 436

Refer to the exhibit.

```
R1#sh logging
Syslog logging: enabled (12 messages dropped, 0 messages rate-limited, 0 flushes, 0 overruns, xml disabled, filtering disabled)

No Active Message Discriminator.

No Inactive Message Discriminator.

  Console logging: level debugging, 28 messages logged, xml disabled, filtering disabled
  Monitor logging: level debugging, 0 messages logged, xml disabled, filtering disabled
  Buffer logging:  level debugging, 7 messages logged, xml disabled, filtering disabled
  Logging Exception size (4096 bytes)
  Count and timestamp logging messages: disabled
  Persistent logging: disabled

No active filter modules.

ESM: 0 messages dropped

%SYS-5-CONFIG_I: Configured from console by console
  Trap logging: level informational, 32 message lines logged

Log Buffer (4096 bytes):

%BGP-5-ADJCHANGE: neighbor 209.165.200.226 Down Interface flap
%LINK-5-CHANGED: Interface FastEthernet0/0, changed state to administratively down
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state to down
%SYS-5-CONFIG_I: Configured from console by console
```

Which statement about the R1 configuration is true?

- A. It supports the service timestamps log uptime command to display time stamps.
- B. The logging buffer command was used to increase the default of the buffer.
- C. The logging of warning messages is disabled.
- D. Log message sequence numbering is disabled.

Correct Answer: D

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

To disable sequence numbers, use the **no service sequence-numbers** global configuration command.

This example shows part of a logging display with sequence numbers enabled.

000019: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)

In this example we see the absence of sequence numbers on the log messages.

Not A. In this example there are no time stamps or uptimes shown in the logs.

Not B. The default buffer size is 4096 bytes.

Not C. The logging level in this example is informational (level 6), which will display levels 0-6 in the logs. Warnings are level 4.

QUESTION 437

Which two statements about class maps are true? (Choose two.)

- A. As many as eight DSCP values can be included in a match dscp statement.
- B. The default parameter on a class map with more than one match command is match-any.
- C. The match class command can nest a class map within another class map.
- D. A policy map can be used to designate a protocol within a class map.

Correct Answer: AC

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Answer A.

Router(config-cmap)# match [ip] dscp <i>dscp-value [dscp-value dscp-value</i> <i>dscp-value dscp-value dscp-value</i> <i>dscp-value dscp-value]</i>	(Optional) Identifies a specific IP differentiated service code point (DSCP) value as a match criterion. Up to eight DSCP values can be included in one match statement.
---	--

Answer C.

Router config-cmap)# match class-map <i>class-name</i>	(Optional) Specifies the name of a traffic class to be used as a matching criterion (for nesting traffic class [nested class maps] within one another).
---	---

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_2/qos/configuration/guide/fqos_c/qcfccli2.html

QUESTION 438

Which IP SLA operation type is enhanced by the use of the IP SLAs Responder?

- A. DNS
- B. HTTP
- C. ICMP Echo
- D. UDP Echo

Correct Answer: D


Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Figure 1. UDP Echo Operation




Response time (round-trip time) is computed by measuring the time taken between sending a UDP echo request message from Device B to the destination device--Device A--and receiving a UDP echo reply from Device A. UDP echo accuracy is enhanced by using the IP SLAs Responder at Device A, the destination Cisco device. If the destination device is a Cisco device, then IP SLAs sends a UDP datagram to any port number that you specified. Using the IP SLAs Responder is optional for a UDP echo operation when using Cisco devices. The IP SLAs Responder cannot be configured on non-Cisco devices.

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

QUESTION 439

Refer to the exhibit.



Router 1 and Router 2 use HSRP to provide first hop redundancy for hosts on the 10.1.2.0/24 network.

Which feature can provide additional failover coverage for the PC?

- A. Cisco Express Forwarding
- B. NetFlow
- C. Accounting
- D. Enhanced Object Tracking

Correct Answer: D

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

HSRP has a mechanism for tracking the interface line-protocol state. The enhanced object tracking feature separates the tracking mechanism from HSRP. It creates a separate, standalone tracking process that can be used by processes other than HSRP. This feature allows tracking of other objects

in addition to the interface line-protocol state. A client process, such as HSRP, can register an interest in tracking objects and request notification when the tracked object changes state. Several clients can track the same object, and can take different actions when the object changes state. This feature increases the availability and speed of recovery of a router system and decreases outages and outage duration.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/blades/3040/software/release/12-2_44_se/configuration/guide/swhsrp.html#wp1083927

QUESTION 440

Refer to the exhibit.

```
R2#show ntp associations
  address      ref clock   st  when poll reach  delay  offset  disp
~10.1.1.1      0.0.0.0    16   61   64   0    0.0   0.00  16000.
* master (syncd), # master (unsyncd), + selected, - candidate, ~ configured

R2#show ip route | include 10.1.1.1
O    10.1.1.1/32 [110/11] via 10.1.12.1, 00:20:28, FastEthernet0/0.12

R2#show run | include ntp
ntp authentication-key 1 md5 110A1016141D 7
ntp authenticate
ntp trusted-key 1
ntp clock-period 17179894
ntp server 10.1.1.1 key 1

R1#show ip route connected
 209.165.200.0/27 is subnetted, 1 subnets
C    209.165.200.224 is directly connected, FastEthernet0/0.112
 10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
C    10.1.13.0/24 is directly connected, FastEthernet0/1.13
C    10.1.12.0/24 is directly connected, FastEthernet0/0.12
C    10.1.1.0/24 is directly connected, Loopback0
```

Which two possible network conditions can you infer from this configuration? (Choose two.)

- A. The authentication parameters on R1 and R2 are mismatched.
- B. R1 is using the default NTP source configuration.
- C. R1 and R2 have established an NTP session.
- D. R2 is configured as the NTP master with a stratum of 7.

Correct Answer: AB

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Answer A. The NTP associations are not synced, it is only listed as a candidate because it was configured. Routing is not the issue, so it must be mismatched authentication parameters.

Answer B. NTP sets the source IP address for all NTP packets based on the address of the interface through which the NTP packets are sent. You can configure NTP to use a specific source IP address.

QUESTION 441

Which three message types are used for prefix delegation in DHCPv6? (Choose three.)

- A. DHCP Discover
- B. Renew
- C. Solicit
- D. DHCP Offer
- E. Advertise
- F. DHCP Ack

Correct Answer: BCE

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

DHCPv6 Message Types

For a client to get an IPv6 address successfully from a DHCPv6 server, the Client-Server Conversation happens using the following messages.

Client-->Server Messages	Server-->Client Messages
<i>Solicit, Request, Confirm, Renew, Rebind, Release, Decline, Information-Request</i>	<i>Advertise, Reply, Reconfigure</i>

Lets look at each message types in detail:

SOLICIT

This is the first step in DHCPv6, where a DHCPv6 client sends a Solicit message to locate DHCPv6 servers.

ADVERTISE

Upon receiving a Solicit Message from the client, the DHCPv6 server sends an Advertise message to indicate that it is available for DHCP service, in response to a Solicit message received from a client.

REQUEST

This message is sent by the DHCPv6 client. Client sends a Request message to request configuration parameters which includes IP addresses or delegated prefixes, from a specific server.

CONFIRM

Confirm message is sent by the client to any available server in the network to confirm that the client is still on the same link or it has to be removed. This message also confirms the IPv6 addresses that are assigned to the link are still valid. This could happen in case when a client detects a change in link-layer connectivity or if the device is powered on and it is found that one or more leases are still valid. Note that only the prefix portion of the addresses are validated and not the actual leases.

RENEW

A client sends a Renew message to the server when it wants to extend the lifetimes on the addresses and other configuration parameters assigned to the client and also to update other configuration parameters.

REBIND

In case of No response from the DHCPv6 Server for the Renew message, the client sends a Rebind message to any available server to extend the lifetimes on the address and to update other configuration parameters.

REPLY

A Reply message is sent by the DHCPv6 Server in response to a Solicit, Request, Renew, Rebind message received from a client. The reply message is sent by the server in response to a confirm message (either confirming or denying) that the addresses assigned to the client are appropriate. In short the server acknowledge receipt of a Release or Decline message by sending a REPLY message.

RELEASE

Release message as the name implies, is sent by the client to the server that has assigned the addresses, to indicate that the client will no longer use the assigned addresses (one or more).

DECLINE

Client sends a Decline message to the DHCPv6 server to tell that the one or more addresses assigned by the server is already in use

RECONFIGURE

The Reconfigure Message is sent by the DHCPv6 server to the client when the server has new or updated information of configuration parameters. It tells the client to initiate a information- request/reply message to the server so as to receive the updated information.

INFORMATION-REQUEST

Information-Request message is sent by the client to the server to update the configuration parameters

Reference. <https://supportforums.cisco.com/blog/153426/implementing-dhcpv6-introduction>

QUESTION 442

Which two statements about static NAT are true? (Choose two.)

- A. An outside local address maps to the same outside global IP address.
- B. An inside local address maps to a different inside global IP address.
- C. An outside local address maps to a different outside global IP address.
- D. An inside local address maps to the same inside global IP address.

Correct Answer: AD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Example found at the reference link below:

Reference. <http://www.cisco.com/c/en/us/support/docs/ip/network-address-translation-nat/4606-8.html>

QUESTION 443

DRAG DROP

Drag and drop the NetFlow Export feature on the left to the NetFlow version that first supported it on the right.

Select and Place:

Drag and drop the NetFlow Export feature on the left to the NetFlow version that first supported it on the right.	
exports data from the aggregation cache only	Version 5
exports data from the main and aggregation caches	
exports data from the main cache only	
supports BGP next-hop	Version 8
supports BGP AS information	
supports IPv6	Version 9

Correct Answer:

Drag and drop the NetFlow Export feature on the left to the NetFlow version that first supported it on the right.

	Version 5
	exports data from the main cache only
	supports BGP AS information
	Version 8
	exports data from the aggregation cache only
	Version 9
	supports IPv6
	exports data from the main and aggregation caches
	supports BGP next-hop

Section: Infrastructure Services

Explanation

Explanation/Reference:

QUESTION 444

Which two options are required parts of an EEM policy? (Choose two.)

- A. event register keyword
- B. body
- C. environment must defines

- D. namespace import
- E. entry status
- F. exit status

Correct Answer: AB

Section: Infrastructure Services

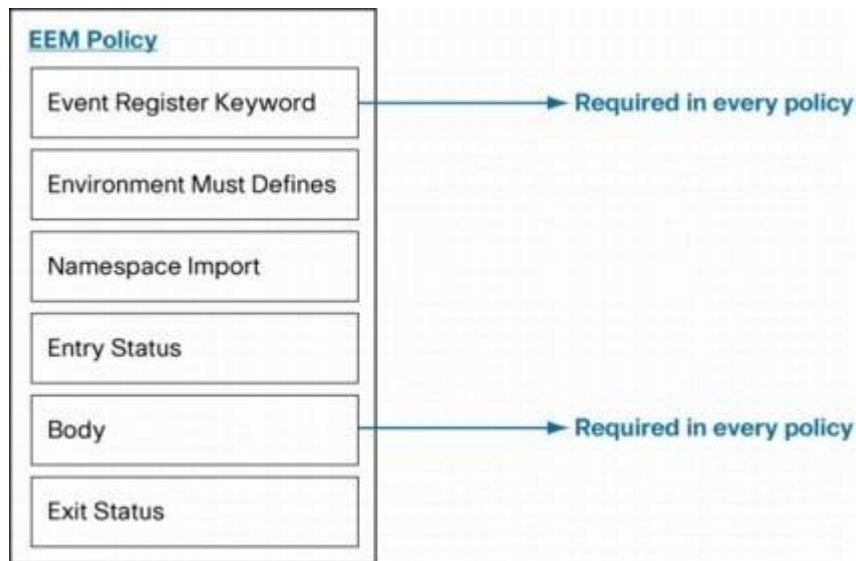
Explanation

Explanation/Reference:

Explanation:

EEM policies require two parts: the event register keyword and the body. The remaining parts of the policy are optional: environment must defines, namespace import, entry status, and exit status (Figure 5).

Figure 5. EEM Policy Parts



Reference. http://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/ios-embedded- event-manager-eem/config_guide_eem_configuration_for_cisco_integrated_services_router_platforms.html

QUESTION 445

Which two actions can you take to allow the greatest number of pertinent packets to be stored in the temporary buffer of Cisco IOS Embedded Packet Capture? (Choose two.)

- A. Specify the sampling interval.

- B. Specify the capture buffer type.
- C. Specify a reflexive ACL.
- D. Specify the minimum packet capture rate.
- E. Specify the packet size.
- F. Store the capture simultaneously onto an external memory card as the capture occurs.

Correct Answer: AB

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Embedded Packet Capture (EPC) provides an embedded systems management facility that helps in tracing and troubleshooting packets. This feature allows network administrators to capture data packets flowing through, to, and from a Cisco device. The network administrator may define the capture buffer size and type (circular, or linear) and the maximum number of bytes of each packet to capture. The packet capture rate can be throttled using further administrative controls. For example, options allow for filtering the packets to be captured using an Access Control List and, optionally, further defined by specifying a maximum packet capture rate or by specifying a sampling interval.

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/epc/configuration/xe-3s/asr1000/epc-xe-3s-asr1000-book/nm-packet-capture-xe.html>

QUESTION 446

Which statement describes Cisco PfR link groups?

- A. Link groups enable Cisco PfR Fast Reroute when NetFlow is enabled on the external interfaces of the border routers.
- B. Link groups define a strict or loose hop-by-hop path preference.
- C. Link groups are required only when Cisco PfR is configured to load-balance all traffic.
- D. Link groups are enabled automatically when Cisco PfR is in Fast Reroute mode.
- E. Link groups set a preference for primary and fallback (backup) external exit interfaces.

Correct Answer: E

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The Performance Routing - Link Groups feature introduced the ability to define a group of exit links as a preferred set of links, or a fallback set of links for PfR to use when optimizing traffic classes specified in an PfR policy. PfR currently selects the best link for a traffic class based on the preferences specified in a policy and the traffic class performance--using parameters such as reachability, delay, loss, jitter or MOS--on a path out of the specified link.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/pfr/configuration/guide/15_1/pfr_15_1_book/pfr-link-group.html

QUESTION 447

Which two statements about NetFlow are true? (Choose two.)

- A. It must be configured on each router in a network.
- B. It supports ATM LAN emulation.
- C. The existing network is unaware that NetFlow is running.
- D. It uses SIP to establish sessions between neighbors.
- E. It provides resource utilization accounting.

Correct Answer: CE

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

NetFlow identifies packet flows for both ingress and egress IP packets. It does not involve any connection-setup protocol, either between routers or to any other networking device or end station. NetFlow does not require any change externally--either to the packets themselves or to any networking device. NetFlow is completely transparent to the existing network, including end stations and application software and network devices like LAN switches. Also, NetFlow capture and export are performed independently on each internetworking device; NetFlow need not be operational on each router in the network.

