Cisco 350-001



CCIE Routing and Switching Written Exam v4.0

Version: 12.1

Topic 1, Implement Layer 2 Technologies

QUESTION NO: 1

Which statement is true about loop guard?

A. Loop guard only operates on interfaces that are considered point-to-point by the spanning tree.
B. Loop guard only operates on root ports.
C. Loop guard only operates on designated ports.
D. Loop guard only operates on edge ports.

Answer: A
Explanation:
Understanding How Loop Guard Works
Unidirectional link failures may cause a root port or alternate port to become designated as root if BPDUs are absent. Some software failures may introduce temporary loops in the network. Loop guard checks if a root port or an alternate root port receives BPDUs. If the port is receiving BPDUs, loop guard puts the port into an inconsistent state until it starts receiving BPDUs again. Loop guard isolates the failure and lets spanning tree converge to a stable topology without the failed link or bridge.

You can enable loop guard per port with the set spantree guard loop command.
Note When you are in MST mode, you can set all the ports on a switch with the set spantree global-defaults loop-guard command.
When you enable loop guard, it is automatically applied to all of the active instances or VLANs to which that port belongs. When you disable loop guard, it is disabled for the specified ports. Disabling loop guard moves all loop-inconsistent ports to the listening state.
If you enable loop guard on a channel and the first link becomes unidirectional, loop guard blocks the entire channel until the affected port is removed from the channel. Figure 8-6 shows loop guard in a triangle switch configuration.
Figure 8-6 Triangle Switch Configuration with Loop Guard

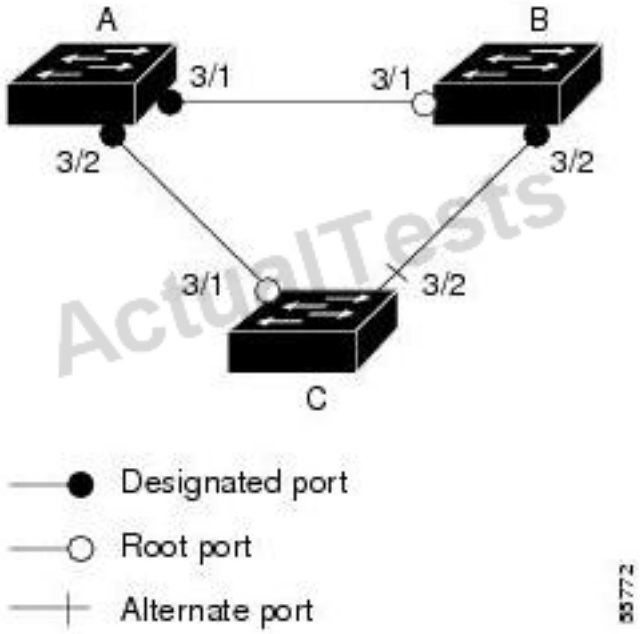


Figure 8-6 illustrates the following configuration:
Switches A and B are distribution switches.
Switch C is an access switch.
Loop guard is enabled on ports 3/1 and 3/2 on Switches A, B, and C.
Use loop guard only in topologies where there are blocked ports. Topologies that have no blocked ports, which are loop free, do not need to enable this feature. Enabling loop guard on a root switch has no effect but provides protection when a root switch becomes a nonroot switch.
Follow these guidelines when using loop guard:

Do not enable loop guard on PortFast-enabled or dynamic VLAN ports.
Do not enable PortFast on loop guard-enabled ports.
Do not enable loop guard if root guard is enabled.
Do not enable loop guard on ports that are connected to a shared link.

Note: We recommend that you enable loop guard on root ports and alternate root ports on access switches.
Loop guard interacts with other features as follows:
Loop guard does not affect the functionality of UplinkFast or BackboneFast.
Root guard forces a port to always be designated as the root port. Loop guard is effective only if the port is a root port or an alternate port. Do not enable loop guard and root guard on a port at the same time.
PortFast transitions a port into a forwarding state immediately when a link is established. Because a PortFast-enabled port will not be a root port or alternate port, loop guard and PortFast cannot be configured on the same port. Assigning dynamic VLAN membership for the port requires that the port is PortFast enabled. Do not configure a loop guard-enabled port with dynamic VLAN

membership.
If your network has a type-inconsistent port or a PVID-inconsistent port, all BPDUs are dropped until the misconfiguration is corrected. The port transitions out of the inconsistent state after the message age expires. Loop guard ignores the message age expiration on type-inconsistent ports and PVID-inconsistent ports. If the port is already blocked by loop guard, misconfigured BPDUs that are received on the port make loop guard recover, but the port is moved into the type-inconsistent state or PVID-inconsistent state.

In high-availability switch configurations, if a port is put into the blocked state by loop guard, it remains blocked even after a switchover to the redundant supervisor engine. The newly activated supervisor engine recovers the port only after receiving a BPDU on that port.
Loop guard uses the ports known to spanning tree. Loop guard can take advantage of logical ports provided by the Port Aggregation Protocol (PAgP). However, to form a channel, all the physical ports grouped in the channel must have compatible configurations. PAgP enforces uniform configurations of root guard or loop guard on all the physical ports to form a channel.
These caveats apply to loop guard:
–Spanning tree always chooses the first operational port in the channel to send the BPDUs. If that link becomes unidirectional, loop guard blocks the channel, even if other links in the channel are functioning properly.
–If a set of ports that are already blocked by loop guard are grouped together to form a channel, spanning tree loses all the state information for those ports and the new channel port may obtain the forwarding state with a designated role.
–If a channel is blocked by loop guard and the channel breaks, spanning tree loses all the state information.
The individual physical ports may obtain the forwarding state with the designated role, even if one or more of the links that formed the channel are unidirectional.
You can enable UniDirectional Link Detection (UDLD) to help isolate the link failure. A loop may occur until
UDLD detects the failure, but loop guard will not be able to detect it.
Loop guard has no effect on a disabled spanning tree instance or a VLAN.

Reference:

http://www.cisco.com/en/US/docs/switches/lan/catalyst4000/8.2glx/configuration/guide/stp_enha.html#wp1048163

QUESTION NO: 2

Which command is used to enable EtherChannel hashing for Layer 3 IP and Layer 4 port-based CEF?

A. mpls ip cef
B. port-channel ip cef
C. mpls ip port-channel cef
D. port-channel load balance
E. mpls ip load-balance
F. ip cef EtherChannel channel-id XOR L4
G. ip cef connection exchange

Answer: D
Explanation:
Port-channel load balance is normally used for enable etherchannel hashing for Layer 3 IP and Layer 4 port based CEF.

QUESTION NO: 3

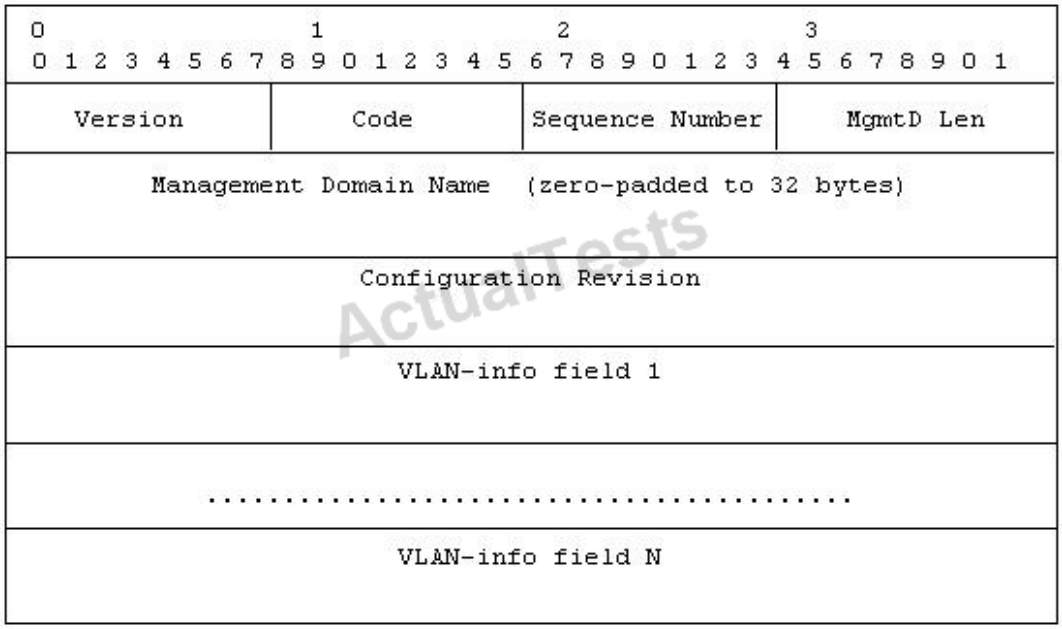
Which two options are contained in a VTP subset advertisement? (Choose two.)

A. followers field
B. MD5 digest
C. VLAN information
D. sequence number

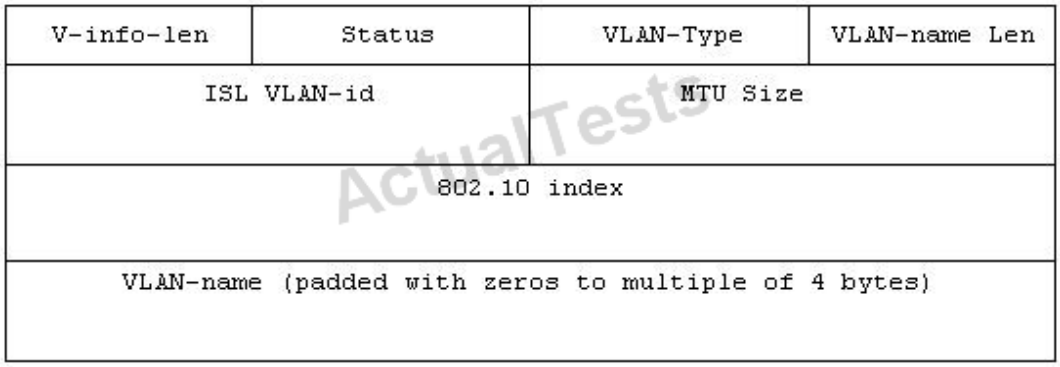
Answer: C,D
Explanation:
Subset Advertisements
When you add, delete, or change a VLAN in a Catalyst, the server Catalyst where the changes are made increments the configuration revision and issues a summary advertisement. One or several subset advertisements follow the summary advertisement. A subset advertisement contains a list of VLAN information.

If there are several VLANs, more than one subset advertisement can be required in order to advertise all the VLANs.

Subset Advertisement Packet Format



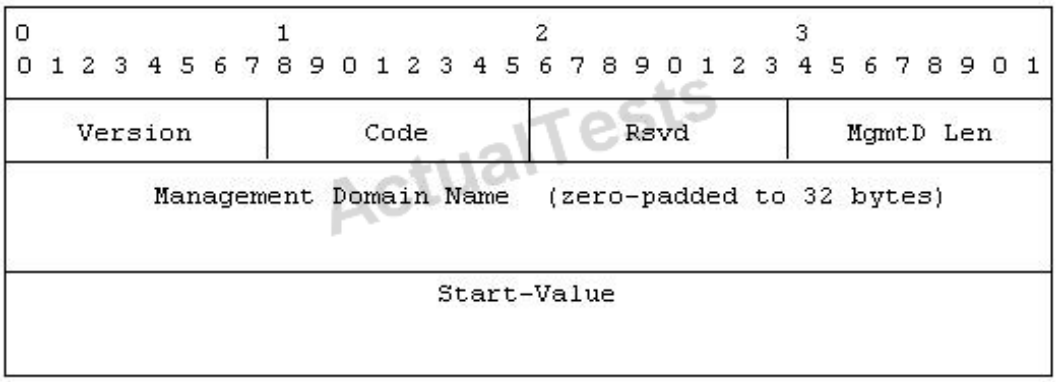
This formatted example shows that each VLAN information field contains information for a different VLAN. It is ordered so that lowered-valued ISL VLAN IDs occur first:



Most of the fields in this packet are easy to understand. These are two clarifications:
Code — The format for this is 0x02 for subset advertisement.
Sequence number — This is the sequence of the packet in the stream of packets that follow a summary advertisement. The sequence starts with 1.

Advertisement Requests
A switch needs a VTP advertisement request in these situations:
The switch has been reset.
The VTP domain name has been changed.
The switch has received a VTP summary advertisement with a higher configuration revision than its own.

Upon receipt of an advertisement request, a VTP device sends a summary advertisement. One or more subset advertisements follow the summary advertisement. This is an example:



Code—The format for this is 0x03 for an advertisement request.
Start-Value—This is used in cases in which there are several subset advertisements. If the first (n) subset advertisement has been received and the subsequent one (n+1) has not been received, the Catalyst only requests advertisements from the (n+1)th one.
Reference

http://www.cisco.com/en/US/tech/tk389/tk689/technologies_tech_note09186a0080094c52.shtml

QUESTION NO: 4

Which two statements are true about traffic shaping? (Choose two.)

A. Out-of-profile packets are queued.
B. It causes TCP retransmits.
C. Marking/remarking is not supported.
D. It does not respond to BECN and ForeSight Messages.
E. It uses a single/two-bucket mechanism for metering.

Answer: A,C
Reference:
https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&ved=0CCwQFjAA&url=http%3A%2F%2Fstaffweb.itsligo.ie%2Fstaff%2Fpflynn%2FTelecoms%25203%2FSlides%2FONT%2520Mod%25204%2520Lesson%25207.ppt&ei=LoDIUfTTGtO3hAeQz4HQCA&usg=AFQjCNGY24UkAfy8tKIHlzEm9gfoIjv6fg&sig2=t4UIzkZ12wnO2988dEDyug&bvm=bv.48293060,d.ZG4 (slide 6)

QUESTION NO: 5

Which three options are features of VTP version 3? (Choose three.)

A. VTPv3 supports 8K VLANs.
B. VTPv3 supports private VLAN mapping.
C. VTPv3 allows for domain discovery.
D. VTPv3 uses a primary server concept to avoid configuration revision issues.
E. VTPv3 is not compatible with VTPv1 or VTPv2.
F. VTPv3 has a hidden password option.

Answer: B,D,F
Explanation:
Key Benefits of VTP Version 3
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).

Brief Background on VTP Version 1 and VTP Version 2
VTP version 1 was developed when only 1k VLANs where available for configuration. A tight internal coupling of the VLAN implementation, the VLAN pruning feature, and the VTP function itself offered an efficient means of implementation. It has proved in the field to reliably support Ethernet, Token Ring, and FDDI networks via VTP.
The use of consistent VLAN naming was a requirement for successful use of VMPS (Vlan Membership Policy Server). VTP ensures the consistency of VLAN names across the VTP domain. Most VMPS implementations are likely to be migrated to a newer, more flexible and feature-rich method.
To add support for Token Ring, VTP version 1 was enhanced and called VTP version 2. Certain other minor changes and enhancements were also added at this time.
The functional base in VTP version 3 is left unchanged from VTP version 2, so backward

compatibility is built in. It is possible, on a per link basis, to automatically discover and support VTP version 2 devices.

VTP version 3 adds a number of enhancements to VTP version 1 and VTP version 2:
Support for a structured and secure VLAN environment (Private VLAN, or PVLAN)
Support for up to 4k VLANs
Feature enhancement beyond support for a single database or VTP instance
Protection from unintended database overrides during insertion of new switches
Option of clear text or hidden password protection
Configuration option on a per port base instead of only a global scheme
Optimized resource handling and more efficient transfer of information
These new requirements made a new code foundation necessary. The design goal was to make VTP version 3 a versatile vehicle. This was not only for the task of transferring a VLAN DB but also for transferring other databases-for example, the MST database.

Reference
http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/solution_guide_c78_508010.html

QUESTION NO: 6

Which three options are considered in the spanning-tree decision process? (Choose three.)

A. lowest root bridge ID
B. lowest path cost to root bridge
C. lowest sender bridge ID
D. highest port ID
E. highest root bridge ID
F. highest path cost to root bridge

Answer: A,B,C
Explanation:
Configuration bridge protocol data units (BPDUs) are sent between switches for each port. Switches use s four step process to save a copy of the best BPDU seen on every port. When a port receives a better BPDU, it stops sending them. If the BPDUs stop arriving for 20 seconds (default), it begins sending them again.

Step 1 Lowest Root Bridge ID (BID)
Step 2 Lowest Path Cost to Root Bridge

Step 3 Lowest Sender BID
Step 4 Lowest Port ID

Reference
Cisco General Networking Theory Quick Reference Sheets

QUESTION NO: 7

In 802.1s, how is the VLAN to instance mapping represented in the BPDU?

A. The VLAN to instance mapping is a normal 16-byte field in the MST BPDU.
B. The VLAN to instance mapping is a normal 12-byte field in the MST BPDU.
C. The VLAN to instance mapping is a 16-byte MD5 signature field in the MST BPDU.
D. The VLAN to instance mapping is a 12-byte MD5 signature field in the MST BPDU.

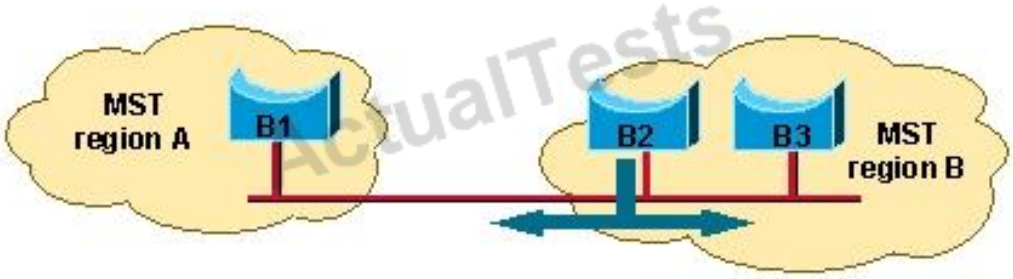
Answer: C
Explanation:
MST Configuration and MST Region
Each switch running MST in the network has a single MST configuration that consists of these three attributes:
1. An alphanumeric configuration name (32 bytes)
2. A configuration revision number (two bytes)
3. A 4096-element table that associates each of the potential 4096 VLANs supported on the chassis to a given instance.

In order to be part of a common MST region, a group of switches must share the same configuration attributes.
It is up to the network administrator to properly propagate the configuration throughout the region. Currently, this step is only possible by the means of the command line interface (CLI) or through Simple Network
Management Protocol (SNMP). Other methods can be envisioned, as the IEEE specification does not explicitly mention how to accomplish that step.
Note: If for any reason two switches differ on one or more configuration attribute, the switches are part of different regions. For more information refer to the Region Boundary section of this document.

Region Boundary
In order to ensure consistent VLAN-to-instance mapping, it is necessary for the protocol to be able to exactly identify the boundaries of the regions. For that purpose, the characteristics of the region

are included in the BPDUs. The exact VLANs-to-instance mapping is not propagated in the BPDU, because the switches only need to know whether they are in the same region as a neighbor. Therefore, only a digest of the VLANs-toinstance mapping table is sent, along with the revision number and the name. Once a switch receives a BPDU, the switch extracts the digest (a numerical value derived from the VLAN-to-instance mapping table through a mathematical function) and compares this digest with its own computed digest. If the digests differ, the port on which the BPDU was received is at the boundary of a region.

In generic terms, a port is at the boundary of a region if the designated bridge on its segment is in a different region or if it receives legacy 802.1d BPDUs. In this diagram, the port on B1 is at the boundary of region A, whereas the ports on B2 and B3 are internal to region B:



MST Instances

According to the IEEE 802.1s specification, an MST bridge must be able to handle at least these two instances:

One Internal Spanning Tree (IST)

One or more Multiple Spanning Tree Instance(s) (MSTIs)

The terminology continues to evolve, as 802.1s is actually in a pre-standard phase. It is likely these names will change in the final release of 802.1s. The Cisco implementation supports 16 instances: one IST (instance 0) and 15 MSTIs.

show vtp status

Cisco switches "show vtp status" Field Descriptions has a MD5 digest field that is a 16-byte checksum of the

VTP configuration as shown below

Router# show vtp status

VTP Version: 3 (capable)

Configuration Revision: 1

Maximum VLANs supported locally: 1005

Number of existing VLANs: 37

VTP Operating Mode: Server

VTP Domain Name: [smartports]

VTP Pruning Mode: Disabled

VTP V2 Mode: Enabled
VTP Traps Generation: Disabled
MD5 digest : 0x26 0xEE 0x0D 0x84 0x73 0x0E 0x1B 0x69
Configuration last modified by 172.20.52.19 at 7-25-08 14:33:43
Local updater ID is 172.20.52.19 on interface Gi5/2 (first layer3 interface fou)
VTP version running: 2

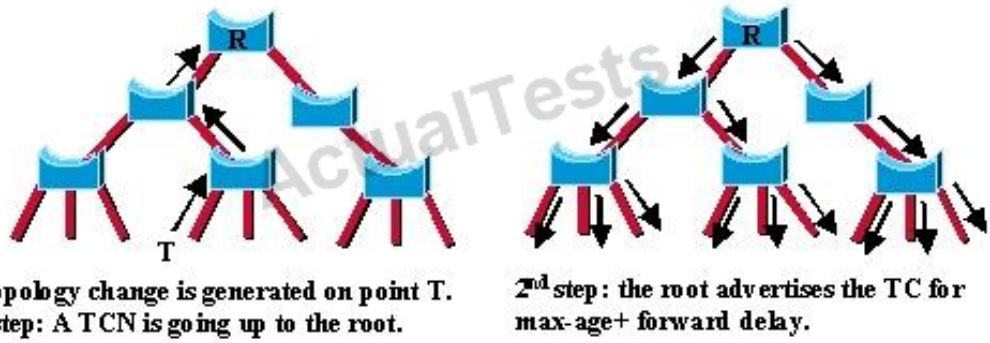
Reference
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_white_paper09186a0080094cfc.shtml
http://www.cisco.com/en/US/docs/ios-xml/ios/lanswitch/command/lsw-cr-book.pdf

QUESTION NO: 8

Which statement is true about TCN propagation?

A. The originator of the TCN immediately floods this information through the network.
B. The TCN propagation is a two step process.
C. A TCN is generated and sent to the root bridge.
D. The root bridge must flood this information throughout the network.

Answer: C
Explanation:
Explanation
New Topology Change Mechanisms
When an 802.1D bridge detects a topology change, it uses a reliable mechanism to first notify the root bridge.
This is shown in this diagram:



Once the root bridge is aware of a change in the topology of the network, it sets the TC flag on the BPDUs it sends out, which are then relayed to all the bridges in the network. When a bridge

receives a BPDU with the TC flag bit set, it reduces its bridging-table aging time to forward delay seconds. This ensures a relatively quick flush of stale information. Refer to Understanding Spanning-Tree Protocol Topology Changes for more information on this process. This topology change mechanism is deeply remodeled in RSTP. Both the detection of a topology change and its propagation through the network evolve.

Topology Change Detection

In RSTP, only non-edge ports that move to the forwarding state cause a topology change. This means that a loss of connectivity is not considered as a topology change any more, contrary to 802.1D (that is, a port that moves to blocking no longer generates a TC). When a RSTP bridge detects a topology change, these occur:

It starts the TC While timer with a value equal to twice the hello-time for all its non-edge designated ports and its root port, if necessary.

It flushes the MAC addresses associated with all these ports.

Note: As long as the TC While timer runs on a port, the BPDUs sent out of that port have the TC bit set.

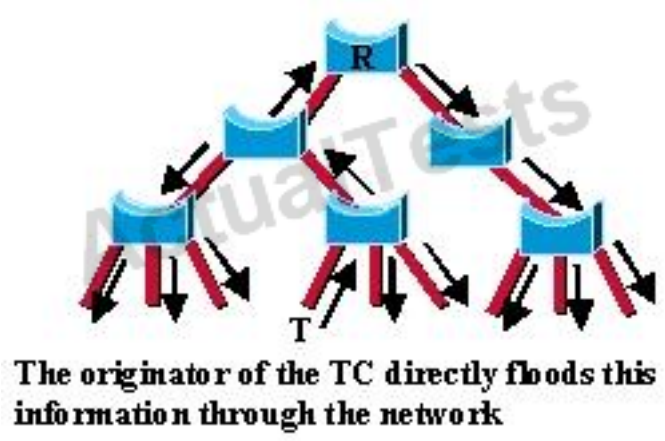
BPDUs are also sent on the root port while the timer is active.

Topology Change Propagation

When a bridge receives a BPDU with the TC bit set from a neighbor, these occur:

It clears the MAC addresses learned on all its ports, except the one that receives the topology change.

It starts the TC While timer and sends BPDUs with TC set on all its designated ports and root port (RSTP no longer uses the specific TCN BPDU, unless a legacy bridge needs to be notified). This way, the TCN floods very quickly across the whole network. The TC propagation is now a one step process. In fact, the initiator of the topology change floods this information throughout the network, as opposed to 802.1D where only the root did. This mechanism is much faster than the 802.1D equivalent. There is no need to wait for the root bridge to be notified and then maintain the topology change state for the whole network for <max age plus forward delay> seconds.



In just a few seconds, or a small multiple of hello-times, most of the entries in the CAM tables of the entire network (VLAN) flush. This approach results in potentially more temporary flooding, but on the other hand it clears potential stale information that prevents rapid connectivity restitution.

Reference
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_white_paper09186a0080094cfa.shtml

QUESTION NO: 9

While you are troubleshooting network performance issues, you notice that a switch is periodically flooding all unicast traffic. Further investigation reveals that periodically the switch is also having spikes in CPU utilization, causing the MAC address table to be flushed and relearned. What is the most likely cause of this issue?

A. a routing protocol that is flooding updates
B. a flapping port that is generating BPDUs with the TCN bit set
C. STP is not running on the switch
D. a user that is downloading the output of the show-tech command
E. a corrupted switch CAM table

Answer: B
Explanation:
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. TCNs are triggered by a port that is transitioning to or from the forwarding state. After the TCN, even if the particular destination MAC address has aged out, flooding should not happen for long in most cases since the address will be relearned. The issue might arise when TCNs are occurring repeatedly with short intervals. The switches will constantly be fast-aging their forwarding tables so flooding will be nearly constant.
Normally, a TCN is rare in a well-configured network. When the port on a switch goes up or down, there is eventually a TCN once the STP state of the port is changing to or from forwarding. When the port is flapping, repetitive TCNs and flooding occurs.
Ports with the STP portfast feature enabled will not cause TCNs when going to or from the forwarding state. Configuration of portfast on all end-device ports (such as printers, PCs, servers,

and so on) should limit TCNs to a low amount. Refer to this document for more information on TCNs:
Understanding Spanning-Tree Protocol Topology Changes Note: In MSFC IOS, there is an optimization that will trigger VLAN interfaces to repopulate their ARP tables when there is a TCN in the respective VLAN. This limits flooding in case of TCNs, as there will be an ARP broadcast and the host MAC address will be relearned as the hosts reply to ARP.

Reference
http://www.cisco.com/en/US/products/hw/switches/ps700/products_tech_note09186a00801d0808.shtml

QUESTION NO: 10

Your network is suffering from regular outages. After troubleshooting, you learn that the transmit lead of a fiber uplink was damaged. Which two features can prevent the same issues in the future? (Choose two.)

A. root guard
B. loop guard
C. BPDU guard
D. UDLD
E. BPDU skew detection

Answer: B,D
Explanation: STP Loop Guard
The STP loop guard feature provides additional protection against Layer 2 forwarding loops (STP loops). An STP loop is created when an STP blocking port in a redundant topology erroneously transitions to the forwarding state. This usually happens because one of the ports of a physically redundant topology (not necessarily the STP blocking port) no longer receives STP BPDUs. In its operation, STP relies on continuous reception or transmission of BPDUs based on the port role. The designated port transmits BPDUs, and the non-designated port receives BPDUs. When one of the ports in a physically redundant topology no longer receives BPDUs, the STP conceives that the topology is loop free. Eventually, the blocking port from the alternate or backup port becomes designated and moves to a forwarding state. This situation creates a loop. The loop guard feature makes additional checks. If BPDUs are not received on a non-designated port, and loop guard is enabled, that port is moved into the STP loop-inconsistent blocking state, instead of the listening / learning / forwarding state. Without the loop guard feature, the port assumes the designated port role. The port moves to the STP forwarding state and creates a loop.

Loop Guard versus UDLD

Loop guard and Unidirectional Link Detection (UDLD) functionality overlap, partly in the sense that both protect against STP failures caused by unidirectional links. However, these two features differ in functionality and how they approach the problem. This table describes loop guard and UDLD functionality:

Functionality	Loop Guard	UDLD
Configuration	Per-port	Per-port
Action granularity	Per-VLAN	Per-port
Autorecover	Yes	Yes, with err-disable timeout feature
Protection against STP failures caused by unidirectional links	Yes, when enabled on all root and alternate ports in redundant topology	Yes, when enabled on all links in redundant topology
Protection against STP failures caused by problems in the software (designated switch does not send BPDU)	Yes	No
Protection against miswiring.	No	Yes

ActualTests

Based on the various design considerations, you can choose either UDLD or the loop guard feature. In regards to STP, the most noticeable difference between the two features is the absence of protection in UDLD against STP failures caused by problems in software. As a result, the designated switch does not send BPDUs.

However, this type of failure is (by an order of magnitude) more rare than failures caused by unidirectional links. In return, UDLD might be more flexible in the case of unidirectional links on EtherChannel. In this case, UDLD disables only failed links, and the channel should remain functional with the links that remain. In such a failure, the loop guard puts it into loop-inconsistent state in order to block the whole channel.

Additionally, loop guard does not work on shared links or in situations where the link has been unidirectional since the link-up. In the last case, the port never receives BPDU and becomes designated. Because this behavior could be normal, this particular case is not covered by loop guard. UDLD provides protection against such a scenario.

As described, the highest level of protection is provided when you enable UDLD and loop guard.

Reference

http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a0080094640.shtml#loop_guard_vs_uld

QUESTION NO: 11

When troubleshooting the issue, you notice the election of a new root bridge with an unknown MAC address. Knowing that all access ports have the PortFast feature enabled, what would be the easiest way to resolve the issue without losing redundant links?

A. Enable bpduguard globally.
B. Enable rootguard.
C. Enable loopguard.
D. Enable spanning tree.
E. Enable UDLD.

Answer: A

Explanation:

Loopguard, spanning tree, and UDLD are obvious red herrings. This leaves enabling rootguard or bpduguard. One key is that enabling bpduguard only affects ports that have portfast enabled; see the following URL under "Configuration."
Reference:
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a008009482f.shtml

QUESTION NO: 12

You are the network administrator of a small Layer 2 network with 50 users. Lately, users have been complaining that the network is very slow. While troubleshooting, you notice that the CAM table of your switch is full, although it supports up to 12, 000 MAC addresses. How can you solve this issue and prevent it from happening in the future?

A. Upgrade the switches
B. Configure BPDU guard
C. Configure VLAN access lists
D. Configure port security
E. Configure Dynamic ARP inspection

Answer: D

Explanation:

Enabling Port Security
Port security is either autoconfigured or enabled manually by specifying a MAC address. If a MAC address is not specified, the source address from the incoming traffic is autoconfigured and secured, up to the maximum number of MAC addresses allowed. These autoconfigured MAC addresses remain secured for a time, depending upon the aging timer set. The autoconfigured MAC addresses are cleared from the port in case of a link-down event.

When you enable port security on a port, any dynamic CAM entries that are associated with the port are cleared. If there are any currently configured static or permanent CAM entrie on that same port, you may not be able to enable the port-security on that port. If this is the case, clear the configured static and permanent earl entries on that port and then enable port-security.
To enable port security, perform this task in privileged mode:

	Task	Command
Step 1	Enable port security on the desired ports. If desired, specify the secure MAC address.	**set port security** *mod_num/port_num* **enable** [*mac_addr*]
Step 2	You can add MAC addresses to the list of secure addresses.	**set port security** *mod_num/port_num* **mac_addr**
Step 3	Verify the configuration.	**show port** [*mod_num*[/*port_num*]]

Actualtests

Reference
http://www.cisco.com/en/US/docs/switches/lan/catalyst4000/8.2glx/configuration/guide/sec_port.html

QUESTION NO: 13

In Layer 2 topologies, spanning-tree failures can cause loops in the network. These unblocked loops can cause network failures because of excessive traffic. Which two Catalyst 6500 features can be used to limit excessive traffic during spanning-tree loop conditions? (Choose two.)

A. loop guard
B. storm control
C. storm suppression
D. broadcast suppression
E. BPDU guard

Answer: B,D
Explanation: Explanation
Traffic Storm Control
A traffic storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. The traffic storm control feature prevents LAN ports from being disrupted by a broadcast, multicast, or unicast traffic storm on physical interfaces.
Traffic storm control (also called traffic suppression) monitors incoming traffic levels over a 1-

second traffic storm control interval and, during the interval, compares the traffic level with the traffic storm control level that you configure. The traffic storm control level is a percentage of the total available bandwidth of the port. Each port has a single traffic storm control level that is used for all types of traffic (broadcast, multicast, and unicast).
Traffic storm control monitors the level of each traffic type for which you enable traffic storm control in 1-second traffic storm control intervals. Within an interval, when the ingress traffic for which traffic storm control is enabled reaches the traffic storm control level that is configured on the port, traffic storm control drops the traffic until the traffic storm control interval ends.
Broadcast suppression Broadcast suppression prevents the switched ports on a LAN from being disrupted by a broadcast storm on one of the ports. A LAN broadcast storm occurs when the broadcast or multicast packets flood the LAN, creating excessive traffic and degrading the network performance. Errors in the protocol-stack implementation or in the network configuration can cause a broadcast storm.
Broadcast suppression uses filtering that measures the broadcast activity on a LAN over a time period (15264 nsec to ~1 sec) that varies based on the type of line card and speed setting on the port, and compares the measurement with a predefined threshold. If the threshold is reached, further broadcast activity is suppressed for the duration of a specified time period. Broadcast suppression is disabled by default.

Reference
http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SXF/native/configuration/guide/storm.html\
http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/catos/8.x/configuration/guide/bcastsup.html

QUESTION NO: 14

Which two statements are true about 802.1s? (Choose two.)

A. 802.1s supports a reduced number of spanning-tree instances.
B. 802.1s has better convergence times than 802.1w.
C. 802.1s does not support load balancing over the same physical topology.
D. The CPU utilization for 802.1s is lower than the CPU utilization for 802.1w.

Answer: A,D
Reference:
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_white_paper09186a0080094cfc.shtml

QUESTION NO: 15

Which configuration is used to enable root guard?

A. interface gig3/1
spanning-tree guard root
B. interface gig3/1
spanning-tree root guard
C. interface gig3/1
spanning-tree root-guard
D. interface gig3/1
spanning-tree root-guard default

Answer: A
Reference:
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a0080094640.shtml

QUESTION NO: 16

Which two statements describe spanning-tree BPDU processing for a blocking port? (Choose two.)

A. BPDUs that enter a blocking port are discarded.
B. BPDUs that enter a blocking port are processed.
C. Loopguard puts an interface into a loop-inconsistent state when BPDUs stop being received on a blocking port.
D. BPDUs are only processed on forwarding ports.

Answer: B,C
Reference: http://www.infraworld.eu/spanning-tree-root-guard-and-loop-guard/

QUESTION NO: 17

When troubleshooting duplex mismatches, which two are errors that are typically seen on the half duplex end? (Choose two.)

A. excessive collisions
B. FCS errors
C. runts
D. late collisions

Answer: A,D
Reference: http://www.netcordia.com/resources/tech-tips/switch-port-duplex-mismatch.asp

QUESTION NO: 18

You are using VTP (version 2) in your network to transport VLAN information between switches. When adding a switch to the network (that has been used in the lab previously), you notice that a lot of the existing VLANs have been deleted or replaced with other names. What can you do to prevent this from happening in the future, without losing all VTP features that you are using today?

A. configure a hard-to-guess VTP domain name
B. use a hard-to-guess VTP password
C. use VTP transparent mode
D. implement VTP version 3

Answer: D
Explanation:
http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/solution_guide_c78_508010.html

QUESTION NO: 19

Which two combinations are valid PAgP configurations that will set up a PAgP channel? (Choose two.)

A. On-Passive
B. On-Auto
C. Passive-Active
D. Desirable-Auto
E. Active-Active
F. Desirable-Desirable

Answer: D,F
Reference:

http://www.cisco.com/en/US/products/hw/switches/ps607/products_configuration_example09186a0080094789.shtml (see port aggregation protocol)

QUESTION NO: 20

Spanning Tree Protocol IEEE 802.1 s defines the ability to deploy which of these?

A. one global STP instance for all VLANs
B. one STP instance for each VLAN
C. one STP instance per set of VLANs
D. one STP instance per set of bridges

Answer: C
Explanation:
The IEEE 802.1s standard is the Multiple Spanning Tree (MST). With MST, you can group VLANs and run one instance of Spanning Tree for a group of VLANs. Other STP types:
Common Spanning Tree (CST), which is defined with IEEE 802.1Q, defines one spanning tree instance for all VLANs.
Rapid Spanning Tree (RSTP), which is defined with 802.1w, is used to speed up STP convergence. Switch ports exchange an explicit handshake when they transition to forwarding.

QUESTION NO: 21

Which two of these are used in the selection of a root bridge in a network utilizing Spanning Tree Protocol IEEE 802.1 D? (Choose two.)

A. Designated Root Cost
B. bridge ID priority
C. max age
D. bridge ID MAC address
E. Designated Root Priority
F. forward delay

Answer: B,D
Explanation: Explanation
The root bridge of the spanning tree is the bridge with the smallest (lowest) bridge ID. Each bridge has a unique identifier (ID) and a configurable priority number; the bridge ID contains both

numbers. To compare two bridge IDs, the priority is compared first. If two bridges have equal priority, then the MAC addresses are compared. For example, if switches A (MAC=0200.0000.1111) and B (MAC=0200.0000.2222) both have a priority of 10, then switch A will be selected as the root bridge. If the network administrators would like switch B to become the root bridge, they must set its priority to be less than 10.