NetFlow data provides fine-grained metering for highly flexible and detailed resource utilization accounting. For example, flow data includes details such as IP addresses, packet and byte counts, timestamps, type-of-service, and application ports. Service providers might utilize the information for billing based on time-of-day, bandwidth usage, application usage, or quality of service. Enterprise customers might utilize the information for departmental chargeback or cost allocation for resource utilization.

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/netflow/configuration/12-4t/nf-12-4t-book/ios-netflow-ov.html>

QUESTION 448

You are installing a new device to replace a device that failed. The configuration of the failed device is stored on a networked server, and the new device has an RXBOOT image installed. Under which condition does the streamlined Setup mode fail?

- A. The last four bits of the configuration register are not equal to the decimal value 0 or 1.
- B. The startup configuration file was deleted.
- C. Bit 6 is set in the configuration register.
- D. The startup configuration is corrupt.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The lowest four bits of the configuration register (bits 3, 2, 1, and 0) form the boot field. The boot field determines if the router boots manually, from ROM, or from Flash or the network. To change the boot field value and leave all other bits set to their default values, follow these guidelines:

- If you set the configuration register boot field value to 0x0, you must boot the operating system manually with the **boot** command.
- If you set the configuration register boot field value to 0x1, the router boots using the default ROM software.
- If you set the configuration register boot field to any value from 0x2 to 0xF, the router uses the boot field value to form a default boot filename for booting from a network server.


For more information about the configuration register bit settings and default filenames, refer to the appropriate router hardware installation guide.

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/12_2/configfun/command/reference/ffun_r/frf010.html

QUESTION 449

Refer to the exhibit.



You are configuring the S1 switch for the switchport connecting to the client computer. Which option describes the effect of the command `mls qos map cos-dscp 0 8 16 24 32 40 46 56`?

- A. Voice traffic is excluded from the default priority queue.
- B. Voice packets are given a class selector of 5.
- C. Video conferencing is marked CS3.
- D. Voice packets are processed in the priority queue.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The default CoS to DSCP mappings are shown below:

Default CoS-to-DSCP Map	
CoS Value	DSCP Value
0	0
1	8
2	16
3	24
4	32
5	40
6	48
7	56

In our example, we see that COS 6 is mapped to DSCP, not the default of DSCP 48 as shown above. DSCP 46 is Expedited Forwarding (EF), which is typically used for voice traffic, and this value has not been included in this class map.

QUESTION 450

Refer to the exhibit.

```
class-map match-any voice
match dscp ef
class-map match-any router
match dscp cs6
class-map match-any gold
match dscp af41
class-map match-any silver
match dscp af31

policy-map egress_queue
class voice
priority percent 25
class gold
bandwidth percent 40
class silver
bandwidth percent 15
class router
bandwidth percent 5
class class-default
bandwidth percent remaining

policy-map egress_queue_2
class class-default
shape average 6000000
service-policy egress_queue

interface GigabitEthernet0/1
service-policy output egress_queue_2
```

If the network switch is configured as shown, which two statements about network traffic are true? (Choose two.)

- A. Traffic enters the shaper on a FIFO basis.
- B. Traffic enters the shaper on a weighted fair queueing basis.
- C. Drop behavior is random for traffic in excess of 6 Mbps.
- D. Voice traffic is given priority until it reaches 1.5 Mbps.
- E. Voice traffic is given priority until it reaches 6 Mbps.

Correct Answer: AD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

A. Serial interfaces at E1 (2.048 Mbps) and below use weighted fair queueing (WFQ) by default. Other interfaces use first-in first-out (FIFO) by default.

D. Voice traffic is given priority up to 25% of the shape average value, which is 6000000, so 25% of 6 Mbps is 1.5 Mbps.

QUESTION 451

Which two options are two characteristics of the HSRPv6 protocol? (Choose two.)

- A. It uses virtual MAC addresses 0005.73a0.0000 through 0005.73a0.0fff.
- B. It uses UDP port number 2029.
- C. It uses virtual MAC addresses 0005.73a0.0000 through 0005.73a0.ffff.
- D. It uses UDP port number 2920.
- E. If a link local IPv6 address is used, it must have a prefix.

Correct Answer: AB

Section: Infrastructure Services

Explanation

Explanation/Reference:

HSRP IPv6 Virtual MAC Address Range

HSRP IPv6 uses a different virtual MAC address block than does HSRP for IP:

0005.73A0.0000 through 0005.73A0.0FFF (4096 addresses)

HSRP IPv6 UDP Port Number

Port number 2029 has been assigned to HSRP IPv6.

Reference. http://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 452

Which statement about VRRP is true?

- A. It supports load balancing.
- B. It can be configured with HSRP on a switch or switch stack.
- C. It supports IPv4 and IPv6.
- D. It supports encrypted authentication.

Correct Answer: B

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

VRRP Limitations

- You can configure both HSRP and VRRP on a switch or switch stack. However, you cannot add a switch model that supports only one protocol to a stack that is configured for both protocols.
- The VRRP implementation on the switch does not support the MIB specified in RFC 2787.
- The VRRP implementation on the switch supports only text -based authentication.
- The switch supports VRRP only for IPv4.

Reference.

http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750x_3560x/software/release/12-2_58_se/configuration/guide/3750xscg/swhsrp.html#pgfId-1107127

QUESTION 453

Refer to the exhibit.

```
ip sla monitor 10
  type echo protocol ipIcmpEcho 10.1.1.1 source-ipaddr 10.1.1.2
  frequency 60

ip sla monitor schedule 10 life 360
```

What is the polling frequency set by this configuration?

- A. 60 seconds
- B. 10 seconds
- C. 360 seconds
- D. 60 milliseconds
- E. 10 milliseconds

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The frequency value lists the polling interval, in seconds.

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

QUESTION 454

Refer to the exhibit.

```
configure terminal
  interface Ethernet 0/0
    ip address 10.1.1.2 255.255.255.0
    ip flow-export destination 10.1.1.1
```

Which additional information must you specify in this configuration to capture NetFlow traffic?

- A. ingress or egress traffic
- B. the number of cache entries
- C. the flow cache active timeout
- D. the flow cache inactive timeout

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Configuring NetFlow

Perform the following task to enable NetFlow on an interface.

SUMMARY STEPS

1. enable
2. configure terminal
3. interface *type number*
4. ip flow {ingress | egress}
5. exit
6. Repeat Steps 3 through 5 to enable NetFlow on other interfaces.
7. end
- 8.

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. • __Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	interface <i>type number</i> Example: Router(config)# interface ethernet 0/0	Specifies the interface that you want to enable NetFlow on and enters interface configuration mode.
Step 4	ip flow {ingress egress} Example: Router(config-if)# ip flow ingress	Enables NetFlow on the interface. • __ingress—Captures traffic that is being received by the interface • __egress—Captures traffic that is being transmitted by the interface
Step 5	exit Example: Router(config-if)# exit	(Optional) Exits interface configuration mode and enters global configuration mode. Note You need to use this command only if you want to enable NetFlow on another interface.
Step 6	Repeat Steps 3 through 5 to enable NetFlow on other interfaces.	This step is optional.
Step 7	end Example: Router(config-if)# end	Exits the current configuration mode and returns to privileged EXEC mod

Reference.

http://www.cisco.com/c/en/us/td/docs/ios/netflow/configuration/guide/12_2sr/nf_12_2sr_book/cfg_nflow_data_expt.html

QUESTION 455

For which three routing protocols can Cisco PfR provide direct route control? (Choose three.)

- A. OSPF
- B. ISIS
- C. BGP
- D. EIGRP
- E. static routing
- F. ODR

Correct Answer: CDE

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Q. Can you elaborate more on the Parent Route and why it's so important to PfR?

A. Yes. For any route that PfR modifies or controls (BGP, Static, PIRO, EIGRP, PBR), having a Parent prefix in the routing table eliminates the possibility of a routing loop occurring. This is naturally a good thing to prevent in routed networks.

Reference. http://docwiki.cisco.com/wiki/Performance_Routing_FAQs#Route_Control

QUESTION 456

DRAG DROP

Drag and drop the NAT operations on the left into the correct sequential order on the right.

Select and Place:

Drag and drop the NAT operations on the left into the correct sequential order on the right.

Check the IP routing table.	step 1
Check the outbound access list.	step 2
Check the inbound access list.	step 3
Inspect CBAC.	step 4
Translate inside local to outside global.	step 5
Check the policy routing	step 6

Correct Answer:

Drag and drop the NAT operations on the left into the correct sequential order on the right.

	Check the inbound access list.
	Check the policy routing
	Check the IP routing table.
	Translate inside local to outside global.
	Check the outbound access list.
	Inspect CBAC.

Section: Infrastructure Services

Explanation

Explanation/Reference:

QUESTION 457

DRAG DROP

Drag and drop each step in the performance-monitoring configuration process on the left into the correct order on the right.

Select and Place:

Drag and drop each step in the performance-monitoring configuration process on the left into the correct order on the right.	
Configure a policy with at least one performance-monitor type flow monitor.	1
Configure a flow record.	2
Configure a class that describes the filtering criteria.	3
Associate a performance-monitor type policy with its corresponding interface.	4
Configure a flow monitor that includes the flow record and flow exporter.	5

Correct Answer:

Drag and drop each step in the performance-monitoring configuration process on the left into the correct order on the right.

	Configure a flow record.
	Configure a flow monitor that includes the flow record and flow exporter.
	Configure a class that describes the filtering criteria.
	Configure a policy with at least one performance-monitor type flow monitor.
	Associate a performance-monitor type policy with its corresponding interface.

Section: Infrastructure Services
Explanation

Explanation/Reference:

QUESTION 458
Refer to the exhibit.

```
R1#show run

ip ssh time-out 30
ip ssh authentication-retries 2

access-list 10 permit 10.1.1.2
no cdp log mismatch duplex

control-plane

line con 0
  exec-timeout 5 30
  logging synchronous
line aux 0
line vty 0 4
  access-class 10 in
  login
  transport input ssh
```

Which statement about the R1 configuration is true?

- A. It permits host 10.1.1.2 to establish a Telnet connection to R1.
- B. It limits remote hosts to two SSH connection attempts.
- C. SSH connections to R1 will log out after a 5-minute idle interval.
- D. Hosts that reside on network 10.0.0.0/8 can SSH to R1.
- E. The R1 timeout for outgoing SSH connection attempts is 30 seconds.

Correct Answer: E

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The timeout for outgoing SSH connection is defined by the "ip sshh time-out" command (in seconds), which is configured here as 30.

QUESTION 459

Which two statements about the default SNMP configuration are true? (Choose two.)

- A. The SNMP agent is enabled.

- B. The SNMP trap receiver is configured.
- C. All SNMP notification types are sent.
- D. SNMPv1 is the default version.
- E. SNMPv3 is the default version.

Correct Answer: CD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Default SNMP Configuration	
Feature	Default Setting
SNMP agent	Enabled
SNMP community strings	Read-Only: Public Read-Write: Private Read-Write-all: Secret
SNMP trap receiver	None configured
SNMP traps	None enabled
SNMP version	If no version keyword is present, the default is version 1.
SNMPv3 authentication	If no keyword is entered, the default is the noauth (noAuthNoPriv) security level.
SNMP notification type	If no type is specified, all notifications are sent.

Reference. http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2950/software/release/12-1_19_ea1/configuration/guide/2950scg/swsnmp.html

QUESTION 460

Which two statements about logging are true? (Choose two.)

- A. Log messages are sent to the console port by default.
- B. Log messages are displayed in a Telnet session by default.
- C. Interface status changes are logged at the Notification level.
- D. Interface status changes are logged at the Informational level.
- E. System restart messages are logged at the Critical level.
- F. Reload requests are logged at the Notification level.

Correct Answer: AC

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

By default, switches send the output from system messages and debug privileged EXEC commands to a logging process. The logging process controls the distribution of logging messages to various destinations, such as the logging buffer, terminal lines, or a UNIX syslog server, depending on your configuration. The process also sends messages to the console.

Table 29-3 Message Logging Level Keywords			
Level Keyword	Level	Description	Syslog Definition
emergencies	0	System unstable	LOG_EMERG
alerts	1	Immediate action needed	LOG_ALERT
critical	2	Critical conditions	LOG_CRIT
errors	3	Error conditions	LOG_ERR
warnings	4	Warning conditions	LOG_WARNING
notifications	5	Normal but significant condition	LOG_NOTICE
informational	6	Informational messages only	LOG_INFO
debugging	7	Debugging messages	LOG_DEBUG

The software generates four other categories of messages:

- . Error messages about software or hardware malfunctions, displayed at levels warnings through emergencies. These types of messages mean that the functionality of the switch is affected. For information on how to recover from these malfunctions, see the system message guide for this release.
- . Output from the debug commands, displayed at the debugging level. Debug commands are typically used only by the Technical Assistance Center.
- . Interface up or down transitions and system restart messages, displayed at the notifications level. This message is only for information; switch functionality is not affected.

- Reload requests and low-process stack messages, displayed at the informational level. This message is only for information; switch functionality is not affected.

References: http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2950/software/release/12-1_9_ea1/configuration/guide/scg/swlog.html
http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960/software/release/12-2_55_se/configuration/guide/scg_2960/swlog.html

QUESTION 461

Refer to the exhibit.

```
mls qos queue-set output 1 threshold 2 80 90 100 100
mls qos queue-set output 1 threshold 3 400 400 100 800
mls qos queue-set output 1 threshold 4 60 100 100 100
```

If the remaining configuration uses default values, what is the expected output of the show mls qos queue-set command?

A.

```
Queueset: 1
Queue      :      1      2      3      4
-----
buffers    :      25      25      25      25
threshold1:     100      80     400      60
threshold2:     100      90     400     100
reserved   :       50     100     100     100
maximum    :     400     100     800     100
```

B.

```
Queueset: 1
Queue      :      1      2      3      4
-----
buffers    :      25      25      25      25
threshold1:     100      80     400      60
threshold2:     100      90     400     100
reserved   :       50     100     100     100
maximum    :     100     100     800     100
```

C.

Queueset: 1					
Queue	:	1	2	3	4
<hr/>					
buffers	:	25	25	25	25
threshold1:		50	80	400	60
threshold2:		50	90	400	100
reserved	:	50	100	100	100
maximum	:	400	100	800	100

D.

Queueset: 1					
Queue	:	1	2	3	4
<hr/>					
buffers	:	25	25	25	25
threshold1:		100	80	400	60
threshold2:		100	90	400	100
reserved	:	100	100	100	100
maximum	:	400	100	800	100

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

mls qos queue-set output <i>qset-id</i> threshold <i>queue-id</i> drop-threshold1 drop-threshold2 reserved-threshold maximum-threshold	<p>Configure the WTD thresholds, guarantee the availability of buffers, and configure the maximum memory allocation for the queue-set (four egress queues per port). By default, the WTD thresholds for queues 1, 3, and 4 are set to 100 percent. The thresholds for queue 2 are set to 200 percent. The reserved thresholds for queues 1, 2, 3, and 4 are set to 50 percent. The maximum thresholds for all queues are set to 400 percent.</p> <ul style="list-style-type: none">• For <i>qset-id</i>, enter the ID of the queue-set specified in Step 2. The range is 1 to 2.• For <i>queue-id</i>, enter the specific queue in the queue-set on which the command is performed. The range is 1 to 4.• For <i>drop-threshold1 drop-threshold2</i>, specify the two WTD thresholds expressed as a percentage of the queue's allocated memory. The range is 1 to 3200 percent.• For <i>reserved-threshold</i>, enter the amount of memory to be guaranteed (reserved) for the queue expressed as a percentage of the allocated memory. The range is 1 to 100 percent.• For <i>maximum-threshold</i>, enable a queue in the full condition to obtain more buffers than are reserved for it. This is the maximum memory the queue can have before the packets are dropped if the common pool is not empty. The range is 1 to 3200 percent
---	--

Reference.

http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750x_3560x/software/release/12-2_55_se/configuration/guide/3750xscg/swqos.html

QUESTION 462

Which two statements about the client-identifier in a DHCP pool are true? (Choose two.)

- A. It specifies a unique identifier that is used only for DHCP requests.
- B. It is specified by appending 01 to the MAC address of a DHCP client.
- C. It specifies a hardware address for the client.
- D. It specifies a unique identifier that is used only for BOOTP requests.
- E. It requires that you specify the hardware protocol.

Correct Answer: AB

Section: Infrastructure Services

Explanation

Explanation/Reference:


Reference.

client- identifier <i>unique- identifier</i> Example: Device(dhcp- config)# client- identifier 01b7.0813.881 1.66	<p>Specifies the unique identifier for DHCP clients.</p> <ul style="list-style-type: none"> This command is used for DHCP requests. DHCP clients require client identifiers. You can specify the unique identifier for the client in either of the following ways: <ul style="list-style-type: none"> A 7-byte dotted hexadecimal notation. For example, 01b7.0813.8811.66, where 01 represents the Ethernet media type and the remaining bytes represent the MAC address of the DHCP client. A 27-byte dotted hexadecimal notation. For example, 7665.6e64.6f72.2d30.3032.342e.3937.6230.2e33.3734.312d.4661.302f.31. The equivalent ASCII string for this hexadecimal value is vendor-0024.97b0.3741-fa0/1, where vendor represents the vendor, 0024.97b0.3741 represents the MAC address of the source interface, and fa0/1 represents the source interface of the DHCP client. See the Troubleshooting Tips section for information about how to determine the client identifier of the DHCP client. <table border="1"> <tr> <td>Note</td><td>The identifier specified here is considered for a DHCP client that sends a client identifier in the packet.</td></tr> </table>	Note	The identifier specified here is considered for a DHCP client that sends a client identifier in the packet.
Note	The identifier specified here is considered for a DHCP client that sends a client identifier in the packet.		

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-mt/dhcp-15-mt-book/config-dhcp-server.html

QUESTION 463

Refer to the exhibit.



If router R1 is functioning as a DHCPv6 server and you enter the command `show ipv6 dhcp binding`, which two options are pieces of information in the output? (Choose two.)

- A. The IA PD
- B. The DUID
- C. The prefix pool
- D. The DNS server
- E. The Rapid-Commit setting

Correct Answer: AB

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

In the following example, the show ipv6 dhcp binding command shows information about two clients, including their DUIDs, IAPDs, prefixes, and preferred and valid lifetimes:

Router# **show ipv6 dhcp binding**

Client: FE80::202:FCFF:FEA5:DC39 (GigabitEthernet2/1/0)

DUID. 000300010002FCA5DC1C

IA PD. IA ID 0x00040001, T1 0, T2 0

Prefix: 3FFE:C00:C18:11::/68

preferred lifetime 180, valid lifetime 12345

expires at Nov 08 2002 02:24 PM (12320 seconds)

Client: FE80::202:FCFF:FEA5:C039 (GigabitEthernet2/1/0)

DUID. 000300010002FCA5C01C

IA PD. IA ID 0x00040001, T1 0, T2 0

Prefix: 3FFE:C00:C18:1::/72

preferred lifetime 240, valid lifetime 54321

expires at Nov 09 2002 02:02 AM (54246 seconds)

Prefix: 3FFE:C00:C18:2::/72

preferred lifetime 300, valid lifetime 54333

expires at Nov 09 2002 02:03 AM (54258 seconds)

Prefix: 3FFE:C00:C18:3::/72

preferred lifetime 280, valid lifetime 51111

Reference. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/xs-3s/dhcp-xe-3s-book/ip6-dhcp-prefix-xe.html

QUESTION 464

Which two statements about NPTv6 are true? (Choose two.)

- A. The translation is invisible to applications that hard code IP information within the application logic.
- B. It is a one-way stateful translation for the IPv6 address.
- C. Translation is 1:1 at the network layer.
- D. It is a two-way stateless translation for the network prefix.

Correct Answer: CD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

This document describes a stateless, transport-agnostic IPv6-to-IPv6 Network Prefix Translation (NPTv6) function that provides the address-independence benefit associated with IPv4-to-IPv4 NAT (NAPT44) and provides a 1:1 relationship between addresses in the "inside" and "outside" prefixes, preserving end-to-end reachability at the network layer.

NPTv6 Translation is stateless, so a "reset" or brief outage of an NPTv6 Translator does not break connections that traverse the translation function, and if multiple NPTv6 Translators exist between the same two networks, the load can shift or be dynamically load shared among them. NPTv6 is defined

to include a two-way, checksum-neutral, algorithmic translation function, and nothing else.

Reference. <https://tools.ietf.org/html/rfc6296>

QUESTION 465

Which three protocols can use enhanced object tracking? (Choose three.)

- A. HSRP
- B. Proxy-ARP
- C. VRRP
- D. GLBP
- E. NTP
- F. DHCP

Correct Answer: ACD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The Enhanced Object Tracking feature separates the tracking mechanism from HSRP and creates a separate standalone tracking process that can be used by other processes and HSRP. This feature allows tracking of other objects in addition to the interface line-protocol state.

A client process such as HSRP, Virtual Router Redundancy Protocol (VRRP), or Gateway Load Balancing Protocol (GLBP), can register its interest in tracking objects and then be notified when the tracked object changes state.

Reference. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/15-mt/iap-15-mt-book/iap-eot.html>

QUESTION 466

What are the three primary components of NetFlow? (Choose three.)

- A. Flow caching
- B. A flow collector
- C. The data analyzer
- D. Flow sequence numbers
- E. Cisco Express Forwarding
- F. Multicast

Correct Answer: ABC

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

NetFlow includes three key components that perform the following capabilities:

- *Flow caching* analyzes and collects IP data flows entering router or switch interfaces and prepares data for export. It enables the accumulation of data on flows with unique characteristics, such as IP addresses, application, and CoS.
- *FlowCollector and Data Analysis* captures exported data from multiple routers and filters and aggregates the data according to customer policies, and then stores this summarized or aggregated data. Users can leverage Cisco NetFlow collector as a flow collector, or they can opt for a variety of third-party partner products. A Graphical user interface displays and analyzes NetFlow data collected from FlowCollector files. This allows users to complete near-real-time visualization or trending analysis of recorded and aggregated flow data. Users can specify the router and aggregation scheme and desired time interval.

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

QUESTION 467

Which two options are actions that EEM can perform after detecting an event? (Choose two.)

- A. Place a port in err-disabled.
- B. Generate an SNMP trap.
- C. Reload the Cisco IOS Software.
- D. Send an SMS.

Correct Answer: BC

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

action snmp-trap

To specify the action of generating a Simple Network Management Protocol (SNMP) trap when an Embedded Event Manager (EEM) applet is triggered, use the **action snmp-trap** command in applet configuration mode.

action reload

To specify the action of reloading the Cisco IOS software when an Embedded Event Manager (EEM) applet is triggered, use the **action reload** command in applet configuration mode.

Reference. http://www.cisco.com/c/en/us/td/docs/ios/12_2s/feature/guide/fs_eem2.html

QUESTION 468

On which three options can Cisco PfR base its traffic routing? (Choose three.)

- A. Time of day
- B. An access list with permit or deny statements

- C. Load-balancing requirements
- D. Network performance
- E. User-defined link capacity thresholds
- F. Router IOS version

Correct Answer: CDE

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Key Advantages of using PfR for Load balancing:

- Utilization based load-balancing: PfR takes real-time link utilization into account when load balancing the links. This will ensure that a link will not go beyond a certain percentage of its maximum capacity (75% by default).
- Application Performance based Load Balancing: PfR does not randomly forward traffic through one link or another. It takes application performance requirements into consideration and then forwards the traffic through a link which meets the performance policy requirements. PfR also load balances the link at the same time.
- Bi-directional Solution: PfR is a bi-directional load balancing solution which influences outbound as well as in-bound traffic.
- Consolidated Centralized View: PfR offers consolidated and centralized view of the state of all external links in the network. At any given time, the network administrator can see the current link utilization (in kbps and percentage of its capacity), maximum link threshold, and the policies applied to the links in the network.

Reference. <http://docwiki.cisco.com/wiki/PfR:Solutions:InternetOutboundLoadBalancing>

QUESTION 469

DRAG DROP

Drag and drop the DSCP PHB on the left to the corresponding binary representation on the right.

Select and Place:

AF31	10111000
AF43	01010000
AF22	1110000
AF13	01101000
EF	00111000
CS7	10011000

Correct Answer:

	EF
	AF22
	CS7
	AF31
	AF13
	AF43

Section: Infrastructure Services
Explanation

Explanation/Reference:

QUESTION 470

Which two routing protocols are not directly supported by Cisco PfR route control, and rely on the Cisco PfR subfeature PIRO? (Choose two.)

- A. BGP
- B. EIGRP
- C. Static routing
- D. OSPF
- E. IS-IS

Correct Answer: DE

Section: Infrastructure Services**Explanation****Explanation/Reference:**

Explanation:

Protocol Independent Route Optimization (PIRO) introduced the ability of Performance Routing (PfR) to search for a parent route--an exact matching route, or a less specific route--in the IP Routing Information Base (RIB), allowing PfR to be deployed in any IP-routed environment including Interior Gateway Protocols (IGPs) such as OSPF and IS-IS.

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/pfr/configuration/guide/15_1/pfr_15_1_book/pfr- piro.html

QUESTION 471

Which two tasks are required for configuring SNMP to send traps on a Cisco IOS device? (Choose two.)

- A. Create access controls for an SNMP community.
- B. Configure SNMP notifications.
- C. Configure the SNMP agent.
- D. Configure SNMP status monitoring and troubleshooting.
- E. Configure SNMP server group names.
- F. Configure the SNMP server engine ID.

Correct Answer: AB

Section: Infrastructure Services**Explanation****Explanation/Reference:**

Explanation:

The best current practices recommend applying Access Control Lists (ACLs) to community strings and ensuring that the requests community strings are not identical to notifications community strings. Access lists provide further protection when used in combination with other protective measures.

This example sets up ACL to community string:

```
access-list 1 permit 1.1.1.1 snmp-server community string1 ro 1
```

SNMP Notifications

A key feature of SNMP is the ability to generate notifications from an SNMP agent. These notifications do not require that requests be sent from the SNMP manager. Unsolicited (asynchronous) notifications can be generated as *traps or inform requests*. Traps are messages alerting the SNMP manager to a condition on the network. Inform requests (informs) are traps that include a request for confirmation of receipt from the SNMP manager. Notifications can indicate improper user authentication, restarts, the closing of a connection, loss of connection to a neighbor router, or other significant events.

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/12_2/configfun/configuration/guide/ffun_c/fcf014.html# wp1007320

QUESTION 472

Which two statements about SNMP traps are true? (Choose two.)

- A. They are sent by an agent after a specified event.
- B. They are sent when solicited after a specified event.
- C. They are equivalent to a community string.
- D. They provide solicited data to the manager.
- E. They are sent by a management station to an agent.
- F. Vendor-specific traps can be configured.

Correct Answer: AF

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The SNMP agent contains MIB variables whose values the SNMP manager can request or change. A manager can get a value from an agent or store a value into the agent. The agent gathers data from the MIB, the repository for information about device parameters and network data. The agent can also respond to a manager's requests to get or set data.

An agent can send unsolicited traps to the manager. Traps are messages alerting the SNMP manager to a condition on the network. Traps can mean improper user authentication, restarts, link status (up or down), MAC address tracking, closing of a TCP connection, loss of connection to a neighbor, or other significant events.

Reference: http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960/software/release/12-2_55_se/configuration/guide/scg_2960/swnmp.html

QUESTION 473

A configuration includes the line `ip nbar port-map SSH tcp 22 23 443 8080`. Which option describes the effect of this configuration line?

- A. It configures NBAR to search for SSH using ports 22, 23, 443, and 8080.
- B. It configures NBAR to allow SSH connections only on ports 22, 23, 443, and 8080.
- C. It enables NBAR to inspect for SSH connections.
- D. It creates a custom NBAR port-map named SSH and associates TCP ports 22, 23, 443, and 8080 to itself.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The `ip nbar-port-map` command configures NBAR to search for a protocol or protocol name using a port number other than the well-known port.

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/12_2/qos/command/reference/fqos_r/qrfcmd10.pdf

QUESTION 474

Which configuration sets a minimum quality of service on a Layer 2 access switch?

- A. mls qos cos override
mls qos cos 2
- B. mls qos cos 2
- C. mls qos trust cos
mls qos cos 2
- D. mls qos trust cos
- E. mls qos trust dscp

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The mls qos cos override interface command must be used to ensure that untrusted CoS values are explicitly set 0 (default).

Reference:

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book/QoSDesign.html

QUESTION 475

Refer to the exhibit.

```
R2#show run | include ntp
ntp server 10.1.1.1 prefer
ntp server 10.3.3.3
ntp server 10.4.4.4
```

If the route to 10.1.1.1 is removed from the R2 routing table, which server becomes the master NTP server?

- A. R2
- B. the NTP server at 10.3.3.3
- C. the NTP server at 10.4.4.4

D. the NTP server with the lowest stratum number

Correct Answer: D

Section: Infrastructure Services

Explanation

Explanation/Reference:


Explanation:

NTP uses a concept called "stratum" that defines how many NTP hops away a device is from an authoritative time source. For example, a device with stratum 1 is a very accurate device and might have an atomic clock attached to it. Another NTP server that is using this stratum 1 server to sync its own time would be a stratum 2 device because it's one NTP hop further away from the source. When you configure multiple NTP servers, the client will prefer the NTP server with the lowest stratum value.