QUESTION NO: 22

If a port configured with STP loop guard stops receiving BPDUs, the port will be put into which state?

A. learning state
B. listening state
C. forwarding state
D. loop-inconsistent state

Answer: D
Explanation: Explanation
STP Loop Guard
Feature Description
The STP loop guard feature provides additional protection against Layer 2 forwarding loops (STP loops). An STP loop is created when an STP blocking port in a redundant topology erroneously transitions to the forwarding state. This usually happens because one of the ports of a physically redundant topology (not necessarily the STP blocking port) no longer receives STP BPDUs. In its operation, STP relies on continuous reception or transmission of BPDUs based on the port role. The designated port transmits BPDUs, and the non-designated port receives BPDUs.

When one of the ports in a physically redundant topology no longer receives BPDUs, the STP conceives that the topology is loop free. Eventually, the blocking port from the alternate or backup port becomes designated and moves to a forwarding state. This situation creates a loop.
The loop guard feature makes additional checks. If BPDUs are not received on a non-designated port, and loop guard is enabled, that port is moved into the STP loop-inconsistent blocking state, instead of the listening / learning / forwarding state. Without the loop guard feature, the port assumes the designated port role. The port moves to the STP forwarding state and creates a loop. When the loop guard blocks an inconsistent port, this message is logged:

CatOS%SPANTREE-2-LOOPGUARDBLOCK: No BPDUs were received on port 3/2 in vlan 3. Moved to loop-inconsistent state.
Cisco IOS%SPANTREE-2-LOOPGUARD_BLOCK: Loop guard blocking port FastEthernet0/24 on

VLAN0050.
Once the BPDU is received on a port in a loop-inconsistent STP state, the port transitions into another STP state. According to the received BPDU, this means that the recovery is automatic and intervention is not necessary. After recovery, this message is logged:
CatOS%SPANTREE-2-LOOPGUARDUNBLOCK: port 3/2 restored in vlan 3.
Cisco IOS%SPANTREE-2-LOOPGUARD_UNBLOCK: Loop guard unblocking port FastEthernet0/24 on VLAN0050.

Reference
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a0080094640.shtml

QUESTION NO: 23

What is the purpose of the STP PortFast BPDU guard feature?

A. enforce the placement of the root bridge in the network
B. ensure that a port is transitioned to a forwarding state quickly if a BPDU is received
C. enforce the borders of an STP domain
D. ensure that any BPDUs received are forwarded into the STP domain

Answer: C
Explanation: Explanation
STP configures meshed topology into a loop-free, tree-like topology. When the link on a bridge port goes up, STP calculation occurs on that port. The result of the calculation is the transition of the port into forwarding or blocking state. The result depends on the position of the port in the network and the STP parameters. This calculation and transition period usually takes about 30 to 50 seconds. At that time, no user data pass via the port. Some user applications can time out during the period.

In order to allow immediate transition of the port into forwarding state, enable the STP PortFast feature.

PortFast immediately transitions the port into STP forwarding mode upon linkup. The port still participates in STP. So if the port is to be a part of the loop, the port eventually transitions into STP blocking mode.

As long as the port participates in STP, some device can assume the root bridge function and affect active STP topology. To assume the root bridge function, the device would be attached to the port and would run STP with a lower bridge priority than that of the current root bridge. If

another device assumes the root bridge function in this way, it renders the network suboptimal. This is a simple form of a denial of service (DoS) attack on the network. The temporary introduction and subsequent removal of STP devices with low (0) bridge priority cause a permanent STP recalculation.

The STP PortFast BPDU guard enhancement allows network designers to enforce the STP domain borders and keep the active topology predictable. The devices behind the ports that have STP PortFast enabled are not able to influence the STP topology. At the reception of BPDUs, the BPDU guard operation disables the port that has PortFast configured. The BPDU guard transitions the port into errdisable state, and a message appears on the console. This message is an example:
2000 May 12 15:13:32 %SPANTREE-2-RX_PORTFAST:Received BPDU on PortFast enable port.

Disabling 2/1
2000 May 12 15:13:32 %PAGP-5-PORTFROMSTP:Port 2/1 left bridge port 2/1

Reference
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a008009482f.shtml#topic1

QUESTION NO: 24

When STP UplinkFast is enabled on a switch utilizing the default bridge priority, what will the new bridge priority be changed to?

A. 8192
B. 16384
C. 49152
D. 65535

Answer: C
Explanation:
Explanation
The STP UplinkFast is used to fast switchover to alternate ports when the root port fails. When STP UplinkFast is enabled on a switch utilizing the default bridge priority (32768), the new bridge priority will be changed to 49152. The reason for the priority being raised is to prevent the switch from becoming the root (recall that lower bridge priority is preferred). To enable UplinkFast feature, use the "set spantree uplinkfast enable" in privileged mode
The set spantree uplinkfast enable command has the following results:

Changes the bridge priority to 49152 for all VLANs (allowed VLANs).
Increases the path cost and portvlancost of all ports to a value greater than 3000.
On detecting the failure of a root port, an instant cutover occurs to an alternate port selected by Spanning Tree Protocol (without using this feature, the network will need about 30 seconds to re-establish the connection.

Reference
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a0080094641.shtml

QUESTION NO: 25

Which of these best describes the actions taken when a VTP message is received on a switch configured with the VTP mode "transparent"?

A. VTP updates are ignored and forwarded out all ports.
B. VTP updates are ignored and forwarded out trunks only.
C. VTP updates are made to the VLAN database and are forwarded out trunks only.
D. VTP updates are ignored and are not forwarded.

Answer: B
Explanation: Explanation
You can configure a switch to operate in any one of these VTP modes:
Server--In VTP server mode, you can create, modify, and delete VLANs and specify other configuration parameters, such as VTP version and VTP pruning, for the entire VTP domain. VTP servers advertise their VLAN configuration to other switches in the same VTP domain and synchronize their VLAN configuration with other switches based on advertisements received over trunk links. VTP server is the default mode.

Client--VTP clients behave the same way as VTP servers, but you cannot create, change, or delete VLANs on a VTP client.

Transparent--VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements, but transparent switches do forward VTP advertisements that they receive out their trunk ports in VTP Version 2. Off (configurable only in CatOS switches)--In the three described modes, VTP advertisements are received and transmitted as soon as the switch enters the management domain state. In the VTP off mode, switches behave the same as in VTP transparent mode with the exception that VTP advertisements are not forwarded.

VTP V2
VTP V2 is not much different than VTP V1. The major difference is that VTP V2 introduces support for Token Ring VLANs. If you use Token Ring VLANs, you must enable VTP V2. Otherwise, there is no reason to use VTP V2. Changing the VTP version from 1 to 2 will not cause a switch to reload.
VTP Password If you configure a password for VTP, you must configure the password on all switches in the VTP domain. The password must be the same password on all those switches. The VTP password that you configure is translated by algorithm into a 16-byte word (MD5 value) that is carried in all summary-advertisement VTP packets.

VTP Pruning
VTP ensures that all switches in the VTP domain are aware of all VLANs. However, there are occasions when VTP can create unnecessary traffic. All unknown unicasts and broadcasts in a VLAN are flooded over the entire VLAN. All switches in the network receive all broadcasts, even in situations in which few users are connected in that VLAN. VTP pruning is a feature that you use in order to eliminate or prune this unnecessary traffic.

Reference
http://www.cisco.com/en/US/tech/tk389/tk689/technologies_tech_note09186a0080094c52.shtml

QUESTION NO: 26

On what type of ports would STP Port Fast BPDU guard be most appropriate?

A. root ports
B. Designated ports
C. Host ports
D. alternate ports

Answer: C
Explanation:
BPDU Guard is a basic feature that will automatically shut down a port when BPDU’s are received on that particular port. It is common to configure PortFast and BPDU Guard on host access ports.

QUESTION NO: 27

The 802.1w protocol is seen as the next evolution beyond the 802.1 D standard protocol. Which of

these statements regarding port states is true of both 802.1 D and 802.1w?

A. All 802.1 D port states (Disabled, Blocking, Listening, Learning, and Forwarding) are identical in 802.1w.
B. The 802.1 D port states Disabled and Blocking have become the 802.1w port state Discarding, and all other 802.1D port states remain the same in 802.1w.
C. The 802.1 D port states Disabled, Blocking, and Listening have become the 802.1w port state Discarding, and all other 802.1D port states remain the same in 802.1w.
D. The 802.1 D port states Disabled, Blocking, and Listening have been removed completely from 802.1w (there is no corresponding port state), and all other 802.1 D port states remain the same in 802.1w.) E. The 802.1 D port state Disabled has been removed from 802.1w, and the 802.1 D port states Blocking and Listening have become the 802.1w port state Discarding; all other 802.1D port states remain the same in 802.1w.

Answer: C
Explanation:

Table 7-3 Port State Comparison

Operational Status	STP Port State	RSTP Port State	Is Port Included in the Active Topology?
Enabled	Blocking	Discarding	No
Enabled	Listening	Discarding	No
Enabled	Learning	Learning	Yes
Enabled	Forwarding	Forwarding	Yes
Disabled	Disabled	Discarding	No

ActualTests

QUESTION NO: 28

The classic Spanning Tree Protocol (802.1 D 1998) uses which sequence of variables to determine the best received BPDU?

A. 1) lowest root bridge id, 2) lowest sender bridge id, 3) lowest port id, 4) lowest root path cost
B. 1) lowest root path cost, 2) lowest root bridge id, 3) lowest sender bridge id, 4) lowest sender port id
C. 1) lowest root bridge id, 2) lowest sender bridge id, 3) lowest root path cost 4) lowest sender port id
D. 1) lowest root bridge id, 2) lowest root path cost, 3) lowest sender bridge id, 4) lowest sender

port id

Answer: D

Explanation: Configuration bridge protocol data units (BPDUs) are sent between switches for each port. Switches use a fourstep process to save a copy of the best BPDU seen on every port. When a port receives a better BPDU, it stops sending them. If the BPDUs stop arriving for 20 seconds (default), it begins sending them again.

Step 1 Lowest Root Bridge ID (BID)
Step 2 Lowest Path cost to Root Bridge
Step 3 Lowest Sender BID
Step 4 Lowest Port ID

Reference
Cisco General Networking Theory Quick Reference Sheets

QUESTION NO: 29

Which three port states are used by RSTP 802.1w? (Choose three.)

A. Listening
B. Learning
C. Forwarding
D. Blocking
E. Discarding
F. Disabled

Answer: B,C,E

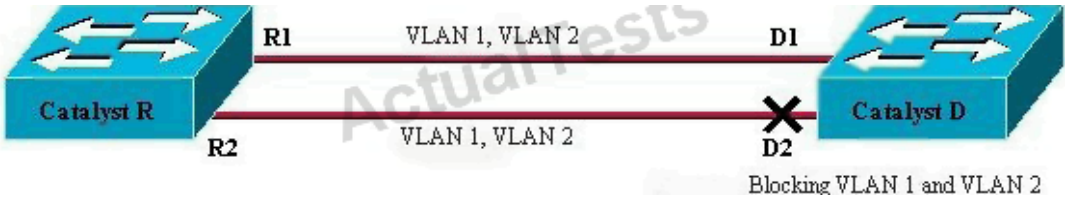
Explanation:

RSTP and STP Port States		
Administrative State	**STP State (802.1d)**	**RSTP State (802.1w)**
Disabled	Disabled	Discarding
Enabled	Blocking	Discarding
Enabled	Listening	Discarding
Enabled	Learning	Learning
Enabled	Forwarding	Forwarding

ActualTests

QUESTION NO: 30

Refer to the exhibit. Catalyst R is the root bridge for both VLAN 1 and VLAN 2. What is the easiest way to load-share traffic across both trunks and maintain redundancy in case a link fails, without using any type of EtherChannel link-bundling?



A. Increase the root bridge priority (increasing the numerical priority number) for VLAN 2 on Catalyst D so that port D2 becomes the root port on Catalyst D for VLAN 2.
B. Decrease the port priority on R2 for VLAN 2 on Catalyst R so that port D1 will be blocked for VLAN 2 and port D2 will remain blocked for VLAN 1.
C. Decrease the path cost on R2 on Catalyst R for VLAN 2 so that port D1 will be blocked for VLAN 2 and port D2 will remain blocked for VLAN 1.
D. Increase the root bridge priority (decreasing the numerical priority number) for VLAN 2 on Catalyst R so that R2 becomes the root port on Catalyst D for VLAN 2.

Answer: B
Explanation:
First we should understand what will happen if nothing is configured (use default values). Because R is the root bridge so all of its ports will forward. D will need to block one of its ports to avoid a bridging loop between the two switches. But how does D select its blocked port? Well, the answer is based on the BPDUs it receives from R. A BPDU is superior than another if it has:

1. A lower Root Bridge ID
2. A lower path cost to the Root
3. A lower Sending Bridge ID
4. A lower Sending Port ID

These four parameters are examined in order. In this case, all the BPDUs sent by R have the same Root Bridge ID, same path cost to the Root and same Sending Bridge ID. The only parameter left to select the best one is the Sending Port ID (Port ID = port priority + port index). If using default values, the default port priority's value is 32 or 128 (128 is much more popular today), so D will compare port index values, which are unique to each port on the switch, and because port R2 is inferior to port R1 (the port's number of R2 is higher than that of R1, for

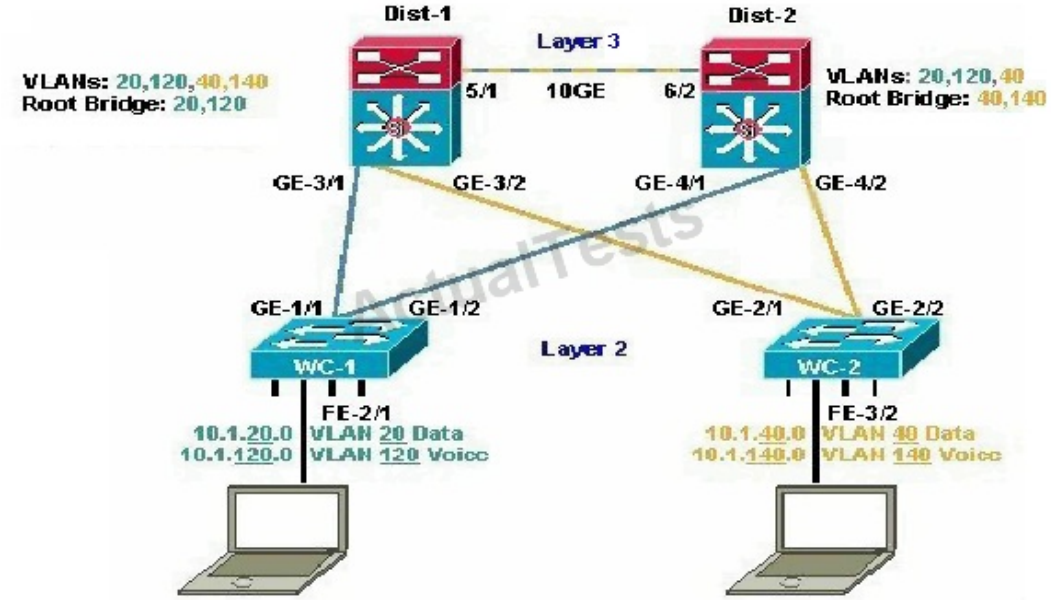
example port Fa0/2 is inferior to port Fa0/1), D will select the port connected with port R1 as its root port and block the other port.

The problem here is port D2 is blocked for both VLAN 1 & 2 and that means we can't use the underneath link for load-sharing. The underneath link is just used in the case the above link fails.

Now as you can guess, the easiest way to load-share traffic across both trunks is decreasing the port priority on R2 for VLAN 2 on Catalyst R so that port D1 will be blocked for VLAN 2. Notice that "decreasing" here means make that port ID superior to the other port.

QUESTION NO: 31

Refer to the exhibit. In the diagram, the switches are running IEEE 802.1s MST. Which ports are in the MST blocking state?



A. GE-1/2andGE2/1
B. GE-1/1 and GE-2/2
C. GE-3/2 and GE 4/1
D. no ports are in the blocking state
E. There is not enough information to determine which ports are in the blocking state.

Answer: A

Explanation:

Switches Dist-1 & 2 will have no ports in a blocking stating. However, switches WC-1 & 2 will have the secondary ports Ge1/2 & Ge2/1 in an Alternate/Blocking state as this will be backup ports for

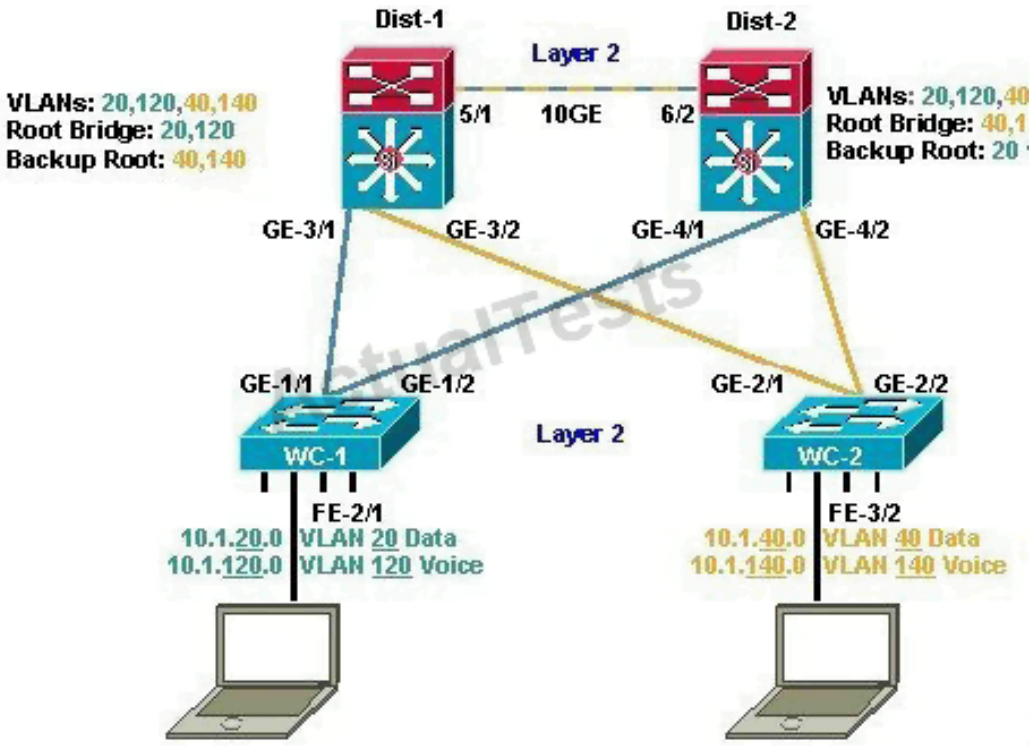
the root port.
Ge1/2 & Ge2/1 will transition to a forwarding state and become root ports if ports Ge1/1 & Ge2/2 go down.
Multiple Spanning Tree (MST) is an IEEE standard inspired from the Cisco proprietary Multiple Instances Spanning Tree Protocol (MISTP) implementation. The main enhancement introduced by MST is that several VLANs can be mapped to a single spanning tree instance. This raises the problem of how to determine which VLAN is to be associated with which instance. More precisely, how to tag BPDUs so that the receiving devices can identify the instances and the VLANs to which each device applies.

MST Configuration and MST Region
Each switch running MST in the network has a single MST configuration that consists of these three attributes:
1. An alphanumeric configuration name (32 bytes)
2. A configuration revision number (two bytes)
3. A 4096-element table that associates each of the potential 4096 VLANs supported on the chassis to a given instance
In order to be part of a common MST region, a group of switches must share the same configuration attributes.
It is up to the network administrator to properly propagate the configuration throughout the region. Currently, this step is only possible by the means of the command line interface (CLI) or through Simple Network Management Protocol (SNMP). Other methods can be envisioned, as the IEEE specification does not explicitly mention how to accomplish that step.
Note: If for any reason two switches differ on one or more configuration attribute, the switches are part of different regions. For more information refer to the Region Boundary section of this document.

Reference
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_white_paper09186a0080094cfc.shtml

QUESTION NO: 32

Refer to the exhibit. In the diagram, the switches are running IEEE 802.1 w RSPT. On which ports should root guard be enabled in order to facilitate deterministic root bridge election under normal and failure scenarios?



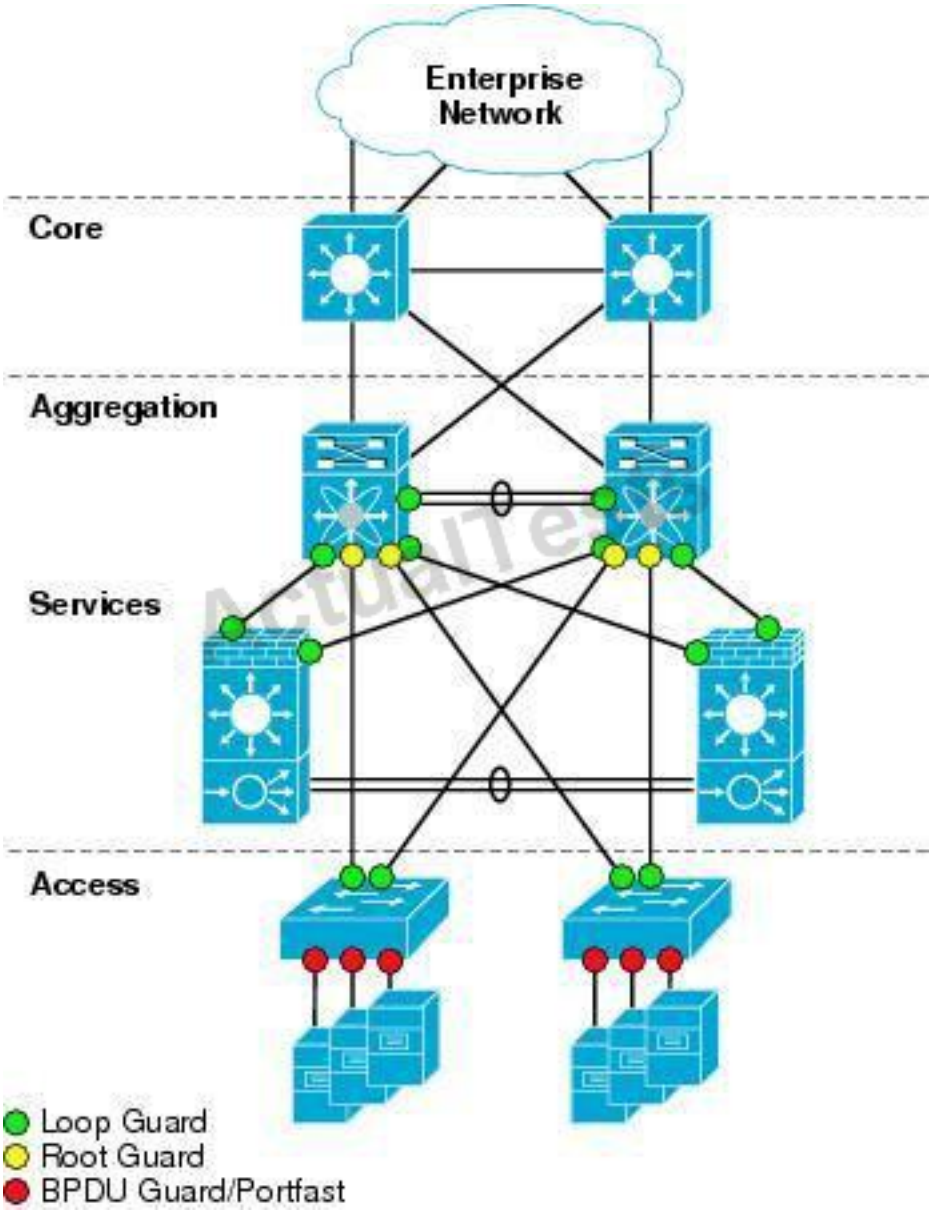
A. GE-3/1, GE-3/2
B. FE-2/1, FE-3/2
C. GE-1/1.GE-1/2
D. GE-4/1, GE-4/2
E. GE-2/1, GE-2/2
F. GE-3/1, GE-3/2, GE-4/1, GE-4/2, FE-2/1, FE-3/2

Answer: F

Explanation:

Root Guard is a Cisco-specific feature that prevents a Layer 2 switched port from becoming a root port. It is enabled on ports other than the root port and on switches other than the root. If a Root Guard port receives a BPDU that might cause it to become a root port, then the port is put into "root-inconsistent" state and does not pass traffic through it. If the port stops receiving these BPDUs, it automatically re-enables itself.

This feature is sometimes recommended on aggregation layer ports that are facing the access layer, to ensure that a configuration error on an access layer switch cannot cause it to change the location of the spanning tree root switch (bridge) for a given VLAN or instance. Below is a recommended port's features should be enabled in a network.



Reference
http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/nx_7000_dc.html
The port FE-2/1 & FE-3/2 should be turned on the Root Guard feature because hackers can try to plug these ports into other switches or try to run a switch-simulation software on these PCs. Imagine a new switch that is introduced into the network with a bridge priority lower than the current root bridge. In a normal STP operation, this new bridge can become the new Root Bridge and disrupt your carefully designed network. The recommended design is to enable Root Guard on all access ports so that a root bridge is not established through this port.
Note: The Root Guard affects the entire port. Therefore it applies to all VLANs on that port. To enable this feature, use the following command in interface configuration:

Switch(config-if)# spanning-tree guard root

QUESTION NO: 33

Loop guard and UniDireclional Link Detection both protect against Layer 2 STP loops. In which two ways does loop guard differ from UDLD in loop detection and prevention? (Choose two.)

A. Loop guard can be used with root guard simultaneously on the same port on the same VLAN while UDLD cannot.
B. UDLD protects against STP failures caused by cabling problems that create one-way links.
C. Loop guard detects and protects against duplicate packets being received and transmitted on different ports.
D. UDLD protects against unidirectional cabling problems on copper and fiber media.
E. Loop guard protects against STP failures caused by problems that result in the loss of BPDUs from a designated switch port.

Answer: B,E

Explanation:

Answers B, D, & E are all correct. However, as UDLD is only enabled on Fiber media by default I have selected B instead of D as you will have to manually configure UDLD if you want it to work on copper media.

The Cisco-proprietary UDLD protocol allows devices connected through fiber-optic or copper (for example, Category 5 cabling) Ethernet cables connected to LAN ports to monitor the physical configuration of the cables and detect when a unidirectional link exists. When a unidirectional link is detected, UDLD shuts down the affected LAN port and alerts the user. Unidirectional links can cause a variety of problems, including spanning tree topology loops.
UDLD is a Layer 2 protocol that works with the Layer 1 protocols to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected LAN ports. When you enable both autonegotiation and UDLD, Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.
Based on the various design considerations, you can choose either UDLD or the loop guard feature. In regards to STP, the most noticeable difference between the two features is the absence of protection in UDLD against STP failures caused by problems in software. As a result, the designated switch does not send BPDUs.
However, this type of failure is (by an order of magnitude) more rare than failures caused by unidirectional links. In return, UDLD might be more flexible in the case of unidirectional links on EtherChannel. In this case,
UDLD disables only failed links, and the channel should remain functional with the links that remain. In such a failure, the loop guard puts it into loop-inconsistent state in order to block the whole channel.

Additionally, loop guard does not work on shared links or in situations where the link has been unidirectional since the link-up. In the last case, the port never receives BPDU and becomes designated. Because this behavior could be normal, this particular case is not covered by loop guard. UDLD provides protection against such a scenario.

QUESTION NO: 34

Which standard supports multiple instances of spanning tree?

A. 802.1 D
B. 802.1s
C. 802.1w
D. 802.1 z

Answer: B
Explanation:
Multiple Spanning Tree Protocol (MSTP) was first specified in IEEE 802.1s and is standardized in IEEE 802.1Q. MSTP enables multiple VLANs to be mapped to the same spanning-tree instance, reducing the number of spanning-tree instances needed to support a large number of VLANs. MSTP provides multiple forwarding paths for data traffic and enables load balancing. It improves the fault tolerance of the network because a failure in one instance, or forwarding path, does not affect other instances

QUESTION NO: 35

Spanning Tree Protocol calculates path cost based on which of these?

A. interface bandwidth
B. interface delay
C. interface bandwidth and delay
D. hop count
E. bridge priority

Answer: A
Explanation:
STP calculates the path cost based on the media speed (bandwidth) of the links between switches and the port cost of each port forwarding frame. Spanning tree selects the root port based on the

path cost. The port with the lowest path cost to the root bridge becomes the root port. The root port is always in the forwarding state.
If the speed/duplex of the port is changed, spanning tree recalculates the path cost automatically. A change in the path cost can change the spanning tree topology.

Data rate and STP path cost
The table below shows the default cost of an interface for a given data rate.

Data rate	STP Cost (802.1D-1998)	RSTP Cost (802.1W-2001)
4 Mbit/s	250	5,000,000
10 Mbit/s	100	2,000,000
16 Mbit/s	62	1,250,000
100 Mbit/s	19	200,000
1 Gbit/s	4	20,000
2 Gbit/s	3	10,000
10 Gbit/s	2	2,000

ActualTests

QUESTION NO: 36

What does the root guard feature provide in a bridged network?

A. it ensures that the bridge is elected as Root Bridge in the network.
B. it enforces the root bridge placement in the network
C. It ensures that BPDUs sent by the root bridge are forwarded in a timely manner.
D. It ensures that all ports receiving BPDUs from the root bridge are in the forwarding state.

Answer: B
Explanation:
Root Guard—Enabled per port; ignores any received superior BPDUs to prevent a switch

connected to this port from becoming root. Upon receipt of superior BPDUs, this switch puts the port in a loop-inconsistent state, ceasing forwarding and receiving frames until the superior BPDUs cease.
The STP topology can be changed based on one of these unexpected and undesired switches being added to the network. For instance, this newly added and unexpected switch might have the lowest bridge ID and become the root. To prevent such problems, BPDU Guard and Root Guard can be enabled on these access ports to monitor for incoming BPDUs.

QUESTION NO: 37

Why does RSTP have a better convergence time than 802.1D?

A. it is newer
B. it has smaller timers
C. it has less overhead
D. it is not timer-based

Answer: D
Explanation:
RSTP identifies certain links as point to point. When a point-to-point link fails, the alternate link can transition to the forwarding state.

Although STP provides basic loop prevention functionality, it does not provide fast network convergence when there are topology changes. STP's process to determine network state transitions is slower than RSTP's because it is timer-based. A device must reinitialize every time a topology change occurs. The device must start in the listening state and transition to the learning state and eventually to a forwarding or blocking state.

When default values are used for the maximum age (20 seconds) and forward delay (15 seconds), it takes 50 seconds for the device to converge. RSTP converges faster because it uses a handshake mechanism based on point-to-point links instead of the timer-based process used by STP.

An RSTP domain running switch has the following components:
A root port, which is the "best path" to the root device.
A designated port, indicating that the switch is the designated bridge for the other switch connecting to this port.
An alternate port, which provides an alternate root port.
A backup port, which provides an alternate designated port.

Port assignments change through messages exchanged throughout the domain. An RSTP device generates configuration messages once every hello time interval. If an RSTP device does not receive a configuration message from its neighbor after an interval of three hello times, it determines it has lost connection with that neighbor. When a root port or a designated port fails on a device, the device generates a configuration message with the proposal bit set. Once its neighbor device receives this message, it verifies that this configuration message is better than the one saved for that port and then it starts a synchronizing operation to ensure that all of its ports are in sync with the new information.

Similar waves of proposal agreement handshake messages propagate toward the leaves of the network, restoring the connectivity very quickly after a topology change (in a well-designed network that uses RSTP, network convergence can take as little as 0.5 seconds). If a device does not receive an agreement to a proposal message it has sent, it returns to the original IEEE 802.D convention. RSTP was originally defined in the IEEE 802.1w draft specification and later incorporated into the IEEE 802.1D-2004 specification.

QUESTION NO: 38

Under which two circumstances would an RSTP bridge flush its CAM table? (Choose two.)

A. upon a port state change
B. upon receiving a topology change notification
C. when transitioning from discarding to forwarding
D. when transitioning from forwarding to discarding
E. only when changing from listening to discarding
F. when CAM resources have been completely used up

Answer: B,C
Explanation: Explanation
First, the goal of RSTP is fast re-convergence. Since ports are assumed to transition to forwarding relatively fast, simply increasing MAC address aging speed is not enough. Thus, when a topology change is detected, RSTP instructs the bridge to flush all MAC address table entries. With Ethernet, this process results in unconstrained flooding until the moment MAC addresses are re-learned. The bridge detecting a topology change sets the TC (Topology Change) bit in all outgoing BPDUs and starts sending BPDUs with the TC bit set upstream through the root port as well. This marking lasts for TCWhile=2xHelloTime seconds and allows the detecting bridge the start the flooding process.

QUESTION NO: 39

Which of these correctly identifies a difference between the way BPDUs are handled by 802.1w and 802.1 D?

A. 802.1 D bridges do not relay
B. 802.1w bridges do not relay BPDUs
C. 802.1D bridges only relay BPDUs received from the root
D. 802.1w bridges only relay BPDUs received from the root.

Answer: C
Explanation:
A bridge sends a BPDU frame using the unique MAC address of the port itself as a source address, and a destination address of the STP multicast address 01:80:C2:00:00:00.
There are three types of BPDUs:

Configuration BPDU (CBPDU), used for Spanning Tree computation
Topology Change Notification (TCN) BPDU, used to announce changes in the network topology
Topology Change Notification Acknowledgment (TCA)

BPDU are Sent Every Hello-Time
BPDU are sent every hello-time, and not simply relayed anymore. With 802.1D, a non-root bridge only generates BPDUs when it receives one on the root port. In fact, a bridge relays BPDUs more than it actually generates them. This is not the case with 802.1w. A bridge now sends a BPDU with its current information every <hello-time> seconds (2 by default), even if it does not receive any from the root bridge.

Reference
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_white_paper09186a0080094cfa.shtml#topic4

QUESTION NO: 40

Which three of these statements about Dynamic Trunking Protocol are correct? (Choose three)

A. It supports autonegotiation for both ISL and IEEE 802.1 Q trunks.
B. It must be disabled on an interface if you do not want the interface to work as a trunk or start negotiation to become a trunk.

C. It is a point-to-multipoint protocol.
D. It is a point-to-point protocol.
E. It is not supported on private VLAN ports or tunneling ports

Answer: A,B,D
Explanation:
By default Cisco states that PVLANs will be forwarded. Keep in mind that if you do not disable DTP it will attempt to negotiate a trunk with any additional switch that it is connected to on the port in question. Switchport mode access - This command puts the interface (access port) into permanent nontrunking mode.

The interface will generate DTP frames, negotiating with the neighboring interface to convert the link into a nontrunk link. The interface becomes a nontrunk interface even if the neighboring interface does not agree to the change.

Switchport mode dynamic desirable - This command makes the interface actively attempt to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to trunk, desirable, or auto mode. This is the default mode for all Ethernet interfaces. If the neighboring interface is set to the access or non-negotiate mode, the link will become a non-trunking link.

Switchport mode dynamic auto - This command makes the interface willing to convert the link to a trunk link if the neighboring interface is set to trunk or desirable mode. Otherwise, the link will become a non-trunking link. Switchport mode trunk - This command puts the interface into permanent trunking mode and negotiates to convert the link into a trunk link. The interface becomes a trunk interface even if the neighboring interface does not agree to the change.

Switchport nonegotiate - Prevents the interface from generating DTP frames. You can use this command only when the interface switchport mode is access or trunk. You must manually configure the neighboring interface as a trunk interface to establish a trunk link, otherwise the link will be a non-trunking link.

Using these different trunking modes, an interface can be set to trunking or nontrunking or even able to negotiate trunking with the neighboring interface. To automatically negotiate trunking, the interfaces must be in the same VTP domain. Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which is a Cisco proprietary Point-to-Point Protocol.

QUESTION NO: 41

You are designing your network to be able to use trunks. As part of this process you are comparing the ISL and 802.1 Q encapsulation options. All of these statements about the two encapsulation options are correct except which one?

A. Both support normal and extended VLAN ranges.
B. ISL is a Cisco proprietary encapsulation method and 802.1 Q is an IEEE standard.
C. ISL encapsulates the original frame
D. Both support native VLANs.
E. 802.1 Q does not encapsulate the original frame.

Answer: D

Explanation: Explanation
ISL is a Cisco proprietary protocol for the interconnection of multiple switches and maintenance of VLAN information as traffic goes between switches. ISL provides VLAN trunking capabilities while it maintains full wire-speed performance on Ethernet links in full-duplex or half-duplex mode. ISL operates in a point- to-point environment and can support up to 1000 VLANs. In ISL, the original frame is encapsulated and an additional header is added before the frame is carried over a trunk link. At the receiving end, the header is removed and the frame is forwarded to the assigned VLAN. ISL uses Per VLAN Spanning Tree (PVST), which runs one
instance of Spanning Tree Protocol (STP) per VLAN. PVST allows the optimization of root switch placement for each VLAN and supports the load balancing of VLANs over multiple trunk links. 802.1Q is the IEEE standard for tagging frames on a trunk and supports up to 4096 VLANs. In 802.1Q, the trunking device inserts a 4-byte tag into the original frame and recomputes the frame check sequence (FCS) before the device sends the frame over the trunk link. At the receiving end, the tag is removed and the frame is forwarded to the assigned VLAN. 802.1Q does not tag frames on the native VLAN. It tags all other frames that are transmitted and received on the trunk. When you configure an 802.1Q trunk, you must make sure that you
configure the same native VLAN on both sides of the trunk. IEEE 802.1Q defines a single instance of spanning tree that runs on the native VLAN for all the VLANs in the network. This is called Mono Spanning Tree (MST). This lacks the flexibility and load balancing capability of PVST that is available with ISL. However, PVST+ offers the capability to retain multiple spanning tree topologies with 802.1Q trunking.