Reference: <https://networklessons.com/network-services/cisco-network-time-protocol-ntp/>

QUESTION 476

Refer to the exhibit.



Which feature can R1 use to fail over from R2 to R3 if the address for R2 becomes unavailable?

- A. object tracking
- B. HSRP
- C. GLBP
- D. LACP

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The object tracking feature allows you to create a tracked object that multiple clients can use to modify the client behavior when a tracked object changes. Several clients register their interest with the tracking process, track the same object, and take different actions when the object state changes.

Clients include the following features:

- · Embedded Event Manager (EEM)
- · Gateway Load Balancing Protocol (GLBP)
- · Hot Standby Redundancy Protocol (HSRP)
- · Virtual port channel (vPC)
- · Virtual Router Redundancy Protocol (VRRP)

The object tracking monitors the status of the tracked objects and communicates any changes made to interested clients. Each tracked object is identified by a unique number that clients can use to configure the action to take when a tracked object changes state.

Reference: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/unicast/configuration/guide/l3_cli_nxos/l3_object.html

QUESTION 477

Refer to the exhibit.

```
Router1#show run
interface FastEthernet0/0
  ip address 10.20.10.1 255.255.255.0
  ip route-cache flow

ip flow-export version 5 origin-as
ip flow-export destination 209.165.200.227 49152
```

Which two options are effects of the given configuration? (Choose two.)

- A. It sets the data export destination to 209.165.200.227 on UDP port 49152.
- B. It enables Cisco Express Forwarding on interface FastEthernet0/0.
- C. It configures the export process to include the BGP peer AS of the router gathering the data.
- D. It enables NetFlow switching on interface FastEthernet0/0.
- E. It sets the data export destination to 209.165.200.227 on TCP port 49152.

Correct Answer: AD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The "ip flow-export destination 209.165.200.227 49152" command specifies that the data export destination server is 209.165.200.227 using UDP port 49152.

The "ip route-cache flow" command under the fastethernet 0/0 interface enable netflow switching on that interface.

QUESTION 478

Which three options are components of an EEM CLI policy? (Choose three.)

- A. Safe-Tcl
- B. applet name
- C. Fast Tcl
- D. event
- E. action
- F. Tcl bytecode

Correct Answer: BDE

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The Embedded Event Manager (EEM) monitors events that occur on your device and takes action to recover or troubleshoot these events, based on your configuration.

EEM consists of three major components:

Event statements -- Events to monitor from another Cisco NX-OS component that might require some action, workaround, or notification.

Action statements -- An action that EEM can take, such as sending an e-mail, or disabling an interface, to recover from an event.

Policies -- An applet name paired with one or more actions to troubleshoot or recover from the event.

Reference:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5500/sw/system_management/6x/b_5500_System_Mgmt_Config_6x/b_5500_System_Mgmt_Config_6x_chapter_010011.html

QUESTION 479

Which two statements best describe the difference between active mode monitoring and passive mode monitoring? (Choose two.)

- A. Passive mode monitoring uses IP SLA to generate probes for the purpose of obtaining information regarding the characteristics of the WAN links.

- B. Active mode monitoring is the act of Cisco PfR gathering information on user packets assembled into flows by NetFlow.
- C. Active mode monitoring uses IP SLA probes for obtaining performance characteristics of the current exit WAN link.
- D. Passive mode monitoring uses NetFlow for obtaining performance characteristics of the exit WAN links.

Correct Answer: CD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Passive and Active Monitoring

Passive monitoring is the act of OER gathering information on user packets assembled into flows by NetFlow. OER, when enabled, automatically enables NetFlow on the managed interfaces on the border routers. By aggregating this information on the border routers and periodically reporting the collected data to the master controller, the network prefixes and applications in use can automatically be learned. Additionally, attributes like throughput, reachability, loading, packet loss, and latency can be deduced from the collected flows.

Active monitoring is the act of generating IP SLA probes to generate test traffic for the purpose of obtaining information regarding the characteristics of the WAN links. Active probes can either be implicitly generated by OER when passive monitoring has identified destination hosts, or explicitly configured by the network manager in the OER configuration.

Reference:

http://products.mcisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/Transport_diversity/Transport_Diversity_PfR.html#wp199209

QUESTION 480

Which option is a core event publisher for EEM?

- A. Timer
- B. Policy Director
- C. Applet
- D. Script

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

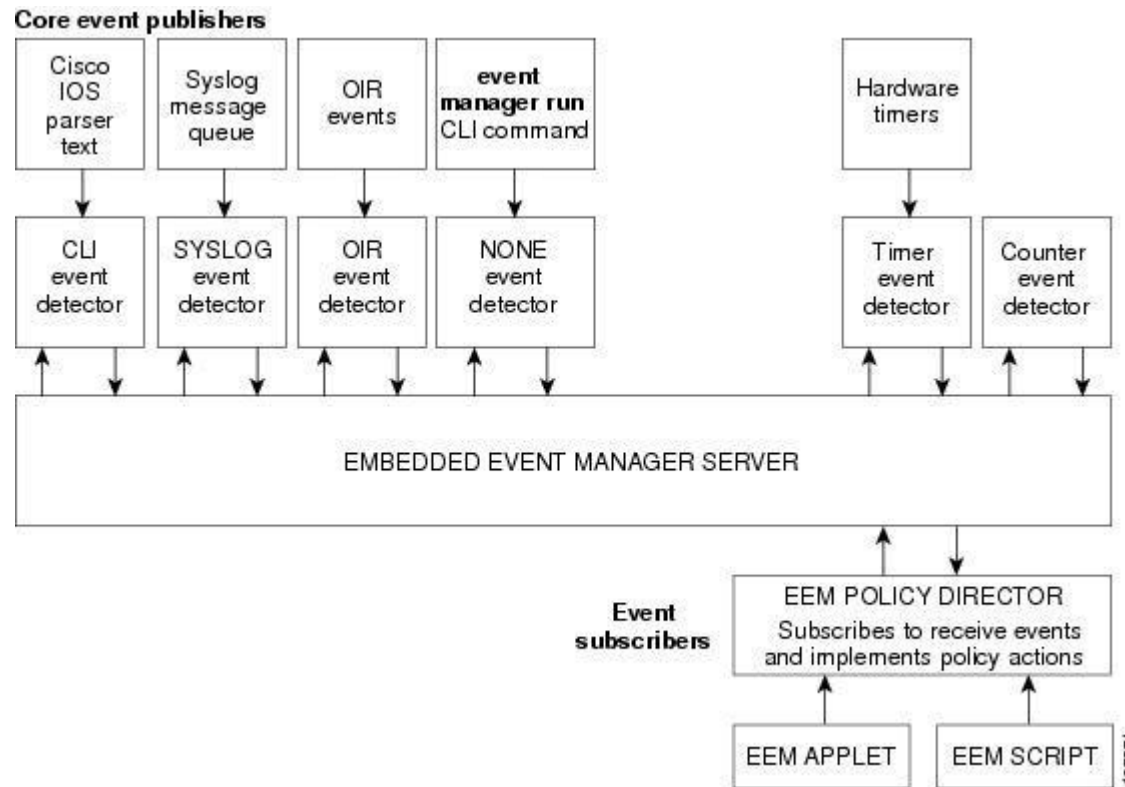
Explanation:

EEM is a flexible, policy-driven framework that supports in-box monitoring of different components of the system with the help of software agents known as event detectors. The figure below shows the relationship between the EEM server, core event publishers (event detectors), and the event subscribers (policies). Basically, event publishers screen events and publish them when there is a match on an event specification that is provided by the event subscriber. Event detectors notify the EEM server when an event of interest occurs. The EEM policies that are configured using the Cisco command-line interface (CLI) then implement recovery on the basis of the current state of the system and the actions specified in the policy for the given event.

EEM offers the ability to monitor events and take informational or corrective action when the monitored events occur or when a threshold is reached. An

EEM policy is an entity that defines an event and the actions to be taken when that event occurs. There are two types of EEM policies: an applet or a script. An applet is a simple form of policy that is defined within the CLI configuration. A script is a form of policy that is written in Tool Command Language (Tcl).

Figure 1. Embedded Event Manager Core Event Detectors



Reference: <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/eem/configuration/15-mt/eem-15-mt-book/eem-overview.html>

QUESTION 481

Which three statements about implementing an application layer gateway in a network are true? (Choose three.)

- A. It allows client applications to use dynamic ports to communicate with a server regardless of whether NAT is being used.
- B. It maintains granular security over application-specific data.
- C. It allows synchronization between multiple streams of data between two hosts.
- D. Application layer gateway is used only in VoIP/SIP deployments.

- E. Client applications require additional configuration to use an application layer gateway.
- F. An application layer gateway inspects only the first 64 bytes of a packet before forwarding it through the network.

Correct Answer: ABC

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:


An ALG may offer the following functions:

- allowing client applications to use dynamic ephemeral TCP/ UDP ports to communicate with the known ports used by the server applications, even though a firewall configuration may allow only a limited number of known ports. In the absence of an ALG, either the ports would get blocked or the network administrator would need to explicitly open up a large number of ports in the firewall -- rendering the network vulnerable to attacks on those ports.
- converting the network layer address information found inside an application payload between the addresses acceptable by the hosts on either side of the firewall/NAT. This aspect introduces the term 'gateway' for an ALG.
- recognizing application-specific commands and offering granular security controls over them
- synchronizing between multiple streams/sessions of data between two hosts exchanging data. For example, an FTP application may use separate connections for passing control commands and for exchanging data between the client and a remote server. During large file transfers, the control connection may remain idle. An ALG can prevent the control connection getting timed out by network devices before the lengthy file transfer completes.

Reference: http://en.wikipedia.org/wiki/Application-level_gateway

QUESTION 482

Refer to the exhibit.



You are configuring the S1 switch for the switch port that connects to the client computer. Which configuration blocks users on the port from using more than 6 Mbps of traffic and marks the traffic for a class of service of 1?

- A.
- ```
class-map match-all cos1
 match any
policy-map cos1
 class cos1
 set cos1
 police cir 6000000 bc 1125000 be 2250000 conform-action
 set-dscp-transmit cs1 exceed-action drop
 violate-action drop
```
- B.
- ```
class-map match-any cos1
  match any
policy-map cos1
  class cos1
    police cir 6000000 bc 1125000 be 2250000 conform-action
      set-dscp-transmit cs1 exceed-action drop
      violate-action drop
```
- C.
- ```
class-map match-all cos1
 match any
policy-map cos1
 class cos1
 set cos1
 policy cir 6000000 conform-action set-dscp-transmit cs1
 exceed-action permit violate-action permit
```
- D.
- ```
class-map match-any cos1
  match any
policy-map cos1
  class cos1
    set cos1
    policy cir 6000000 conform-action transmit exceed-action
      permit violate-action drop
```

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Only option A specified that the exceed and violate actions are set to drop for traffic over the CIR of 6 Mbps, and is also configured to set all traffic with a COS of 1 using the "set cos1" command.

QUESTION 483

Which EIGRP packet types are sent as unicast packets?

- A. hello, update, query
- B. query, SIA query, reply
- C. SIA query, reply, ACK
- D. query, SIA query, SIA reply

Correct Answer: C

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 484

What is a reason for an EIGRP router to send an SIA reply to a peer?

- A. to respond to an SIA query with the alternative path requested
- B. to respond to a query reporting that the prefix has gone stuck-in-active
- C. to respond to an SIA query that the router is still waiting on replies from its peers
- D. to respond to a reply reporting that the prefix has gone stuck-in-active

Correct Answer: C


Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 485

Refer to the exhibit.



Which prefixes will appear in the EIGRP topology table?

- A. 10.0.0.0/8, 172.16.1.0/24, 192.168.0.0/16
- B. 10.1.1.0/24, 10.1.2.0/24, 172.16.1.0/26, 192.168.1.0/26, 192.168.2.0/26
- C. 10.1.1.0/24, 10.1.2.0/24, 172.16.1.0/26, 172.16.2.0/26, 192.168.1.0/26, 192.168.2.0/26
- D. 10.1.1.1/24, 10.1.2.1/24, 172.16.1.1/26, 172, 192.168.1.1/26, 192.168.2.1/26

Correct Answer: B

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 486

What is the most common use for route tagging in EIGRP?

- A. to determine the route source for management purposes

- B. to change the metric of a prefix
- C. to filter routes in order to prevent routing loops
- D. to modify path selection for certain classes of traffic

Correct Answer: C

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 487

Refer to the exhibit.

```
RP1#show ipv6 mroute
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF7E:200:2A02:B111:FC02:1:11FF:11EE), 00:00:35/never, RP
2A02:B111:FC02:1::1, flags: SCLJ
  Incoming interface: Null
  RPF nbr: ::
  Immediate Outgoing interface list:
    Ethernet2/0, Forward, 00:00:35/never
```

Which statement is true?

- A. The output shows an IPv6 multicast address with link-local scope.
- B. The output shows an IPv6 multicast address that is used for unique local sources only.
- C. The output shows an IPv6 multicast address that can be used for BIDIR-PIM only.
- D. The output shows an IPv6 multicast address with embedded RP.

Correct Answer: D

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 488

Which two statements about the max-age time in IS-IS are true? (Choose two.)

- A. The IS-IS max-age time is 20 minutes by default.
- B. The IS-IS max-age time is 60 minutes by default.
- C. The IS-IS max-age time increments from zero to max-age.
- D. The IS-IS max-age time decrements from max-age to zero.

Correct Answer: AD

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 489

Which two statements about the default behavior of IS-IS are true? (Choose two.)

- A. The default IS-IS router type is L1/L2.
- B. The default IS-IS metric type is wide.
- C. The default IS-IS interface circuit type is L1/L2.
- D. By default, two IS-IS routers must use the same hello interval and hold timer in order to become neighbors.

Correct Answer: AC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 490

Which two statements about BPDU guard are true? (Choose two.)

- A. The global configuration command spanning-tree portfast bpduguard default shuts down interfaces that are in the PortFast-operational state when a

BPDU is received on that port.

- B. The interface configuration command spanning-tree portfast bpduguard enable shuts down only interfaces with PortFast enabled when a BPDU is received.
- C. BPDU guard can be used to prevent an access port from participating in the spanning tree in the service provider environment.
- D. BPDU guard can be used to protect the root port.
- E. BPDU guard can be used to prevent an invalid BPDU from propagating throughout the network.

Correct Answer: AC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 491

Which two 802.1D port states are expected in a stable Layer 2 network? (Choose two.)

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

Correct Answer: AD

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 492

Which three fields are part of a TCN BPDU? (Choose three.)

- A. protocol ID
- B. version
- C. type
- D. max-age
- E. flags
- F. message age

Correct Answer: ABC
Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 493

With AutoInstall, which mechanism allows for automatic addressing of the serial interface using HDLC?

- A. ARP
- B. BOOTP
- C. DHCP
- D. SLARP

Correct Answer: D
Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 494

Which two protocols does the Management Plane Protection feature support? (Choose two.)


- A. ARP
- B. HTTPS
- C. TFTP
- D. OSPF

Correct Answer: BC
Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 495

Refer to the exhibit.



Which configuration reduces CPU utilization on R2 while still advertising the connected routes of R2 to R1?

- A. Configure eigrp stub connected on R2.
- B. Configure eigrp stub receive-only on R1.
- C. Configure eigrp stub static on R2.
- D. Configure eigrp stub summary on R1.

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 496

Which authentication types does OSPF support?

- A. null and clear text
- B. MD5 only
- C. MD5 and clear text
- D. null, clear text, and MD5
- E. clear text only

Correct Answer: D

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 497

Which ICMP message type is used to assist path MTU discovery?

- A. destination unreachable
- B. redirect message
- C. source quench
- D. time exceeded

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 498

A configuration includes the line `ip route 10.0.0.0 255.0.0.0 172.16.10.10 permanent`. Which option is a benefit of configuring this static route as permanent?

- A. It allows the route to be redistributed into the network even if the outgoing interface is down.
- B. It allows the route to be saved in the running configuration of the device.
- C. It places a hidden tag on the route that can be matched on other devices.
- D. It allows the route to have a tracking status even if no tracking object is configured.

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 499

Refer to the exhibit.

```
R1#sh run | i mpls

mpls ldp session protection
mpls ldp discovery targeted-hello accept
no mpls ldp advertise-labels
mpls ldp advertise-labels for LOOPBACK-ONLY
no mpls ip propagate-ttl forwarded
mpls label protocol ldp
```

Which two statements about the R1 configuration are true? (Choose two.)

- A. The IP TTL value is copied to the MPLS field during label imposition.
- B. The structure of the MPLS network is hidden in a traceroute.
- C. The LDP session interval and hold times are configured for directly connected neighbors.
- D. R1 protects the session for 86400 seconds.
- E. All locally assigned labels are discarded.

Correct Answer: BD

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 500

Which two statements about IPsec VTI implementation are true? (Choose two.)

- A. The IKE SA can be bound to the VTI and the crypto map.
- B. The transform set can be configured only in tunnel mode.
- C. SVTIs support only a single IPsec SA.
- D. SVTIs support IPv4 packets that carry IPv6 packets.

Correct Answer: BC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 501

Which command sets the maximum segment size for a TCP packet initiated from a router?

- A. ip mtu
- B. ip tcp adjust-mss
- C. ip tcp mss
- D. ip tcp window-size

Correct Answer: C

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 502

Which circumstance can cause TCP starvation and UDP dominance to occur?

- A. Too few queues are available.
- B. UDP is comprised of smaller packets than TCP.
- C. Retransmitted TCP packets are on the network.
- D. UDP and TCP data are assigned to the same service-provider class.

Correct Answer: D


Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 503

Refer to the exhibit.



Which statement about the topology is true?

- A. It provides a transparent LAN service.
- B. It provides only point-to-multipoint connections between UNIs.
- C. It uses port-based connections at the hub.
- D. It provides point-to-point connections between UNIs.

Correct Answer: D

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 504


Which two statements about reverse ARP are true? (Choose two.)

- A. Its servers require static mappings.
- B. It works with AutoInstall to configure new devices.
- C. It provides IP addresses for subnet masks.
- D. It provides IP addresses for default gateways.
- E. It requires less maintenance than DHCP.

Correct Answer: AB
Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 505
Refer to the exhibit.



This network is configured with PIM, and the RPF check has failed toward the multicast source. Which two configuration changes must you make to router R3 to enable the RPF check to pass? (Choose two.)

- A. Configure a static multicast route to the multicast source through the tunnel interface.
- B. Configure a static multicast route to the multicast source LAN through the tunnel interface.
- C. Configure a static multicast route to the multicast source LAN through the Ethernet interface.
- D. Remove the command `ip pim bidir-enable` from the R3 configuration.

Correct Answer: AB
Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 506

In which two situations is an EIGRP hello packet sent as unicast? (Choose two.)

- A. during neighbor discovery
- B. when link costs change
- C. when the neighbor command is used
- D. when an ACK is sent

Correct Answer: CD

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 507

Which three options are results of the command no mpls ip propagate-ttl? (Choose three.)

- A. It prevents the TTL from being copied from the IP header to the MPLS header.
- B. It prevents the MPLS hops from being visible to a CE router when you perform a traceroute.
- C. A fixed TTL value of 255 is used for the first label of the IP packet.
- D. It prevents the TTL from being copied from the MPLS header back into the IP header.
- E. MPLS hops remain visible on a CE router when you perform a traceroute.
- F. A fixed TTL value of 1 is used for the first label of the IP packet.

Correct Answer: ABC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 508

Which statement about how a CE router is used in an MPLS VPN is true?

- A. It is located on the customer premises, where it peers and exchanges routes with the provider edge router.

- B. It is located on the provider premises, where it peers and exchanges routes with the customer edge router.
- C. It is located on the customer premises, but it is fully controlled by the provider, which provides a full routing table to the customer.
- D. It is located on the provider premises, and it routes only MPLS label traffic.

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 509

Which three options are three benefits of an MPLS VPN? (Choose three.)

- A. It allows IP address space overlap by maintaining customer routes in a private routing table.
- B. It offers additional security by preventing intrusions directly into the customer routing table.
- C. It offers a transparent virtual network in which all customer sites appear on one LAN.
- D. It offers additional security by allowing only dynamic routing protocols between CE and PE routers.
- E. It allows IP address space overlap by maintaining customer routes in the global routing table with unique BGP communities.
- F. Providers can send only a default route for Internet access into the customer VPN.

Correct Answer: ABC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 510

Into which two pieces of information does the LISP protocol split the device identity? (Choose two.)

- A. Routing Locator
- B. Endpoint Identifier
- C. Resource Location
- D. Enterprise Identifier
- E. LISP ID
- F. Device ID

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 511

Which two protocols are used to establish IPv6 connectivity over an MPLS network? (Choose two.)

- A. 6PE
- B. 6VPE
- C. RSVP
- D. ISATAP
- E. LDP
- F. IPv6IP

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 512

Which three types of traffic are protected when you implement IPsec within an IPv6-in-IPv4 tunnel? (Choose three.)

- A. IPv6 link-local traffic
- B. IPv6 multicast traffic
- C. IPv6 unicast traffic
- D. IPv4 tunnel control traffic
- E. IPv4 broadcast traffic
- F. IPv6 broadcast traffic

Correct Answer: ABC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 513

Which three features does GETVPN support to improve deployment and scalability? (Choose three.)

- A. configuration of multiple key servers to work cooperatively
- B. allowing traffic to be discarded until a group member registers successfully
- C. local exceptions in the traffic classification ACL
- D. GDOI protocol configuration between group members and the key server
- E. redundant IPsec tunnels between group members and the key server
- F. redundant multicast replication streaming through the use of a bypass tunnel

Correct Answer: ABC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 514

Refer to the exhibit.

```
aaa new-model
aaa authentication login default group tacacs+ local
aaa authentication enable default group tacacs+ enable
aaa authorization console
aaa authorization exec default group tacacs+ if-authenticated
aaa authorization commands 1 default group tacacs+ local if-authenticated
aaa authorization commands 4 default group tacacs+ if-authenticated
aaa authorization commands 15 default group tacacs+ local if-authenticated
aaa accounting exec default start-stop group tacacs+
aaa accounting commands 1 default start-stop group tacacs+
aaa accounting commands 15 default start-stop group tacacs+
aaa session-id common

line con 0
  exec-timeout 120 0
  logging synchronous
  transport input ssh telnet
  password cisco

line vty 0 4
  exec-timeout 120 0
  logging synchronous
  transport input ssh telnet
  password cisco
```

Which two configuration changes enable you to log in to the router? (Choose two.)

- A. Configure a user name and password on the device.
- B. Modify the default login authentication group to use the terminal line password.
- C. Remove the terminal line password on the console line.
- D. Modify the terminal lines to include transport input none.
- E. Configure the terminal lines to use the local user database.

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 515

Refer to the exhibit.

```
R1#show logging
Syslog logging: enabled (0 messages dropped, 1 messages rate-limited, 0 flushes, 0 overruns, xml disabled, filtering disabled)

No Active Message Discriminator.

No Inactive Message Discriminator.

  Console logging: level informational, 47 messages logged, xml disabled, filtering disabled
  Monitor logging: level debugging, 0 messages logged, xml disabled, filtering disabled
  Buffer logging:  level debugging, 47 messages logged, xml disabled, filtering disabled
  Exception Logging: size (8192 bytes)
  Count and timestamp logging messages: disabled
  Persistent logging: disabled

No active filter modules.

  Trap logging: level informational, 51 message lines logged

Log Buffer (4096 bytes):
```

Which log levels are enabled for the console?

- A. informational only
- B. informational and debugging
- C. informational, debugging, notifications, warnings, errors, critical, alerts, and emergencies
- D. informational, notifications, warnings, errors, critical, alerts, and emergencies

Correct Answer: D

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 516

External EIGRP route exchange on routers R1 and R2 was failing because the routers had duplicate router IDs. You changed the eigrp router-id

command on R1, but the problem persists. Which additional action must you take to enable the routers to exchange routes?

- A. Change the corresponding loopback address.
- B. Change the router ID on R2.
- C. Reset the EIGRP neighbor relationship.
- D. Clear the EIGRP process.

Correct Answer: D

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 517

Which two BGP path attributes are visible in Wireshark? (Choose two.)

- A. weight
- B. AS path
- C. local preference
- D. route maps

Correct Answer: BC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 518

Refer to the exhibit.

```
%CFIB-7-CFIB_EXCEPTION: FIB TCAM exception, Some entries will be software switched
```

If a Layer 3 switch running OSPF in a VRF-lite configuration reports this error, which action can you take to correct the problem?

- A. Set mls cef maximum-routes in the global configuration.

- B. Add the vrf-lite capability to the OSPF configuration.
- C. Upgrade the Layer 3 switch to a model that can support more routes.
- D. Configure the control plane with a larger memory allocation to support the Cisco Express Forwarding Information Base.

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 519

Which statement describes the effect of the configuration line redistribute maximum-prefix 1500 90 withdraw?

- A. After the 1500th route is redistributed, a warning is posted in the log file and 90 more routes are redistributed before further routes are discarded.
- B. After the 1350th route is redistributed, a warning is posted in the log file until the 1500th route is redistributed, and then further routes are discarded.
- C. After the 1500th route is redistributed, further routes are discarded only if the CPU is above 90%.
- D. The routing protocol receives 1500 routes. After the routing process has redistributed 90% of the routes, the process supernets routes and injects a NULL route to prevent black-hole routing.

Correct Answer: B

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 520

Which option is a correct match criterion for policy-based routing?

- A. length
- B. interface type
- C. interface
- D. cost

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 521

Which EIGRP configuration results in subsecond failover outside of the basic routing protocol convergence?

- A. bfd all-interfaces
- B. timers active-time disabled
- C. timers active-time 1
- D. timers nsf route-hold 20

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 522

Which two statements about BGP loop prevention are true? (Choose two.)

- A. Advertisements from PE routers with per-neighbor SOO configured include a Site of Origin value that is equal to the configured value of the BGP peering.
- B. If the configured Site of Origin value of a BGP peering is equal to the Site of Origin value on a route it receives, route advertisement is blocked to prevent a route loop.
- C. AS-override aids BGP loop prevention, but alternate loop prevention mechanisms are also necessary.
- D. Advertisements from the neighbors a BGP peering include a Site of Origin value that is separate from the configured value of the BGP peering.
- E. If the configured Site of Origin value of a BGP peering is greater than the Site of Origin value on a route it receives, route advertisement is blocked to prevent a route loop.
- F. If the configured Site of Origin value of a BGP peering is equal to the Site of Origin value on a route it receives, route advertisement is permitted.

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 523

Which option is the default point of insertion for the BGP cost community?

- A. before best path calculation
- B. after best path calculation
- C. after the IGP metric comparison
- D. after the router ID comparison

Correct Answer: C


Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 524

Refer to the exhibit.



Which BGP feature allows R1 to send R2 a list of prefixes that R2 is prevented from advertising to R1?

- A. route refresh
- B. Prefix-Based Outbound Route Filtering
- C. distribute lists
- D. prefix lists

Correct Answer: B

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 525

Refer to the exhibit.

```
R1#show ip bgp neighbor 192.168.1.1

BGP neighbor is 192.168.1.1, remote AS 64496, external link
BGP version 4, remote router ID 192.168.1.1
Neighbor under common administration
BGP state = Established, up for 00:00:10
Last read 00:00:10, last write 00:00:10, hold time is 180, keep alive interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received(old & new)
  Address family IPv4 Unicast: advertised and received
Message statistics:
  InQ depth is 0
  OutQ depth is 0

      Sent      Rcvd
Opens:          1          1
Notifications:  0          0
Updates:        0          0
Keepalives:     1          1
Route Refresh:  0          0
Total:          2          2
Default minimum time between advertisement runs is 30 seconds
```

Which type of BGP peer is 192.168.1.1?

- A. route reflector client
- B. iBGP
- C. confederation
- D. VPNv4

Correct Answer: C

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 526

Which option describes what the default RT filter indicates when you implement the BGP RT constrained route distribution feature?

- A. A peer receives only a default route for each VRF.
- B. A peer receives all routes, regardless of the RT value.
- C. A peer receives routes only for RTs that are used on that router.
- D. A peer receives no routes, regardless of the RT value.

Correct Answer: B

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 527

Refer to the exhibit.

```
SW1# show interface | include Vlan10 is  
interface Vlan10 is down, line protocol is down
```

Which two issues can cause the interface VLAN10 to be down/down? (Choose two.)

- A. The VLAN is inactive or has been removed from the VLAN database.
- B. STP is in a forwarding state on the port.
- C. A Layer 2 access port is configured with VLAN10, but is in a down/down state.
- D. The autostate exclude feature was used on interface VLAN10.

Correct Answer: AC


Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 528

Refer to the exhibit.



SW1

```
vlan 10
vlan 200
```

```
interface vlan 10
  ip address 172.29.1.1 255.255.255.0

interface fa0/21
  switchport trunk encapsulation dot1q
  switchport mode trunk
  switchport trunk allowed vlan add 200
```

```
interface fa0/22
  switchport trunk encapsulation dot1q
  switchport mode trunk
  switchport trunk allowed vlan add 200
```

SW2

```
vlan 10
vlan 200
```

```
interface vlan 10
  ip address 172.29.1.2 255.255.255.0

interface fa0/21
  switchport trunk encapsulation dot1q
  switchport mode trunk
  switchport trunk allowed vlan add 200, 10
```

```
interface fa0/22
  switchport trunk encapsulation dot1q
  switchport mode trunk
  switchport trunk allowed vlan add 200, 10
```

Which two statements about this configuration are true? (Choose two.)