QUESTION NO: 42

The network administrator is trying to add Switch1 to the network, but the 802.1 Q trunk is not coming up. Switch1 was previously tested in the laboratory and its trunk configuration worked fine. What are three possible causes of this problem? (Choose three.)

A. The trunking configuration mode on Switch1 is set to Off.

B. The trunking configuration mode on the other end is set to On.
C. The trunking configuration mode on the other end is set to Desirable.
D. Cisco Discovery Protocol is not running on the other end.
E. There is a VTP domain name mismatch.
F. Switch1 does not support 802.1Q.

Answer: A,E,F

Explanation:

There are 5 possible trunking modes for a switch port:

Auto: this is the default mode. In this mode, a port will become a trunk port if the device the port is connected to is set to the on or desirable mode.

Desirable: allows the port to become a trunk port if the device the port is connected to is set to the on, desirable, or auto mode

On: sets the port to permanent trunking mode.

Nonegotiate: sets the port to permanent trunking mode without sending Dynamic Trunking Protocol (DTP)

Frame Off: sets the port to permanent non-trunking mode

In this case, we can guess the trunking mode of Switch 1 is auto (default mode). When in the laboratory, the trunking mode of the other end is set to On or Desirable so 2 switches can negotiate and the link becomes trunk with no problem. But when plugging to the network, other switches may have the trunking mode set to auto so the 802.1Q trunk is not coming up

Of course these switches need to be in the same VTP domain so that they can talk with each other. When trying to configure a trunk negation with a mismatched VTP domain you will receive the following error %DTP-5-DOMAINMISMATCH: Unable to perform trunk negotiation on port Gig0/1 because of VTP domain mismatch.

Note that the following combinations are possible resulting in the trunk status shown:

	OFF	ON	AUTO	DESIR.	NONEG.
OFF	No Trunk	No Trunk	No Trunk	No Trunk	No Trunk
ON	No Trunk	Trunk	Trunk	Trunk	Trunk
AUTO	No Trunk	Trunk	No Trunk	Trunk	No Trunk
DESIR.	No Trunk	Trunk	Trunk	Trunk	Trunk
NONEG.	No Trunk	Trunk	No Trunk	Trunk	Trunk

ActualTests

Reference

Cisco General Networking Theory Quick Reference Sheets

QUESTION NO: 43

You are the network administrator of a large Layer 2 network. At certain times during the day, users complain that the network is responding very slowly.

When troubleshooting the issue, you notice the election of a new root bridge with an unknown MAC address. Knowing that all access ports have the PortFast feature enabled, what should be done to resolve the issue without losing redundant links?

A. Enable bpduguard globally.
B. Enable rootguard.
C. Enable loopguard.
D. Enable spanning tree.
E. Enable UDLD.

Answer: A

Explanation:

Loopguard, spanning tree, and UDLD are obvious red herrings. This leaves enabling rootguard or bpduguard. One key is that enabling bpduguard only affects ports that have portfast enabled; see the following URL under "Configuration."
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a008009482f.shtml

QUESTION NO: 44

An 802.1Q trunk is not coming up between two switches. The ports on both switches are configured as "switchport mode desirable." Assuming that there is no physical issue, choose two possible causes. (Choose two.)

A. Incorrect VTP domain
B. Incorrect VTP password
C. Incorrect VTP mode
D. Incorrect VTP configuration revision

Answer: A,B

Explanation:

Table 2-4. VTP Configuration Options	
Option	**Meaning**
domain	Sends domain name in VTP updates. Received VTP update is ignored if it does not match a switch's domain name. One VTP domain name per switch is allowed.
password	Used to generate an MD5 hash that is included in VTP updates. Received VTP updates are ignored if the passwords on the sending and receiving switch do not match.
mode	Sets server, client, or transparent mode on the switch.
version	Sets version 1 or 2. Servers and clients must match version to exchange VLAN configuration data. Transparent mode switches at version 2 forward version 1 or version 2 VTP updates.
pruning	Enables VTP pruning, which prevents flooding on a per-VLAN basis to switches that do not have any ports configured as members of that VLAN.
interface	Specifies the interface whose IP address is used to identify this switch in VTP updates.

Reference
CCIE Routing and Switching Certification Guide, Fourth Edition

QUESTION NO: 45

Refer to the exhibit. Look at the command output. What would be the most probable reason for this port-ID mismatch?

```
Switch1#show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
Switch2          Gig 1/0/3         160          S I       WS-C3750G- Gig 1/0/11
Switch1#

Switch2#show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
Switch1          Gig 1/0/11        173          R S I     WS-C3750G- Gig 1/0/4
Switch2#
```

A. spanning-tree misconfiguration
B. speed mismatch configuration

C. cabling problem
D. configuration problem

Answer: C

Explanation:

As CDP has different information for Switch1 it is obvious that there is a cabling issue

Verifying Layer 2 Connectivity

Verifying Layer 2 connectivity involves checking switch ports and each connected device at an Ethernet level, ensuring Layer 2 connectivity has been established and also checking errors for a connection. Although all indications might be that a Layer 2 connection is operational, it is also good to actually confirm that Layer 2 traffic is being sent and received between two devices. Cisco Discovery Protocol (CDP) provides such functionality, verifying connectivity to locally connected Cisco devices.

Checking Network Cabling

Many physical layer connectivity problems are caused by faulty network cabling. All cabling between your switch and the rest of the network should be thoroughly tested prior to implementation using a professional cable tester. Ensure that you always use compatible cables for the protocol you are using.

A common source of cabling issue is the incorrect use of straight-thru and crossover RJ-45 cables. Any device to-switch connection should use a straight-thru RJ-45 cable, while any device-to-device or switch-to-switch connection requires the use of a crossover RJ-45 cable.

QUESTION NO: 46

Refer to the exhibit. Look at the command output. Assume that there is no other path, and the configuration is correct. What would be the consequences of this situation?

```
Switch1#show cdp neighbor

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID        Local Intrfce     Holdtme   Capability   Platform   Port ID
Switch2          Gig 1/0/3         160          S I       WS-C2955C-Fas 0/13

Switch2#show cdp neighbor

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID      Local Intrfce    Holdtme   Capability  Platform  Port ID
Switch1        Fas 0/13         173       R S I     WS-C3750G-Gig 1/0/4
```

ActualTests

v

A. Users in SW1 can ping SW2 but not vice versa.
B. Users in SW2 can ping SW1 but not vice versa.
C. Users in SW1 and SW2 can ping each other.
D. Users in SW1 and SW2 cannot ping each other.

Answer: D

Explanation: References
http://www.cisco.com/en/US/docs/ios/12_1/configfun/configuration/guide/fcd301c.html
http://www.cisco.com/en/US/docs/wireless/access_point/12.2_11_JA/configuration/guide/s11cdp.html
http://www.cisco.com/en/US/docs/ios/12_1/configfun/command/reference/frd3001b.html

QUESTION NO: 47

Refer to the exhibit. Look at the command output. What can you use to prevent this behavior?

```
Switch1#show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
Switch2          Gig 1/0/3         160          S I       WS-C2955C- Fas 0/13
Switch1#

Switch2#show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
Switch1          Fas 0/13          173          R S I     WS-C3750G- Gig 1/0/4
Switch2#
```

A. udld
B. spanning-tree loopguard
C. VTP mode transparent
D. switchport mode desirable

Answer: A

Explanation:

UDLD Overview

The Cisco-proprietary Unidirectional Link Detection (UDLD) protocol allows ports that are connected through fiber optics or copper (for example, Category 5 cabling) Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. When the switch detects a unidirectional link, UDLD shuts down the affected LAN port and alerts the user. Unidirectional links can cause a variety of problems, including spanning tree topology loops.

UDLD is a Layer 2 protocol that works with the Layer 1 protocols to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected LAN ports. When you enable both autonegotiation and UDLD, Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

A unidirectional link occurs whenever traffic transmitted by the local device over a link is received by the neighbor but traffic transmitted from the neighbor is not received by the local device. If one of the fiber strands in a pair is disconnected, as long as autonegotiation is active, the link does not stay up. In this case, the logical link is undetermined, and UDLD does not take any action. If both fibers are working normally at Layer 1, then UDLD at Layer 2 determines whether those fibers are connected correctly and whether traffic is flowing bidirectionally between the correct neighbors. This check cannot be performed by autonegotiation, because autonegotiation operates at Layer 1.

A Cisco Nexus 5000 Series switch periodically transmits UDLD frames to neighbor devices on LAN ports with UDLD enabled. If the frames are echoed back within a specific time frame and they lack a specific acknowledgment (echo), the link is flagged as unidirectional and the LAN port is shut down. Devices on both ends of the link must support UDLD in order for the protocol to successfully identify and disable unidirectional links.

STP Loop Guard
The STP loop guard feature provides additional protection against Layer 2 forwarding loops (STP loops). An STP loop is created when an STP blocking port in a redundant topology erroneously transitions to the forwarding state. This usually happens because one of the ports of a physically redundant topology (not necessarily the STP blocking port) no longer receives STP BPDUs. In its operation, STP relies on continuous reception or transmission of BPDUs based on the port role. The designated port transmits BPDUs, and the non-designated port receives BPDUs.

When one of the ports in a physically redundant topology no longer receives BPDUs, the STP conceives that the topology is loop free. Eventually, the blocking port from the alternate or backup port becomes designated and moves to a forwarding state. This situation creates a loop.
The loop guard feature makes additional checks. If BPDUs are not received on a non-designated port, and loop guard is enabled, that port is moved into the STP loop-inconsistent blocking state, instead of the listening / learning / forwarding state. Without the loop guard feature, the port assumes the designated port role. The port moves to the STP forwarding state and creates a loop

Reference
http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli_rel_4_0_1a/BasicEthernet.html#wp1267151

http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a0080094640.shtml

QUESTION NO: 48

Before inserting a new switch in the network, the network administrator checks that the VTP domain name is correct, the VTP mode is set to server, and revision is lower than the switches in the network. The administrator then configures interfaces and trunks, erases existing VLANs, and connects the switch to the network. Following that procedure, there is no connectivity in the network. What is a possible cause of this problem?

A. Because the configuration revision of the new switches is lower than the rest of the network, it can change the VLAN database of the other switches.

B. As a VTP server, the new switch deleted all VLANs of the network.
C. Erasing VLANs increases the VTP configuration revision.
D. Since the configuration revision of the network is higher than the new switch, the VLAN database was automatically synchronized.

Answer: C

Explanation:

Reset the Configuration Revision Number

You can easily reset the configuration revision number by either of the two procedures provided in this section.

Reset the Configuration Revision using Domain Name Complete these steps in order to reset the configuration revision number with the change of the domain name:

1. Issue "show vtp domain " in order to see that the configuration is empty
2. Configure the VTP Domain name
3. Change the VTP Domain back
4. Change the VTP Domain to what it was in step 2

Reset the Configuration Revision using VTP Mode

Complete these steps in order to reset the configuration revision number with the change of the domain name:

1. Issue "show vtp domain " in order to see that the configuration is empty
2. Configure the VTP Domain name
3. Change the VTP mode from server to transparent
4. Change the VTP mode from transparent to server or client.

Reference

http://www.cisco.com/en/US/tech/tk389/tk689/technologies_tech_note09186a0080890613.shtml#topic9

QUESTION NO: 49

The network administrator wants to enable an EtherChannel between two switches in "on" mode. The administrator connects the cables and enables the interfaces, but while configuring the EtherChannel in the first switch, a spanning-tree loop was detected. Which two of these procedures can avoid this problem? (Choose two.)

A. Configure the EtherChannel as "desirable" first.
B. Assign all interfaces to the same VLAN.
C. Disable PortFast on the interfaces in the EtherChannels.
D. Disable all interfaces first.

E. Fast Ethernet and Gigabit Ethernet ports cannot be assigned to the same EtherChannel.
F. Fix cabling problems.

Answer: A,D

Explanation:

If a workstation or a server is connected with a single Network Interface Card (NIC) to a switch port, this connection cannot create a physical loop. These connections are considered leaf nodes. There is no reason to make the workstation wait 30 seconds while the switch checks for loops when the workstation cannot cause a loop. With the addition of the PortFast or fast-start feature, the STP for this port assumes that the port is not part of a loop. In this case, the port immediately moves to the forwarding state, and skips the blocking, listening, or learning states. This command does not turn STP off. This command makes STP skip a few steps in the beginning on the selected port, although unnecessary in this circumstance.

Note:
The PortFast feature must never be used on switch ports that connect to other switches, hubs, or routers. These connections can cause physical loops, and it is very important that Spanning Tree go through the full initialization procedure in these situations. A Spanning Tree loop can bring the network down. If the PortFast feature is turned on for a port that is part of a physical loop, it can cause packets to be continuously forwarded and even multiply in such a way that the network cannot recover.

Reference
https://supportforums.cisco.com/docs/DOC-4897

QUESTION NO: 50

Customer X has a hub-and-spoke Frame Relay network, with a central office and two branch offices (RemoteA and RemoteB). Each location has only one physical link to the Frame Relay cloud and RemoteB has a router that is not a Cisco router. Since the installation, there is no connectivity between RemoteB and the central office. What is a possible solution to this issue?

A. Because Frame Relay IETF encapsulation is only configurable at interface level, you must use IETF encapsulation on all routers.
B. This is not a possible scenario. A dedicated Frame Relay link to RemoteB is mandatory at the central office.
C. The router at RemoteB must be replaced by a Cisco router.
D. Use Frame Relay IETF encapsulation on a per-VC basis on the central office router.
E. There is a problem in the Frame Relay cloud, because Cisco routers are compatible with IETF Frame Relay.

Answer: D

Explanation:

QUESTION NO: 51

You are configuring an 802.1Q trunk between a Layer 2 switch and a firewall. You read in the documentation that the best way to set up a trunk is to set the port as dynamic desirable. The trunk is not coming up. Which one of these options would be a valid explanation?

A. The firewall does not support Cisco Discovery Protocol. You should set the switchport trunk mode ON.
B. The firewall does not support Cisco Discovery Protocol. You should set the switchport trunk mode to OFF.
C. The firewall does not support Cisco Discovery Protocol. You should set the switchport trunk mode as auto.
D. The firewall does not support DTP. You should set the switchport trunk mode to ON.

Answer: D

Explanation:

Explanation

PortFast, Channeling, and Trunking

By default, many switches, such as Cisco switches that run the Catalyst operating system (OS), are designed to be plug-and-play devices. As such, many of the default port parameters are not desirable when a PIX is plugged into the switch. For example, on a switch that runs the Catalyst OS, default channeling is set to Auto, trunking is set to Auto, and PortFast is disabled. If you connect a PIX to a switch that runs the Catalyst OS, disable channeling, disable trunking, and enable PortFast.

Channeling, also known as Fast EtherChannel or Giga EtherChannel, is used to bind two or more physical ports in a logical group in order to increase the overall throughput across the link. When a port is configured for automatic channeling, it sends out Port Aggregation Protocol (PAgP) frames as the link becomes active in order to determine if it is part of a channel. These frames can cause problems if the other device tries to autonegotiate the speed and duplex of the link. If channeling on the port is set to Auto, it also results in an additional delay of about 3 seconds before the port starts to forward traffic after the link is up.

Note: On the Catalyst XL Series Switches, channeling is not set to Auto by default. For this reason, you should disable channeling on any switch port that connects to a PIX.

Trunking, also known by the common trunking protocols Inter-Switch Link (ISL) or Dot1q, combines multiple virtual LANs (VLANs) on a single port (or link). Trunking is typically used between two switches when both switches have more than one VLAN defined on them. When a port is configured for automatic trunking, it sends out Dynamic Trunking Protocol (DTP) frames as the link comes up in order to determine if the port that it connects to wants to trunk. These DTP

frames can cause problems with autonegotiation of the link. If trunking is set to Auto on a switch port, it adds an additional delay of about 15 seconds before the port starts to forward traffic after the link is up. PortFast, also known as Fast Start, is an option that informs the switch that a Layer 3 device is connected out of a switch port. The port does not wait the default 30 seconds (15 seconds to listen and 15 seconds to learn); instead, this action causes the switch to put the port into forwarding state immediately after the link comes up. It is important to understand that when you enable PortFast, spanning tree is not disabled. Spanning tree is still active on that port. When you enable PortFast, the switch is informed only that there is not another switch or hub (Layer 2-only device) connected at the other end of the link. The switch bypasses the normal 30-second delay while it attempts to determine if a Layer 2 loop results if it brings up that port. After the link is brought up, it still participates in spanning tree. The port sends out bridge packet data units (BPDUs), and the switch still listens for BPDUs on that port. For these reasons, it is recommended that you enable PortFast on any switch port that connects to a PIX.
Note: Catalyst OS releases 5.4 and later include the set port host <mod>/<port> command that allows you to use a single command to disable channeling, disable trunking, and enable PortFast.

Reference
http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_tech_note09186a008009491c.shtml#portfastchanneltrunk

QUESTION NO: 52

Refer to the exhibit. You are setting up a 2-gigabit EtherChannel. Following IEEE standards, the exhibit shows your configuration in a local switchl. However, EtherChannel is not coming up. Which one of these statements could be a possible reason?

```
!
interface Port-channel1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 50
 switchport trunk allowed vlan 50,100,150,200
 !
interface GigabitEthernet5/1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 50
 switchport trunk allowed vlan 50,100,150,200
 channel-group 1 mode desirable
!
interface GigabitEthernet5/2
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 50
 switchport trunk allowed vlan 50,100,150,200
 channel-group 1 mode desirable
```

A. EtherChannel is only available in Cisco equipment.
B. The customer side is supposed to be running PAgP, which is a Cisco standard.
C. PAgP is not an IEEE standard. VRRP should be used.
D. The configuration on switch1 needs to be modified to use LACP.

Answer: D
Explanation:
Background Theory
LACP trunking supports four modes of operation, as follows:
On: The link aggregation is forced to be formed without any LACP negotiation .In other words, the switch will neither send the LACP packet nor process any incoming LACP packet. This is similar to the on state for PAgP.

Off: The link aggregation will not be formed. We do not send or understand the LACP packet. This is similar to the off state for PAgP.

Passive: The switch does not initiate the channel, but does understand incoming LACP packets. The peer (in active state) initiates negotiation (by sending out an LACP packet) which we receive and reply to, eventually forming the aggregation channel with the peer. This is similar to the auto mode in PAgP.
Active: We are willing to form an aggregate link, and initiate the negotiation. The link aggregate will be formed if the other end is running in LACP active or passive mode. This is similar to the desirable mode of PAgP.

There are only three valid combinations to run the LACP link aggregate, as follows:

Switch	Switch	Comments
active	active	Recommended.
active	passive	Link aggregation occurs if negotiation is successful.
on	on	Link aggregation occurs without LACP. Although this works, it is not recommended.

Note: By default, when an LACP channel is configured, the LACP channel mode is passive.

Reference
http://www.cisco.com/en/US/tech/tk389/tk213/technologies_configuration_example09186a0080094470.shtml

QUESTION NO: 53

Where should frame-relay traffic shaping be applied?

A. on the physical interface
B. on the subinterface
C. under the frame-relay map class
D. on any of the above

Answer: A
Reference: http://blog.ine.com/2008/01/22/legacy-frts/ (first bullet on the page)

QUESTION NO: 54

Prior to 802.1w, Cisco implemented a number of proprietary enhancements to 802.1D to improve convergence in a Layer 2 network. Which statement is correct?

A. Only UplinkFast and BackboneFast are specified in 802.1w; PortFast must be manually configured.
B. Only PortFast is specified in 802.1w; UplinkFast and BackboneFast must be manually configured.

C. None of the proprietary Cisco enhancements are specified in 802.1w.
D. PortFast, UplinkFast, and BackboneFast are specified in 802.1w.

Answer: D

Explanation:

Spanning-tree PortFast causes a spanning-tree port to enter the forwarding state immediately, bypassing the listening and learning states. You can use PortFast on switch ports connected to a single workstation or server to allow those devices to connect to the network immediately, rather than waiting for spanning tree to converge.

UplinkFast provides fast convergence after a spanning-tree topology change and achieves load balancing between redundant links using uplink groups. An uplink group is a set of ports (per VLAN), only one of which is forwarding at any given time. Specifically, an uplink group consists of the root port (which is forwarding) and a set of blocked ports, except for self-looping ports. The uplink group provides an alternate path in case the currently forwarding link fails.

BackboneFast is initiated when a root port or blocked port on a switch receives inferior BPDUs from its designated bridge. An inferior BPDU identifies one switch as both the root bridge and the designated bridge. When a switch receives an inferior BPDU, it indicates that a link to which the switch is not directly connected (an indirect link) has failed (that is, the designated bridge has lost its connection to the root bridge). Under normal spanning-tree rules, the switch ignores inferior BPDUs for the configured maximum aging time, as specified by the aging time variable of the "set spantree MaxAge" command. The switch tries to determine if it has an alternate path to the root bridge. If the inferior BPDU arrives on a blocked port, the root port and other blocked ports on the switch become alternate paths to the root bridge. (Self-looped ports are not considered alternate paths to the root bridge.) If the inferior BPDU arrives on the root port, all blocked ports become alternate paths to the root bridge. If the inferior BPDU arrives on the root port and there are no blocked ports, the switch assumes that it has lost connectivity to the root bridge, causes the maximum aging time on the root to expire, and becomes the root switch according to normal spanning-tree rules.

If the switch has alternate paths to the root bridge, it uses these alternate paths to transmit a new kind of PDU called the Root Link Query PDU. The switch sends the Root Link Query PDU out all alternate paths to the root bridge. If the switch determines that it still has an alternate path to the root, it causes the maximum aging time on the ports on which it received the inferior BPDU to expire. If all the alternate paths to the root bridge indicate that the switch has lost connectivity to the root bridge, the switch causes the maximum aging times on the ports on which it received an inferior BPDU to expire. If one or more alternate paths can still connect to the root bridge, the switch makes all ports on which it received an inferior BPDU its designated ports and moves them out of the blocking state (if they were in blocking state), through the listening and learning states, and into the forwarding state.

QUESTION NO: 55

As a network administrator, can you tell me what the root guard feature provides in a bridged network?

A. It ensures that BPDUs sent by the root bridge are forwarded in a timely manner
B. It enforces the root bridge placement in the network
C. It ensures that all ports receiving BPDUs from the root bridge are in the forwarding state.
D. It ensures that the bridge is elected as root bridge in the network.

Answer: B

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.

QUESTION NO: 56

Refer to the following descriptions, which three are true about Cisco spanning-tree features? (Choose three.)

A. RPVST+ converges faster than RSTP during a topology change.
B. STP BPDUs are relayed by all non-root bridges and RSTP BPDUs are generated by each bridge.
C. RSTP can only achieve rapid transition to Forwarding on edge ports and on point-to-point links.
D. RPVST+ and RSTP are both based upon the IEEE 802.1w specification.

Answer: B,C,D

Explanation:

PVST+ is per-VLAN spanning tree (which is the default for most cisco switches). It means that you will run an spanning-tree instance per VLAN. This is useful when you need different layer 2 behaviors per VLAN, for example you can have different root bridge on different VLANs (so that spanning tree does not have to run as a whole on the layer 2 domain, but can run a different instance per- VLAN) RSTP is rapid STP. It is an enhancement to STP. RSTP does not work with timers as regular STP (which takes up to 30-50 seconds to converge due to the transition to all its states) Regular STP can use port-fast for ports not connected to other switches, but all ports

connected to other switches need to transition from blocking to listening, learning and finally forwarding. RSTP optimizes this by using P2P links and taking up to only 2 seconds to converge. RPVST+ Is a mix of PVST+ and RSTP. You have an instance of rapid STP running per VLAN. Also, some use MST which is another variance of STP which can group several VLANs to be part of a single MST region (and behave like RSTP inside that region). MST is useful because if you have 1000 VLANs, normally you don't need to have 1000 STP/RSTP instances! You can instead have one instance with VLAN 1-500 and another instance with VLANs 501-1000 (just to give you an example)

QUESTION NO: 57

Which switch port error is an indication of duplex mismatches on 10/100/1000 IEEE 802.3u Gigabit Ethernet ports?

A. FCS errors
B. Runts
C. Multiple collisions
D. Alignment errors

Answer: C
Explanation:
Communication is possible over a connection in spite of a duplex mismatch. Single packets are sent and acknowledged without problems. As a result, a simple ping command fails to detect a duplex mismatch because single packets and their resulting acknowledgments at 1-second intervals do not cause any problem on the network. A terminal session which sends data slowly (in very short bursts) can also communicate successfully. However, as soon as either end of the connection attempts to send any significant amount of data, the network suddenly slows to very low speed. Since the network is otherwise working, the cause is not so readily apparent.
A duplex mismatch causes problems when both ends of the connection attempt to transfer data at the same time. This happens even if the channel is used (from a high-level or user's perspective) in one direction only, in case of large data transfers. Indeed, when a large data transfer is sent over a TCP, data is sent in multiple packets, some of which will trigger an acknowledgment packet back to the sender. This results in packets being sent in both directions at the same time.
In such conditions, the full-duplex end of the connection sends its packets while receiving other packets; this is exactly the point of a full-duplex connection. Meanwhile, the half-duplex end cannot accept the incoming data while it is sending -- it will sense it as a collision. The half-duplex device ceases its current transmission and then retries later as per CSMA/CD. As a result, when both devices are attempting to transmit at the same time, packets sent by the full-duplex end will be lost and packets sent by the half duplex device will be delayed or lost.

The lost packets force the TCP protocol to perform error recovery, but the initial (streamlined) recovery attempts fail because the retransmitted packets are lost in exactly the same way as the original packets. Eventually, the TCP transmission window becomes full and the TCP protocol refuses to transmit any further data until the previously-transmitted data is acknowledged. This, in turn, will quiescence the new traffic over the connection, leaving only the retransmissions and acknowledgments. Since the retransmission timer grows progressively longer between attempts, eventually a retransmission will occur when there is no reverse traffic on the connection, and the acknowledgments are finally received. This will restart the TCP traffic, which in turn immediately causes lost packets as streaming resumes. The end result is a connection that is working but performs extremely poorly because of the duplex mismatch. Symptoms of a duplex mismatch are connections that seem to work fine with a ping command, but "lock up" easily with very low throughput on data transfers; the effective data transfer rate is likely to be asymmetrical, performing much worse in one direction than the other. In normal half-duplex operations late collisions do not occur. However, in a duplex mismatch the collisions seen on the half-duplex side of the link are often late collisions. The full-duplex side usually will register frame check sequence errors, or runt frames. Viewing these standard Ethernet statistics can help diagnose the problem. Contrary to what one might reasonably expect, both sides of a connection need to be identically configured for proper operation. In other words, setting one side to automatic (either speed or duplex or both) and setting the other to be fixed (either speed or duplex or both) will result in a speed mismatch, a duplex mismatch or both. A duplex mismatch can be fixed by either enabling autonegotiation (if available and working) on both ends or by forcing the same settings on both ends (availability of a configuration interface permitting). If there is no option but to have a locked setting on one end and autonegotiation the other (for example, an old device with broken autonegotiation connected to an unmanaged switch) half duplex must be used. All modern LAN equipment comes with autonegotiation enabled and the various compatibility issues have been resolved. The best way to avoid duplex mismatches is to use autonegotiation and to replace any legacy equipment that does not use autonegotiation or does not autonegotiate correctly.

QUESTION NO: 58

Which one of the following potential issues is eliminated by using split horizon?

A. Joined horizons
B. Packet forwarding loops
C. cisco Express Forwarding load-balancing inconsistency
D. Asymmetric routing throughout the network

Answer: B
Explanation:

Split horizon is a base technique used to reduce the chance of routing loops. Split horizon states that it is never useful to send information about a route back in the direction from which the information came and therefore routing information should not be sent back to the source from which it came. In fact, only the interfaces are considered for the direction, not the neighbors. Note that this rule works well not only for routes learned via a distance vector routing protocol but also for routes installed in a routing table as directly connected networks.
As they reside on the same network, the neighbors do not need any advertisements on a path to that shared network.
The split horizon rule helps prevent two-node (two-neighbor) routing loops and also improves performance by eliminating unnecessary updates.

QUESTION NO: 59

Phase I and Phase II DMVPN differ in terms of which of these characteristics?

A. Utilization of spoke-to-spoke dynamic tunnels
B. Utilization of multipoint GRE tunnels at the hub site
C. Utilization of hub-to-spoke dynamic tunnels
D. Support for multicast

Answer: A
Explanation: Explanation
DMVPN Phases
Phase 1: Hub and spoke functionality
Phase 2: Spoke-to-spoke functionality
Phase 3: Architecture and scaling
DMVPN Phase 1 Benefits Simplified and Smaller Comfig's for Hub and Spoke Zero touch provisioning for adding spokes to the VPN Easily supports dynamically addressed CPEs
DMVPN Phase 2 Benefits Future Functionality On-demand spoke-to-spoke tunnels avoids dual encrypts/ decrypts smaller spoke CPE can participate in the virtual full mesh

QUESTION NO: 60

Which statement is true about shaping?

A. Shaping supports queuing of excess traffic.

B. Shaping can be applied both input and output on interfaces.
C. Shaping does not introduce delay in voice packet handling in the event of congestion.
D. Shaping makes instantaneous packet drop decisions.

Answer: A
Explanation:
Traffic shaping retains excess packets in a queue and then schedules the excess for later transmission over increments of time. The result of traffic shaping is a smoothed packet output rate. In contrast, traffic policing propagates bursts. When the traffic rate reaches the configured maximum rate, excess traffic is dropped (or remarked). The result is an output rate that appears as a saw-tooth with crests and troughs.

QUESTION NO: 61

When two bridges are competing for the root bridge of an IEEE 802.1D spanning tree and both have the same bridge priority configured, which parameter determines the winner?

A. highest-numbered IP interface
B. MAC address
C. device uptime
D. root port cost

Answer: B
Explanation:
Select a root bridge. The root bridge of the spanning tree is the bridge with the smallest (lowest) bridge ID.
Each bridge has a unique identifier (ID) and a configurable priority number; the bridge ID contains both numbers. To compare two bridge IDs, the priority is compared first. If two bridges have equal priority, then the MAC addresses are compared. For example, if switches A (MAC=0200.0000.1111) and B (MAC=0200.0000.2222) both have a priority of 10, then switch A will be selected as the root bridge. If the network administrators would like switch B to become the root bridge, they must set its priority to be less than 10.

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

QUESTION NO: 62

For the following ports, which port is on every bridge in a Spanning Tree Protocol IEEE 802.1w network except the root bridge?

A. root port
B. backup port
C. designated port
D. alternate port

Answer: A

Explanation:

Explanation The root bridge does not have a root port as this is the port on all non-root bridges that is used to communicate with the root bridge. All ports on the root bridge are designed ports.

Reference
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_configuration_example09186a008009467c.shtml#maintask1

QUESTION NO: 63

IEEE 802.1w is a Rapid Spanning Tree Protocol (RSTP) that can be seen as an evolution of the 802.1 standard. What are the port roles described by 802.1w?

A. root port, designated port, alternate port, backup port, and disabled
B. standby port, alternate port, root port, and disabled
C. standby port, designated port, backup port, and disabled
D. root port, designated port, alternate port, and standby port

Answer: A

Explanation:

Port Roles
The role is now a variable assigned to a given port. The root port and designated port roles remain, while the blocking port role is split into the backup and alternate port roles. The Spanning Tree Algorithm (STA) determines the role of a port based on Bridge Protocol Data Units (BPDUs). In order to simplify matters, the thing to remember about a BPDU is there is always a method to compare any two of them and decide whether one is more useful than the other. This is based on the value stored in the BPDU and occasionally on the port on which they are received.

Root Port Roles
The port that receives the best BPDU on a bridge is the root port. This is the port that is the closest to the root bridge in terms of path cost. The STA elects a single root bridge in the whole

bridged network (per-VLAN). The root bridge sends BPDUs that are more useful than the ones any other bridge sends. The root bridge is the only bridge in the network that does not have a root port. All other bridges receive BPDUs on at least one port.

Designated Port

A port is designated if it can send the best BPDU on the segment to which it is connected. 802.1D bridges link together different segments, such as Ethernet segments, to create a bridged domain. On a given segment, there can only be one path toward the root bridge. If there are two, there is a bridging loop in the network. All bridges connected to a given segment listen to the BPDUs of each and agree on the bridge that sends the best BPDU as the designated bridge for the segment. The port on that bridge that corresponds is the designated port for that segment.

Alternate and Backup Port Roles

These two port roles correspond to the blocking state of 802.1D. A blocked port is defined as not being the designated or root port. A blocked port receives a more useful BPDU than the one it sends out on its segment.

Remember that a port absolutely needs to receive BPDUs in order to stay blocked. RSTP introduces these two roles for this purpose.

An alternate port receives more useful BPDUs from another bridge and is a port blocked.

A backup port receives more useful BPDUs from the same bridge it is on and is a port blocked.

This distinction is already made internally within 802.1D. This is essentially how Cisco UplinkFast functions.

The rationale is that an alternate port provides an alternate path to the root bridge and therefore can replace the root port if it fails. Of course, a backup port provides redundant connectivity to the same segment and cannot guarantee an alternate connectivity to the root bridge. Therefore, it is excluded from the uplink group.

As a result, RSTP calculates the final topology for the spanning tree that uses the same criteria as 802.1D.

There is absolutely no change in the way the different bridge and port priorities are used. The name blocking is used for the discarding state in Cisco implementation. CatOS releases 7.1 and later still display the listening and learning states. This gives even more information about a port than the IEEE standard requires. However, the new feature is now there is a difference between the role the protocol determines for a port and its current state. For example, it is now perfectly valid for a port to be designated and blocking at the same time. While this typically occurs for very short periods of time, it simply means that this port is in a transitory state towards the designated forwarding state.

Reference
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_white_paper09186a0080094cfa.shtml#roles

QUESTION NO: 64

What is the STP root guard feature designed to prevent?

A. a root port being transitioned to the blocking state
B. a port being assigned as a root port
C. a port being assigned as an alternate port
D. a root port being transitioned to the forwarding state

Answer: B
Reference: http://ericleahy.com/index.php/bpdu-guard-bpdu-filter-root-guard-loop-guard-udld/

QUESTION NO: 65

Which two statements are true about the role of split horizon? (Choose two.)

A. It is a function used by routing protocols to install routes into routing table
B. It is a function that prevents the advertising of routes over an interface that the router is using to reach a route
C. Its function is to help avoid routing loops.
D. It is a redistribution technique used by routing protocols

Answer: B,C
Explanation:
Split horizon is a method of preventing a routing loop in a network. The basic principle is simple: Information about the routing for a particular packet is never sent back in the direction from which it was received.
Split horizon can be achieved by means of a technique called poison reverse. This is the equivalent of route poisoning all possible reverse paths - that is, informing all routers that the path back to the originating node for a particular packet has an infinite metric. Split horizon with poison reverse is more effective than simple split horizon in networks with multiple routing paths, although it affords no improvement over simple split horizon in networks with only one routing path.

QUESTION NO: 66

How many bytes make up the spanning-tree bridge ID?

A. 4
B. 8
C. 12
D. 16

Answer: B
Explanation:
The IEEE 802.1D standard requires that each switch has an unique bridge identifier (bridge ID), which controls the selection of the root switch. Because each VLAN is considered as a different logical bridge with PVST+ and rapid PVST+, the same switch must have a different bridge IDs for each configured VLAN. Each VLAN on the switch has a unique 8-byte bridge ID. The 2 most-significant bytes are used for the switch priority, and the remaining 6 bytes are derived from the switch MAC address.

QUESTION NO: 67

In which two spanning-tree port states is the port learning MAC addresses? (Choose two.)

A. disabled
B. blocking
C. listening
D. learning
E. forwarding

Answer: D,E
Explanation:
When the spanning-tree algorithm determines that a port should be placed in the forwarding state, the following occurs:

QUESTION NO: 68

In an 802.1s BPDU, what is the size of the configuration revision number?

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

Answer: B
Reference: http://blog.ine.com/tag/8021s/ (see implementing MSTP, 5th bullet)

QUESTION NO: 69

What is the BPDU protocol version for 802.1w?

A. 0
B. 1
C. 2
D. 3

Answer: C
Reference:
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_white_paper09186a0080094cfa.shtml
(See New BPDU format, second para)

QUESTION NO: 70

Which two are natively included by the IEEE 802.1w standard? (Choose two.)

A. instances can control a selection of VLANs
B. load balancing
C. fast transition to forwarding state
D. backbone, uplink, and portfast (or equivalent)
E. root, loop, and BPDU guard (or equivalent)

Answer: C,D
Explanation:
802.1D Spanning Tree Protocol (STP) has a drawback of slow convergence. Cisco Catalyst switches support three types of STPs, which are PVST+, rapid-PVST+ and MST. PVST+ is based

on IEEE802.1D standard and includes Cisco proprietary extensions such as BackboneFast, UplinkFast, and PortFast. Rapid-PVST+ is based on IEEE 802.1w standard and has a faster convergence than 802.1D. RSTP (IEEE 802.1w) natively includes most of the Cisco proprietary enhancements to the 802.1D Spanning Tree, such as BackboneFast and UplinkFast. Rapid-PVST+ has these unique features:
Rapid-PVST uses RSTP to provide faster convergence. When any RSTP port receives legacy 802.1D BPDU, it falls back to legacy STP and the inherent fast convergence benefits of 802.1w are lost when it interacts with legacy bridges.

QUESTION NO: 71

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

A. LACP packets are sent with multicast group MAC address 01-80-c2-00-00-02.
B. The Type/Field value is 0x8808.
C. During detection, LACP packets are transmitted every second.
D. The timeout for a failed LACP channel is 30 seconds by default.