- A. Pings from SW2 to SW1 fail because SW1 is pruning VLAN 10.
- B. VLANs 10 and 200 are added to the SW2 allowed list on interface fa0/22.
- C. Pings from SW2 to SW1 are successful.
- D. Only VLAN 200 is added to the SW1 allowed list on interface fa0/22.

Correct Answer: BC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 529

Refer to the exhibit.

```
%PM-4-ERR_DISABLE: channel-misconfig error detected on Po1, putting fa0/12 in err-disable state
```

Which two conditions can cause this error message to be displayed on the console? (Choose two.)

- A. The EtherChannel is configured as desirable on both ends.
- B. The port-channel on the adjacent device is misconfigured.
- C. There is a speed and duplex mismatch on interface fa0/12.
- D. The EtherChannel is configured as auto on one of the interfaces.

Correct Answer: BC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 530

Which three statements about RIPvng are true? (Choose three.)

- A. It supports route tags.
- B. It sends updates on FF02::9.
- C. Its RTE last byte is 0XFF.
- D. It supports authentication.
- E. It sends updates on UDP port 520.
- F. It can be used on networks of greater than 15 hops.

Correct Answer: ABC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 531

Which option is the result if two adjacent routers are configured for OSPF with different process IDs?

- A. The routers are unable to establish an adjacency.
- B. The routers establish an adjacency, but route exchange fails.
- C. The routers establish an adjacency and exchange routes, but the routes are unreachable.
- D. The routers establish an adjacency and exchange routes, and the routes are reachable.

Correct Answer: D

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 532

Which two commands enable OSPF graceful shutdown? (Choose two.)

- A. nsf cisco
- B. ip ospf shutdown
- C. shutdown
- D. nsf ietf helper disable

Correct Answer: BC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 533

Which object tracking function tracks the combined states of multiple objects?

- A. application
- B. interface
- C. stub-object
- D. list

Correct Answer: D

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 534

Which two options are EEM policies? (Choose two.)

- A. applets
- B. event detectors
- C. scripts
- D. syslogs
- E. actions

Correct Answer: AC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 535

Which two metrics are measured with active probes when PfR voice traffic optimization is in use? (Choose two.)

- A. MOS
- B. cost
- C. jitter
- D. bandwidth

Correct Answer: AC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 536

Which two actions can you take to recover an interface in a errdisable state? (Choose two.)

- A. Enable UDLD on the switch.
- B. Enable errdisable recovery on the switch.

- C. Execute the shutdown command on the interface, followed by the no shutdown command.
- D. Remove the related commands from the configuration and reenter them.
- E. Enable loop guard on the switch.

Correct Answer: BC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 537

Which three protocols support SSM? (Choose three.)

- A. IGMPv2
- B. IGMPv3
- C. IGMP v3lite
- D. URD
- E. CGMP
- F. IGMPv1

Correct Answer: BCD

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 538

Which additional feature must be enabled on a switch to allow PIM snooping to function correctly?

- A. IGMP snooping
- B. port security
- C. storm control
- D. dynamic ARP inspection

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 539

Which protocol uses a proprietary 2-byte Type field for multiple protocol support?

- A. HDLC
- B. PPP
- C. CHAP
- D. PAP

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 540

How many address families can a single OSPFv3 instance support?

- A. 1
- B. 2
- C. 5
- D. 10

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 541

Which two statements about LDP advertising when Explicit Null is in effect are true? (Choose two.)

- A. Penultimate hop popping is disabled.
- B. Penultimate hop popping is enabled.
- C. It is the default behavior for LDP.
- D. It is used for the advertisement of static routes.

E. It is used for the advertisement of connected routes.

Correct Answer: AE

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 542

You are configuring a DMVPN hub to perform CBWFQ on a per-spoke basis. Which information is used to identify the spoke?

- A. the NHRP network ID
- B. the spoke tunnel source IP
- C. the spoke tunnel interface IP address
- D. the NHRP group

Correct Answer: D

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 543

Which option is true about output policing for the control plane?

- A. It improves router performance by limiting traffic sent to the control plane.
- B. It improves router performance by limiting traffic sent from the control plane.
- C. It improves router performance by limiting traffic sent to and from the control plane.
- D. It controls traffic originated from the router.

Correct Answer: D

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 544

You are configuring a DHCPv6 client for a DHCPv6 server with the prefix delegation feature. Which option is a result of the interface configuration when

you enter the command `ipv6 address autoconfig default`?

- A. a static IPv6 default route pointing to the upstream DHCP server
- B. a static IPv6 default route pointing to the upstream DHCP relay
- C. a static IPv6 default route pointing to the upstream router
- D. a temporary stateless address, formed from the EUI-64 bit address and the prefix from the route advertisement of the upstream router

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 545

Refer to the exhibit.

```
!  
interface Loopback25  
 ip address 25.25.25.1 255.255.255.255  
!  
interface Ethernet0/0  
 ip address 192.168.12.1 255.255.255.252  
 mpls ip  
!
```

You are bringing a new MPLS router online and have configured only what is shown to bring LDP up. Assume that the peer has been configured in a similar manner. You verify the LDP peer state and see that there are no neighbors. What will the output of `show mpls ldp discovery` show?

- A. Interfaces:
Ethernet0/0 (ldp): xmit
- B. Interfaces:
Ethernet0/0 (ldp): xmit/recv
LDP Id: 25.25.25.2:0; IP addr: 192.168.12.2
- C. Interfaces:
Ethernet0/0 (ldp): xmit/recv
LDP Id: 192.168.12.2:0; no route
- D. Interfaces:
Ethernet0/0 (ldp): xmit/recv

LDP Id: 25.25.25.2:0; no route

Correct Answer: D

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 546

Which three features are common to OSPF and IS-IS? (Choose three.)

- A. They both maintain a link-state database from which a Dijkstra-based SPF algorithm computes a shortest path tree.
- B. They both use DR and BDR in the broadcast network.
- C. They both use hello packets to form and maintain adjacencies.
- D. They both use NSSA and stub type areas to scale the network design.
- E. They both have areas to form a two-level hierarchical topology.

Correct Answer: ACE

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 547

Which three parameters must match to establish OSPF neighbor adjacency? (Choose three.)

- A. the process ID
- B. the hello interval
- C. the subnet mask
- D. authentication
- E. the router ID
- F. the OSPF interface priority

Correct Answer: BCD

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 548

Which IP SLA operation type uses IP to measure the round-trip time between a router and a device?

- A. HTTP
- B. ICMP Echo
- C. ICMP Path Jitter
- D. UDP Jitter for VoIP

Correct Answer: B

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 549

What are two reasons to use the ip ospf database filter all out command? (Choose two.)

- A. to maintain a centralized OSPF database on a single master device
- B. to avoid flooding LSAs on low-speed links
- C. to ensure a consistent OSPF database across the network
- D. to selectively filter OSPF routes without disrupting the SPF algorithm
- E. to filter only type 7 LSAs from an OSPF area
- F. to enable OSPF to send triggered updates

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 550

On a broadcast interface, which two OSPF states support BFD sessions? (Choose two.)

- A. DR
- B. BDR
- C. DROTHER

- D. 2WAY
- E. FULL
- F. ACTIVE

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 551

Which two statements about BGP best-path selection are true? (Choose two.)

- A. The route with the highest local preference is preferred.
- B. The weight attribute is advertised to peers.
- C. The route with the lowest MED is preferred.
- D. A route that originates from iBGP peers is preferred.
- E. A route that originates from a router with a higher BGP router ID is preferred.
- F. The lowest weight advertised is preferred.

Correct Answer: AC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 552

The no ip unreachable command is configured on interfaces to protect the control plane of a router. Which mechanism is impacted by using this command?

- A. ICMP redirects
- B. path MTU discovery
- C. source routing
- D. ICMP router discovery protocol

Correct Answer: B


Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 553

Refer to the exhibit.



The customer wants to use IP SLA to create a failover to ISP2 when both Ethernet connections to ISP1 are down. The customer also requires that both connections to ISP1 are utilized during normal operations.

Which IP route configuration accomplishes these requirements for the customer?

- A. ip route 0.0.0.0 0.0.0.0 192.168.0.1 track 1
ip route 0.0.0.0 0.0.0.0 192.168.1.1 track 2
ip route 0.0.0.0 0.0.0.0 192.168.2.1 track 3
- B. ip route 0.0.0.0 0.0.0.0 192.168.0.1 track 1
ip route 0.0.0.0 0.0.0.0 192.168.1.1 track 2
ip route 0.0.0.0 0.0.0.0 192.168.2.1 track 4 100
- C. ip route 0.0.0.0 0.0.0.0 192.168.0.1 track 1
ip route 0.0.0.0 0.0.0.0 192.168.1.1 track 2
ip route 0.0.0.0 0.0.0.0 192.168.2.1 track 3 100
- D. ip route 0.0.0.0 0.0.0.0 192.168.0.1 track 1 1
ip route 0.0.0.0 0.0.0.0 192.168.1.1 track 2 2
ip route 0.0.0.0 0.0.0.0 192.168.2.1 track 3 3

Correct Answer: C

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 554

Which statement describes what it means if a router has an OSPF priority set to 0?

- A. A router with the OSPF priority set to 0 is one that can participate in the election of a DR. It has the highest priority.
- B. A router with the OSPF priority set to 0 is one that cannot participate in the election of a DR, but it can become a BDR
- C. A router with the OSPF priority set to 0 is one that cannot participate in the election of a DR. It can become neither a DR nor a BDR.
- D. A router with the OSPF priority set to 0 is one that cannot participate in the election of a BDR, but it can become a DR

Correct Answer: C

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 555

What is the maximum number of classes that MQC can support in a single policy map?

- A. 512
- B. 256
- C. 128
- D. 64

Correct Answer: B

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 556

DRAG DROP

Drag each IPv6 extension header on the left to its corresponding description on the right.

Select and Place:

AH	Specifies the path for a datagram.
Destination	Carries encrypted data.
ESP	Specifies the parameters used to split datagrams.
Fragment	Carries authentication information.
Hop-by-Hop	Specifies options to be examined only at the final device.
Routing	Specifies options to be examined by all devices.

Correct Answer:

Routing
ESP
Fragment
AH
Destination
Hop-by-Hop

Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 557

DRAG DROP

Drag each traceroute text character on the left to its meaning on the right.

Select and Place:

*	The port is unreachable.
?	The probe timed out.
A	The protocol is unreachable.
P	Unknown packet type.
Q	The destination is too busy.
U	Prohibited.

Correct Answer:

	U
	*
	P
	?
	Q
	A

Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 558

DRAG DROP

Drag and drop each IPv6 neighbor discovery message type on the left to the corresponding description on the right.

Select and Place:

neighbor redirect	The message a node uses to share its link-layer address
router solicitation	The message a node uses to notify hosts on the link of a better first-hop for a destination
router advertisement	The message a node uses to discover the link-local addresses of other nodes on the link
neighbor advertisement	The message a node uses to share information about its status and its local prefixes
neighbor solicitation	The message a host sends when it starts up, requesting local routers to transmit information

Correct Answer:

	neighbor advertisement
	neighbor redirect
	neighbor solicitation
	router advertisement
	router solicitation

Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 559

DRAG DROP

Drag and drop each BGP feature on the left to the corresponding function it performs on the right.

Select and Place:

peer session templates	Applies configuration commands to a group of neighbors
peer policy templates	Separates updates from configurations, allowing groups to
peer groups	Supports the configuration of a group of neighbors by defining
BGP Dynamic Update Peer-Groups	Applies configuration commands to a group of neighbors
BGP dynamic neighbors	Creates a group of neighbors in the same address family that

Correct Answer:

	peer session templates
	BGP Dynamic Update Peer-Groups
	BGP dynamic neighbors
	peer policy templates
	peer groups

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 560

DRAG DROP

Drag and drop each step of the Unicast RPF process on the left into the correct order on the right.

Select and Place:

Unicast RPF performs a CEF table lookup for packet forwarding.

1

Unicast RPF performs a reverse lookup of the return path in

2

The packet is received on the interface.

3

The output ACL is checked on the forwarding interface.

4

The packet is forwarded.

5

The packet is checked against the inbound ACL.

6

Correct Answer:

The packet is received on the interface.

The packet is checked against the inbound ACL.

Unicast RPF performs a reverse lookup of the return path in

Unicast RPF performs a CEF table lookup for packet forwarding.

The output ACL is checked on the forwarding interface.

The packet is forwarded.

Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 561

DRAG DROP

Drag and drop the SNMP element on the left to the corresponding definition on the right.

Select and Place:

Get	an inquiry for a MIB-leaf variable
GetNext	a reply to a request
GetBulk	a router request for a single variable's value
Response	a single inquiry for multiple, consecutive MIB variables

Correct Answer:

	GetNext
	Response
	Get
	GetBulk

Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 562

DRAG DROP

Drag and drop each PHB on the left to the functionality it performs on the right.

Select and Place:

Default	Provides low-latency, low-loss, low-jitter, and assured bandwidth service.
Expedited Forwarding	Provides backward compatibility with IP-precedence.
Assured Forwarding	Assigns best effort service to the packet.
Class Selector	Defines classes for traffic allocation

Correct Answer:

	Expedited Forwarding
	Class Selector
	Default
	Assured Forwarding

Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 563

DRAG DROP

Drag and drop each policy command on the left to the function it performs on the right.

Select and Place:

bandwidth	Enables CBWFQ.
random-detect	Configures the queuing of excess traffic for later transmission.
service-policy	Configures the dropping of excess traffic when a maximum rate
police	Enables WFQ.
shape	Enables a traffic policy on an interface.
fair-queue	Enabled WRED or DWRED.

Correct Answer:

	bandwidth
	shape
	police
	fair-queue
	service-policy
	random-detect

Section: Mix Questions**Explanation****Explanation/Reference:****QUESTION 564**

Which topology allows the split-horizon rule to be safely disabled when using EIGRP?

- A. full mesh
- B. partial mesh
- C. hub and spoke
- D. ring

Correct Answer: C

Section: Mix Questions**Explanation**

Explanation/Reference:

QUESTION 565

Which two issues is TCP Sequence Number Randomization designed to prevent? (Choose two.)

- A. DDOS attacks
- B. OS fingerprinting
- C. man-in-the-middle attacks
- D. ARP poisoning
- E. Smurf attack

Correct Answer: BC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 566

Which value is the maximum segment size if you start with an MTU of 1500 bytes and then remove the overhead of the Ethernet header, IP header, TCP header, and the MAC frame check sequence?

- A. 1434 bytes
- B. 1460 bytes
- C. 1458 bytes
- D. 1464 bytes

Correct Answer: B

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 567

Which three improvements does Cisco IOS XE Software offer over traditional IOS Software? (Choose three.)

- A. It can run applications as separate processes on multicore CPUs.

- B. It supports drivers for data plane ASICs outside of the operating system.
- C. It allows platform-dependent code to be compiled into a single image.
- D. It supports multiple IOS instances simultaneously, sharing resources and internal infrastructure for scalability.
- E. It allows platform-independent code to be abstracted into a single microkernel for portability across platforms.
- F. It uses a QNX Neutrino-based environment underneath the IOS Software.

Correct Answer: ABC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 568

Refer to the exhibit.

```
Switch(config-if)#switchport block unicast
```

Which two benefits result from using this command on a switch? (Choose two.)

- A. The port cannot forward unknown unicast packets.
- B. Network security is increased on the configured port.
- C. The port cannot forward unknown multicast packets.
- D. The port cannot forward unknown broadcast packets.
- E. Network security is increased on all ports of the switch.
- F. Unknown packets of all types, except unicast, are blocked.

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 569

Which two statements about IPv4 and IPv6 networks are true? (Choose two.)

- A. In IPv6, hosts perform fragmentation.

- B. IPv6 uses a UDP checksum to verify packet integrity.
- C. In IPv6, routers perform fragmentation.
- D. In IPv4, fragmentation is performed by the source of the packet.
- E. IPv4 uses an optional checksum at the transport layer.
- F. IPv6 uses a required checksum at the network layer.

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 570

Refer to the exhibit.

```
R1#show monitor capture status
Load for five secs: 0%/0%; one minute: 0%; five minutes: 0%

capture state      : ON
  [running for 00:03:26.860]
capture mode       : Circular [wrap count = 0 ]
Number of packets
  captured : 201
  dropped  : 0
  received : 591
Capture will stop after 00:01:33
```

Which two statements about this capture are true? (Choose two.)

- A. It is set to run for five minutes.
- B. It continues to capture data after the buffer is full.
- C. It is set to run for a period of 00:03:26.
- D. It captures data only until the buffer is full.
- E. It is set to use the default buffer type.

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 571

In which way does the Bridge Assurance mechanism modify the default spanning-tree behavior in an effort to prevent bridging loops?

- A. Received BPDUs are looped back toward the sender to ensure that the link is bidirectional.
- B. If BPDUs are no longer received on a port, the switch immediately sends out a TCN BPDU.
- C. Extended topology information is encoded into all BPDUs.
- D. BPDUs are sent bidirectional on all active network ports, including blocked and alternate ports.

Correct Answer: D

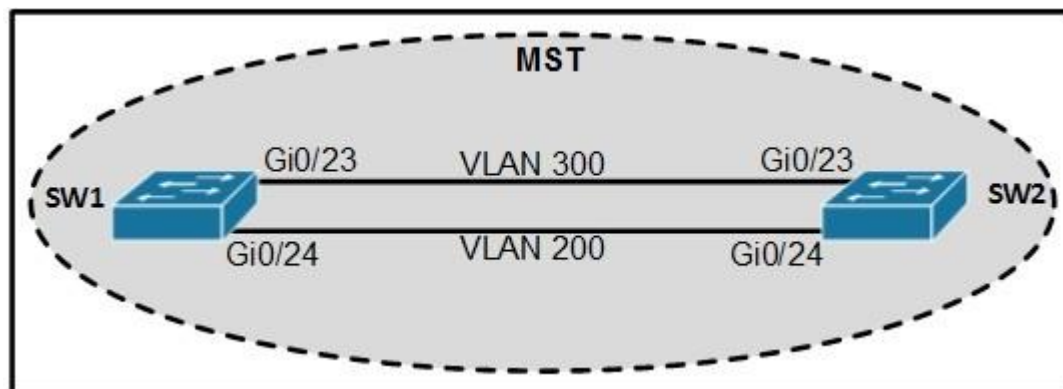
Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 572

Refer to the exhibit.



The VLAN-to-MST mapping is shown. (Assume SW1 acts as root for all possible MST instances.)

spanning-tree mst configuration name MST

revision 2

instance 0 vlan 1-200,301-4094 instance 1 vlan 201-300

!

If this topology is deployed, which action is required for traffic to flow on VLAN 200 and 300?

- A. Map VLAN 300 to instance 0.
- B. Map VLAN 200 to instance 2.
- C. Move instance 0 root to SW2.
- D. Move instance 1 root to SW2.
- E. Map both VLANs to instance 2.

Correct Answer: B

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 573

Which statement about UDLD is true?

- A. The udd reset command resets ports that have been error-disabled by both UDLD and Fast UDLD.
- B. Fast UDLD is configured in aggressive mode.
- C. Only bidirectional link failures can be detected in normal mode.
- D. Each switch in a UDLD topology can send and receive packets to and from its neighbors.

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 574

What is the VLAN ID range of VLANs that are eligible for pruning?

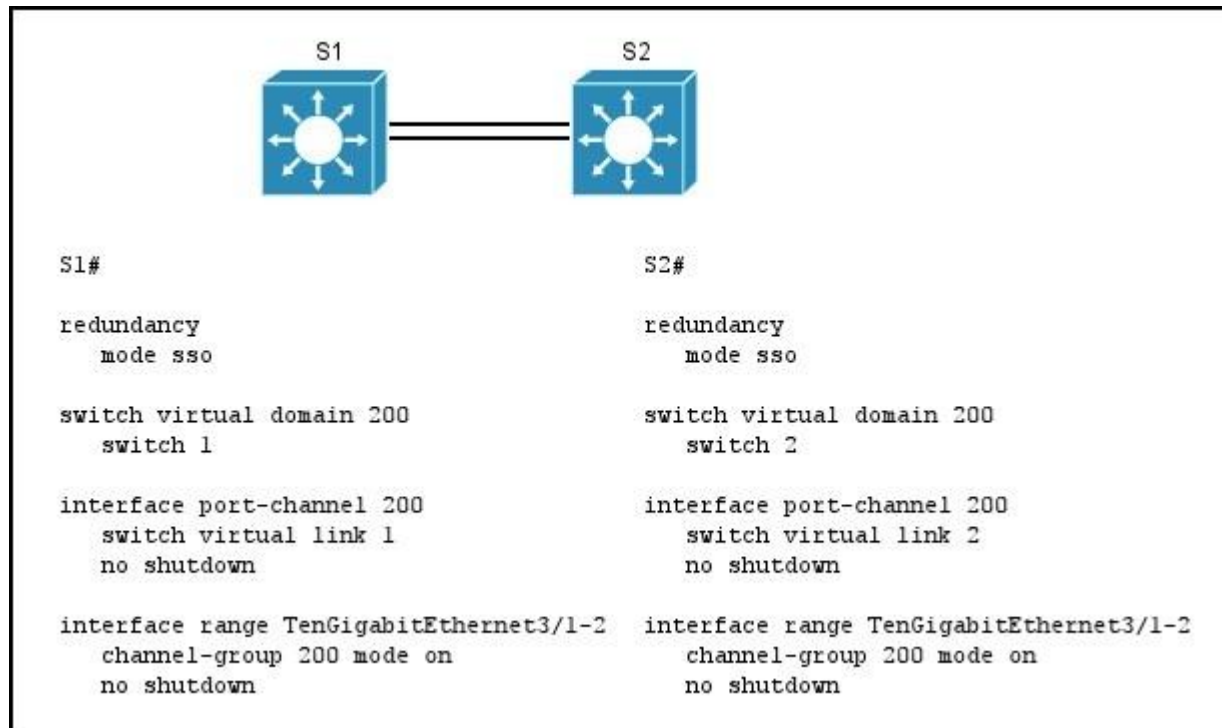
- A. 2 through 1001
- B. 1 through 1005
- C. 1 through 4096
- D. 2 through 1005

Correct Answer: A
Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 575

Refer to the exhibit.



The two standalone chassis are unable to convert into a VSS. What can you do to correct the problem?

- A. Set a different port channel number on each chassis.
- B. Set a different virtual domain ID on each chassis.
- C. Set the redundancy mode to rpr on both chassis.
- D. Add two ports to the port channel group.

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 576

Which two statements about the MAC address table space are true? (Choose two.)

- A. You can disable learning on a VLAN to reduce table-space requirements.
- B. When you disable learning on a VLAN with an SVI, IP packet flooding in the Layer 2 domain is also disabled.
- C. Unicast, multicast, and broadcast MAC address filtering is configured globally and disabled by default.
- D. The default setting for static MAC addresses to age out of the MAC address table is 300 seconds.
- E. Turning off MAC learning on VLANs 900 through 1005 disables learning on VLANs 900 through 1001.

Correct Answer: AE

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 577

Which two statements about 802.1Q tunneling are true? (Choose two.)

- A. It requires a system MTU of at least 1504 bytes.
- B. The default configuration sends Cisco Discovery Protocol, STP, and VTP information.
- C. Traffic that traverses the tunnel is encrypted.
- D. It is supported on private VLAN ports.
- E. MAC-based QoS and UDLD are supported on tunnel ports.
- F. Its maximum allowable system MTU is 1546 bytes.

Correct Answer: AE

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 578

Which two options about PIM-DM are true? (Choose two.)

- A. PIM-DM initially floods multicast traffic throughout the network.
- B. In a PIM-DM network, routers that have no upstream neighbors prune back unwanted traffic.
- C. PIM-DM supports only shared trees.
- D. PIM-DM uses a pull model to deliver multicast traffic.
- E. PIM-DM cannot be used to build a shared distribution tree.

Correct Answer: AE

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 579

Which two statements about PIM-DM are true? (Choose two.)

- A It forwards multicast packets on a source tree.
- A. It requires an RP.
- B. It forwards multicast packets on a shared distribution tree.
- C. It floods multicast packets to neighbors that have requested the data.
- D. It floods multicast packets throughout the network.
- E. It forwards multicast packets to neighbors that have requested the data.

Correct Answer: AE


Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 580

Refer to the exhibit.



R1 and R2 both advertise 10.50.1.0/24 to R3 and R4 as shown. R1 is the primary path. Which path does traffic take from the R4 data center to the file server?

- A. All traffic travels from R4 to R2 to the file server.
- B. All traffic travels from R4 to R3 to R1 to the file server.
- C. Traffic is load-balanced from R4 to R2 and R3. Traffic that is directed to R3 then continues to R1 to the file server. Traffic that is directed to R2 continues to the file server.
- D. All traffic travels from R4 to R2 to R1 to the file server.

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 581

Which two statements about path selection are true? (Choose two.)

- A. If there are multiple equal matches between OSPF processes, the path with the lowest OSPF PID is chosen.
- B. If the backdoor command is configured on a BGP network, the route is advertised with an AD of 20.
- C. If an OSPF E2 route has an AS of 90, that path is preferred over an OSPF IA route with an AD of 110.
- D. If there are multiple equal matches between the same protocols on an EIGRP network, the preferred path will be EIGRP with the highest AS.
- E. If IS-IS has multiple routes with the same prefix-length, it will prefer Level 1 routes over Level 2 routes.

Correct Answer: AE

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 582

Which two statements about BGP prefix-based outbound filtering are true? (Choose two.)

- A. It must be configured per address family.
- B. It can use prefix lists and route maps for filtering.
- C. It can be configured under the global BGP routing process.
- D. It can be configured for external peering sessions only.
- E. It can increase the processing load on the router.
- F. It supports IP multicast routes.

Correct Answer: AD

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 583

Which statement about OSPF loop prevention is true?

- A. The discard route is generated automatically on the ABR to prevent routing loops.
- B. The ASBR uses type 3 LSAs from non-backbone areas to prevent control-plane routing loops.
- C. The ABR can filter type 3 LSPs to prevent routing loops.
- D. The DN bit ignores LSA types 2, 3, and 5 to prevent routing loops.

Correct Answer: A
Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 584

Which three options must be configured when deploying OSPFv3 for authentication? (Choose three.)

- A. security parameter index
- B. crypto map
- C. authentication method
- D. IPsec peer
- E. encryption algorithm
- F. encryption key
- G. IPsec transform-set
- H. authentication key

Correct Answer: ACH
Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 585

Which two improvements do SIA-Query and SIA-Reply messages add to EIGRP? (Choose two.)

- A. Stuck-in-active conditions are solved faster.
- B. They prevent a route from going into the stuck-in-active state.
- C. They help in the localization of the real failure in the network.
- D. The EIGRP adjacency between two neighbors never goes down.

Correct Answer: AC
Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 586

Which two statements about EIGRP load balancing are true? (Choose two.)

- A. EIGRP supports 6 unequal-cost paths.
- B. A path can be used for load balancing only if it is a feasible successor.
- C. EIGRP supports unequal-cost paths by default.
- D. Any path in the EIGRP topology table can be used for unequal-cost load balancing.
- E. Cisco Express Forwarding is required to load-balance across interfaces.

Correct Answer: AB


Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 587

Refer to the exhibit.



All routers are running EIGRP and the network has converged. R3 and R4 are configured as EIGRP Stub. If the link between R1 and R3 goes down,

which statement is true?

- A. R1 sends traffic destined to 192.168.0.100 via R2.
- B. R2 does not have a route to 192.168.0.0/24 in the routing table.
- C. The prefix 192.168.0.0/24 becomes stuck-in-active on R4.
- D. R3 does not advertise 192.168.0.0/24 to R4 anymore.

Correct Answer: B

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 588

An NSSA area has two ABRs connected to Area 0. Which statement is true?

- A. Both ABRs translate Type-7 LSAs to Type-5 LSAs.
- B. The ABR with the highest router ID translates Type-7 LSAs to Type-5 LSAs.
- C. Both ABRs forward Type-5 LSAs from the NSSA area to backbone area.
- D. No LSA translation is needed.

Correct Answer: B

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 589

Which two OSPF network types require the use of a DR and BDR? (Choose two.)

- A. non-broadcast networks
- B. point-to-point networks
- C. point-to-multipoint networks
- D. broadcast networks
- E. point-to-multipoint non-broadcast networks

Correct Answer: AD

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 590

Which packet does a router receive if it receives an OSPF type 4 packet?

- A. hello packet
- B. database descriptor packet
- C. link state update packet
- D. link state request packet
- E. link state acknowledge packet

Correct Answer: C

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 591

Which two statements about OSPFv3 are true? (Choose two.)

- A. It supports unicast address families for IPv4 and IPv6.
- B. It supports unicast address families for IPv6 only.
- C. It supports only one address family per instance.
- D. It supports the use of a cluster ID for loop prevention.
- E. It supports multicast address families for IPv4 and IPv6.
- F. It supports multicast address families for IPv6 only.