Answer: A,C
Reference: http://www.cisco.com/en/US/docs/ios/cether/configuration/guide/ce_lnkbndl.html

QUESTION NO: 72

Which statement is true about TCN propagation in RSTP (802.1w)?

A. The originator of the TCN immediately floods this information through the network.
B. The TCN propagation is a two step process.
C. A TCN is generated and sent to the root bridge.
D. The root bridge must flood this information throughout the network.

Answer: A
Explanation:
The TCN propagation in RSTP is unique in the way that the originator of the TCN floods information through the network.

QUESTION NO: 73

When using extended system ID in 802.1d, how many bits are reserved for this field?

A. 6
B. 8
C. 10
D. 12

Answer: D
Explanation: The priority value is divided into 4 bit priority and 12 bit extended system id. This extended system id is usually equal to the VLAN id. The idea behind it is to conserve MAC addresses - PVST+ requires a different bridge id per VLAN, and this would mean different MACs -- unless you ensure that the priorities are different. By re-mapping some of the priority bits, this is ensured and all VLANs can use the same MAC address.

http://www.informit.com/library/content.aspx?b=CCIE_Practical_Studies_II&seqNum=13

QUESTION NO: 74

What are two ways to force the selection of a root bridge in a network that is running the 802.1D protocol? (Choose two.)

A. spanning tree vlan all root
B. spanning-tree vlan vlan-id priority 65535
C. spanning-tree vlan vlan-id root
D. spanning-tree vlan vlan-id priority 0
E. spanning-tree vlan vlan-id force root

Answer: C,D
Explanation:
The IEEE 802.1 standard (STP) is used to create a loop-free Layer 2 network. This protocol uses the bridge ID (a field inside BPDU packets) to elect root bridge. It is 8 bytes in length. The first two bytes are the Bridge Priority, which is an integer in the range of 0 – 65,535 (default is 32,768). The last six bytes are a MAC address supplied by the switch.
In STP, lower bridge ID values are preferred. To compare two bridge IDs, the priority is compared first. If two bridges have equal priority, then the MAC addresses are compared. For example, if switch A (MAC=0600.0000.1111) and B (MAC=0600.0000.2222) both have a priority of 10, then switch A will be selected as the root bridge because it has lower MAC.

QUESTION NO: 75

Which three can be achieved by Cisco PVST+? (Choose three.)

A. instances can control a selection of VLANs
B. load balancing
C. fast transition to forwarding state
D. backbone, uplink, and portfast
E. root, loop, and BPDU guard

Answer: B,D,E
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_example09186a00807b0670.shtml (see introduction)

QUESTION NO: 76

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

A. PaGP packets are sent with multicast group MAC address 01-80-c2-00-00-02.
B. PAgP uses the same multicast group MAC address as Cisco Discovery Protocol.
C. The PAgP protocol value is 0x0104.
D. During detection, PAgP packets are transmitted every 2 seconds.

Answer: B,C
Reference: http://ericleahy.com/index.php/etherchannels/ (see PAgP)

QUESTION NO: 77

Which two are contained in a VTP summary advertisement? (Choose two.)

A. configuration revision number
B. VTP domain name
C. VLAN information
D. sequence number

E. VLAN type

Answer: A,B
Reference: http://ericleahy.com/index.php/vlan-trunking-protocol-vtp/ (see VTP advertisements)

QUESTION NO: 78

Which two combinations are valid LACP configurations that will set up an LACP channel? (Choose two.)

A. on-passive
B. on-auto
C. passive-active
D. desirable-auto
E. active-active
F. desirable-desirable

Answer: C,E
Reference:
http://www.cisco.com/en/US/tech/tk389/tk213/technologies_configuration_example09186a0080094470.shtml (background theory, see the table)

QUESTION NO: 79

Which three options are used in the spanning-tree decision process? (Choose three.)

A. lowest root bridge ID
B. lowest path cost to root bridge
C. lowest sender bridge ID
D. highest port ID
E. highest root bridge ID
F. highest path cost to root bridge

Answer: A,B,C
Reference: http://www.ccnpguide.com/ccnp-switch-642-813-spanning-tree-basics/ (see STP convergence)

QUESTION NO: 80

When troubleshooting duplex mismatches, which two are errors that are seen on the half duplex end? (Choose two.)

A. excessive collisions
B. FCS errors
C. runts
D. late collisions

Answer: A,D
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a00800a7af0.shtml

QUESTION NO: 81

Which three combinations are valid PAgP configurations that will set up a channel? (Choose three.)

A. On-On
B. On-Auto
C. Passive-Active
D. Desirable-Auto
E. Active-Active
F. Desirable-Desirable

Answer: A,D,F
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps607/products_configuration_example09186a0080094789.shtml (see port aggregation protocol)

QUESTION NO: 82

Which three combinations are valid LACP configurations that will set up a channel? (Choose three.)

A. On/On
B. On/Auto

C. Passive/Active
D. Desirable/Auto
E. Active/Active
F. Desirable/Desirable

Answer: A,C,E
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps663/products_configuration_example09186a0080094aec.shtml (pagp and lacp modes, see the table)

QUESTION NO: 83

Network A has a spanning-tree problem in which the traffic is selecting a longer path. How is the path cost calculated?

A. number of hops
B. priority of the bridge
C. interface bandwidth
D. interface delay
E. None of the above

Answer: C
Explanation:
STP Path Cost Automatically Changes When a Port Speed/Duplex Is Changed STP calculates the path cost based on the media speed (bandwidth) of the links between switches and the port cost of each port forwarding frame. Spanning tree selects the root port based on the path cost. The port with the lowest path cost to the root bridge becomes the root port. The root port is always in the forwarding state.

If the speed/duplex of the port is changed, spanning tree recalculates the path cost automatically. A change in the path cost can change the spanning tree topology.

Reference
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_configuration_example09186a008009467c.shtml

QUESTION NO: 84
For the following LMI types, which three can be configured for use with Frame Relay on a Cisco

router? (Choose three.)

A. Cisco
B. ANSI - Annex D
C. Q.931 - Annex B
D. Q.933 - Annex A

Answer: A,B,D

Explanation:

ANSI-617d (ANSI or annex D) LMI type, DLCI 0
Serial1(in): Status, myseq 3
RT IE 1, length 1, type 0
KA IE 3, length 2, yourseq 4 , myseq 3
PVC IE 0×7 , length 0×3 , dlci 100, status 0×0
PVC IE 0×7 , length 0×3 , dlci 200, status 0×0
Q933a (CCITT or annex A) LMI type, DLCI 0
Serial1(in): Status, myseq 1
RT IE 51, length 1, type 0
KA IE 53, length 2, yourseq 2 , myseq 1
PVC IE 0×57, length 0×3 , dlci 100, status 0×0
PVC IE 0×57, length 0×3 , dlci 200, status 0×0
Cisco LMI type, DLCI 1023
Serial1(in): Status, myseq 68
RT IE 1, length 1, type 0
KA IE 3, length 2, yourseq 68, myseq 68
PVC IE 0×7 , length 0×6 , dlci 100, status 0×2 , bw 0
PVC IE 0×7 , length 0×6 , dlci 200, status 0×2 , bw 0

QUESTION NO: 85

For the following items, what is the mathematical relationship between the committed information rate (CIR), committed burst (Bc), and committed rate measurement interval (Tc)?

A. CIR = TC / Bc
B. CIR = Be / Tc
C. Tc = CIR / Bc
D. Tc = Bc / CIR

Answer: D

Explanation:

Explanation
Terminologies: The term CIR refers to the traffic rate for a VC based on a business contract.
Tc is a static time interval, set by the shaper.

Committed burst (Bc) is the number of bits that can be sent in each Tc. Be is the excess burst size, in bits. This is the number of bits beyond Bc that can be sent after a period of inactivity.

QUESTION NO: 86

If a Cisco switch is configured with VTPv1 in transparent mode, what is done with received VTP advertisements?

A. They are discarded
B. The contents are altered to reflect the switch's own VTP database and then they are forward out all trunking ports
C. The changes within the advertisements are made to the switch's VTP database.
D. The contents are ignored and they are forwarded out all trunking ports.

Answer: D
Explanation:
VTPv1 & VTPv2 are the same in regards to Transparent mode VTP advertisements. Therefore the Transparent mode switch will NOT update it's local VTP database but WILL forward the VTP advertisement out all of it's trunk ports.

QUESTION NO: 87

Which of these potential issues is eliminated by the use of split horizon?

A. asymmetric routing throughout the network
B. packet forwarding loops
C. joined horizons
D. Cisco Express Forwarding load-balancing inconsistency

Answer: B
Explanation:
Distance-vector routing protocols employ the split horizon rule which prohibits a router from advertising a route back out the interface from which it was learned. Split horizon is one of the

methods used to prevent routing loops due to the slow convergence times of distance-vector routing protocols.

QUESTION NO: 88

In Frame Relay, FECN messages indicating congestion are sent or received by which of following?

A. Sent by the destination
B. Received by the sender
C. Received by the destination
D. Sent by the sender

Answer: C
Explanation:
Congestion control
The Frame Relay network uses a simplified protocol at each switching node. It achieves simplicity by omitting link-by-link flow-control. As a result, the offered load has largely determined the performance of Frame Relay networks. When offered load is high, due to the bursts in some services, temporary overload at some Frame Relay nodes causes a collapse in network throughput. Therefore, frame-relay networks require some effective mechanisms to control the congestion. Congestion control in frame-relay networks includes the following elements:
Admission Control provides the principal mechanism used in Frame Relay to ensure the guarantee of resource requirement once accepted. It also serves generally to achieve high network performance. The network decides whether to accept a new connection request, based on the relation of the requested traffic descriptor and the network's residual capacity. The traffic descriptor consists of a set of parameters communicated to the switching nodes at call set-up time or at service-subscription time, and which characterizes the connection's statistical properties.

The traffic descriptor consists of three elements:
Committed Information Rate (CIR) - The average rate (in bit/s) at which the network guarantees to transfer information units over a measurement interval T. This T interval is defined as: T = Bc/CIR.

Committed Burst Size (BC) - The maximum number of information units transmittable during the interval T. Excess Burst Size (BE) - The maximum number of uncommitted information units (in bits) that the network will attempt to carry during the interval.
Once the network has established a connection, the edge node of the Frame Relay network must monitor the connection's traffic flow to ensure that the actual usage of network resources does not exceed this specification. Frame Relay defines some restrictions on the user's information rate. It

allows the network to enforce the end user's information rate and discard information when the subscribed access rate is exceeded.
Explicit congestion notification is proposed as the congestion avoidance policy. It tries to keep the network operating at its desired equilibrium point so that a certain Quality of Service (QoS) for the network can be met. To do so, special congestion control bits have been incorporated into the address field of the Frame Relay:
FECN and BECN. The basic idea is to avoid data accumulation inside the network.
FECN means Forward Explicit Congestion Notification. The FECN bit can be set to 1 to indicate that congestion was experienced in the direction of the frame transmission, so it informs the destination that congestion has occurred. BECN means Backwards Explicit Congestion Notification. The BECN bit can be set to 1 to indicate that congestion was experienced in the network in the direction opposite of the frame transmission, so it informs the sender that congestion has occurred.

QUESTION NO: 89

Which two attributes need to match for two switches to become members of the same MST region? (Choose two.)

A. the table of 4096 elements that map the respective VLAN to STP instance number
B. VTP version
C. configuration revision number
D. native VLAN ID

Answer: A,C
Explanation:
The final implementation adopted by the IEEE 802.1s standard made this mechanics more elegant and simple. Before we process with discussing IEEE’s implementation, let’s define **MSTP region** as a collection of switches, sharing the same view of physical topology partitioning into set of logical topologies. For two switches to become members of the same region, the following attributes must match:

QUESTION NO: 90

An administrator of a peer-to-peer server application reports that the maximum bandwidth that his application receives is 90 Mb/s. You have an 8-port, 100-Mb/s EtherChannel bundle on the switching infrastructure between the two servers, resulting in a bidirectional throughput of more

than what is recorded. Which two solutions would allow for more bandwidth for the application? (Choose two.)

A. change the EtherChannel default hashing algorithm from XOR to use source port Layer 4 port load balancing to better load balance the traffic
B. change the EtherChannel default hashing algorithm from XOR to use destination Layer 4 port load balancing to better load balance the traffic
C. have the application recoded to use multiple connections instead of a single connection so EtherChannel can better load balance the traffic
D. upgrade the CPU and memory of the compute devices so they can better process traffic
E. upgrade the EtherChannel bundle to a single gigabit link because EtherChannel is not well suited for single-connection traffic

Answer: C,E
Reference:
http://www.cisco.com/en/US/prod/collateral/routers/ps5855/prod_white_paper0900aecd806f698a.pdf

QUESTION NO: 91

A user has no network connectivity. A check of the associated port indicates that the interface is up, the line protocol is down. Which item would most likely cause this problem?

A. Speed mismatch
B. Incorrect encapsulation
C. MTU set too low
D. Duplex mismatch

Answer: A
Explanation:
If there is duplex mismatch, there will be connectivity, but there will be errors and late collisions.
Incorrect encapsulation - This is a user connectivity, so it is an ethernet connection, no encapsulation configuration.
If MTU is different on each device or too low, there might problems in data transfer but the line protocol will not go down

QUESTION NO: 92

A router is connected to an HDLC circuit via a T1 physical interface. The SLA for this link only allows for a sustained rate of 768 kb/s. Bursts are allowed for up to 30 seconds at up to line rate, with a window Tc of 125 ms. What should the Bc and Be setting be when using generic traffic shaping?

A. Be = 46320000 , Bc = 96000
B. Be = , 768000 Be = 32000
C. Be = , 128000 Be = 7680
D. Be = , 0 Be = 96000

Answer: A
Explanation:

Tc= 125
CIR = 768
What is the Be
T1 = 1.544 Mbps
Bursts are allowed for 30 seconds
Seconds * Bandwidth in bps = Be
30 * 1544000 = Be
30 * 1544000 = 46320000
Be = 46320000
What is Bc?
Bc = Tc * CIR
Bc = 125 * 768
Bc = 96000
Traffic Shaping Parameters
We can use the following traffic shaping parameters:
CIR = committed information rate (= mean time)
EIR = excess information rate
TB = token bucket (= Bc + Be)
Bc = committed burst size (= sustained burst size)
Be = excess burst size
DE = discard eligibility
Tc = measurement interval
AR = access rate corresponding to the rate of the physical interface (so if you use a T1, the AR is approximately 1.5 Mbps).

Committed Burst Size (Bc)
The maximum committed amount of data you can offer to the network is defined as Bc. Bc is a measure for the volume of data for which the network guarantees message delivery under normal conditions. It is measured during the committed rate Tc.

Excess Burst Size (Be)
The number of noncommitted bits (outside of CIR) that are still accepted by the Frame Relay switch but are marked as eligible to be discarded (DE). The token bucket is a 'virtual' buffer. It contains a number of tokens, enabling you to send a limited amount of data per time interval. The token bucket is filled with Bc bits per Tc.

The maximum size of the bucket is Bc + Be. If the Be is very big and, if at T0 the bucket is filled with Bc + Be tokens, you can send Bc + Be bits at the access rate. This is not limited by Tc but by the time it takes to send the Be. This is a function of the access rate.
Committed Information Rate (CIR)
The CIR is the allowed amount of data which the network is committed to transfer under normal conditions. The rate is averaged over a increment of time Tc. The CIR is also referred to as the minimum acceptable throughput. Bc and Be are expressed in bits, Tc in seconds, and the access rate and CIR in bits per second.
Bc, Be, Tc and CIR are defined per datalink connection identifier (DLCI). Due to this, the token bucket filter controls the rate per DLCI. The access rate is valid per usernetwork interface. For Bc, Be and CIR incoming and outgoing values can be distinguished. If the connection is symmetrical, the values in both directions are the same. For permanent virtual circuits, we define incoming and outgoing Bc, Be and CIR at subscription time.

Peak = DLCI's maximum speed. The bandwidth for that particular DLCI.
Tc = Bc / CIR
Peak = CIR + Be/Tc = CIR (1 + Be/Bc)
If the Tc is one second then:
Peak = CIR + Be = Bc + Be

Reference
http://www.cisco.com/warp/public/125/21.pdf

QUESTION NO: 93

The EtherChannel between your LAN switch and the Internet router is not load-balancing efficiently. On the switch, there are several workstations with valid IP ranges. Which load-balance algorithms can you use in the switch in order to optimize this load balancing? (Choose four.)

A. source IP address
B. destination IP address
C. per-packet load balance
D. destination MAC address

E. source MAC address

Answer: A,B,D,E

Explanation:

EtherChannel load balancing can use MAC addresses, IP addresses, or Layer 4 port numbers with a Policy Feature Card 2 (PFC2) and either source mode, destination mode, or both. The mode you select applies to all EtherChannels that you configure on the switch. Use the option that provides the greatest variety in your configuration. For example, if the traffic on a channel only goes to a single MAC address, use of the destination MAC address results in the choice of the same link in the channel each time. Use of source addresses or IP addresses can result in a better load balance. Issue the port-channel load- balance {src-mac | dst-mac | srcdst-mac | src-ip | dst-ip | src-dst-ip | src-port | dst-port | src-dst-port | mpls} global configuration command in order to configure the load balancing.

Issue the show etherchannel load-balance command in order to check the frame distribution policy. You can determine which interface in the EtherChannel forwards traffic, with the frame distribution policy as a basis.

Issue the remote login switch command to log in remotely to the Switch Processor (SP) console in order to make this determination. Then, issue the test etherchannel load-balance interface port-channel number {ip |

l4port | mac} [source_ip_add | source_mac_add | source_l4_port] [dest_ip_add | dest_mac_add | dest_l4_port] command.

QUESTION NO: 94

Which three combinations are valid LACP or static configurations that will set up a channel? (Choose three.)

A. On/On
B. On/Auto
C. Passive/Active
D. Desirable/Auto
E. Active/Active
F. Desirable/Desirable

Answer: A,C,E

Explanation:

There are only three valid combinations to run the LACP link aggregate, as this table shows:

Switch

Switch

Comments

active

active

Recommended.

active

passive

Link aggregation occurs if negotiation is successful.

on

on

Link aggregation occurs without LACP. Although this works, this combination is not recommended.

Note: By default, with the configuration of an LACP channel, the LACP channel mode is passive.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps663/products_configuration_example09186a0080094aec.shtml

QUESTION NO: 95

Which three combinations are valid PAgP or static configurations that will set up a channel? (Choose three.)

A. On-On
B. On-Auto
C. Passive-Active
D. Desirable-Auto
E. Active-Active
F. Desirable-Desirable

Answer: A,D,F

Explanation:

PAgP modes are off, auto, desirable, and on. Only the combinations auto-desirable, desirable-desirable, and on-on allow the formation of a channel.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps663/products_configuration_example09186a0080094aec.shtml

Topic 2, Implement IPv4

QUESTION NO: 96

A new backup connection is being deployed on a remote site router. The stability of the connection has been a concern. In order to provide more information to EIGRP regarding this interface, you wish to incorporate the "reliability" cost metric in the EIGRP calculation with the command metric weights 1 0 1 0 1.

What impact will this modification on the remote site router have for other existing EIGRP neighborships from the same EIGRP domain?

A. Existing neighbors will immediately begin using the new metric.
B. Existing neighbors will use the new metric after clearing the EIGRP neighbors.
C. Existing neighbors will resync, maintaining the neighbor relationship.
D. All existing neighbor relationships will go down.

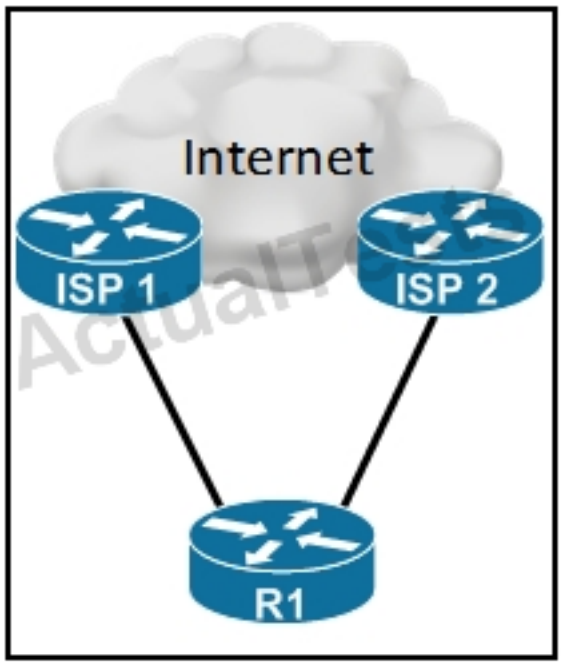
Answer: D

Explanation:

For EIGRP Neighbor relationship to form, K values must match on both routers.

QUESTION NO: 97

Refer to the exhibit.



R1 has an EBGP session to ISP 1 and an EBGP session to ISP 2. R1 receives the same prefixes

through both links.

Which configuration should be applied so that the link between R1 and ISP 2 will be preferred for outgoing traffic (R1 to ISP 2)?

A. Increase local preference on R1 for received routes
B. Decrease local preference on R1 for received routes
C. Increase MED on ISP 2 for received routes
D. Decrease MED on ISP 2 for received routes

Answer: A

Explanation: Explanation
Local preference is an indication to the AS about which path has preference to exit the AS in order to reach a certain network. A path with higher local preference is preferred more. The default value of preference is 100.

Reference
http://www.cisco.com/en/US/tech/tk872/technologies_configuration_example09186a0080b82d1f.shtml?

QUESTION NO: 98

Policy-based routing allows network administrators to implement routing policies to allow or deny paths based on all of these factors except which one?

A. End system
B. Protocol
C. Application
D. Throughput

Answer: D

Explanation:
Policy-based routing (PBR) provides a mechanism for expressing and implementing forwarding/routing of data packets based on the policies defined by the network administrators. It provides a more flexible mechanism for routing packets through routers, complementing the existing mechanism provided by routing protocols.

Policy-based routing allows network administrators to determine and implement routing policies to allow or deny paths based on the following:
• Identity of a particular end system

• Application
• Protocol
• Size of packets

QUESTION NO: 99

Refer to the exhibit.

```
BGP table version is 2, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 160.0.0.0/8      4.4.4.1                                0 300 {200,100} i
```

Which statement is correct about the prefix 160.0.0.0/8?

A. The prefix has encountered a routing loop.
B. The prefix is an aggregate with an as-set.
C. The prefix has been aggregated twice, once in AS 100 and once in AS 200.
D. None of these statements is true.

Answer: B
Explanation:
The as-set keyword within an aggregate-address will set all AS which the more-specific routes have passed to the point where the aggregate-address has been configured. Without this keyword, the aggregate address would only show the passed AS' from the point where the aggregate-route has been configured.

QUESTION NO: 100

Which two options does Cisco PfR use to control the entrance link selection with inbound optimization? (Choose two.)

A. Prepend extra AS hops to the BGP prefix.
B. Advertise more specific BGP prefixes (longer mask).
C. Add (prepend) one or more communities to the prefix that is advertised by BGP.
D. Have BGP dampen the prefix.

Answer: A,C

Explanation: PfR Entrance Link Selection Control Techniques
The PfR BGP inbound optimization feature introduced the ability to influence inbound traffic. A network advertises reachability of its inside prefixes to the Internet using eBGP advertisements to its ISPs. If the same prefix is advertised to more than one ISP, then the network is multihoming. PfR BGP inbound optimization works best with multihomed networks, but it can also be used with a network that has multiple connections to the same ISP. To implement BGP inbound optimization, PfR manipulates eBGP advertisements to influence the best entrance selection for traffic bound for inside prefixes. The benefit of implementing the best entrance selection is limited to a network that has more than one ISP connection.

To enforce an entrance link selection, PfR offers the following methods:
BGP Autonomous System Number Prepend When an entrance link goes out-of-policy (OOP) due to delay, or in images prior to Cisco IOS Releases 15.2(1) T1 and 15.1(2)S, and PfR selects a best entrance for an inside prefix, extra autonomous system hops are prepended one at a time (up to a maximum of six) to the inside prefix BGP advertisement over the other entrances. In Cisco IOS Releases 15.2(1)T1, 15.1(2)S, and later releases, when an entrance link goes out-of policy (OOP) due to unreachable or loss reasons, and PfR selects a best entrance for an inside prefix, six extra autonomous system hops are prepended immediately to the inside prefix BGP advertisement over the other entrances. The extra autonomous system hops on the other entrances increase the probability that the best entrance will be used for the inside prefix. When the entrance link is OOP due to unreachable or loss reasons, six extra autonomous system hops are added immediately to allow the software to quickly move the traffic away from the old entrance link. This is the default method PfR uses to control an inside prefix, and no user configuration is required.

BGP Autonomous System Number Community Prepend
When an entrance link goes out-of-policy (OOP) due to delay, or in images prior to Cisco IOS Releases 15.2

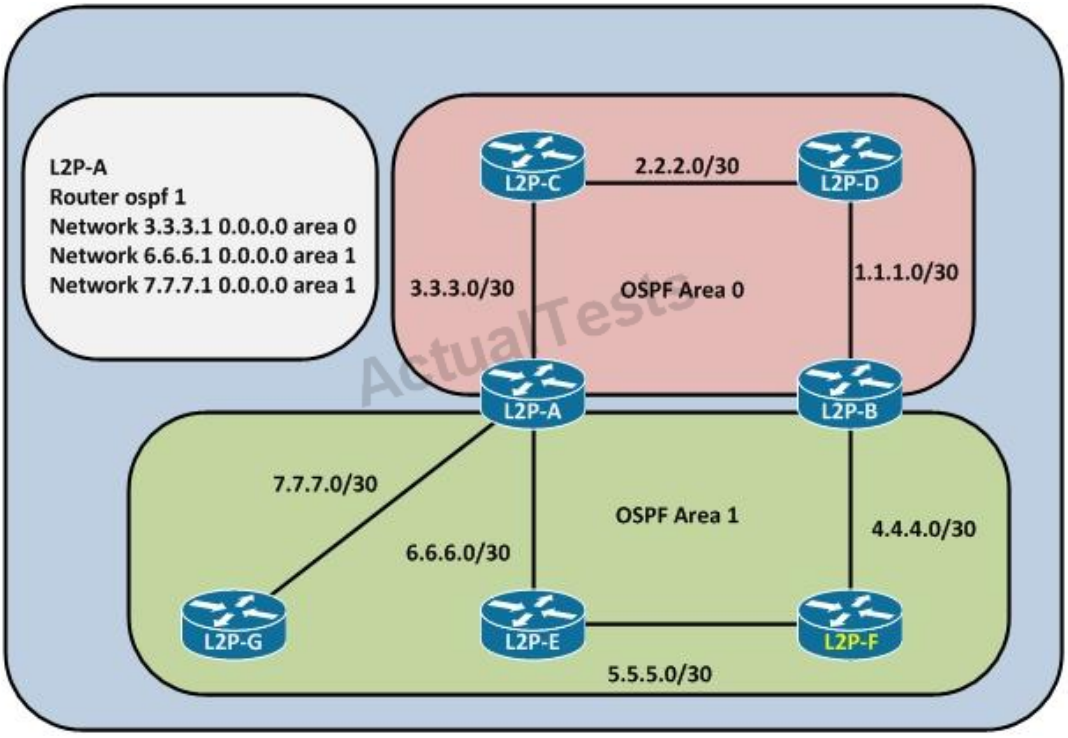
(1)T1 and 15.1(2)S, and PfR selects a best entrance for an inside prefix, a BGP prepend community is attached one at a time (up to a maximum of six) to the inside prefix BGP advertisement from the network to another autonomous system such as an ISP. In Cisco IOS Releases 15.2(1)T1, 15.1(2)S, and later releases, when an entrance link goes out-of-policy (OOP) due to unreachable or loss reasons, and PfR selects a best entrance for an inside prefix, six BGP prepend communities are attached to the inside prefix BGP advertisement. The BGP prepend community will increase the number of autonomous system hops in the advertisement of the inside prefix from the ISP to its peers. Autonomous system prepend BGP community is the preferred method to be used for PfR BGP inbound optimization because there is no risk of the local ISP filtering the extra autonomous system hops. There are some issues, for example, not all ISPs support the BGP prepend community, ISP policies may ignore or modify the autonomous

system hops, and a transit ISP may filter the autonomous system path. If you use this method of inbound optimization and a change is made to an autonomous system, you must issue an outbound reconfiguration using the "clear ip bgp" command.

Reference
http://www.cisco.com/en/US/docs/ios-xml/ios/pfr/configuration/15-2s/pfr-bgp-inbound.html#GUID-F8A59E24-
1D59-4924-827D-B23B43D9A8E0
http://www.cisco.com/en/US/products/ps8787/products_ios_protocol_option_home.html

QUESTION NO: 101

Refer to the exhibit.



What is the potential issue with this configuration?

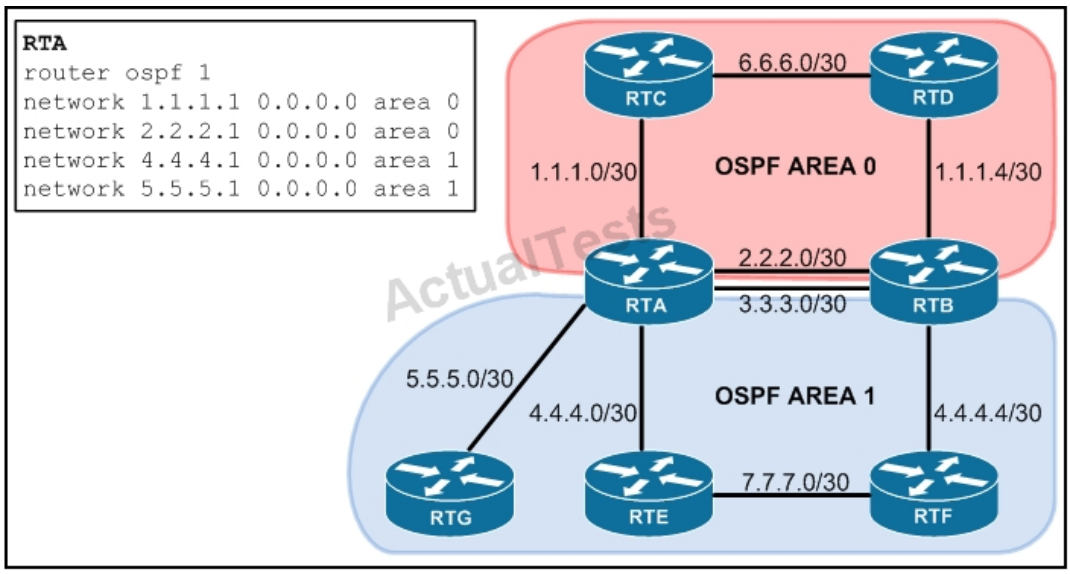
A. There is no potential issue; OSPF will work fine in any condition.
B. Sub-optimal routing may occur since there is no area 1 adjacency between the ABRs.
C. This is a wrong OSPF configuration because all routers must be in area 0 only.
D. This is a wrong OSPF configuration because /30 requires 0.0.0.3 wild card.

Answer: B

Explanation:
The primary difference is that the GRE tunnel hides the real recipients under the outer IP header. The "transit" area, in that case, does not need to actually know all the prefixes. What it needs, though, is a default route that points towards the backbone. Suboptimal routing may occur, then, because the traffic from the "transit" area must first reach the backbone in order to get GRE-encapsulated and carried again through the "transit" area to the disconnected area.

QUESTION NO: 102

Refer to the exhibit.



A packet from RTD with destination RTG, is reaching RTB. What is the path this packet will take from RTB to reach RTG?

A. RTB - RTA - RTG
B. RTB - RTD - RTC - RTA - RTG
C. RTB - RTF - RTE - RTA - RTG
D. RTB will not be able to reach RTG since the OSPF configuration is wrong.

Answer: A
Explanation: RTB takes the intra area route in Area 1, and RTA is directly connected.

RTB#sh ip cef exact-route 1.1.1.5 5.5.5.2
1.1.1.5 -> 5.5.5.2 (next hop 3.3.3.1)

RTA#sh ip cef exact-route 1.1.1.5 5.5.5.2
1.1.1.5 -> 5.5.5.2 : Fal0/0 (attached)

QUESTION NO: 103

Refer to the exhibit.

```
R1# show ip bgp 10.1.0.1
BGP routing table entry for  10.1.0.0/16, version 11
Paths: (2 available, best #?, table  Default-IP-Routing-Table)
  Advertised to non peer-group peers:
    2
  65000
    10.168.30.4 (metric 74) from 3.3.3.3  (3.3.3.3)
      Origin IGP, metric 100, localpref 100, valid,  internal
  65000
    10.168.20.4 from 192.168.20.4  (4.4.4.4)
      Origin IGP, metric 200, localpref 100, valid,  external
```

ActualTests

Which path is selected as best path?

A. path 1, because it is learned from IGP
B. path 1, because the metric is the lowest
C. path 2, because it is external
D. path 2, because it has the higher router ID

Answer: B
Explanation:
Metrics is a property of a route in computer networking, consisting of any value used by a routing protocol to determine whether one particular route should be chosen over another. The routing table stores only the best possible routes, while link-state or topological databases may store all other information as well. For example, Routing Information Protocol uses hopcount (number of hops) to determine the best possible route. The route will go in the direction of the gateway with the lowest metric (default gateway).
Reference: http://en.wikipedia.org/wiki/Metrics_(networking)

QUESTION NO: 104

What action will a BGP route reflector take when it receives a prefix marked with the community attribute NO ADVERTISE from a client peer?

A. It will advertise the prefix to all other client peers and non-client peers.
B. It will not advertise the prefix to EBGP peers.
C. It will only advertise the prefix to all other IBGP peers.
D. It will not advertise the prefix to any peers.

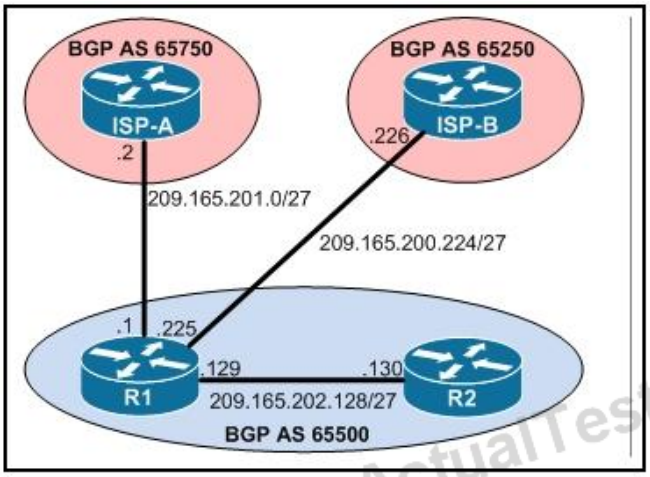
Answer: D

Explanation:

BGP requires that all BGP peers in the same autonomous system form an iBGP session with all peers in the autonomous system. This is too difficult in many environments. Route reflectors are fully functional iBGP speakers that form iBGP sessions with other iBGP speakers, and they also perform a second function - they forward routes from other iBGP speakers to route reflector clients. The route reflector clients and clients form a cluster.

QUESTION NO: 105

Refer to the exhibit.




```
R1#show ip route 172.16.0.0
Routing entry for 172.16.0.0/16
  Known via "bgp 65500", distance 200, metric 0, type internal
  Last update from 209.165.202.130 00:03:15 ago
  Routing Descriptor Blocks:
  * 209.165.202.130, from 209.165.202.130, 00:03:15 ago
      Route metric is 0, traffic share count is 1
      AS Hops 0

R2#show ip route 172.16.0.0
Routing entry for 172.16.0.0/24, 1 known subnets
  Redistributing via eigrp 100, bgp 65500
  Advertised by bgp 65500

D 172.16.10.0 [90/156160] via 209.165.202.158, 00:09:19, FastEthernet0/0
```

ActualTests

R1 is not learning about the 172.16.10.0 subnet from the BGP neighbor R2 (209.165.202.130). What can be done so that R1 will learn about this network?

A. Disable auto-summary on R2.
B. Configure an explicit network command for the 172.16.10.0 subnet on R2.
C. Subnet information cannot be passed between IBGP peers.
D. Disable auto-summary on R1.

Answer: B

Explanation:

By default, BGP does not accept subnets redistributed from IGP. To advertise and carry subnet routes in BGP, use an explicit network command or the no auto-summary command. If you disable auto-summarization and have not entered a network command, you will not advertise network routes for networks with subnet routes unless they contain a summary route.

Reference
http://www.cisco.com/en/US/docs/ios/11_3/np1/command/reference/1rbgp.html

QUESTION NO: 106

Refer to the exhibit.

```
r13#show ip eigrp nei det
IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   192.168.1.1             Et0/0            9999 00:20:26    9   200  0  9
   Version 12.4/1.2, Retrans: 0, Retries: 0, Prefixes: 1
4   192.168.3.7             Et0/2             10 00:21:07   25   200  0  27
   Version 12.4/1.2, Retrans: 0, Retries: 0
   Stub Peer Advertising ( STATIC ) Routes
   Suppressing queries
3   192.168.3.8             Et0/2             12 00:21:26   26   200  0  25
   Version 12.4/1.2, Retrans: 0, Retries: 0
   Stub Peer Advertising ( SUMMARY ) Routes
   Suppressing queries
2   192.168.3.6             Et0/2             14 00:33:41   16   200  0  19
   Restart time 00:33:14
   Version 12.4/1.2, Retrans: 0, Retries: 0, Prefixes: 1
0   192.168.2.1             Et0/1            9999 00:43:06   17   200  0  6
   Restart time 00:33:14
   Version 12.4/1.2, Retrans: 2, Retries: 0, Prefixes: 1
5   192.168.3.9             Et0/2             11 00:33:41   16   200  0  19
   Restart time 00:33:14
   Version 12.4/1.2, Retrans: 0, Retries: 0, Prefixes: 1
```

ActualTests

After a link flap in the network, which two EIGRP neighbors will not be queried for alternative paths? (Choose two.)

A. 192.168.1.1
B. 192.168.3.7
C. 192.168.3.8
D. 192.168.3.6
E. 192.168.2.1
F. 192.168.3.9

Answer: B,C

Explanation:

Explanation

Both 192.168.3.7 & 192.168.3.8 are in an EIGRP Stub area

The Enhanced Interior Gateway Routing Protocol (EIGRP) Stub Routing feature improves network stability, reduces resource utilization, and simplifies stub router configuration.

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.

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. Any neighbor that receives a packet informing it of the stub status will not query the stub router for any routes, and a router that has a stub peer will not query that peer. The stub router will depend on the distribution router to send the proper updates to all peers.

Reference

http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/eigrpstb.html#wp1021949

QUESTION NO: 107

Refer to the exhibit.

```
BGP table version is 11, local router ID is 192.168.3.2
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
*  172.16.1.0/24    172.16.1.1               0    100      0 (65333) 62000 ?
*>i                 192.168.2.1              0    100      0 62000 ?
```

Why is AS 65333 in parentheses?

A. It is an external AS.
B. It is a confederation AS.
C. It is the AS of a route reflector.
D. It is our own AS.
E. A route map has been applied to this route.
F. The BGP next hop is unreachable.

Answer: B

Explanation:

The AS numbers that are part of the confederation are grouped between parentheses and are replaced by the confederation identifier (the real AS number) in nonconfederation eBGP sessions.

.

QUESTION NO: 108

Refer to the exhibit.

```
C1#show ip ospf statistics
            OSPF Router with ID (11.100.1.11) (Process ID 100)
  Area 0: SPF algorithm executed 2 times
  Summary OSPF SPF statistic
  SPF calculation time
Delta T   Intra D-Intra Summ    D-Summ  Ext     D-Ext   Total   Reason
00:05:12  0     0       0       0       0       0       0       R, SN, X
00:05:02  0     0       0       0       0       0       0       R, SN, X
00:02:57  0     0       0       0       0       0       0       X
```

What triggered the first SPF recalculation?

A. changes in a router LSA, subnet LSA, and external LSA

B. changes in a router LSA, summary network LSA, and external LSA
C. changes in a router LSA, summary network LSA, and summary ASBR LSA
D. changes in a router LSA, summary ASBR LSA, and external LSA

Answer: B

Explanation:

OSPFv2

Is built around links, and any IP prefix change in an area will trigger a full SPF. It advertises IP information in Router and Network LSAs. The routers thus, advertise both the IP prefix information (or the connected subnet information) and topology information in the same LSAs. This implies that if an IP address attached to an interface changes, OSPF routers would have to originate a Router LSA or a Network LSA, which btw also carries the topology information. This would trigger a full SPF on all routers in that area, since the same LSAs are flooded to convey topological change information. This can be an issue with an access router or the one sitting at the edge, since many stub links can change regularly.

Only changes in interarea, external and NSSA routes result in partial SPF calculation (since type 3, 4, 5 and 7 LSAs only advertise IP prefix information) and thus IS-IS's PRC is more pervasive than OSPF's partial SPF.

This difference allows IS-IS to be more tolerant of larger single area domains whereas OSPF forces hierarchical designs for relatively smaller networks. However with the route leaking from L2 to L1 incorporated into IS-IS the apparent motivation for keeping large single area domains too goes away.

SPF is calculated in three phases. The first is the calculation of intra-area routes by building the shortest path tree for each attached area. The second phase calculates the inter-area routes by examining the summary LSAs and the last one examines the AS-External-LSAs to calculate the routes to the external destinations.

Reference
http://routingfreak.wordpress.com/2008/03/04/shortest-path-first-calculation-in-ospf-and-is-is/

QUESTION NO: 109

Which two orders in the BGP Best Path Selection process are correct? (Choose two.)

A. Higher local preference, then lowest MED, then eBGP over iBGP paths
B. Higher local preference, then highest weight, then lowest router ID
C. Highest weight, then higher local preference, then shortest AS path
D. Lowest origin type, then higher local preference, then lowest router ID

E. Highest weight, then higher local preference, then highest MED

Answer: A,C

Explanation:

Weight is the first attribute BGP uses in the route selection process. Route with a higher weight is preferred when multiple routes exist to the same destination.

QUESTION NO: 110

Refer to the exhibit.

```
R2#show ip ospf interface ethernet 1/0
Ethernet1/0 is up line protocol is up
 Internet Address 10.10.23.1/30 Area0
 Process ID1.Router ID 10.10.10.2. Network Type POINT_TO_POINT,Cost:10
 Transmit Delay is 1 sec,State POINT_TO_POINT,
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
   Hello due in 00:00:06
 Index 1/1, flood queue length 0
 Next 0x0(0)/0x0(0)
 Last flood scan length is 1, maximum is 7
 Last flood scan time is 0 msec, maximum is 0 msec
 Neighbor Count is 0, Adjacent neighbor count is 0
 Suppress hello for 0 neighbor(s)

R2#debug ip packet detail
00:22:57:IP:s*10.10.23.1 (local),d=224.0.0.5 (Ethernet1/0),len63,sending b
road/multicast, proto*89
00:23:07:IP:s*10.10.23.1 (local),d=224.0.0.5 (Ethernet1/0),len63,sending b
road/multicast, proto*89

R3#show interfaces ethernet1/0
Ethernet2/0 is up, line protocol is up
 Internet address is 10.10.23.2/30
 MTU 1500 bytes, BW 10000 Kbit, DLY1000 usec,
 Encapsulation ARPA, loopback not set
 Keepalive set (10 sec)
 ARP type: ARPA, ARP Timeout 04:00:00
 Last input 00:00:00 output 00:00:01, output hand never
 Last clearing of "show interface" counters 00:03:58
 Input queue: 0/75/0/0(size/max/drops/flashes): Total output drops:0
 5 minute input rate 0 bit/sec, 0 packets/sec
 5 minute output rate 0 bits/sec, 0 packets/sec
  93 packets input, 2604 bytes, 0 no buffer
  Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overruns, 0 ignored
  0 input packets with dribble condition detected
  115 packets output, 9122 bytes, 0 underruns
  0 output errors, 0 collisions, 0 interface resets
  0 babbles, 0 late collision, 0 deferred
  0 lost carrier, 0 no carrier
```

```
R3#show ip ospf interface ethernet 2/0
Ethernet2/0 is up line protocol is up
 Internet Address 10.10.23.2/30 Area0
 Process ID2.Router ID 10.10.10.3. Network Type POINT_TO_POINT,Cost:20
 Transmit Delay is 1 sec,State POINT_TO_POINT,
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
   Hello due in 00:00:06
 Index 1/1, flood queue length 0
 Next 0x0(0)/0x0(0)
 Last flood scan length is 1, maximum is 7
 Last flood scan time is 0 msec, maximum is 0 msec
 Neighbor Count is 0, Adjacent neighbor count is 0
 Suppress hello for 0 neighbor(s)

R3#debug ip packet detail
00:28:36:IP:s*10.10.23.2 (local),d=224.0.0.5 (Ethernet2/0),len68,sending b
road/multicast, proto*89
00:28:46:IP:s*10.10.23.2 (local),d=224.0.0.5 (Ethernet2/0),len68,sending b
road/multicast, proto*89

R2#show interfaces ethernet1/0
Ethernet1/0 is up, line protocol is up
 Internet address is 10.10.23.1/30
 MTU 1500 bytes, BW 20000 Kbit, DLY 1000 usec,
 Encapsulation ARPA, loopback not set
 Keepalive set (10 sec)
 ARP type: ARPA, ARP Timeout 04:00:00
 Last input 00:00:02 output 00:00:02, output hand never
 Last clearing of "show interface" counters 00:04:44
 Input queue: 0/75/0/0(size/max/drops/flashes): Total output drops:0
 5 minute input rate 0 bit/sec, 0 packets/sec
 5 minute output rate 0 bits/sec, 0 packets/sec
  101 packets input, 8346 bytes, 0 no buffer
  Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overruns, 0 ignored
  0 input packets with dribble condition detected
  132 packets output, 10044 bytes, 0 underruns
  0 output errors, 0 collisions, 0 interface resets
  0 babbles, 0 late collision, 0 deferred
  0 lost carrier, 0 no carrier
```

R2 and R3 are routers connected using Ethernet services from a service provider and can receive pings from each other. OSPF is configured as the routing protocol but adjacency is not happening. According to the output of the show commands in the exhibit, what could be the most likely cause of the problem?

A. Ethernet interfaces were configured as point-to-point.
B. Process IDs are not matching.
C. Configured bandwidths do not match on both interfaces.
D. Broadcasts and multicast are not being propagated over the Ethernet services.
E. OSPF cost does not match on both interfaces.

Answer: D

Explanation:

OSPF Adjacencies
Occurs through exchange of Hello packets
After adjacency established, link-state databases (LSDBs) are synched
Two OSPF neighbors on point-to-point link form full adjacency with each other
In LANs, all routers form adjacency with the DR and BDR; updates need to be sent only to DR, which updates all other routers; all other routers on LAN are called DROTHERS and maintain a partial neighbor relationship with each other
Once adjacencies are established, LSAs are exchanged through a reliable mechanism. LSAs are flooded to ensure topological awareness. LSAs have a sequence number and a lifetime value. LSAs convey the cost of links used for the SPF calculation. The cost metric is based on interface bandwidth. The LSA aging timer is 30-minute default.
Here are the details of the exchange process between two routers on a LAN (Router 1 and Router 2) and the OSPF adjacency states involved:
Step 1 Router 1 begins in the down state because it is not exchanging OSPF information with any other router.
It sends Hello packets via multicast address 224.0.0.5(all SPF).
Step 2 Router 2 receives the OSPF Hello and adds Router 1 in its list of neighbors. This is the beginning of the init state.
Step 3 Router 2 sends a unicast Hello packet response to Router 1.
Step 4 Router 1 receives the Hello and notes that it is listed in the packet. It ads Router 2 to its list of neighbors. Router 1 knows that it has bidirectional communication with Router 2. This is known as the Two-Way State.
Step 5 In LAN environment, DR and BDR elected
Step 6 In LAN environment, Hello packets function as keepalive mechanism every 10 seconds.

Reference
Cisco General Network Theory Quick Reference Sheets

QUESTION NO: 111

The core of a network has four routers connected in a square design with Gigabit Ethernet links using /30 subnets. The network is used to carry voice traffic and other applications. Convergence time is taking more than expected. Which three actions would you take to improve OSPF convergence time? (Choose three.)

A. Increase MTU of the interfaces to accommodate larger OSPF packets
B. Change the network type to point-to-point on those links.
C. Reduce SPF initial timer.
D. Increase hello interval to avoid adjacency flapping.

E. Enable OSPF.

Answer: B,C,D

Explanation:

The OSPF SPF Throttling is configured using the command timers throttle spf OSPF router configuration command.

spf-start: Initial delay to schedule an SPF calculation after a topology change. Range is 1 to 600000 milliseconds.spf-hold: Minimum hold-time between two SPF calculations. Range is 1 to 600000 milliseconds.spf-max-wait: Maximum wait between two SPF calculations. Range is 1 to 600000 milliseconds.

This command is used to delay the SPF algorithm being executed during network instability. SPF Is very CPU intensive.

QUESTION NO: 112

Refer to the exhibit.

```
R3#show ip route
…
Gateway of last resort is not set

     209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C       209.165.200.224/27 is directly connected, Serial0/0/0
D       209.165.200.0/24 is a summary, 00:01:27, Null0
     209.165.201.0/27 is subnetted, 1 subnets
D       209.165.201.0 [90/20514560] via 209.165.200.242, 20:41:09, Serial0/0/0
     209.165.202.0/24 is variably subnetted, 2 subnets, 2 masks
C       209.165.202.128/27 is directly connected, FastEthernet0/0
D       209.165.202.0/24 is a summary, 00:01:27, Null0

R3#show ip eigrp neighbors
IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   209.165.202.155         Fa0/0             10 00:23:31    1   200  0  7
0   209.165.200.242         Se0/0/0          158 20:41:25  211  1266  0  67
```

What can be done to remove the summary routes to Null0 on R3?

A. Configure the EIGRP routing subcommand no auto-summary on 209.165.202.155.
B. Configure the EIGRP routing subcommand no auto-summary on 209.165.202.24.
C. Configure the EIGRP routing subcommand no auto-summary on both 209.165.202.155 and 209.165.202.242.
D. Configure the EIGRP routing subcommand no auto-summary on R3.

Answer: D

Explanation:

Disabling automatic summarization will remove the Null0 summary route and allow EIGRP to look for a supernet or default route when an EIGRP child route does not match a destination packet.

QUESTION NO: 113

Refer to the exhibit.

```
R4#show ip route
…
Gateway of last resort is not set

     209.165.202.0/27 is subnetted, 1 subnets
C       209.165.202.128 is directly connected, FastEthernet0/0

R4#show ip eigrp neighbor detail
IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   209.165.202.139         Fa0/0             13 00:00:29    1   200  0  99
   Version 12.4/1.2, Retrans: 0, Retries: 0
   Receive-Only Peer Advertising ( No ) Routes

R3#show ip route eigrp
*   209.165.201.0/27 is subnetted, 1 subnets
D*      209.165.201.0 [90/20514560] via 209.165.200.242, 00:01:36, Serial0/0/0
```

R4 is configured as a receive-only EIGRP stub, and is adjacent with 209.165.202.139 (R3). However, R4 is not learning about network 209.165.201.0/27 from R3. What could be the cause of this issue?

A. R4 should learn this route from 209.165.200.242, and not from R3.
B. R3 is configured as a receive-only EIGRP stub.
C. R3 and R4 may be using different EIGRP process numbers.
D. R3 and R4 are asymmetrically adjacent neighbors.

Answer: B

Explanation:

Configuring EIGRP Stub Routing

To configure a remote or spoke router for EIGRP stub routing, use the following commands beginning in router configuration mode:

	Command	Purpose
Step 1	`router(config)# router eigrp as-number`	Configures a remote or distribution router to run an EIGRP process.
Step 2	`router(config-router)# network network-number`	Specifies the network address of the EIGRP distribution router.
Step 3	`router(config-router)# eigrp stub [receive-only | connected | static | summary]`	Configures a remote router as an EIGRP stub router.

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

QUESTION NO: 114

Refer to the exhibit.

```
R4#show ip eigrp neighbors
IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   209.165.202.139         Fa0/0             14 00:00:22    6   200  0  23

R4(config)#router eigrp 100
R4(config-router)#neighbor 209.165.202.139 fa0/0

*Nov 11 02:01:05.354: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 209.165.202.139

R4#show ip eigrp neighbors
IP-EIGRP neighbors for process 100
```

R4 is a remote office router that is running EIGRP; the decision has been made to change EIGRP to use static EIGRP adjacencies. However, once the configuration change was applied, the adjacency between R4 and 209.165.202.139 (HQ) seems to disappear. What could be the cause of this issue?

A. Static EIGRP neighbor configuration is symmetric; it causes the interface to stop processing inbound multicast packets and stop sending multicast packets.
B. Static EIGRP neighbors are not displayed with the show ip eigrp neighbors command.
C. A distance (internal 90 or external 170) must also be configured for the static neighbor.
D. The neighbor 209.165.202.139 should be changed to run under EIGRP autonomous system 0.

Answer: A

Explanation:

As you can see that static EIGRP neighbor configuration is symmetric. This causes the interface to halt processing inbound packets and stop sending multicast packets.

QUESTION NO: 115

Refer to the exhibit.

```
*Nov  2 00:47:06.246: OSPF: Rcv hello from 209.165.202.140 area 209.165.202.128
 from FastEthernet0/0 209.165.202.140
*Nov  2 00:47:06.246: OSPF: Hello from 209.165.202.140 with mismatched
Stub/Transit area option bit
```

What problem does the debug ip ospf event output from R3 indicate?

A. 209.165.202.140 and R3 are not both configured as OSPF stubs.
B. 209.165.202.140 and R3 are not configured in the same OSPF area.
C. 209.165.202.140 is configured as a no-summary stub.
D. Transit area OSPF hello packets are not processed by design.

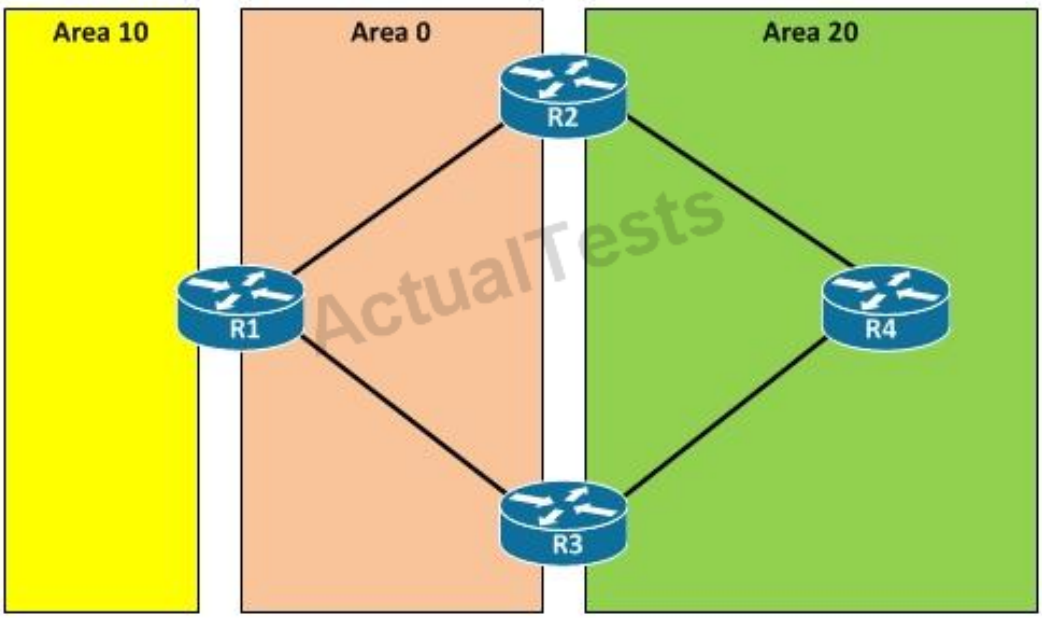
Answer: A

Explanation:

As you can see that the hello packets are mismatched. This means that 209.165.202.140 and R3 are not configured as OSPF stubs.

QUESTION NO: 116

Refer to the exhibit.



R4 is configured as an OSPF stub; however, R4 should still be learning the OSPF LSA type 3 interarea routes from 209.165.202.130. Which action will solve this issue?

A. Remove any route maps from R4 that are filtering the incoming OSPF updates.

B. Enable sending summary LSA's by removing no-summary from the stub command on the ABR.
C. Enable sending summary LSA's by removing no-summary from the stub command on the ASBR.
D. Control of interarea route propagation is best handled with EIGRP.

Answer: B

Explanation:

By default the no-summary command should NOT be enabled on the ABR as this would make Area 20 a Totally Stubby Area not a Stub area. And R4 should receive routes fro 209.165.202.130 as it is directly connected to Area 0 based off of the exhibit.

Table 5. OSPF Stubby Area Types

Area Type	LSA 1	LSA 2	LSA 3	LSA 4	LSA 5	LSA 7
Backbone	Yes	Yes	Yes	Yes	Yes	No
Nonbackbone	Yes	Yes	Yes	Yes	Yes	No
Stub	Yes	Yes	Yes	Yes	No	No
Totally stubby	Yes	Yes	No	No	No	No
Not-so-stubby area (NSSA)	Yes	Yes	Yes	Yes	No	Yes
Totally NSSA	Yes	Yes	No	No	No	Yes

ActualTests

Table 6. Stub Area Configuration Options

Stub Type	Router OSPF Subcommand
NSSA	**area** *area-id* **nssa**
Totally NSSA	**area** *area-id* **nssa no-summary**
Stub	**area** *area-id* **stub**
Totally stubby	**area** *area-id* **stub no-summary**

QUESTION NO: 117

Refer to the exhibit.

```
Router1# show run
!
!
router ospf 1
network 192.168.1.0 0.0.0.255 area 1
network 172.16.1.0 0.0.0.255 area 1
area 1 stub
redistribute static subnets
!
ip route 10.10.10.0 255.255.255.0 172.16.2.1
!
```

Router 1 has three OSPF neighbors, but none of them have routes to the 10.10.10.0 network. Why?

A. A metric is not defined on the redistribute statement.
B. Redistribution cannot be done in a non-zero area.
C. The subnets keyword should not be used.
D. Area 1 is a stub area.
E. Area 1 is a not-so-stubby area.
F. A route map must be used on the redistribute statement to permit the 10.10.10.0/24 route.

Answer: D
Reference:
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094819.shtml (see the problem - mismatched area ID)

QUESTION NO: 118

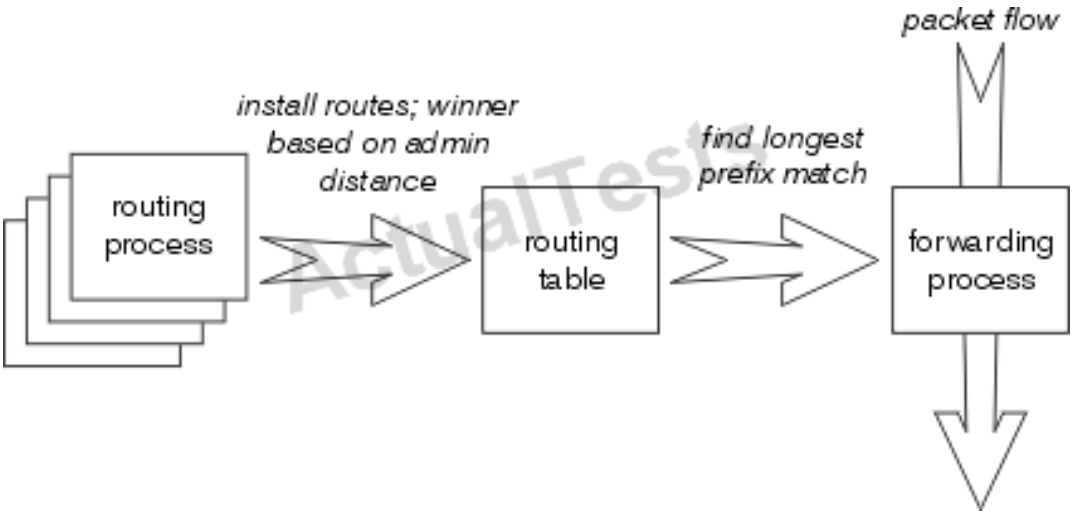
When a router makes a forwarding decision, which of these routes in the routing table always wins?

A. administrative distance
B. router ID
C. longest prefix match
D. routing process ID

Answer: C
Explanation:

Making a forwarding decision actually consists of three sets of processes: the routing protocols, the routing table, and the actual process, which makes a forwarding decision and switches

packets. These three sets of processes are illustrated, along with their relationship, below:



The longest prefix match always wins among the routes actually installed in the routing table, while the routing protocol with the lowest administrative distance always wins when installing routes into the routing table.

Reference: www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094823.shtml

QUESTION NO: 119

Refer to the exhibit.

```
CE1#show ip eigrp topology active
IP-EIGRP Topology Table for AS(65000)/ID(11.1.3.1)

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

A 55.55.55.55/32, 1 successors, FD is Inaccessible, tag is 1
    1 replies, active 00:00:43, query-origin: Local origin
        via Redistributed (40256000/0)
      Remaining replies:
         via 11.1.3.2, r, Ethernet1/0
```

Which statement is correct?

A. Router CE1 is waiting for the query for the prefix 55.55.55.55/32 to be acknowledged.
B. Router CE1 is waiting for the reply for the prefix 55.55.55.55/32 to be received.
C. Router CE1 still has to send out a query for the prefix 55.55.55.55/32.
D. Router CE1 has sent out a reply that has not been acknowledged yet.

Answer: B

Explanation:
If you see the remaining replies, you will know that CE1 is waiting for the reply from 55.55.55.55/32.

QUESTION NO: 120

When a BGP router is not capable of understanding 4-byte AS numbers, it will see 4-byte AS numbers as a special, reserved, 2-byte AS number in the AS path. Which 2-byte AS number is this reserved one?

A. 00000
B. 12345
C. 23456
D. 65000
E. 99999

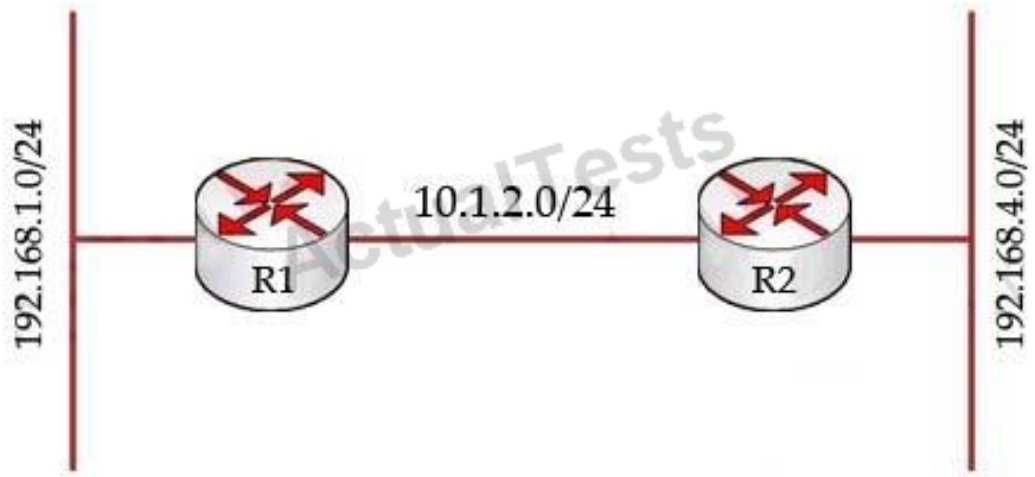
Answer: C
Explanation:
What's New with 4-byte AS Number
The new AS number is 4-bytes and split into two 2-byte values, in X.Y syntax. The support for the 4-byte AS is advertised via BGP capability negotiation. In order to ensure interoperability with existing BGP peers that do not support 4-byte AS, encoding of BGP OPEN message is reserved and 4-byte AS support is exchanged between the BGP peers via the capability field.
In this whitepaper, we will refer to the BGP speaker that supports 4-byte AS as NEW speaker, and the BGP speaker that does not support 4-byte AS as OLD speaker.
When BGP attempts to establish a session with its peer, the OPEN message may include an optional parameter, called Capabilities. A NEW speaker will include the NEW (4-byte AS) capability when it attempts to OPEN a session with its peer. An OLD speaker should simply ignore the NEW capability advertised by its peer and continue to operate in OLD mode, as detailed in RFC 3392.
If the NEW speaker advertises and receives the 4-byte AS capability from its peer, it will just encode the 4-byte AS number in its AS_PATH or AGGREGATOR attributes when exchanging information with this peer.
If the NEW speaker does not receive the 4-byte AS capability from a particular peer, it indicates this peer is an OLD speaker. Two new attributes are introduced, namely AS4_PATH and AS4_AGGREGATOR. Both attributes are optional transitive. These new attributes use the same encoding as the original ASPATH and AGGREGATOR except the AS Number used is 4-bytes instead of 2-bytes. The NEW speaker will substitute a reserved 2-byte AS number (called AS_TRANS with AS # 23456) for each 4-byte AS so that ASPATH and AGGREGATOR is still 2-

byte in length and ASPATH length is still preserved, and at the same time insert the new AS4_PATH and AS4_AGGREGATOR, which will contain the 4-byte encoded copy of the attributes. The NEW speaker will then advertise ASPATH and/or AGGREGATOR together with the AS4_PATH and/or AS4_AGGREGATOR. The OLD speaker that receives these new attributes will preserve and blindly pass them along even though it does not understand them.

Reference
http://www.cisco.com/web/about/security/intelligence/4byte-as.html

QUESTION NO: 121

Refer to the exhibit. In this network, R1 has been configured to advertise a summary route, 192.168.0.0/22, to R2. R2 has been configured to advertise a summary route, 192.168.0.0/21, to R1. Both routers have been configured to remove the discard route (the route to null created when a summary route is configured) by setting the administrative distance of the discard route to 255. What will happen if R1 receives a packet destined to 192.168.3.1 ?



A. The packet will loop between R1 and R2.
B. It is not possible to set the administrative distance on a summary to 255.
C. The packet will be forwarded to R2, where it will be routed to null0.
D. The packet will be dropped by R1, since there is no route to 192.168.3.1.

Answer: A
Explanation:
Indeed, when you change the administrative distance of the discard route to 255, this prevents the route to be installed in the routing table but this does not prevent the route to be advertised to the other peer!

From the Cisco website:
"You can configure a summary aggregate address for a specified interface. If there are any more specific routes in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes"

And don't forget YOU HAVE a directly connected route 192.168.1/24 in R1 and .4/24 in R2.
Then your summary is advertised, and the packet loop between R1 and R2.
If you do the test with GNS3, you will see via "show ip eigrp topology" that you FD of your local subnet 192.168.1.0/24 in R1 become the metric of the summary route "192.168.0.0/22 advertise to R2 ! And viceversa. Numerically, an administrative distance is an integer from 0 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored.
http://www.cisco.com/en/US/docs/ios/12_3/iproute/command/reference/ip2_c1g.html
A route with an AD of 255 would never be installed in the RIB. It is much better for the router to forward a packet to null0 following the default route than for the router to find out it has no route to destination, drop the packet and probably to send an ICMP unreachable message back to the source of that packet.

https://supportforums.cisco.com/thread/192416
You should not use the ip summary-address eigrp summarization command to generate the default route (0.0.0.0) from an interface. This causes the creation of an EIGRP summary default route to the null 0 interface with an administrative distance of 5. The low administrative distance of this default route can cause this route to displace default routes learned from other neighbors from the routing table. If the default route learned from the neighbors is displaced by the summary default route, or if the summary route is the only default route present, all traffic destined for the default route will not leave the router, instead, this traffic will be sent to the null 0 interface where it is dropped.

http://www.cisco.com/en/US/docs/ios/12_0/np1/configuration/guide/1ceigrp.html#wp4937

Default Administrative Distances

Route Source	Default Distance
Connected interface	0
Static route	1
Enhanced Interior Gateway Routing Protocol (EIGRP) summary route	5
External Border Gateway Protocol (eBGP)	20
Internal EIGRP	90
Open Shortest Path First (OSPF)	110
Intermediate System-to-Intermediate System (IS-IS)	115
Routing Information Protocol (RIP)	120
EIGRP external route	170
Internal BGP	200
Unknown	255

ActualTests

Here are the configurations and CLI output to prove the answer is A. This is done on IOS 12.4T train of code.

R1
--
interface Loopback0
ip address 192.168.1.1 255.255.255.0
!
interface FastEthernet0/0
ip address 10.1.2.1 255.255.255.0
ip summary-address eigrp 1 192.168.0.0 255.255.252.0 255
!
router eigrp 1
network 10.0.0.0
network 192.168.1.0
no auto-summary
R2
--
interface Loopback0
ip address 192.168.4.2 255.255.255.0
!
interface FastEthernet0/0
ip address 10.1.2.2 255.255.255.0
ip summary-address eigrp 1 192.168.0.0 255.255.248.0 255
!
router eigrp 1
network 10.0.0.0
network 192.168.4.0
no auto-summary
Show Output

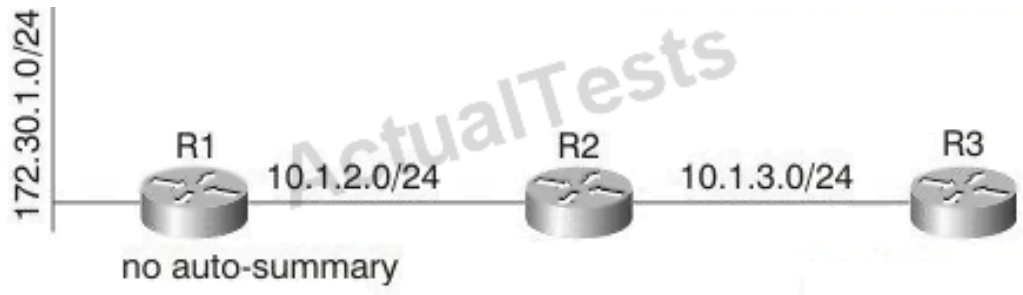
Look at R1 s routing table...notice it does NOT have the summary route to null0 for 192.168.0.0/22 because we
have disabled it. It does have a summary route to 192.168.0.0/21 from R2 it is receiving...
R1#show ip route | b Gateway
Gateway of last resort is not set
10.0.0.0/24 is subnetted, 1 subnets
C 10.1.2.0 is directly connected, FastEthernet0/0

C 192.168.1.0/24 is directly connected, Loopback0
D 192.168.0.0/21 [90/156160] via 10.1.2.2, 00:03:45, FastEthernet0/0
Now look at R2 s routing table. Notice it does NOT have a summary route to null0 for the 192.168.0.0/21 because we disabled it. It does however have a route for 192.168.0.0/22 received from R1.
R2#sh ip route | b Gateway
Gateway of last resort is not set
C 192.168.4.0/24 is directly connected, Loopback0
10.0.0.0/24 is subnetted, 1 subnets
C 10.1.2.0 is directly connected, FastEthernet0/0
D 192.168.0.0/22 [90/156160] via 10.1.2.1, 00:05:36, FastEthernet0/0
Now, what will happen? As I said -- When R1 pings 192.168.3.1 it will look in its routing table, and it will see the
EIGRP route 192.168.0.0/21 from R2 so it will send the packet to R2 since 192.168.3.1 is part of that summary.
R2 will get the packet and look in its routing table. It will find an EIGRP route for 192.168.0.0/22 from R1 and route the packet back to R1 since 192.168.3.1 is within that summary range.

This continues in a loop until the IP TTL expires.
Traceroute from R1. Notice it goes to R2 then right back to R1...Hope this helps! ----------
R1#trace 192.168.3.1
Type escape sequence to abort.
Tracing the route to 192.168.3.1
1 10.1.2.2 0 msec 0 msec 4 msec
2 10.1.2.1 0 msec 0 msec 4 msec
3 * * *

QUESTION NO: 122

Refer to the exhibit. In this network, R1 is configured not to perform auto summarization within EIGRP. What routes will R3 learn from R2 through EIGRP?



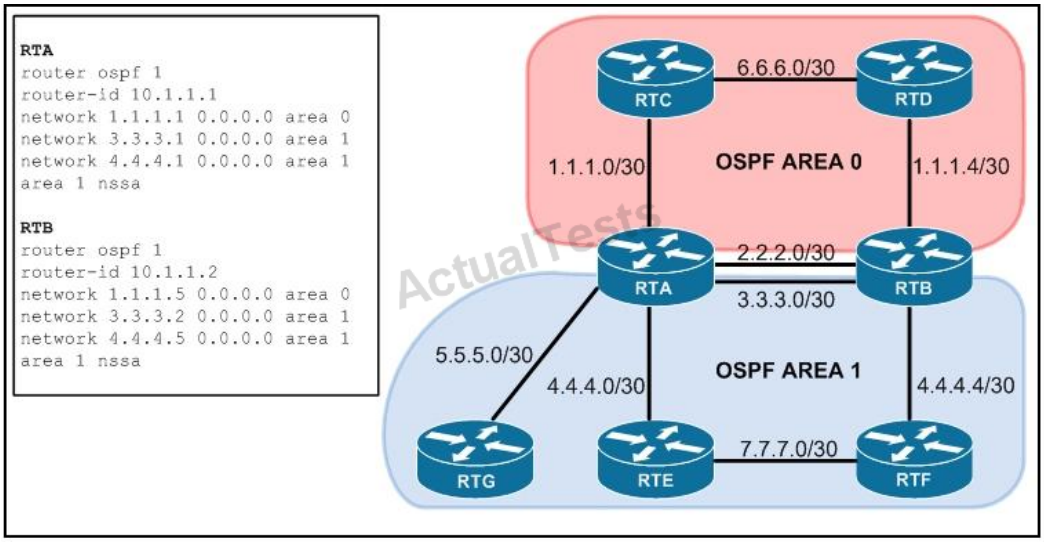
A. 172.30.1.0/24 and 10.1.2.0/24; EIGRP only performs autosummarization at the edqe between two major networks.
B. 172.30.0.0/16 and 10.1.2.0/24; R2 will perform autosummarization, although R1 will not.
C. Since R2 is configured without autosummarization, it will not propagate the 172.30.1.0/24 route.
D. 172.30.0.0/8 and 10.0.0.0/8

Answer: B

Explanation: Reference http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094cb7.shtml#summarizati on

QUESTION NO: 123

Which router is doing LSA type 7 to type 5 translation, and why?



A. Both RTA and RTB will do LSA7/5 translation, since they are both NSSA ABRs.
B. RTA will do LSA7/5 translation, since it has the lower router ID.
C. RTB will do LSA7/5 translation, since it has the higher router ID.
D. The redistributing router will do LSA7/5 translation, since it is within the NSSA.

Answer: C

Reference: http://www.hh.se/download/18.4cf286ee134f03ddb7b800015/1326882212358/Chapter3_VT2012.pdf (see slide 36)

QUESTION NO: 124

Refer to the exhibit.

```
R1#sh ip bgp 0.0.0.0
BGP routing table entry for 0.0.0.0/0, version 2
Paths: (4 available, best #?, table Default-IP-Routing-Table)
Multipath: eBGP
Flag: 0x820
 Advertised to update-groups:
    2
 38264
   10.27.203.9 from 10.27.203.9 (10.27.191.13)
     Origin IGP, metric 0, localpref 150, valid, external, multipath
 38264
   10.27.203.9 from 10.27.203.9 (10.27.191.13)
     Origin IGP, metric 0, localpref 100, valid, external
 9541
   10.124.26.250 from 10.124.26.250 (10.163.120.33)
     Origin IGP, localpref 150, valid, external, multipath
 9541
   10.124.26.250 from 10.124.26.250 (10.163.120.33)
     Origin IGP, localpref 100, valid, external
```

ActualTests

Which path is selected as best path?

A. path 1, because it has a metric of 0 and a localpref of 150
B. paths 1 and 2, because they are multipath from the same AS with the same metric of 0
C. paths 1 and 3, because they are multipath with the same localpref of 150
D. the show command output is wrong, because multipath cannot happen for different eBGP peers

Answer: C

Explanation:

Since both has same localpref 150, path 1 and 3 will be selected as best path.
Reference: http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_bgp/configuration/15-sy/irg-multi-load.html

QUESTION NO: 125

Refer to the exhibit.

```
R1#show ip route 10.0.0.0
O        10.0.0.2/32 [110/101] via 10.1.2.2, 00:02:05, Ethernet0/0
i L1     10.0.0.3/32 [115/20] via 10.1.3.3, Ethernet0/1

R1#show bgp 10.168.0.1/32
BGP routing table entry for 10.168.0.1/32, version 3
Paths: (2 available, best #?, table default)
Flag: 0x820
 Not advertised to any peer
 4
   10.0.0.2 (metric 101) from 10.0.0.2 (10.0.0.2)
     Origin IGP, metric 0, localpref 100, valid, internal
     mpls labels in/out nolabel/2004
 4
   10.0.0.3 (metric 20) from 10.0.0.3 (10.0.0.3)
     Origin IGP, metric 0, localpref 100, valid, internal
```

The next hops are learned via OSPF and IS-IS. Which path is selected as the best path for 10.168.0.1?

A. path 1, because it is an MPLS labeled path
B. path 1, because the next hop is learned via OSPF with an AD of 110, compared to 115 for IS-IS
C. path 2, because it has the highest router ID
D. path 2, because it has the lowest IGP metric

Answer: D

Explanation:

Path 1 has a metric of 101 and path 2 has a metric of 20; therefore, path 2 will be the path that is used as the best path to reach 10.168.0.1

QUESTION NO: 126

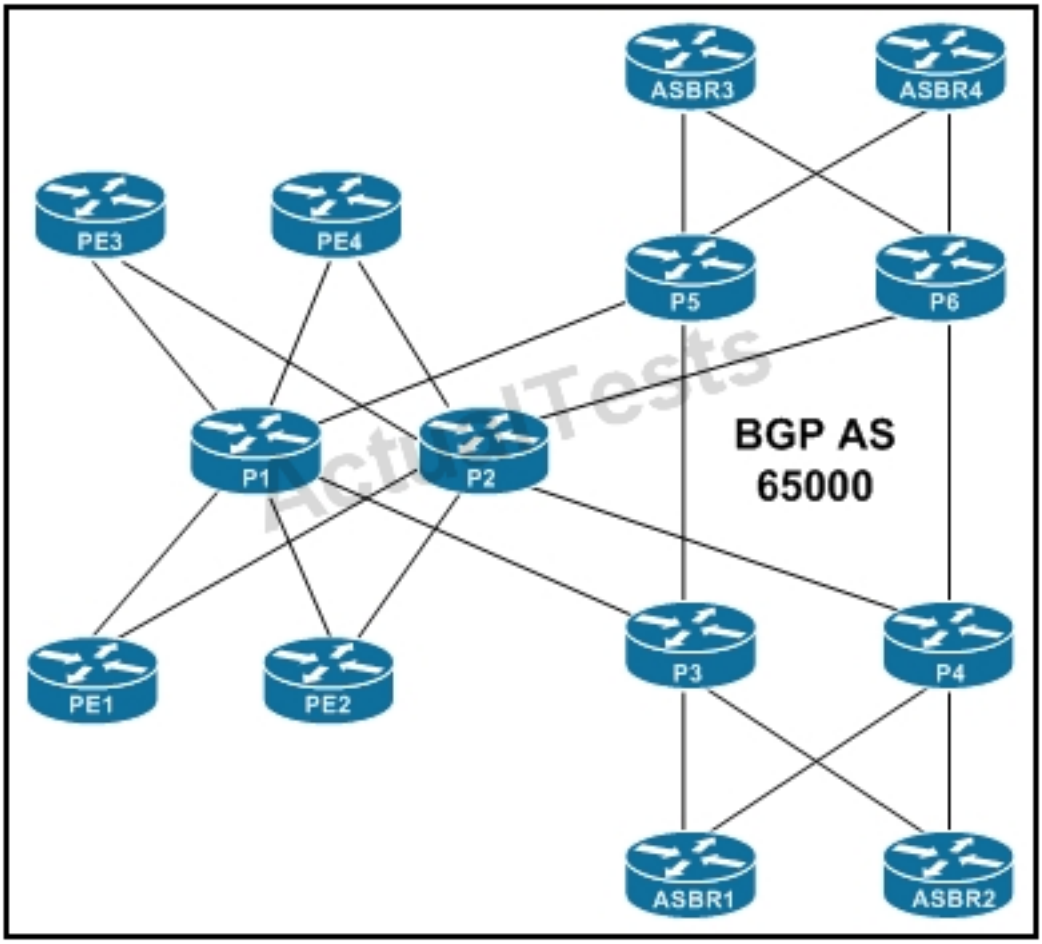
Which statement is correct about an OSPF sham-link?

A. A sham-link is a logical link between PE routers that provides an OSPF adjacency between the PE routers and forwards customer traffic across it.
B. A sham-link is a logical link between PE routers that provides an OSPF adjacency between the PE routers and carries OSPF LSAs.
C. A sham-link is a logical link between PE routers that carries OSPF LSAs and forwards customer traffic across it.
D. A sham-link is a point-to-multipoint link that provides an OSPF adjacency between the PE routers and carries OSPF LSAs.
E. A sham-link is a point-to-multipoint link that provides an OSPF adjacency between the PE routers and forwards customer traffic across it.

Answer: B
Reference: http://www.juniper.net/techpubs/en_US/junose11.0/information-products/topic-collections/swconfig-bgp-mpls/jd0e45628.html (search for sham links)

QUESTION NO: 127

Refer to the exhibit.



AS65000 has core network P1-P6. The eBGP peers to another AS are through ASBR1-2 and ASBR3-4. All business and residential customer POPs currently are connected to the P1-P2 core.

Which pair of routers is the best positioned to become a BGP route reflector?

A. P1 and P2, as they have the same view of the exit point from the AS
B. ASBR1 and ASBR2, as they have eBGP peers to another AS
C. ASBR3 and ASBR4, as they have more eBGP peers
D. a new pair of routers (RR1 and RR2) that are dedicated as a BGP route reflector that is connected to P1 and P3

Answer: A
Reference: http://tools.ietf.org/html/draft-ietf-idr-bgp-optimal-route-reflection-05 (see introduction, 6th para)

QUESTION NO: 128

Which three factors have the biggest influence on OSPF scalability? (Choose three.)

A. Flooding paths and redundancy
B. Amount of routing information in the OSPF area or routing domain
C. Number of routers with Cisco Express Forwarding enabled
D. Number of neighbor adjacencies
E. Other routing protocols in use
F. OSPF timer reconstruction negotiation
G. Redistribution with BGP neighbors
H. Redistribution with other IGP routing protocols, such as RIP or EIGRP

Answer: A,B,D
Reference: http://www.ciscopress.com/articles/article.asp?p=1763921&seqNum=6 (see factors influencing OSPF scalability)

QUESTION NO: 129

Based on the exhibit presented.



What will be the objective of this route map when applied to traffic passing through a router?

A. Take any packet sourced from any address in the 10.2.0.0/16 network or destined to 10.1.14.25 and set the next hop to 10.1.1.1

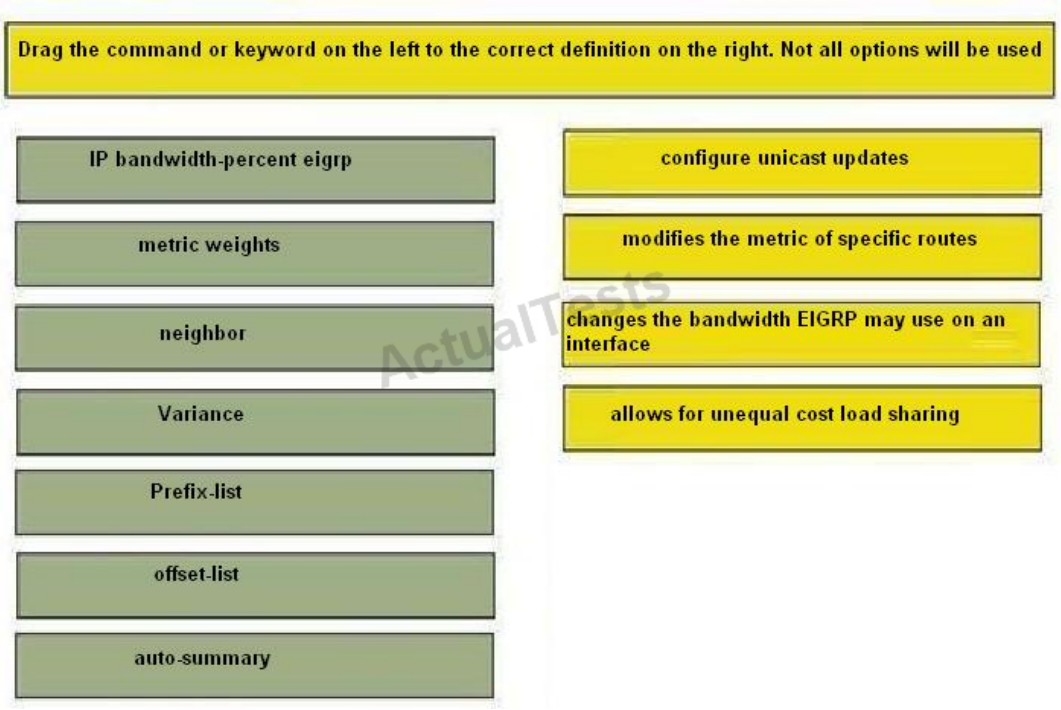
B. Take any packet sourced from any address in the 10.2.0.0/16 network and destined to 10.1.14.25 and set the next hop to 10.1.1.1
C. Nothing; extended access lists are not allowed in route maps used for policy-based routing
D. Drop any packet sourced from 10.2.0.0/16

Answer: A

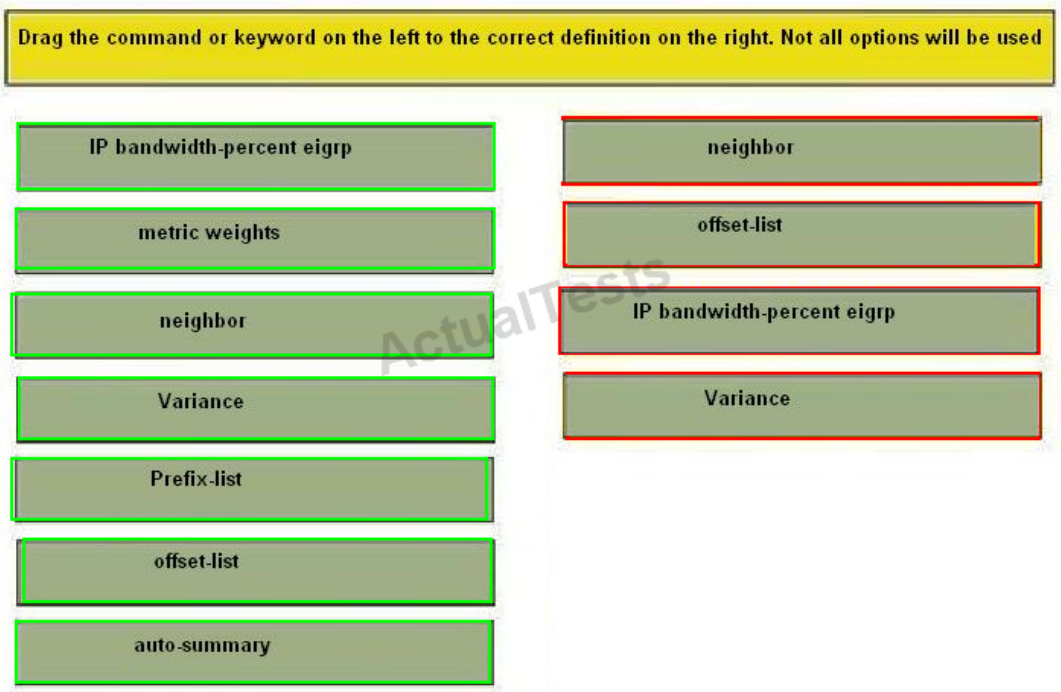
Explanation:

In this configuration example, any traffic matching access list 100 will have their next hop set to 10.1.1.1 overriding the normal behavior of the routing table. Access list 100 has two entries, so any traffic matching either will be policy routed.

QUESTION NO: 130 DRAG DROP



Answer:



Explanation:

ip bandwidth-percent eigrp – changes the bandwidth EIGRP may use on an interface
metric weights –
neighbor – configures unicast updates
variance – allows for unequal cost load sharing
prefix-list –
offset-list – modifies the metric of specific routes
auto-summary –

QUESTION NO: 131

Which of these statements best describes how neighbor adjacencies are formed in a multi-access OSPF network?

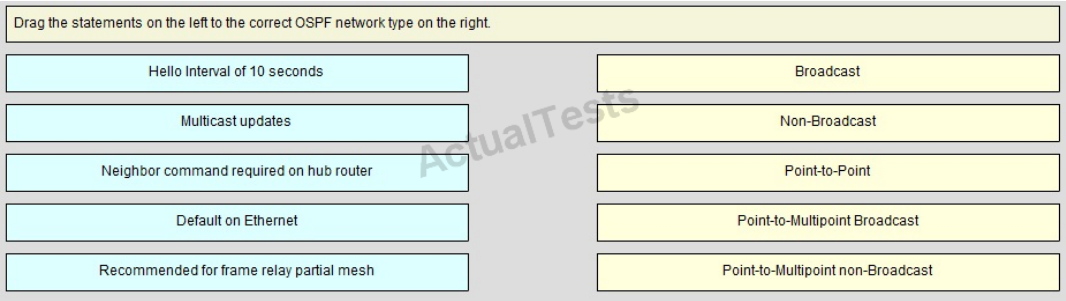
A. The router with the highest priority will become the DR
B. Only those routers with the Cisco default priority of 0 are eligible to become the DR or BDR.
C. The router with the highest loop-back address will become the DR if two or more routers have the same priority.
D. The router with the lowest Router ID will become the DR and the router with the next lowest Router ID will become the BDR.
E. Election of the DR and BDR begins only after a router that wants to become either the DR or BDR enters the ExStart state.

Answer: A

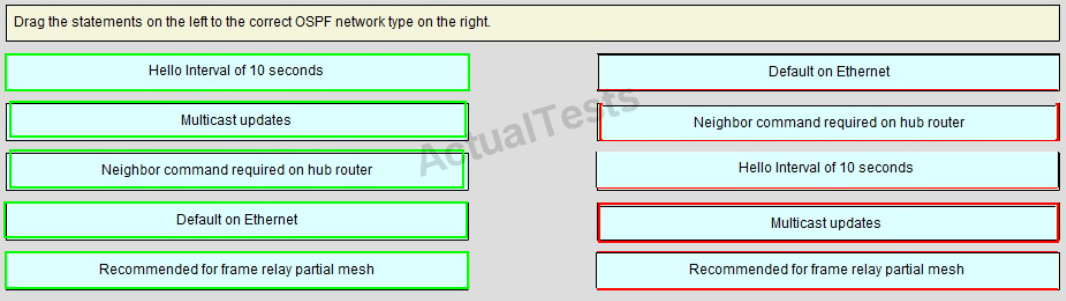
Explanation:

The router with the highest priority is elected the DR on a multiaccess network. A router with a priority of 0 is ineligible to become a DR or BDR. In the event of a tie in priority, the router with the highest router ID is elected the DR. If no router ID has been manually configured on a router, the router uses its numerically highest loopback address as its router ID. If no loopback interfaces have been configured, the router uses its numerically highest IP address of any physical interface. Cisco use a default priority of 1 (higher is better). On a tie, the highest router id wins. If there is still a tie, the highest loopback address wins. Still a tie, the highest active interface
IP wins.

QUESTION NO: 132 DRAG DROP



Answer:

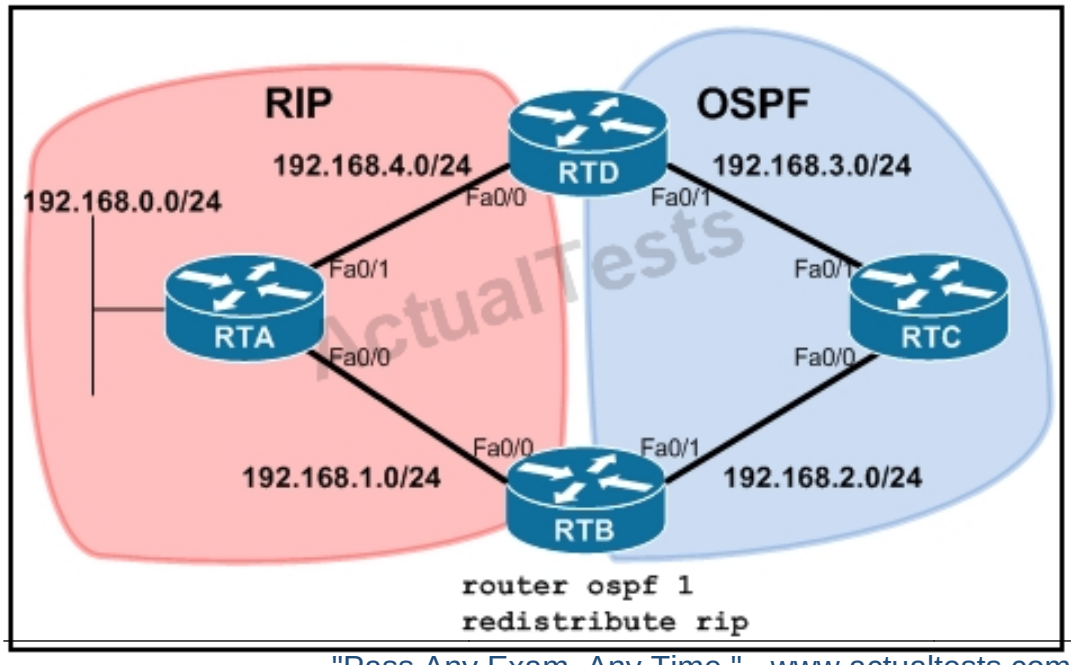


Explanation:

Default on Ethernet
Neighbor command required on hub router
Hello Interval of 10 seconds
Multicast updates
Recommended for frame relay partial mesh

QUESTION NO: 133

Refer to the exhibit.



Router RTB is performing one-way redistribution from RIP to OSPF. Which outgoing interface will router RTD choose for packets to the 192.168.0.0/24 network, and why?

A. Fa0/1, because OSPF is a link-state routing protocol
B. Fa0/0, because RIP is a distance vector protocol
C. Fa0/0, because RIP has a higher administrative distance
D. Fa0/0, because OSPF has a lower administrative distance
E. Fa0/1, because OSPF has a lower administrative distance
F. Fa0/1, because RIP has a lower administrative distance

Answer: E

Explanation:

RTD has information in both topologies RIP and OSPF, and it has to elect which one is better, and the tie-breaker is AD (adminitrative Distance)

QUESTION NO: 134

Refer to the exhibit.

```
PE1#show ip ospf 100 neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
11.100.1.9        0   FULL/  -        00:00:37    11.1.1.4        Serial4/0
11.100.1.2        1   FULL/DR         00:00:38    11.1.3.2        Ethernet0/0
```

ActualTests

Which statement is correct?

A. Router PE1 is not the DR on interface Ethernet0/0 and interface Serial4/0 is an OSPF point-to-point link.
B. Router PE1 is the DR on interface Ethernet0/0 and interface Serial4/0 is an OSPF point-to-point link.
C. Router PE1 is the DR on interface Ethernet0/0 and the DR election has yet to complete on the interface Serial 4/0.
D. Interface Ethernet0/0 does not participate in a DR/BDR election.

Answer: A

Reference: http://vcielka.darksun.sk/~jnx/download/network/19053_MPLS%20Fundamentals.pdf

QUESTION NO: 135

By default, EIGRP will use which percentage of bandwidth on an interface?

A. 10%
B. 25%
C. 50%
D. 75%
E. 80%
F. 100%

Answer: C
Explanation:
The enhanced implementation uses the configured interface bandwidth in order to determine how much EIGRP data to transmit in a given amount of time. By default, EIGRP will limit itself to using no more than 50% of the interface bandwidth. The primary benefit of controlling EIGRP's bandwidth usage is to avoid losing EIGRP packets, which could occur when EIGRP generates data faster than the interface line can absorb it. This is of particular benefit on Frame Relay networks, where the access interface bandwidth and the PVC capacity may be very different. A secondary benefit is to allow the network administrator to ensure that some bandwidth remains for passing user data, even when EIGRP is very busy.

QUESTION NO: 136

Refer to the exhibit.

```
C1#show ip ospf border-routers

OSPF Process 100 internal Routing Table

Codes: i - Intra-area route, I - Inter-area route

i 11.200.1.1 [128] via 11.1.4.4, Serial2/0, ABR/ASBR, Area 0, SPF 14
i 11.100.1.2 [128] via 11.1.5.10, Serial3/0, ABR/ASBR, Area 0, SPF 14
```

ActualTests

Which statement is correct?

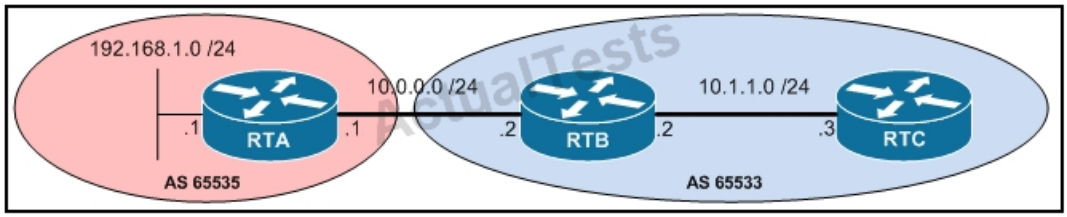
A. The priority of both OSPF border routers is 128.
B. Both routers are border routers for areas 128 and 0.
C. The OSPF cost to both border routers is 128.
D. The age of both border routers is 128.

Answer: C

Explanation:
As per the exhibit, both OSPF border routers has a cost of 128.

QUESTION NO: 137

Refer to the exhibit.



Considering that RTB does not use next-hop-self, what will be the next hop for the route 192.168.1.0/24 on RTC?

A. 192.168.1.1
B. 10.0.0.1
C. 10.0.0.2
D. 10.1.1.2
E. 10.1.1.3
F. the BGP router ID of RTA

Answer: B
Explanation: The next hop will be 10.0.0.1 because RTB is not using next-hop-self.
Reference: http://steve.blogs.exetel.com.au/index.php?/archives/30-next-hop-self.html

QUESTION NO: 138

Which information will the Cisco IOS command show ip ospf rib display?

A. only the local OSPF routes
B. only the OSPF routes installed in the routing table
C. only the remotely learned OSPF routes
D. all the OSPF routes from the OSPF database that are eligible to be put in the routing table

Answer: D

Explanation:

Show ip ospf rib

To display information for the OSPF local Routing Information Base (RIB) or locally redistributed routes, use the show ip ospf rib command in privileged EXEC mode.

Show ip ospf process-id rib [redistribution] [network-prefix] [network-mask] [detail]

Reference

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_osp3.html

QUESTION NO: 139

Refer to the exhibit.

```
CE1#show ip eigrp neighbors detail
IP-EIGRP neighbors for process 65000
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   11.1.3.2                Et1/0            12 00:00:44    1  5000  2  63
   Version 12.4/1.2, Retrans: 9, Retries: 9, Waiting for Init Ack
    UPDATE seq 68 ser 0-0 Sent 44148 Init Sequenced
    UPDATE seq 69 ser 2-80 Sequenced
1   11.1.1.2                Et0/0            11 00:09:16   20   200  0  31
   Version 12.4/1.2, Retrans: 0, Retries: 0
```

Based on the output shown, what will happen?

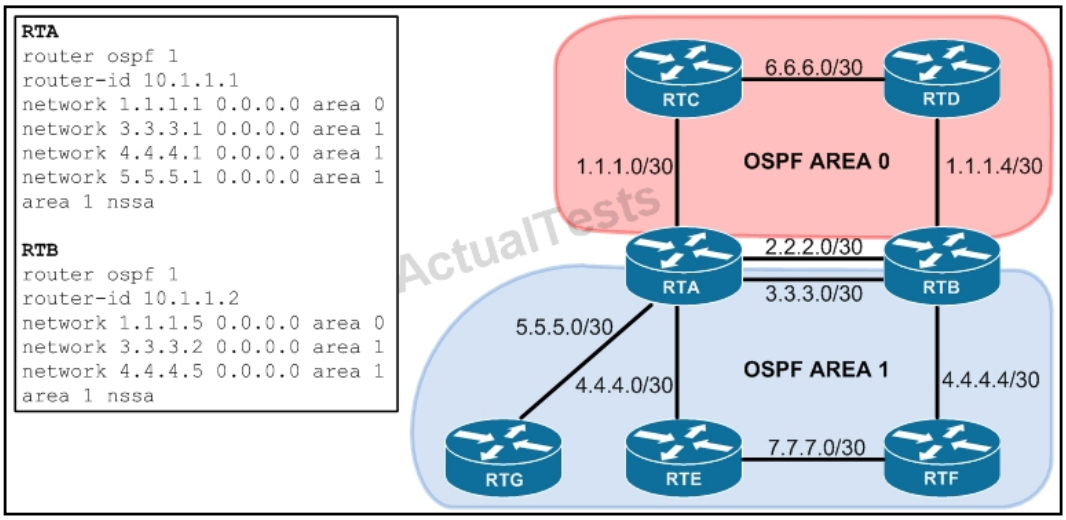
A. EIGRP keeps on retransmitting the reliable EIGRP packets forever.
B. EIGRP will retransmit the reliable EIGRP packets up to 16 times and then delete the related prefixes.
C. EIGRP will retransmit the reliable EIGRP packets up to 16 times and then reset the EIGRP neighbor 11.1.3.2.
D. The EIGRP neighbor 11.1.3.2 goes down when the hold time reaches 0, which is 12 seconds from now.

Answer: C

Explanation:

QUESTION NO: 140

Refer to the exhibit.



RTE is redistributing some static routes into OSPF as E2. Which LSA will be seen in RTA and RTB for those routes?

A. LSA1
B. LSA5
C. LSA7
D. RTA/RTB will not see any LSA for those routes, since this is NSSA and no external redistribution is allowed

Answer: C
Reference:
http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094e9e.shtml#t32

QUESTION NO: 141

Refer to the exhibit.

```
R1#show ip route 10.0.0.0
O        10.0.0.2/32 [110/101] via 10.1.2.2, 00:02:05, Ethernet0/0
i L1     10.0.0.3/32 [115/20] via 10.1.3.3, Ethernet0/1

R1#show bgp 10.168.0.1/32
BGP routing table entry for 10.168.0.1/32, version 3
Paths: (2 available, best #?, table default)
Flag: 0x820
 Not advertised to any peer
 4
   10.0.0.2 (metric 101) from 10.0.0.2 (10.0.0.2)
     Origin IGP, metric 0, localpref 100, valid, internal
     mpls labels in/out nolabel/2004
 4
   10.0.0.3 (metric 20) from 10.0.0.3 (10.0.0.3)
     Origin IGP, metric 0, localpref 100, valid, internal
```

ActualTests

The next hops are learned via OSPF and IS-IS. Which path is selected as the best path for 10.168.0.1?

A. the path via 10.0.0.2, because it is an MPLS labeled path
B. the path via 10.0.0.2, because the next hop is learned via OSPF with an AD of 110, compared to 115 for IS-IS
C. the path via 10.0.0.3, because it has the highest router ID
D. the path via 10.0.0.3, because it has the lowest IGP metric

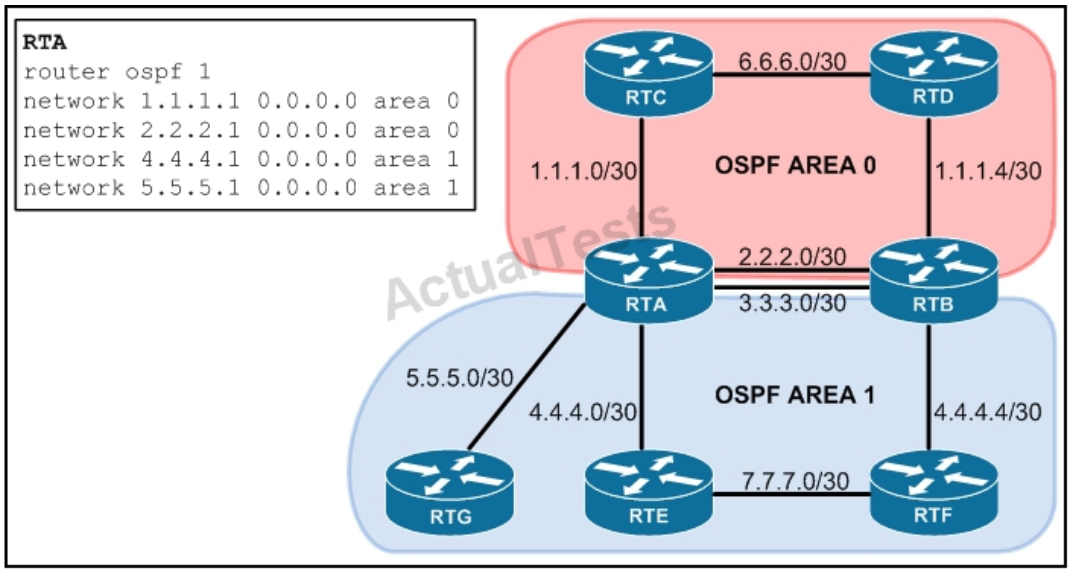
Answer: D

Explanation:

10.0.0.3 has the lowest IGP metric therefore, it will be selected as the best path for 10.168.0.1.

QUESTION NO: 142

Refer to the exhibit.



What will happen if the link between RTA and RTE fails?

A. Area 1 will be split into two between RTB-RTF-RTE and RTG-RTA, so these separate areas will not be able to communicate with each other.
B. Area 1 will be split into two between RTB-RTF-RTE and RTG-RTA, and to avoid communication issues a virtual link between RTA and RTB needs to be configured.
C. Area 1 will be split into two between RTB-RTF-RTE and RTG-RTA, and duplicate LSAs from Area 1 will flood OSPF Area 0.
D. Area 1 will be split into two between RTB-RTF-RTE and RTG-RTA, and OSPF will work just fine.

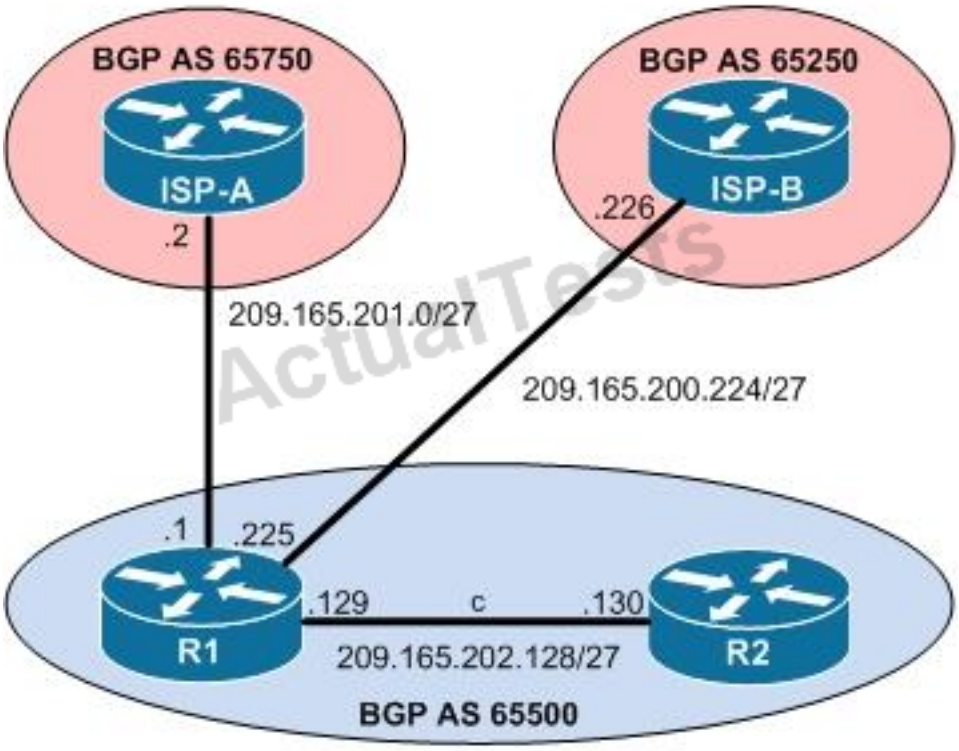
Answer: D

Explanation: Explanation

As RTE will no longer have a direct connection to RTA Area 1 will be split into two groups. Routers RTA & RTG in one group and Routers RTB, RTE, & RTF in the other group. For router RTE to access RTA it will now need to traverse Area 0.

QUESTION NO: 143

Refer to the exhibit.




```
R1#show ip bgp
BGP table version is 47, local router ID is 209.165.202.129
Status codes: s suppressed, * valid, > best, i - internal,r RIB-failure
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i10.0.0.0         209.165.202.158     156160    100      0 ?
*>i172.17.0.0       209.165.202.158     156160    100      0 ?
*>i172.18.0.0       209.165.202.158     156160    100      0 ?
*>i192.168.60.0     209.165.202.158     156160    100      0 ?
*>i192.168.70.0     209.165.202.158     156160    100      0 ?
r>i209.165.202.128/27
                    209.165.202.130          0    100      0 ?

R1#show ip bgp summary
BGP router identifier 209.165.202.129, local AS number 65500

Neighbor        V    AS MsgRcvd MsgSent   InQ OutQ Up/Down  State/PfxRcd
209.165.200.226 4 65250    1356    1376     0    0 21:08:07        0
209.165.201.2   4 64750    1374    1397     0    0 21:08:09        0
209.165.202.130 4 65500    1373    1354     0    0 21:08:09       14
```

All iBGP routes should have the iBGP peer as the next hop address. Why is this not the case for BGP routes learned between R1 and R2?

A. R2 is missing the next-hop-self option under the neighbor command for R1
B. ISP-A is missing the next-hop-self option under the neighbor command for R1
C. ISP-B is missing the next-hop-self option under the neighbor command for R1

D. R2, ISP-A, and ISP-B are missing the next-hop-self option under the neighbor command for R1

Answer: A

Explanation:

When you have an e-bgp peer, the next hop is always changed to your peering IP address (next-hop-self, or setting the IP to its own). This is because ebgp peers don't have any need to know about all the details inside another ASN. Next hops are measured by an AS Path.

In i-bgp peers, information is not changed. The theory is that everyone know about everything inside the same AS! Now, this is nice as long as that's true!

What happens when the link between two e-bgp peers (obviously known to the edge router because it's connected) is NOT exchanged via any IGP? Suddenly other internal routers learn i-bgp routes that appear to be unreachable. They show an unknown IP as the next hop.

Ascii digram:

(AS 100) -- 10.100.1.1----------10.100.1.2 (AS200-R1)10.200.1.1-------10.200.1.2(AS200-R2)

AS 200 routers are i-bgp peers (who could be multiple hops away from each other but this is a short example!). AS100-200 exchange routes over 10.100.1.0 network.

R1 shows all routes with 10.100.1.1 as the next hop. This is simple and per spec, and is reachable because it's connected.

UNLESS R2 is aware of the 10.100.1.0 network, it will learn all of AS100's routes and will view them all as unreachable.

If R1 sets "next-hop-self" when pointing to R2s neighbor statement, all of the routes will be changed to reflect 10.200.1.1 as the next hop which is reachable.

QUESTION NO: 144

Refer to the exhibit.

```
r13#show ip bgp 0.0.0.0/0
BGP routing table entry for 0.0.0.0/0, version 6703
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  65000 65335 65535, (received & used)
    10.31.222.5 from 10.31.222.5 (192.168.1.1)
      Origin IGP, localpref 100, valid, external, best
      Community: 22345643
```

What does "(192.168.1.1)" indicate in the output?

A. the BGP router ID of the eBGP peer
B. the advertising route reflector cluster-id
C. the router originating this prefix
D. the BGP router ID of the iBGP peer
E. our local router ID
F. the BGP next hop for this route

Answer: A
Reference: http://www.cisco.com/en/US/docs/ios/12_2/iproute/command/reference/1rfbgp1.html

QUESTION NO: 145

Refer to the exhibit.

```
r14#show ip bgp
BGP table version is 874, local router ID is 172.16.5.5
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
*>i10.13.0.0/16     192.168.1.1              0    100      0 i
*> 172.16.5.5/32    0.0.0.0                  0         32768 ?
   172.16.5.5/32    192.168.1.1              0    100      0 i
*> 172.16.15.0/24   192.168.2.1              0             0 62000 i
*> 192.168.1.0      0.0.0.0                  0         32768 ?
*  192.168.2.0      192.168.2.1              0             0 62000 ?
*>                  0.0.0.0                  0         32768 ?
```

By default, why will BGP choose a locally originated route over any iBGP or eBGP learned route?

A. The IGP metric to the next hop is always shorter.
B. Locally originated routes have a better AD than eBGP or iBGP routes.
C. All locally originated routes have a better origin code.

D. Locally originated routes have a higher metric.
E. Locally originated routes have a weight of 32, 768.

Answer: E

Explanation:

If you have two routers at the same location, and need to advertise the same internal networks, quite often you will put the same network statements on both routers.
In this case, you would likely have external BGP peers on the routers, but you would also have an internal BGP connection between the two routers.
This would result in each router locally injecting the network via the network statement, but you would also receive an iBGP advertisement from your neighbor within your same AS.

QUESTION NO: 146

Which two EIGRP options will limit a query domain? (Choose two.)

A. configuring EIGRP stubs
B. configuring a second EIGRP AS and redistributing
C. configuring summary addresses
D. configuring an offset-list
E. configuring a prefix-list
F. configuring unicast neighbors

Answer: A,C

Explanation:

Summarization in EIGRP also solves a problem unique to EIGRP: query scoping. EIGRP queries do not stop until they come to a dead-end; because of this, they can loop. Route convergence cannot proceed until all replies have been received. Summarization limits queries by stating: "only these routes are found past this point." Summarizing at logical points in the network is important to prevent stuck in active (SIA)

QUESTION NO: 147

Which three options are valid ways to filter routes from a BGP peer? (Choose three.)

A. filter-list
B. inbound ACL

C. route-map
D. outbound ACL
E. prefix-list
F. neighbor 10.0.0.1 deny-map my-map command

Answer: A,C,E

Explanation:

Configure BGP Route Filtering by Neighbor

You can filter BGP advertisements in two ways:

Use AS-path filters, as with the ip as-path access-list global configuration command and the neighbor filterlist command

Use access or prefix lists, as with the neighbor distribute-list command.

Filtering using prefix lists is described in "Configuring BGP Filtering Using Prefix Lists".

If you want to restrict the routing information that the Cisco IOS software learns or advertises, you can filter BGP routing updates to and from particular neighbors. To do this, you can either define an access list or a prefix list and apply it to the updates.

Note Distribute-list filters are applied to network numbers and not autonomous system paths.

To filter BGP routing updates, use the following command in router configuration mode:

Command	Purpose
neighbor {ip-address | peer-group-name} distribute-list {access-list-number | name} {in | out}	Filter BGP routing updates to/from neighbors as specified in an access list. **Note** The **neighbor prefix-list** command can be used as an alternative to the **neighbor distribute-list** command, but you cannot use both commands in configuring the same BGP peer.

ActualTests

Configuring BGP Filtering Using Prefix Lists

Prefix lists can be used as an alternative to access lists in many BGP route filtering commands.

"How the System Filters Traffic by Prefix List" describes the way prefix list filtering works.

The advantages of using prefix lists are:

Significant performance improvement in loading and route lookup of large lists

Support for incremental updates Filtering using extended access lists does not support incremental updates.

More user-friendly command-line interface

The command-line interface for using access lists to filter BGP updates is difficult to understand and use, since it uses the packet filtering format.

Greater flexibility

Before using a prefix list in a command, you must set up a prefix list, and you may want to assign sequence numbers to the entries in the prefix list.

Enable BGP Routing

To enable BGP routing, establish a BGP routing process by using the following commands beginning in global configuration mode:

Step	Command	Purpose
1	router bgp autonomous-system	Enable a BGP routing process, which places you in router configuration mode.
2	network network-number [mask network-mask] [route-map route-map-name]	Flag a network as local to this autonomous system and enter it to the BGP table.

Note For exterior protocols, a reference to an IP network from the network router configuration command controls only which networks are advertised. This is in contrast to Interior Gateway Protocols (IGP), such as IGRP, which also use the network command to determine where to send updates.

Note
The network command is used to inject IGP routes into the BGP table. The network-mask portion of the command allows supernetting and subnetting. The router's resources, such as configured NVRAM or RAM, determine the number of network commands you can use. Alternatively, you could use the redistribute command to achieve the same result.
Configure Advanced BGP Features The tasks in this section are for configuring advanced BGP features.

Use Route Maps to Modify Updates
You can use a route map on a per-neighbor basis to filter updates and modify various attributes. A route map can be applied to either inbound or outbound updates. Only the routes that pass the route map are sent or accepted in updates.

On both the inbound and the outbound updates, we support matching based on autonomous system path, community, and network numbers. Autonomous system path matching requires the as-path access-list command, community based matching requires the community-list command and network-based matching requires the ip access-list command. Use the following command in router configuration mode:

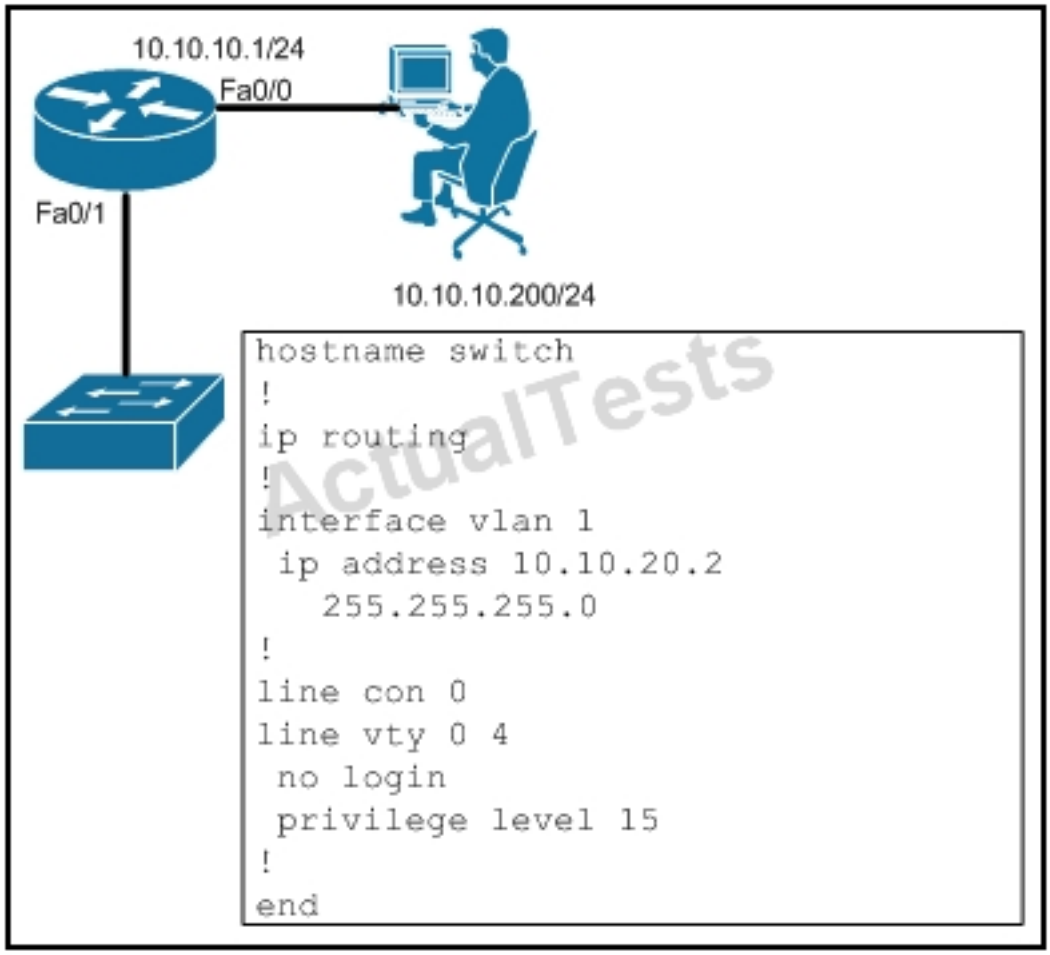
Command	Purpose
neighbor {ip-address | peer-group-name} route-map route-map-name {in | out}	Apply a route map to incoming or outgoing routes.

Reference

http://www.cisco.com/en/US/docs/ios/12_0/np1/configuration/guide/1cbgp.html#wp7487

QUESTION NO: 148

Refer to the exhibit.



A user with IP address 10.10.10.200 fails to telnet to a switch with IP address 10.10.20.2. What is the cause of the issue?

A. The switch is not configured with a default gateway.
B. The HTTP server is not enabled on the switch.
C. STP is blocking the connection from switch to router.
D. IP routing is enabled on the switch, but no route pointing back to the client is configured.
E. The switch is configured with an IP address from the wrong subnet.

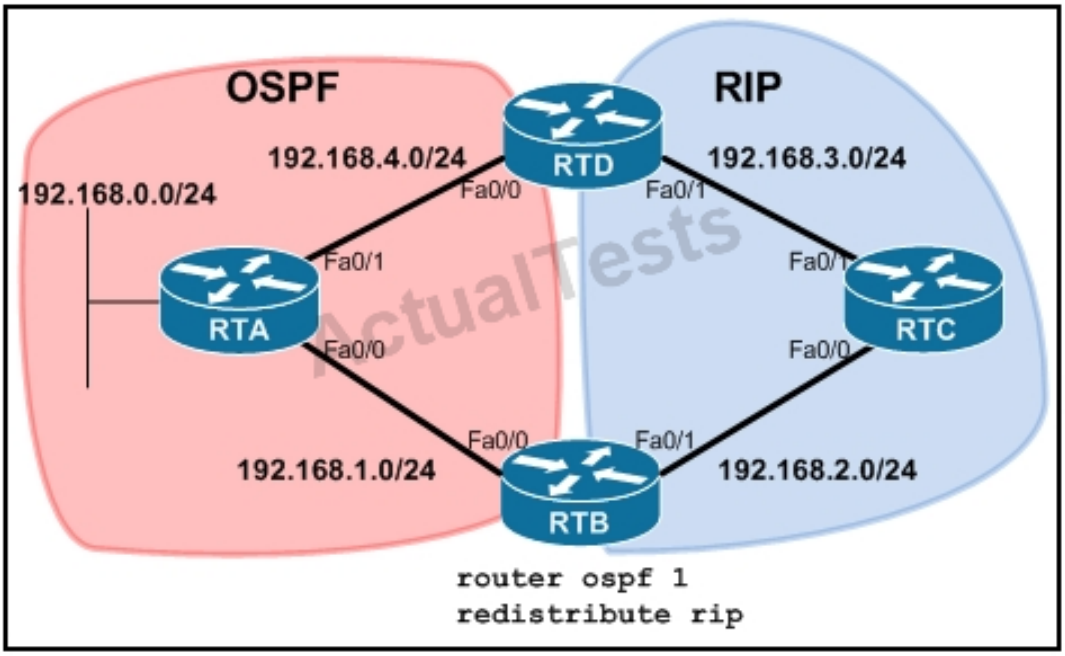
Answer: D

Explanation:

As you can see in the exhibit that IP routing is enabled on the switch but no route is pointing back to the client. The ip address of interface vlan 1 is given but no default route.

QUESTION NO: 149

Refer to the exhibit.



Router RTB is performing one-way redistribution from RIP to OSPF. Which outgoing interface will router RTD choose for packets to the 192.168.0.0/24 network, and why?

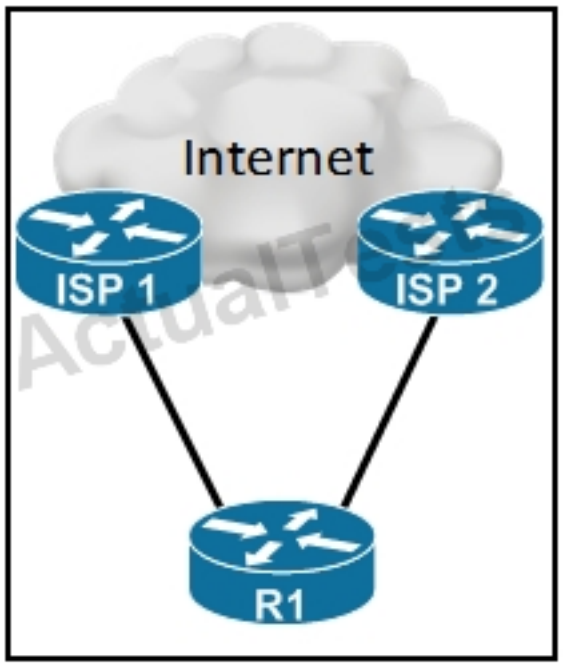
A. Fa0/1, because OSPF is a link-state routing protocol
B. Fa0/0, because RIP is a distance vector protocol
C. Fa0/0, because RIP has a higher administrative distance
D. Fa0/0, because OSPF has a lower administrative distance
E. Fa0/1, because OSPF has a lower administrative distance
F. Fa0/1, because RIP has a lower administrative distance

Answer: D
Explanation:
RTD has information in both topologies RIP and OSPF, and it has to elect which one is better, and the tie-breaker is AD (adminitrative Distance)
Rip pointing to the RIP domain (left i dont remember the interfaces please post the picture)
OSPF pointing to the OSPF domain (to the right)
RIP –> AD = 120

OSPF –> AD = 110
Lower is Better in AD so the correct answer depends on the exibit but i remeber that's the one to the right (to the OSPF domain and the longer path)

QUESTION NO: 150

Refer to the exhibit.



R1 has an EBGP session to ISP 1 and an EBGP session to ISP 2. R1 receives the same prefixes through both links.

Which configuration should be applied so that the link between R1 and ISP 2 will be preferred for outgoing traffic (R1 to ISP 2)?

A. Increase local preference on R1 for routes received from ISP2.
B. Decrease local preference on R1 for routes received from ISP2.
C. Increase MED on ISP 2 for routes received from R1.
D. Decrease MED on ISP 2 for routes received from R1.

Answer: A
Reference:
https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=9&ved=0CGkQFjAI&url=http%3A%2F%2Fwww.cs.cmu.edu%2F~srini%2F15-441%2FS05%2Flectures%2F11-

BGP.ppt&ei=YCbMUYKIFIfOhAeSg4Ew&usg=AFQjCNE2ioClS5UbKXhqIdnDQUt_9yNaUQ&sig2=THFuUNjA-RZEEFIsI97xHw

QUESTION NO: 151

Which individual metrics can be used to calculate the composite EIGRP metric?

A. total delay, minimum bandwidth, reliability, load, MTU
B. total delay, minimum bandwidth, reliability, load, MTU, hop count
C. total delay, minimum bandwidth, reliability, load, hop count
D. total delay, minimum bandwidth, reliability, load
E. total delay, minimum bandwidth

Answer: D
Reference: http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_eigrp/configuration/15-sy/ire-wid-met.html#GUID-736131DE-0B64-46F4-A19A-B0526D24F95B (See table 1)

QUESTION NO: 152

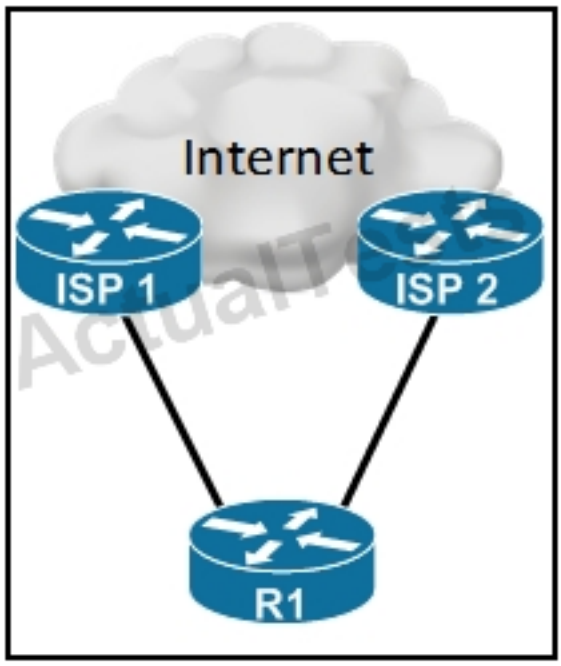
Which four are possible states in the BGP FSM? (Choose four.)

A. Idle
B. Established
C. Wait
D. Active
E. OpenSent
F. Nonconnected

Answer: A,B,D,E
Reference: http://www.freesoft.org/CIE/RFC/1771/52.htm (See the table)

QUESTION NO: 153

Refer to the exhibit.



R1 has two eBGP sessions to ISP1 and ISP2 (one to each ISP router), and R1 receives the same prefixes through both links.

Which configuration should be applied for the link between R1 and ISP2 to be preferred for incoming traffic (ISP2 to R1)?

A. increase local preference on R1 for sent routes
B. decrease local preference on R1 for sent routes
C. increase MED on ISP2 for sent routes
D. decrease MED on ISP2 for sent routes

Answer: D
Reference:
http://books.google.com.pk/books?id=V7MIJorqe_kC&pg=PA344&lpg=PA344&dq=decrease+MED+on+ISP2+for+sent+routes&source=bl&ots=pJvKb3dxua&sig=43gSTo8fhtMwu6DET75T3fJw9cU&hl=en&sa=X&ei=mFDLUZDzAcS47AbU3ICgCQ&ved=0CGgQ6AEwCA#v=onepage&q=decrease%20MED%20on%20ISP2%20for%20sent%20routes&f=false

QUESTION NO: 154

Based on the output provided in the exhibit, to which address or location will the router forward a packet sent to 192.168.32.1?

D 192.168.32.0/26 [90/25789217] via 10.1.1.1

R 192.168.32.0/24 [120/4] via 10.1.1.2

O 192.168.32.0/19 [110/229840] via 10.1.1.3

A. 10.1.1.1
B. 10.1.1.2
C. 10.1.1.3
D. The default gateway

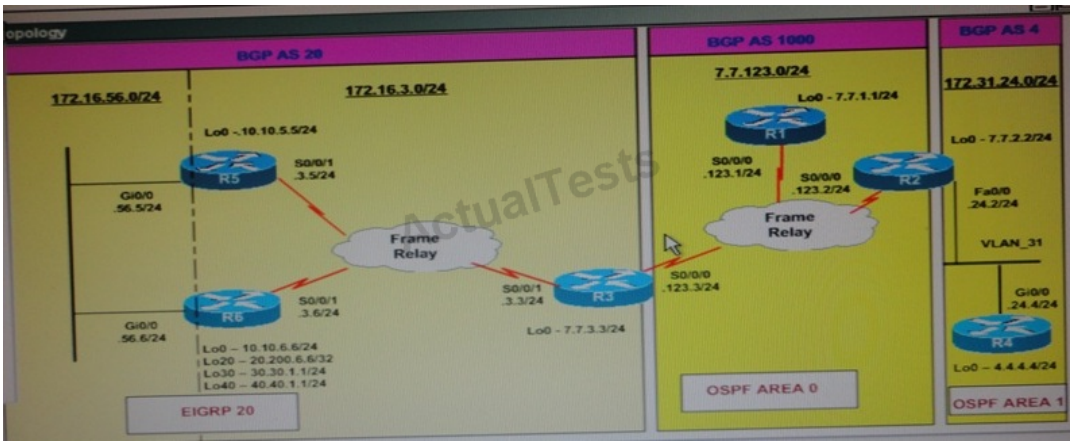
Answer: A

Explanation:

A router forwards the traffic based on the longer prefix match and the shortest administrative distance. In this case the route via 10.1.1.1 has a /26 subnet mask and an AD of 90 as this is using EIGRP. Therefore the router will use this route to connect to 192.168.32.1.

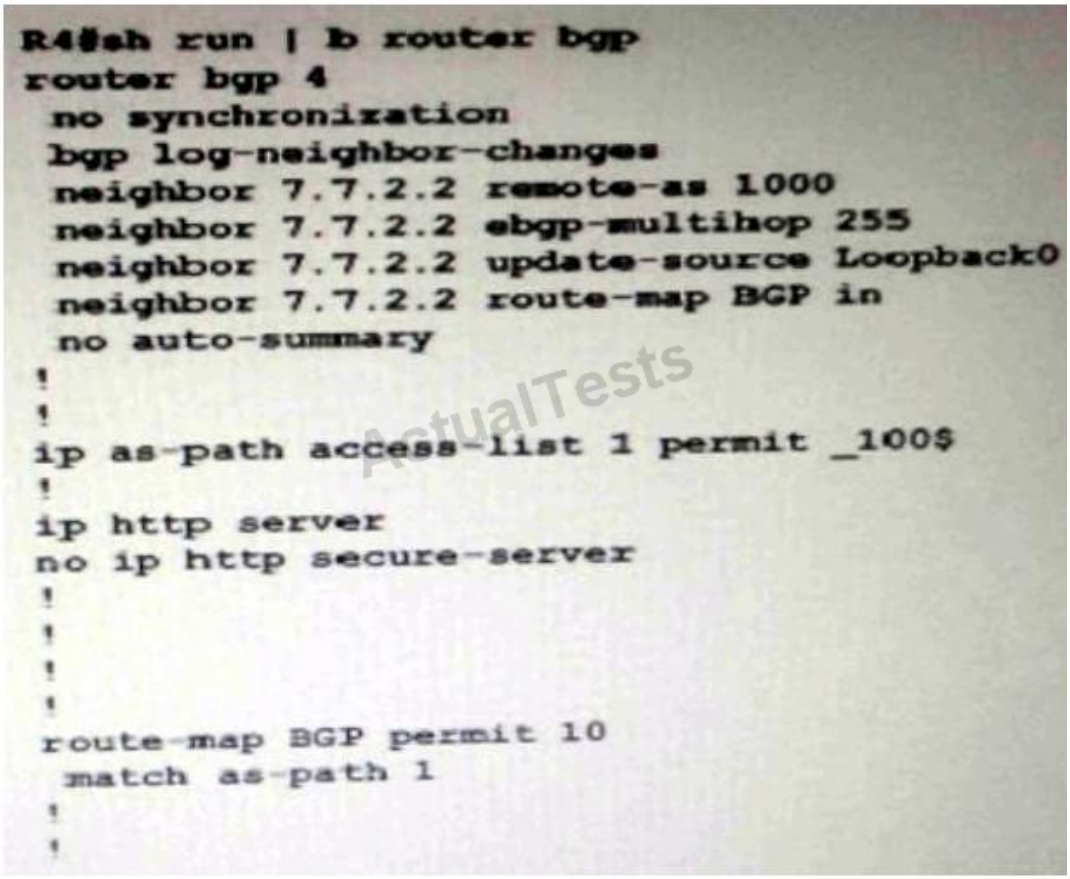
QUESTION NO: 155

Refer to the exhibit.




```
R4#sh ip bgp
BGP table version is 16, local router ID is 4.4.4.4
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
*> 30.30.1.0/24     7.7.2.2                                0 1000 20 i
*> 40.40.1.0/24     7.7.2.2                                0 1000 20 i
*> 200.200.6.6/32   7.7.2.2                                0 1000 20 i
```



What effect will the as-path filter command that is configured on R4 create BGP routing table?

A. It will have all three routes on the R4 BGP routing table
B. It will have none of the three routes on the R4 BGP routing table.
C. It will have only the route 30.30.1.0/24.
D. It will have routes 40.40.1.0/24 and 200.200.6.6/32.
E. It will have routes 30.30.1.0/24 and 200.200.6.6/32.
F. It will have routes 30.30.1.0/24 and 40.40.1.0/24.

Answer: B

Explanation:

Configure BGP Route Filtering by Neighbor

You can filter BGP advertisements in two ways:

Use AS-path filters, as with the ip as-path access-list global configuration command and the neighbor filterlist command

Use access or prefix lists, as with the neighbor distribute-list command.

Filtering using prefix lists is described in "Configuring BGP Filtering Using Prefix Lists".

If you want to restrict the routing information that the Cisco IOS software learns or advertises, you can filter BGP routing updates to and from particular neighbors. To do this, you can either define an access list or a prefix list and apply it to the updates.

Note Distribute-list filters are applied to network numbers and not autonomous system paths.

To filter BGP routing updates, use the following command in router configuration mode:

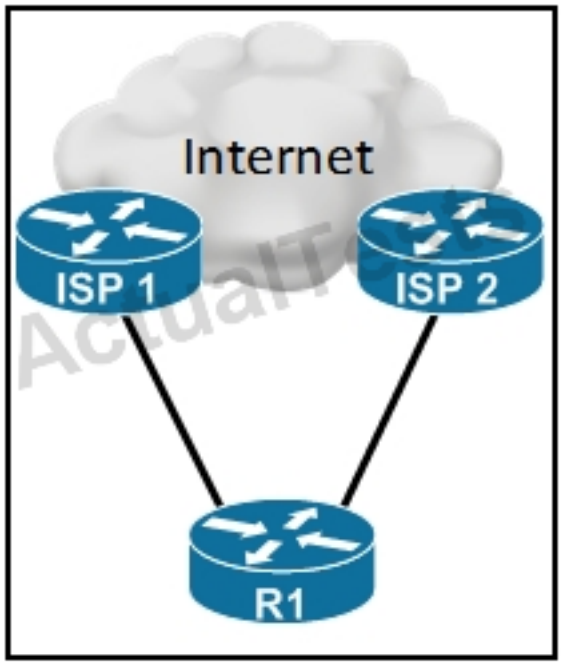
Note Although neighbor prefix-list can be used as an alternative to the neighbor distribute-list

command, do not use attempt to apply both neighbor prefix list and neighbor distribute-list filtering to the same neighbor.

Reference
http://www.cisco.com/en/US/docs/ios/12_0/np1/configuration/guide/1cbgp.html

QUESTION NO: 156

Refer to the exhibit.

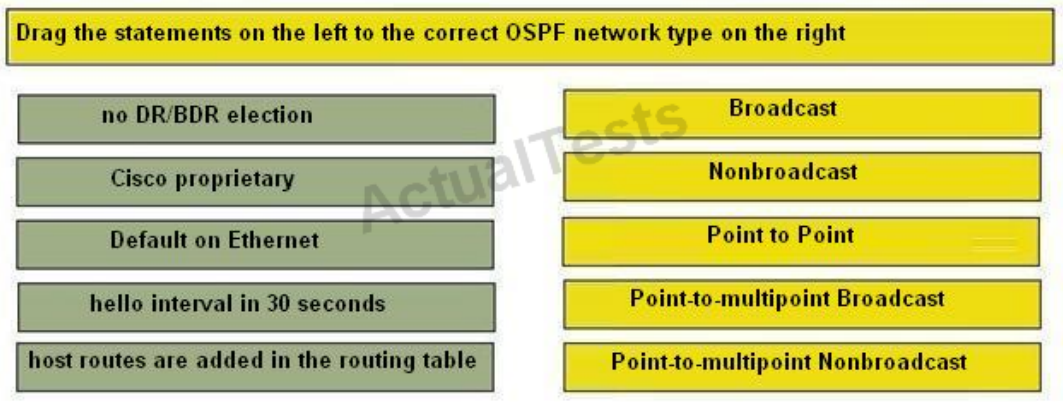


R1 has two eBGP sessions to ISP1 and ISP2 (one to each ISP router), and R1 receives the same prefixes through both links.

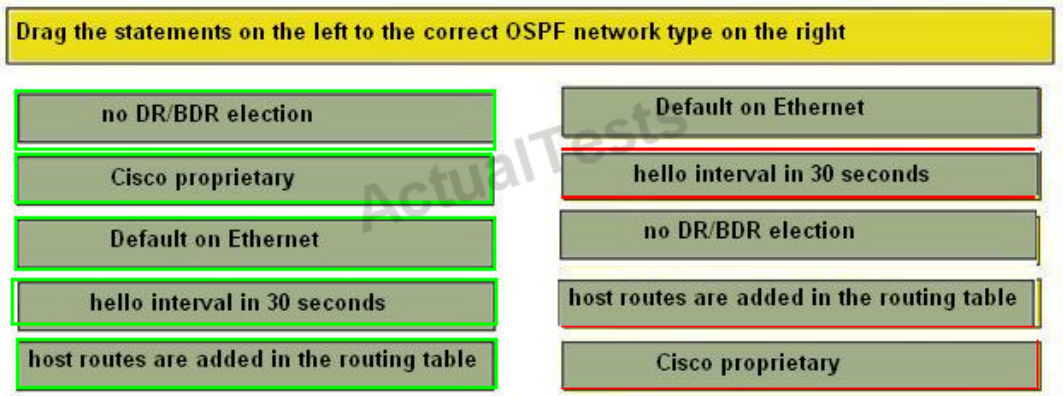
Which configuration should be applied for the link between R1 and ISP2 to be preferred for incoming traffic (ISP2 to R1)?

A. increase local preference on R1 for routes advertised to ISP2
B. decrease local preference on R1 for routes advertised to ISP2
C. increase MED on ISP2 for routes advertised to R1
D. decrease MED on ISP2 for routes advertised to R1

QUESTION NO: 157 DRAG DROP



Answer:



Explanation:

No DR/BDR election – Point to Point

Cisco proprietary – Point-to-Multipoint Nonbroadcast

Default on Ethernet – Broadcast

Hello interval is 30 seconds – Nonbroadcast

Host routes are added in the routing table - Point to Multipoint Broadcast

Reference: OSPF network types

http://blog.sgicc.com/?p=353

QUESTION NO: 158

Two routers configured to run BGP have been connected to a firewall, one on the inside interface and one on the outside interface. BGP has been configured so the two routers should peer, including the correct BGP session endpoint addresses and the correct BGP session hop-count limit (EBGP multihop). What is a good first test to see if BGP will work across the firewall?

A. Attempt to TELNET from the router connected to the inside of the firewall to the router connected to the outside of the firewall. If telnet works, BGP will work, since telnet and BGP both use TCP to transport data.
B. Ping from the router connected to the inside interface of the firewall to the router connected to the outside interface of the firewall. If you can ping between them, BGP should work, since BGP uses IP to transport packets.
C. There is no way to make BGP work across a firewall without special configuration, so there is no simple test that will show you if BGP will work or not, other than trying to start the peering session.
D. There is no way to make BGP work across a firewall.

Answer: C
Explanation:
1. The question doesn't say that you are passing the port parameter to the telnet session. In the answer cisco says "since telnet and BGP both use TCP to transport data." Meaning that TELNET and BGP share TCP, no mention of ports.
2. If you telnet to Port 179 you are testing the path only in 1 direction from the inside to the outside. Yes stateful firewalls will allow return traffic from outside, but they won't allow the outside neighbor to initiate a session.
3. If the Firewall is using NAT for outgoing traffic, which is common, you will be able to telnet to the BGP peer, but the peer won't be able to reach your router back if it needs to initiate a session.
4. The Firewall can translate port 179 to 23 or anything else that will give you a false positive on your Telnet test.
5. Answer C says that
A. "There is no way to make BGP work across a firewall without special configuration" Special configuration refers to the Firewall, since in the question they explicitly say that BGP has been properly configured.
B. "Trying to start the peering session." will provide you with a definitive answer.
C. Therefore correct answer is C.

QUESTION NO: 159

You have two EBGP peers connected via two parallel serial lines. What should you do to be able to load-balance between two EBGP speakers over the parallel serial lines in both directions?

A. nothing, BGP automatically load-balances the traffic between different autonomous systems on all available links
B. peer between the eBGP speaker's loopbacks, configuring eBGP multihop as required, and use an IGP to load-share between the two equal-cost paths between the loopback addresses
C. configure a loopback as update source for both EBGP peers and have on each AS an IGP to introduce two equal-cost paths to reach the EBGP peer loopback address; it is also necessary to

use the next-hop-self command
D. use the ebgp-load-balance command on the neighbor statement on both sides
E. configure a loopback as update source for both EBGP peers and have on each AS an IGP to introduce two equal-cost paths to reach the peer loopback address; it is also necessary to use the ebgp-multihop and next-hop-self commands

Answer: E

Explanation:




```
RTA#
int loopback 0
ip address 150.10.1.1 255.255.255.0
router bgp 100
neighbor 160.10.1.1 remote-as 200
neighbor 160.10.1.1 ebgp-multihop
neighbor 160.10.1.1 update-source loopback 0
network 150.10.0.0

ip route 160.10.0.0 255.255.0.0 1.1.1.2
ip route 160.10.0.0 255.255.0.0 2.2.2.2

RTB#
int loopback 0
ip address 160.10.1.1 255.255.255.0
router bgp 200
neighbor 150.10.1.1 remote-as 100
neighbor 150.10.1.1 update-source loopback 0
neighbor 150.10.1.1 ebgp-multihop
network 160.10.0.0

ip route 150.10.0.0 255.255.0.0 1.1.1.1
ip route 150.10.0.0 255.255.0.0 2.2.2.1
```

This example illustrates the use of loopback interfaces, update-source, and ebgp-multihop. The example is a workaround in order to achieve load balancing between two eBGP speakers over

parallel serial lines. In normal situations, BGP picks one of the lines on which to send packets, and load balancing does not happen. With the introduction of loopback interfaces, the next hop for eBGP is the loopback interface. You use static routes, or an IGP, to introduce two equal-cost paths to reach the destination. RTA has two choices to reach next hop 160.10.1.1: one path via 1.1.1.2 and the other path via 2.2.2.2. RTB has the same choices.

Load balancing with parallel EBGP sessions

Establishing parallel EBGP sessions across parallel links between two edge routers (EBGP peers), as displayed in Figure 1, is the most versatile form of EBGP load balancing. It does not require static routing or extra routing protocol (like the design running EBGP between routers' loopback interfaces), IOS- specific tricks (configuring the same IP address on multiple interfaces) or specific layer-2 encapsulation (like Multilink PPP). It even allows proportional load-balancing across unequal-bandwidth links and combinations of various layer-2 technologies (for example, load-balancing between a serial line and an Ethernet interface). The only drawback of this design is the increased size of the BGP table, as every BGP prefix is received from the EBGP neighbor twice.

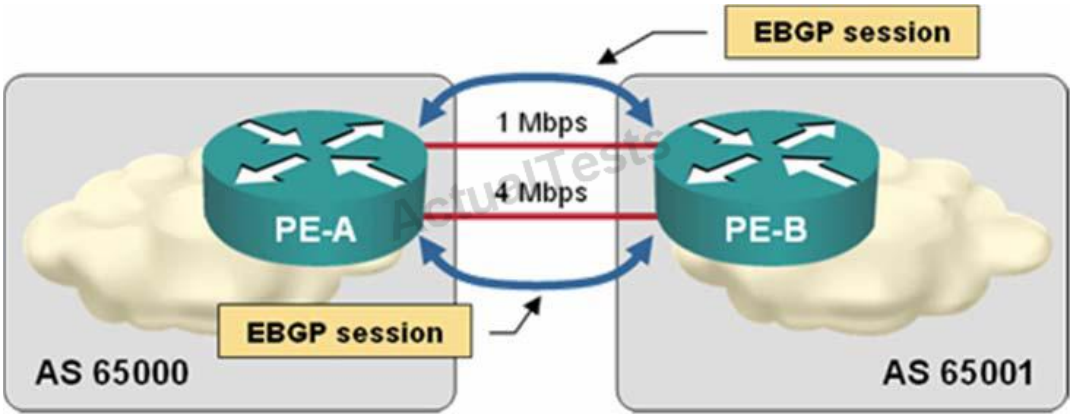


Figure 1: Parallel EBGP sessions

Benefits	Drawbacks
• No extra routing protocol between EBGP neighbors • Reliably detects link and node failures • No IP addressing tricks • No dependency on layer-2 transport technology • Load balancing proportional to link bandwidth	• Duplicate entries in the BGP table increase memory requirements • Additional BGP session and larger BGP table increases CPU utilization

Basic configuration

To implement parallel EBGP sessions, configure multiple neighbors on both EBGP routers, one for each IP subnet (parallel link between the EBGP peers) and enable EBGP multipath load balancing with the maximum paths router configuration command. A sample configuration is shown in the following table:

PE-A (AS 65000)	PE-B (AS 65001)
hostname PE-A ! interface Loopback0 ip address 10.0.1.1 255.255.255.255 ! interface Serial1/4 description Link #1 to PE-B bandwidth 1000 ip address 10.0.7.37 255.255.255.252 encapsulation ppp ! interface Serial1/5 description Link #2 to PE-B bandwidth 4000 ip address 10.0.7.45 255.255.255.252 encapsulation ppp ! router bgp 65000 no synchronization bgp log-neighbor-changes neighbor 10.0.7.38 remote-as 65001 neighbor 10.0.7.46 remote-as 65001 maximum-paths 8	hostname PE-B ! interface Loopback0 ip address 10.0.1.5 255.255.255.255 ! interface Serial1/0 description Link #1 to PE-A bandwidth 1000 ip address 10.0.7.38 255.255.255.252 encapsulation ppp ! interface Serial1/1 description Link #2 to PE-A bandwidth 4000 ip address 10.0.7.46 255.255.255.252 encapsulation ppp ! router bgp 65001 no synchronization bgp log-neighbor-changes neighbor 10.0.7.37 remote-as 65000 neighbor 10.0.7.45 remote-as 65000 maximum-paths 8

Reference
http://wiki.nil.com/Load_balancing_with_parallel_EBGP_sessions
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800c95bb.shtml#ebgpmulithoploadbal

QUESTION NO: 160

Which of these best identifies the types of prefixes a router running BGP will advertise to an EBGP peer?

A. prefixes received from any other BGP peer and prefixes locally originated via network statements or redistributed to BGP
B. all prefixes in its IP routing table.
C. only prefixes received from EBGP peers and prefixes locally originated via network statements or redistributed.
D. only prefixes received from EBGP peers and prefixes received from route reflectors.
E. all prefixes in its routing table except the prefixes received from other EBGP peers.
F. all prefixes in its routing table except the prefixes received from other IBGP peers.

Answer: A

Explanation:

eBGP peers will advertise all known eBGP routes to all other eBGP peers.

iBGP peers will only advertise their own internal routes to other iBGP peers.

A BGP speaking router will never advertise another iBGP peer's routes to any other iBGP peer.

QUESTION NO: 161

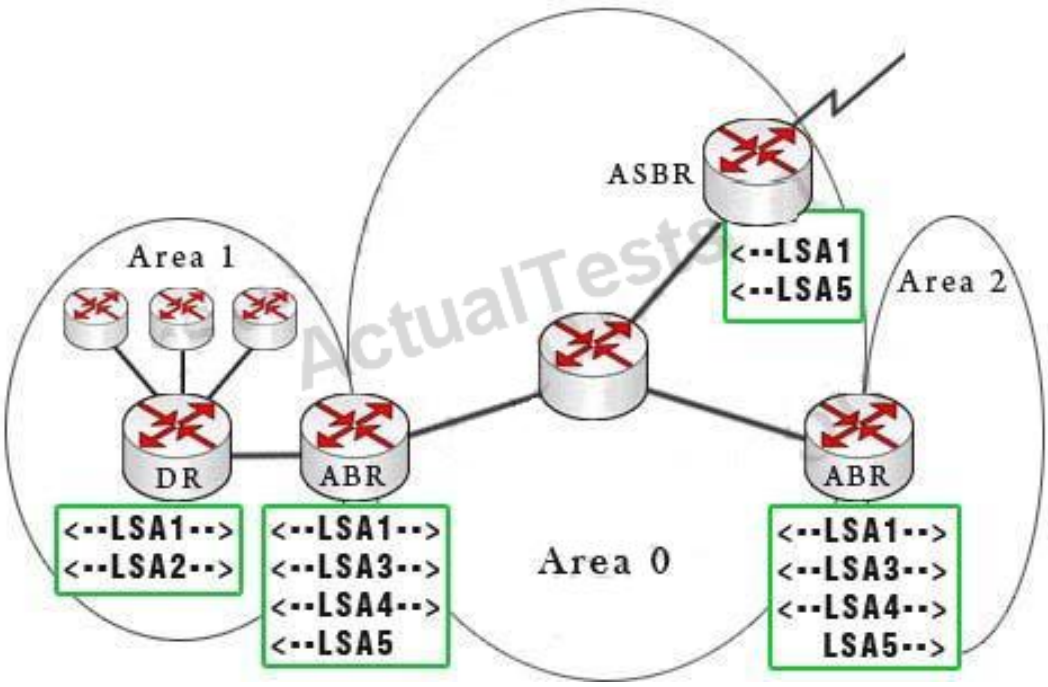
Which one of these statements is true of OSPF type 5 LSAs?

A. They are used to summarize area routes to other areas.
B. They are used in not-so-stubby areas to propagate external routes.
C. They are used to notify areas of the ASBR.
D. They are flooded to all areas except stub areas (external route).

Answer: D

Explanation:

Type 5 external link LSAs are used to advertise external routes originated from an ASBR. They are flooded through the whole OSPF domain.



Note: The dashed arrows show the directions of LSAs in this example

Below is a summary of OSPF Link-state advertisements (LSAs)

Router link LSA (Type 1) - Each router generates a Type 1 LSA that lists its neighbors and the cost to each.

LSA Type 1 is only flooded inside the router's area, does not cross ABR. Network link LSA (Type 2) - is sent out by the designated router (DR) and lists all the routers on the segment it is adjacent to. Types 2 are flooded within its area only; does not cross ABR. Type 1 & type 2 are the basis of SPF path selection. Summary link LSA (Type 3) - ABRs generate this LSA to send between areas (so type 3 is called inter-area link). It lists the networks inside other areas but still belonging to the autonomous system and aggregates routes. Summary links are injected by the ABR from the backbone into other areas and from other areas into the backbone. Summary LSA (Type 4) - Generated by the ABR to describe routes to ASBRs. In the above example, the only ASBR belongs to area 0 so the two ABRs send LSA Type 4 to area 1 & area 2 (not vice versa). This is an indication of the existence of the ASBR in area 0. Note: Type 4 LSAs contain the router ID of the ASBR.

External Link LSA (LSA 5) - Generated by ASBR to describe routes redistributed into the area (which means networks from other autonomous systems). These routes appear as E1 or E2 in the routing table. E2 (default) uses a static cost throughout the OSPF domain as it only takes the cost into account that is reported at redistribution. E1 uses a cumulative cost of the cost reported into the OSPF domain at redistribution plus the local cost to the ASBR. Type 5 LSAs flood throughout the entire autonomous system but notice that Stubby Area and Totally Stubby Area do not accept Type 5. Multicast LSA (Type 6) are specialized LSAs that are used in multicast OSPF applications. NSSA External LSA (Type 7) - Generated by an ASBR inside a NSSA to describe routes redistributed into the NSSA. LSA 7 is translated into LSA 5 as it leaves the NSSA. These routes appear as N1 or N2 in the ip routing table inside the NSSA. Much like LSA 5, N2 is a static cost while N1 is a cumulative cost that includes the cost up to the ASBR

Reference
http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094e9e.shtml#appa1

QUESTION NO: 162

Which OSPF LSA type does an ASBR use to originate a default route into an area?

A. LSA1
B. LSA3
C. LSA4
D. LSA 5

E. LSA7

Answer: D

Explanation:

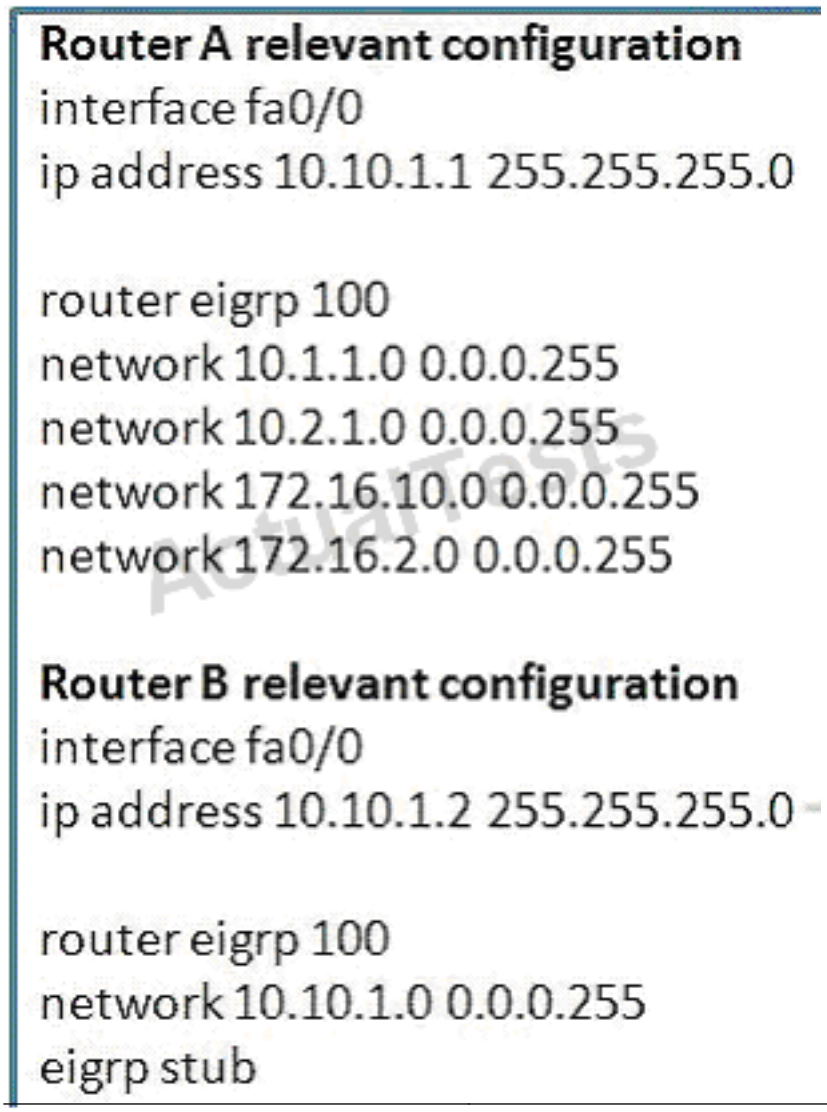
By default, the OSPF router does not generate a default route into the OSPF domain. In order for OSPF to generate a default route, you must use the default-information originate command. With this command, the router will advertise type 5 LSA with a link ID of 0.0.0.0.

Reference

http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a00801ec9f0.shtml

QUESTION NO: 163

Refer to the exhibit. Routers A and B are directly connected. Given the configuration, how many EIGRP routes will router B see in its routing table?



A. 0
B. 1
C. 2
D. 3
E. 4

Answer: A

Explanation:

In this question, router A does not advertise its "network 10.10.1.0 0.0.0.255 in the EIGRP process (the network connected with router B) so no EIGRP neighbor relationship is established between two routers. If we use the "show ip route" command on both routers, we just see a directly connected network

10.10.1.0/24 like this:

```
Router_B#sh ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

Router_B#sh ip eigrp nei
IP-EIGRP neighbors for process 100

Router_B#sh ip eigrp top
IP-EIGRP Topology Table for AS(100)/ID(0.0.0.0)

Router_B#deb eigrp pack
EIGRP Packets debugging is on
    (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP,
     PROBE, ACK, STUB, SIAQUERY, SIAREPLY)

Router_B#sh cdp nei
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID      Local Intrfce   Holdtme   Capability  Platform  Port ID
Router_A       Fas 0/0         176       R           2620XM    Fas 0/0
```

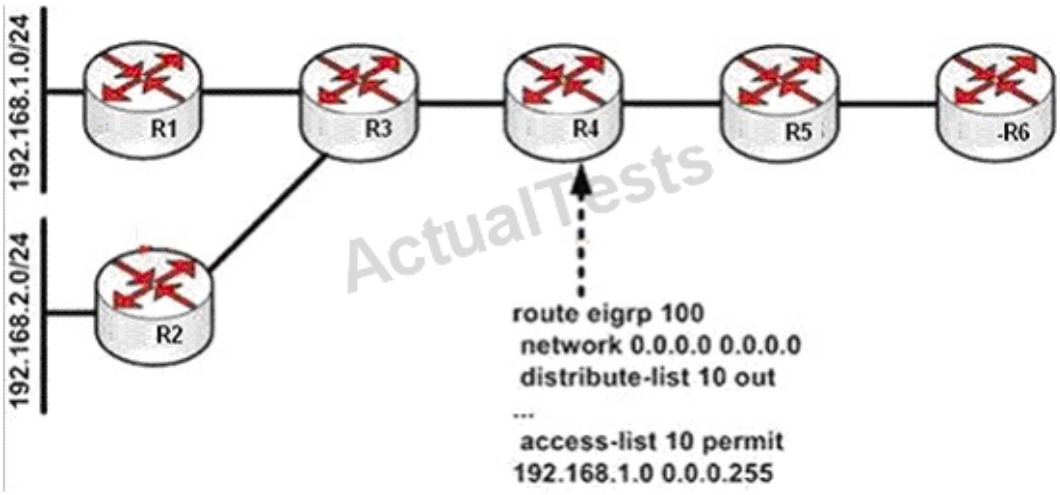
ActualTests

In the EIGRP process of router A
For your information, even if we use the "network 10.10.1.0 0.0.0.255 we still don't see any EIGRP route because router A does not have any interfaces belonging to networks

10.1.1.0/24, 10.2.1.0/24, 172.16.2.0/24 -> it will not advertise these networks to router B.

QUESTION NO: 164

Observe the following network presented in this exhibit carefully. Assume that all routers are running EIGRP in AS 100 on all connected links. If the link between R3 and R4 is down, how many queries will R5 and R6 receive?



A. R5 will receive two queries, one for 192.168.1.0/24 and one for 192.168.2.0/24. R6 will receive one query, for 192.168.2.0/24.
B. R5 will receive one query, for 192.168.1.0/24, and R6 will receive no queries
C. Both R5 and R6 will receive two queries, one for 192.168.1.0/24 and one for 192.168.2.0/24
D. Neither R5 nor R6 will receive any queries for either 192.168.1.0/24 or 192.168.2.0/24.

Answer: A
Explanation:

R4 will go into an active state and send query about the both routes to R5. R5 at that moment have no record in his eigrp topology table for the route 192.168.2.0/24 as this route has been filtered by distribute list. So R5 will reply immediately to the router R4 telling him that route 192.168.2.0/24 is unreachable. As for the query for the route 192.168.1.0/24, R5 at that moment

has the record in his eigrp topology database for this route pointing back to the router R4. This is why R4 will send only one query to router R6 asking for the path to the route 192.168.2.0.

Incorrect answers
B: R5 will receive two queries, for 192.1 68.1.0/24, and for 192 1682 0/24 as distribute list filters only EIGRP updates messages.

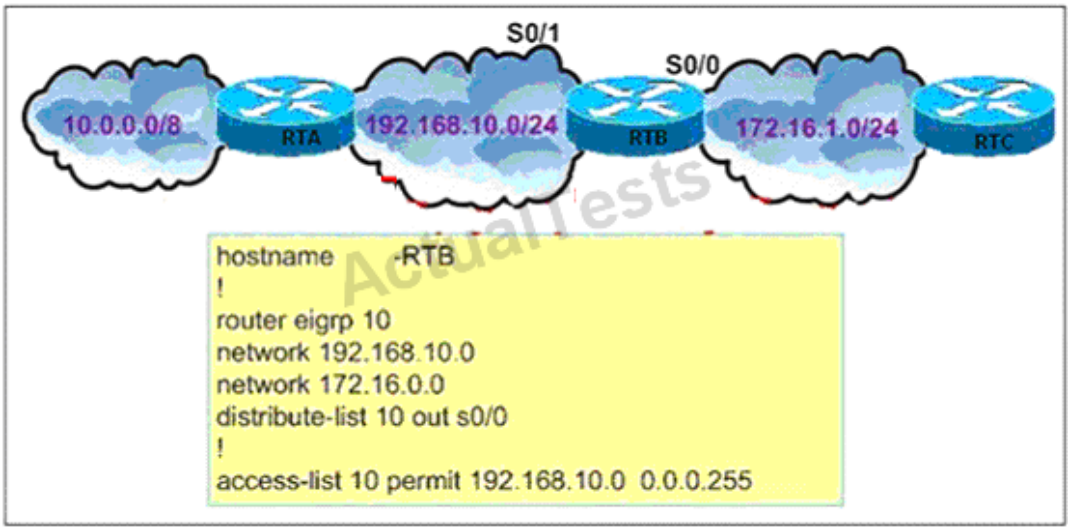
C: Only R5 will receive the one message with two queries, for 192.168.1.0/24 and for 192 1682 0/24 R5 will not send query for 192 1682 0/24 to the router R6 as R5 had no record forthis route at the moment of receiving the query from R4.

D: R5 will receive two queries, one for 192.168.1.0/24 and one for 192 1682 0/24 R6 will receive one query, for 192 1682 0/24.

Reference
http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094cb7.shtml#queries

QUESTION NO: 165

Which description of the following is true according to the information shown in the figure?



A. RTC will not have the 192.168.10.0 network in its routing table.
B. RTC will not have the 10.0.0.0 network in its routing table.
C. RTB will not have the 10.0.0.0 network in its routing table.

D. RTB and RTC will not have the 10.0.0.0 network in their routing tables.

Answer: B

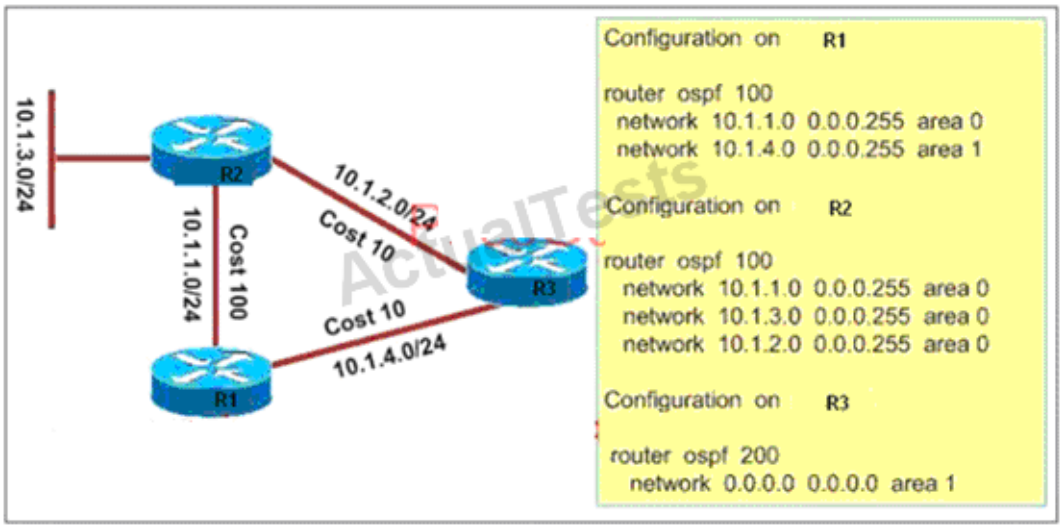
Explanation:

Distribute lists are used to control routing updates. A distribute list allows an access-list to be applied to routing updates. When you configure an access list and use it with a distribute list routing updates can be controlled no matter what their source is. Access list should permit the networks that you want advertised or redistributed and deny all networks that you want to remain hidden. Router then applies the access list to routing updates for that protocol.

RTC will not be able to access the 10.0.0.0 network, because will not have the 10.0.0.0 network in its routing table. With ACL 10 RTB permits routing updates to be sent to RTC for network 192.168.10.0/24 and because of the implicit deny rule at the end of any ACL, update for 10.0.0.0/8 network will not be allowed by the ACL.

QUESTION NO: 166

Observe the following exhibit. Which path will be preferred by traffic destined to 10.1.3.1 and arriving at R1?



A. through R3, because R1 will only have a summary (type 3) LSA from R2
B. through R2, since it is the path through Area 0
C. through R3, since that is the lowest cost path (10+10 = 20, which is lower than 100)
D. through R2; this is the only path available for R1 to reach 10.1.3.0/24, since R3 is in a different autonomous system than R1 and R2

Answer: B

Explanation:

Here is the configuration in a lab environment along with the results. All routers are 3725 routers

R1

```
interface FastEthernet0/0
ip address 10.1.1.1 255.255.255.0
ip ospf cost 100
duplex auto
speed auto
!
interface FastEthernet0/1
ip address 10.1.4.1 255.255.255.0
ip ospf cost 10
duplex auto
speed auto
!
router ospf 100
log-adjacency-changes
network 10.1.1.0 0.0.0.255 area 0
network 10.1.4.0 0.0.0.255 area 1
```

R2

```
interface Loopback0
ip address 10.1.3.2 255.255.255.0
!
interface FastEthernet0/0
ip address 10.1.1.2 255.255.255.0
ip ospf cost 100
duplex auto
speed auto
!
interface FastEthernet0/1
ip address 10.1.2.2 255.255.255.0
ip ospf cost 10
duplex auto
speed auto
!
router ospf 100
log-adjacency-changes
network 10.1.1.0 0.0.0.255 area 0
network 10.1.2.0 0.0.0.255 area 0
```

network 10.1.3.0 0.0.0.255 area 0

R3
interface FastEthernet0/0
ip address 10.1.2.3 255.255.255.0
ip ospf cost 10
duplex auto
speed auto
!
interface FastEthernet0/1
ip address 10.1.4.3 255.255.255.0
ip ospf cost 10
duplex auto
speed auto
!
router ospf 200
log-adjacency-changes
network 0.0.0.0 255.255.255.255 area 1

R1 Additional information
R1#sh ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route
Gateway of last resort is not set
10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O 10.1.2.0/24 [110/20] via 10.1.4.3, 00:07:44, FastEthernet0/1
O 10.1.3.2/32 [110/101] via 10.1.1.2, 00:04:48, FastEthernet0/0
C 10.1.1.0/24 is directly connected, FastEthernet0/0
C 10.1.4.0/24 is directly connected, FastEthernet0/1
R1#sh ip ospf nei
Neighbor ID Pri State Dead Time Address Interface
10.1.3.2 1 FULL/BDR 00:00:36 10.1.1.2 FastEthernet0/0
10.1.4.3 1 FULL/BDR 00:00:34 10.1.4.3 FastEthernet0/1
R1#ping 10.1.3.2

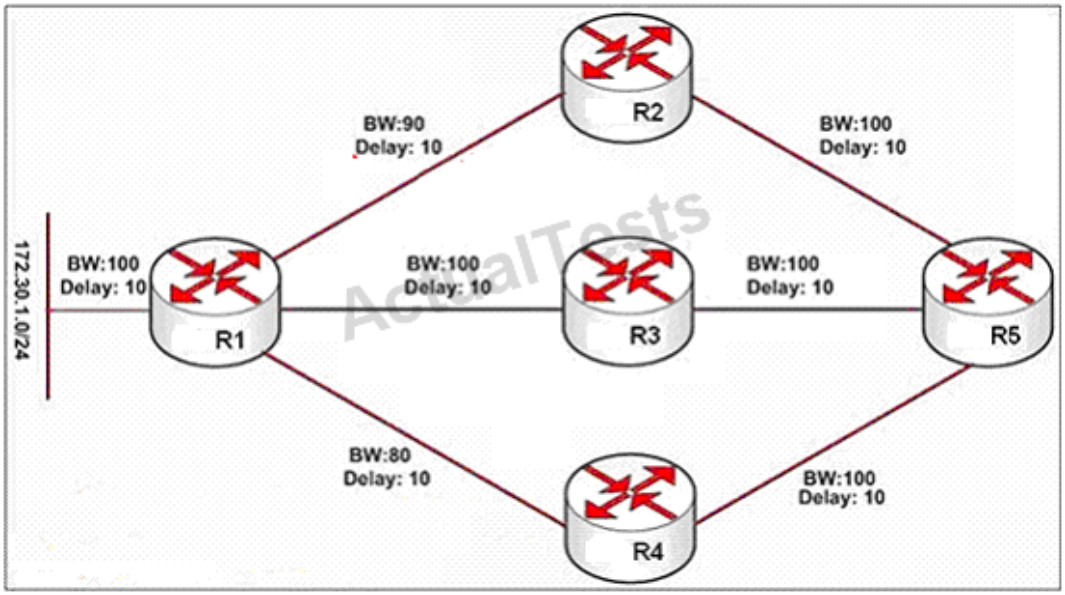
Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.1.3.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 20/39/72 ms
R1#traceroute 10.1.3.2

Type escape sequence to abort.
Tracing the route to 10.1.3.2
1 10.1.1.2 44 msec * 40 msec

QUESTION NO: 167

You are a network engineer for a company, study the exhibit carefully. The company's network is running EIGRP and you want to change the path R5 uses to reach 172.30.1.0/24 to R4. How could you achieve this goal?



A. Change the bandwidth on the link between R2 and R5 to 70, and change the bandwidth on the link between R3 and R5 to 70.
B. Change the bandwidth on the link between R4 and R5 to 110.
C. Change the bandwidth on the link between R3 and R5 to 70.
D. Do nothing, the best path to 172.30.1.0/24 from R5 is already through R4.

Answer: A

Explanation:

Here is an example from Cisco on how and why to change the EIGRP bandwidth along with the verification steps

Change the Bandwidth at R2
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.

Changing the delay does not impact other protocols nor does it cause EIGRP to throttle back.
Check the EIGRP topology table for R1 before you make any changes.

R1# show ip eigrp topology 10.1.3.0 255.255.255.0
IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2221056
Routing Descriptor Blocks:
10.1.1.2 (Serial0.201), from 10.1.1.2, Send flag is 0x0
Composite metric is (2221056/307200), Route is Internal
Vector metric:
Minimum bandwidth is 1544 Kbit
Total delay is 22000 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 2
Check the starting values for the ethernet0 interface on R2.
R2# show interface ethernet0
Ethernet0 is up, line protocol is up
Hardware is Lance, address is 0010.7b3c.6786 (bia 0010.7b3c.6786)
Internet address is 10.1.2.2/24
MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec, rely 255/255, load 1/255
Encapsulation ARPA, loopback not set, keepalive set (10 sec)
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:00:01, output 00:00:02, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
1938 packets input, 165094 bytes, 0 no buffer
Received 1919 broadcasts, 0 runts, 0 giants, 0 throttles

0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
0 input packets with dribble condition detected
1482 packets output, 124222 bytes, 0 underruns
0 output errors, 0 collisions, 18 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier
0 output buffer failures, 0 output buffers swapped out
Decrease the bandwidth to see the impact on R1.
R2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# interface ethernet0
R2(config-if)# bandwidth 5000
R2(config-if)# end
R2#
Confirm the change.
R2# show interface ethernet0
Ethernet0 is up, line protocol is up
Hardware is Lance, address is 0010.7b3c.6786 (bia 0010.7b3c.6786)
Internet address is 10.1.2.2/24
MTU 1500 bytes, BW 5000 Kbit, DLY 1000 usec, rely 255/255, load 1/255
Encapsulation ARPA, loopback not set, keepalive set (10 sec)
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:00:02, output 00:00:01, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
1995 packets input, 169919 bytes, 0 no buffer
Received 1969 broadcasts, 0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
0 input packets with dribble condition detected
1525 packets output, 127831 bytes, 0 underruns
0 output errors, 0 collisions, 18 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier
0 output buffer failures, 0 output buffers swapped out
Confirm that it also changed in the EIGRP topology table.
R2# show ip eigrp topology 10.1.3.0 255.255.255.0
IP-EIGRP (AS 1): topology entry for 10.1.3.0/24

State is Passive, Query origin flag is 1, 2 Successor(s), FD is 563200
Routing Descriptor Blocks:
10.1.2.4 (Ethernet0), from 10.1.2.4, Send flag is 0x0
Composite metric is (563200/281600), Route is Internal
Vector metric:
Minimum bandwidth is 5000 Kbit
Total delay is 2000 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 1
10.1.2.3 (Ethernet0), from 10.1.2.3, Send flag is 0x0
Composite metric is (563200/281600), Route is Internal
Vector metric:
Minimum bandwidth is 5000 Kbit
Total delay is 2000 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 1
Check the impact on the EIGRP topology table for R1.
R1# show ip eigrp topology 10.1.3.0 255.255.255.0
IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2221056
Routing Descriptor Blocks:
10.1.1.2 (Serial0.201), from 10.1.1.2, Send flag is 0x0
Composite metric is (2221056/563200), Route is Internal
Vector metric:
Minimum bandwidth is 1544 Kbit
Total delay is 22000 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 2

There is no change, as the Frame Relay connection between R1 and R2 is still the lowest speed link. You see a change only if you decrease that bandwidth on the ethernet0 interface for R2 to a value less than 1544.

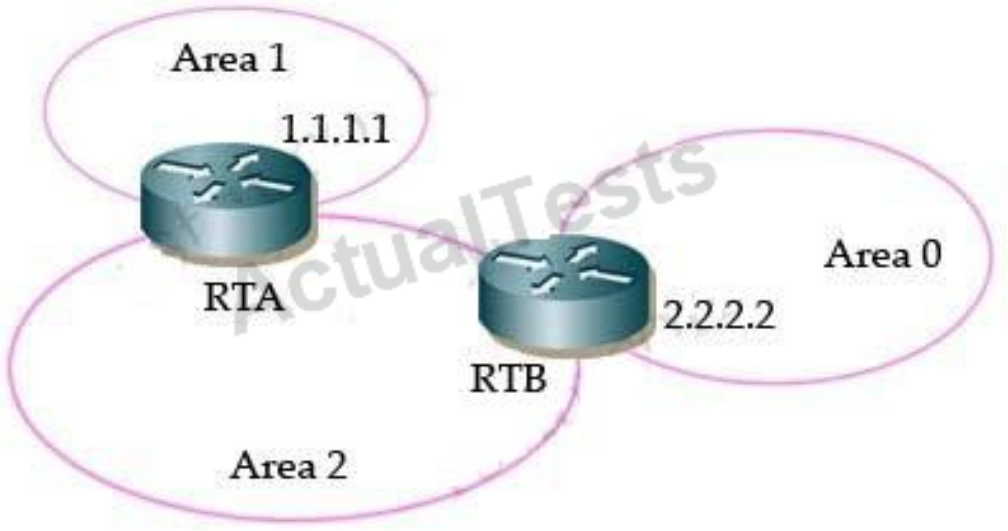
Decrease the bandwidth on the ethernet0 interface for R2 to 1000.
R2# configure terminal

Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# interface ethernet 0
R2(config-if)# bandwidth 1000
R2(config-if)# end
R2#
Check the impact on the EIGRP topology table for R1.
R1# show ip eigrp topology 10.1.3.0 255.255.255.0
IP-EIGRP (AS 1): Topology entry for 10.1.3.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 312320
Routing Descriptor Blocks:
10.1.1.2 (Serial0.201), from 10.1.1.2, Send flag is 0x0
Composite metric is (3123200/2611200), Route is Internal
Vector metric:
Minimum bandwidth is 1000 Kbit
Total delay is 22000 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 2

Reference
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800c2d96.shtml#changebandwidthR2

QUESTION NO: 168

Refer to the exhibit. How would you get the 1.1.1.1 network into the OSPF database?



A. Configure RTA as an ASBR.
B. Redistribute connected routes on RTA into OSPF.
C. Set up a virtual link between area 1 and area 0.
D. Set up a virtual link between area 1 and area 2.
E. Add a static route into RTB and enter it into OSPF.
F. Place a network 1.1.1.0 0.0.0.0 command into RTB.
G. Set up a unique router ID on RTA using an RFC 1918 address
H. Change area 0 on RTB to area 1

Answer: C

Explanation:

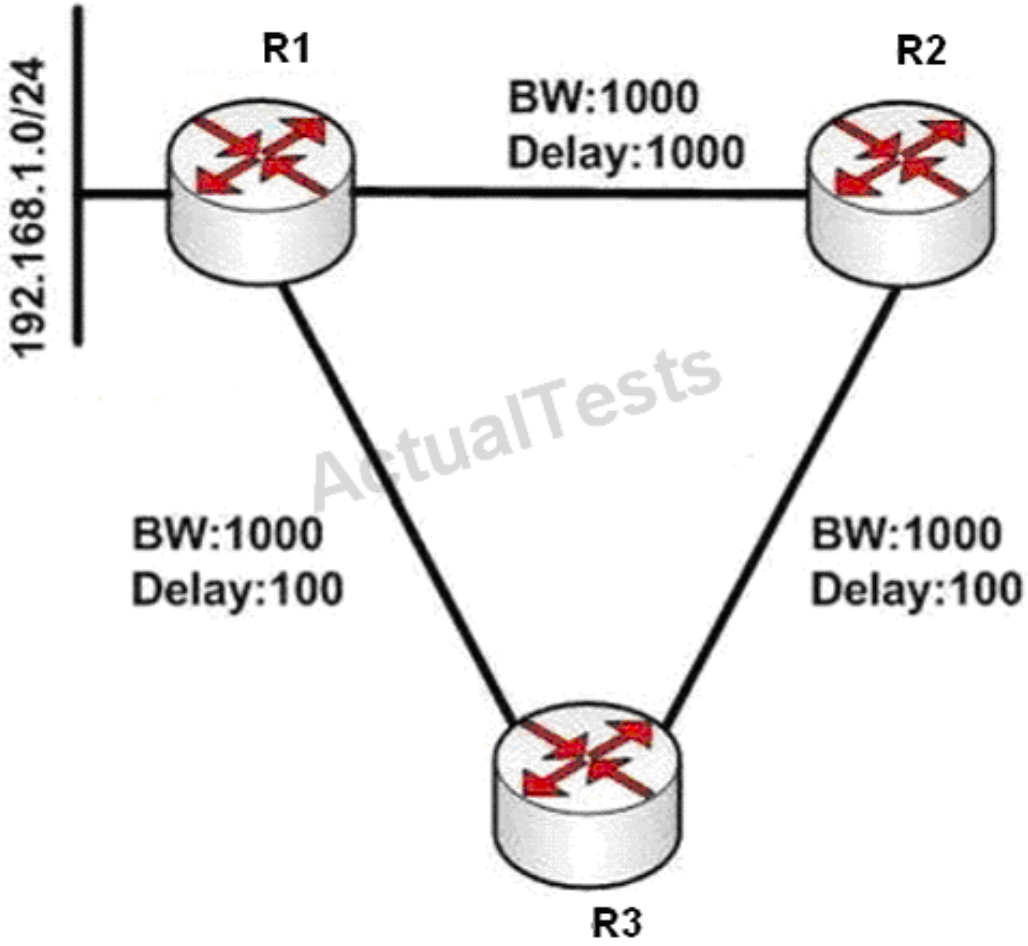
Recall that in OSPF, area 0 is called backbone area and all other areas connect directly to it. In the exhibit above, area 1 is not directly connected with area 0 so we need to set up a virtual link between area 1 & area 0 so that the networks in area 1 can be recognized in area 0. The virtual-link configuration is shown below:

RTB(config)#router ospf 1
RTB(config-router)#area 2 virtual-link 1.1.1.1
RTA(config)#router ospf 1
RTA(config-router)#area 2 virtual-link 2.2.2.2

Notice that the router-id in the "area ... virtual-link " command is the router-id of the neighboring router.

QUESTION NO: 169

Based on the network displayed in the exhibit, both R1 and R2 are configured as EIGRP stub routers. If the link between R1 and R3 is down, will R3 still be able to reach 192.168.1.0/24, and why or why not?



A. No. R3 would remove its route to 192.168.1.0/24 through R1, but would not query R2 for an alternate route, since R2 is a stub.
B. No. The path through R2 would always be considered a loop at R3.
C. Yes. When a directly connected link fails, a router is allowed to query all neighbors, including stub neighbors, for an alternate route.
D. Yes, because R3 would know about both routes, through R1 and R2, before the link between R1 and R3 failed.

Answer: A

Explanation:

When an EIGRP stub is configured a query is actually sent to the neighbor, however, he just replies with the prefix being unreachable. So R2 responds to the query of R3 with the prefix for 192.168.1.0/24 as unreachable. Answer A is the only good answer really.

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.

Reference

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

QUESTION NO: 170

Which statement best describes OSPF external LSAs (type 5)?

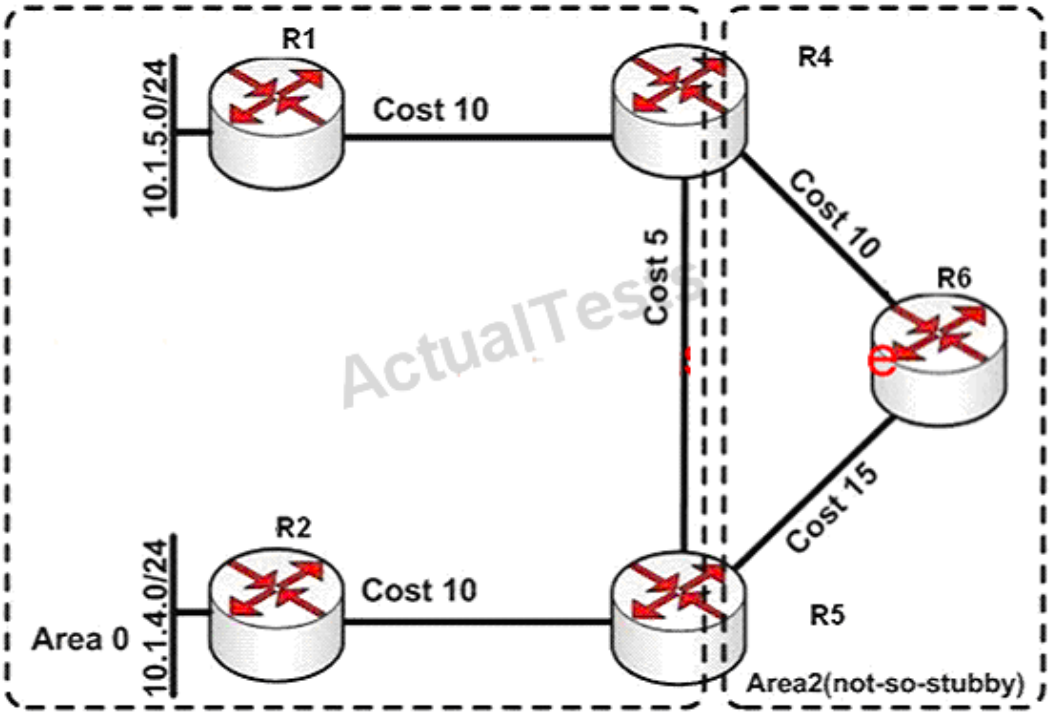
A. OSPF external LSAs are automatically flooded into all OSPF areas, unlike type 7 LSAs, which require that redistribution be configured.
B. External LSAs (type 5) are automatically changed to type 1 LSAs at ASBRs.
C. Type 5 LSAs are route summaries describing routes to networks outside the OSPF Autonomous System.
D. External network LSAs (type 5) redistributed from other routing protocols into OSPF are not permitted to flood into a stub area.

Answer: D
Explanation:
Type 5 - External LSA - these LSAs contain information imported into OSPF from other routing processes.
They are flooded to all areas (except stub areas). For "External Type 1" LSAs routing decisions are made by adding the OSPF metric to get to the ASBR and the external metric from there on, while for "External Type 2" LSAs only the external metric is used. The link-state ID of the type 5 LSA is the external network number

Reference
http://en.wikipedia.org/wiki/Link-state_advertisement

QUESTION NO: 171

You work as a network engineer for a Company. Study the exhibit carefully. In this network, R1 is redistributing 10.1.5.0/24 into OSPF, and R2 is originating 10.1.4.0/24 as an internal route. R6 has received packets destined to 10.1.4.1 and 10.1.5.1. Which statement is true about the path or paths these two packets will take?



A. The packet destined to 10.1.5.1 will follow the optimum path through the network, R4 to R1, while the packet destined to 10.1.4.1 will follow a suboptimal path through the network.
B. There is not enough information provided to determine which packet will take an optimal or suboptimal path through the network.
C. Both packets will follow optimal paths through the network to their destinations
D. The packet destined to 10.1.4.1 will follow an optimal path through the network, while the packet destined to 10.1.5.1 will follow a suboptimal path through the network

Answer: A

Explanation:

Below is the configuration information for each router in the example

R1
interface Loopback0
ip address 10.1.5.1 255.255.255.0
!
interface FastEthernet0/0
ip address 10.1.14.1 255.255.255.0
ip ospf cost 10
router ospf 1
network 10.1.5.0 0.0.0.255 area 