Correct Answer: AC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 592

DRAG DROP

Drag each OSPF route-type identifier on the left to its description on the right.

Select and Place:

E1	A translated, redistributed route whose metric increments along the path.
E2	A route from another area.
N1	A route within an area.
N2	A redistributed route whose metric remains the same along the path.
O	A redistributed route whose metric increments along the path.
O IA	A translated, redistributed route whose metric remains the same along the path.

Correct Answer:

	N1
	O IA
	O
	E2
	E1
	N2

Section: Mix Questions**Explanation****Explanation/Reference:****QUESTION 593**

Which three options are characteristics of a Type 10 LSA? (Choose three.)

- A. It is an area-local, opaque LSA.
- B. Data is flooded to all routers in the LSA scope.
- C. It is used for traffic-engineering extensions to OSPF.
- D. It is a link-local, opaque LSA.
- E. Data is flooded only to the routers in the LSA scope that understand the data.
- F. It is used for traffic-engineering extensions to LDP.

Correct Answer: ABC

Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 594

Which two BGP attributes are optional, non-transitive attributes? (Choose two.)

- A. AS path
- B. local preference
- C. MED
- D. weight
- E. cluster list

Correct Answer: CE
Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 595

Which two statements about IBGP multipath are true? (Choose two.)

- A. The IGP metric of the BGP next hop can be different from the best-path IGP metric if you configure the router for unequal-cost IBGP multipath.
- B. The IGP metric of the BGP next hop must be the same as the best-path IGP metric.
- C. The equivalent next-hop-self is performed on the best path from among the IBGP multipaths before it is forwarded to external peers.
- D. The path should be learned from an external neighbor.
- E. The router BGP process must learn the path from a confederation-external or external neighbor.
- F. The router BGP process must learn the path from an internal neighbor.

Correct Answer: AF
Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 596

Refer to the exhibit.

```
%TCP-6-BADAUTH: Invalid MD5 digest from 172.16.129.4(179) to 172.16.129.7(43766) tableid - 0
BGP: tbl IPv4 Unicast:base Service reset requests
BGP: tbl IPv4 MDT:base Service reset requests
BGP: tbl VPNv4 Unicast:base Service reset requests
BGP: tbl IPv4 Multicast:base Service reset requests
```

Which three statements about the device with this configuration are true? (Choose three.)

- A. Multiple AFI are configured on the device.
- B. The authentication on 172.16.129.7 is configured incorrectly.
- C. The device is configured to support MPLS VPNs.
- D. This device is configured with a single AFI.
- E. The authentication on 172.16.129.4 is configured incorrectly.
- F. The device is configured to support L2VPNs.

Correct Answer: ABC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 597

Which two attributes were introduced with the Cisco IOS BGP 4-byte ASN feature? (Choose two.)

- A. AS4_AGGREGATOR
- B. AS4_PATH
- C. AS4_PLAIN
- D. AS4_DOT
- E. AS4_TRANS

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 598

Which three statements about the default behaviour of eBGP sessions are true? (Choose three.)

- A. eBGP sessions between sub-ASs in different confederations transmit the next hop unchanged.
- B. The next hop in an eBGP peering is the IP address of the neighbor that announced the route.
- C. When a route reflector reflects a route to a client, it transmits the next hop unchanged.
- D. The next hop in an eBGP peering is the loopback address of the interface that originated the route.
- E. The next hop in an eBGP peering is the loopback address of the neighbor that announced the route.
- F. When a route reflector reflects a route to a client, it changes the next hop to its own address.

Correct Answer: ABC

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 599

DRAG DROP

Drag each IS-IS command on the left to its effect on the right.

Select and Place:

isis csnp-interval	Sets the delay between successive LSP transmissions.
isis lsp-interval	Enables graceful restart.
isis retransmit-throttle interval	Sets the delay between point-to-point LSP transmissions.
isis three-way-handshake ietf	Configures point-to-point adjacencies.
nsf cisco	Maintains database synchronization.
nsf ietf	Enables nonstop routing.

Correct Answer:

	isis lsp-interval
	nsf ietf
	isis retransmit-throttle interval
	isis three-way-handshake ietf
	isis csnp-interval
	nsf cisco

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 600

Which statement about a P router in a Layer 3 MPLS VPN is true?

- A. It is unaware of VPN routes.
- B. It connects to customer edge routers.
- C. It participates in MPLS VPN routing.
- D. It uses the running IGP to share VPN routes.

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 601

Which statement about LISP encapsulation in an EIGRP OTP implementation is true?

- A. OTP uses LISP encapsulation for dynamic multipoint tunneling.
- B. OTP maintains the LISP control plane.
- C. OTP uses LISP encapsulation to obtain routes from neighbors.
- D. LISP learns the next hop.

Correct Answer: A


Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 602

Refer to the exhibit.



Your organization has two offices, Site 1 and Site 2, which are connected by a provider backbone, as shown. Where must you configure an attachment circuit to allow the two sites to connect over a Layer 2 network using L2TPv3?

- A. PE Site 1 Fa1/0 and PE Site 2 Fa0/0

- B. CE Site 1 Fa0/0 and CE Site 2 Fa0/0
- C. PE Site 1 Se0/0 and PE Site 2 Se0/0
- D. CE Site 1 Fa0/0 and PE Site 2 Se0/0

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 603

What are two benefits of Per-Tunnel QoS for DMVPN? (Choose two.)

- A. The administrator can configure criteria that, when matched, can automatically set up QoS for each spoke as it comes online.
- B. Traffic from each spoke to the hub can be regulated individually.
- C. When traffic exceeds a configurable threshold, the spokes can automatically set up QoS with the hub.
- D. The hub can send large packets to a spoke during allotted timeframes.
- E. The hub can be regulated to prevent overloading small spokes.

Correct Answer: AE

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 604

Which two statements about VPLS are true? (Choose two.)

- A. Split horizon is used on PE devices to prevent loops.
- B. Spanning tree is extended from CE to CE.
- C. IP is used to switch Ethernet frames between sites.
- D. PE routers dynamically associate to peers.
- E. VPLS extends a Layer 2 broadcast domain.

Correct Answer: AE

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 605

DRAG DROP

Drag each MPLS term on the left to the matching statement on the right.

Select and Place:

Downstream	The transit direction of prefix updates.
LIB	The component that replaces the top label in a label stack.
LSP	A route determined by IGP routing protocols.
LSR	Data traveling to a destination.
Upstream	A table with labels received from peers.

Correct Answer:



Upstream

LSR

LSP


Downstream

LIB

Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 606
Refer to the exhibit.



Which two actions can you take to enable CE-1 at site A to access the Internet? (Choose two.)

- A. Create a default route for site A on PE-1 with the next hop set to the PE-2 interface to the Internet.
- B. Originate a default route in site B with the next hop set to the PE-2 Internet interface, and import the default route into site A.
- C. Create a default route on CE-1 with the next hop set to the PE-1 upstream interface.
- D. Originate a default route in site A with the next hop set to the PE-2 interface to CE-1.
- E. Create a static default route on CE-1 with the next hop set to the PE-2 interface to the Internet.

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 607

Which attribute is transported over an MPLS VPN as a BGP extended community?

- A. route target
- B. route distinguisher
- C. NLRI
- D. origin
- E. local preference

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 608

Which Layer 2 tunneling technique eliminates the need for pseudowires?

- A. OTV
- B. L2TPv3
- C. AToM
- D. VPLS

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 609

Refer to the exhibit.


```
crypto isakmp keepalive 15 periodic
crypto gdoi group testgroup
server local
address ipv4 10.1.1.1
redundancy
local priority 200
peer address ipv4 10.1.2.2
```

Which statement about this GETVPN configuration is true?

- A. Co-operative key servers are configured.
- B. Redundant peers are configured.
- C. The key server uses multicast mode to propagate rekey messages.
- D. PSK authentication is configured.

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 610

DRAG DROP

Drag each GETVPN component on the left to its function on the right.

Select and Place:

Group Domain of Interpretation Protocol	Authenticates group members.
Group Member	Serves as the crypto device.
Key Encryption Key	Secures data during transmission.
Key Server	Manages the group keys.
Traffic Encryption Key	Secures rekey messages.

Correct Answer:

	Key Server
	Group Member
	Traffic Encryption Key
	Group Domain of Interpretation Protocol
	Key Encryption Key

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 611

Refer to the exhibit.

Router 1	Router 2
<pre>crypto isakmp keepalive 15 periodic crypto gdoi group TESTGROUP server local address ipv4 10.1.1.1 redundancy local priority 250 peer address ipv4 10.0.1.2</pre>	<pre>crypto isakmp keepalive 15 periodic crypto gdoi group TESTGROUP server local address ipv4 10.0.1.2 redundancy local priority 200 peer address ipv4 10.1.1.1</pre>

Which statement about this GET VPN configuration is true?

- A. Router 1 acts as the primary key server because it has a higher priority.
- B. An RSA key has been imported into the configuration.
- C. The GDOI group configuration generated a key.
- D. DPD is disabled.

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 612

When you implement CoPP on your network, what is its default action?

- A. permit all traffic
- B. rate-limit bidirectional traffic to the control plane
- C. drop management ingress traffic to the control plane

- D. monitor ingress and egress traffic to the control plane by using access groups that are applied to the interface
- E. block all traffic

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 613

DRAG DROP

Drag each SNMP term on the left to the matching definition on the right.

Select and Place:

Agent	An operation that retrieves object variables.
Get	An operation that retrieves unsolicited information from an agent.
Manager	A system that monitors and controls the activities of network hosts.
MIB	An operation that modifies object variables.
Notification	A software component that maintains and reports data.
Set	A virtual storage area for managed objects.

Correct Answer:

	Get
	Notification
	Manager
	Set
	Agent
	MIB

Section: Mix Questions
Explanation

Explanation/Reference:

QUESTION 614
Refer to the exhibit.

```
username admin password 0 notsecure
aaa new-model
aaa authentication login default group tacacs+ local
aaa authorization exec default group tacacs+ local
banner motd ^C
Authorized users only.
^C
line con 0
  exec-timeout 5 0
  privilege level 15
  no vacant-message
  activation-character 124
```

If a console port is configured as shown, which response is displayed when you connect to the console port?

- A. a blinking cursor
- B. the message "Authorized users only"
- C. the username prompt
- D. three username name prompts followed by a timeout message
- E. the message "Connection refused"

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 615

Which map is locally defined?

- A. DSCP-to-DSCP-mutation
- B. CoS-to-DSCP
- C. IP-precedence-to-DSCP
- D. DSCP-to-CoS

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 616

Which two statements about Layer 2 Frame Prioritization bits are true? (Choose two.)

- A. 802.1Q frame headers carry the CoS value in the three most-significant bits of the 2-byte Tag Control Information field.
- B. ISL frame headers carry an IEEE 802.1P CoS value in the three least-significant bits of the 2- byte User field.
- C. ISL frame headers carry an IEEE 802.1P CoS value in the three most-significant bits of the 1- byte User field.
- D. On 802.1Q trunks, traffic in the native VLAN is carried in 802.1Q frames.
- E. Only 802.1Q and ISL frame types can carry CoS information.
- F. On 802.1Q trunks, traffic in the native VLAN is carried in 802.1P frames.

Correct Answer: AE

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 617

Which two application protocols require application layer gateway support when using NAT on a Cisco router? (Choose two.)

- A. SIP
- B. HTTP
- C. FTP
- D. SMTP
- E. POP3

Correct Answer: AC


Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 618

Refer to the exhibit.



Which configuration can you implement on PE-1 to allow CE-1 to receive delegated IPv6 prefixes? A)

```
ipv6 local pool CE-1 2001:db8:4:8888::/48 56
ipv6 dhcp pool CE-1-DHCP
  prefix-delegation pool CE-1 lifetime infinite infinite
interface GigabitEthernet0/0
  ipv6 address 2001:db8:4:822::1/64
  ipv6 dhcp server CE-1-DHCP
```

B)

```
ipv6 local pool CE-1 2001:db8:4:8888/49 56
ipv6 dhcp pool CE-1-DHCP
interface GigabitEthernet0/0
  ipv6 address 2001:db8:4:822::1/64
  ipv6 dhcp server CE-1
```

C)


```
ipv6 local pool CE-1 2001:db8:4:8888/56 48
ipv6 dhcp pool CE-1-DHCP
  prefix-delegation pool CE-1 lifetime infinite infinite
interface GigabitEthernet0/0
  ipv6 address 2001:db8:4:822::1/64
  ipv6 dhcp server CE-1-DHCP
```

D)

```
ipv6 local pool CE-1 2001:db8:4:8888::/48 32
ipv6 dhcp pool CE-1-DHCP
  prefix-delegation pool CE-1 lifetime infinite infinite
interface GigabitEthernet0/0
  ipv6 address 2001:db8:4:822::1/64
  ipv6 dhcp server CE-1-DHCP
```

E)

```
ip local pool CE-1 2001:db8:4:8888::/64 48
ipv6 dhcp pool CE-1-DHCP
  prefix-delegation pool CE-1 lifetime infinite infinite
interface GigabitEthernet0/0
  ipv6 address 2001:db8:4:822::1/64
  ipv6 dhcp server CE-1-DHCP
```

- A. Exhibit A
- B. Exhibit B
- C. Exhibit C
- D. Exhibit D
- E. Exhibit E

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 619

Which two statements about VRRP are true? (Choose two.)

- A. It is assigned multicast address 224.0.0.18.
- B. The TTL for VRRP packets must be 255.
- C. It is assigned multicast address 224.0.0.9.
- D. Its IP protocol number is 115.
- E. Three versions of the VRRP protocol have been defined.
- F. It supports both MD5 and SHA1 authentication.

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 620

What are two benefits of NVI? (Choose two.)

- A. It provides scalability by maintaining a NAT table on every interface.
- B. It can dynamically create a static route to the NAT pool for translation.
- C. It supports the use of route maps for policy-based NAT.
- D. It supports the use of a single interface for translations.
- E. It injects a route into the existing routing protocol that directs translation to the NAT pool.

Correct Answer: AB

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 621

An IP SLA fails to generate statistics. How can you fix the problem?

- A. Add the verify-data command to the router configuration.

- B. Reload the router configuration.
- C. Remove the ip sla schedule statement from the router configuration and re-enter it.
- D. Add the debug ip sla error command to the router configuration.
- E. Add the debug ip sla trace command to the router configuration.

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 622

Which cache aggregation scheme is supported by NetFlow ToS-based router aggregation?

- A. prefix-port
- B. AS
- C. protocol port
- D. destination prefix

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

QUESTION 623

When you implement PfR, which IP SLA probe is used to determine the MOS?

- A. jitter
- B. latency
- C. packet loss
- D. throughput

Correct Answer: A

Section: Mix Questions

Explanation

Explanation/Reference:

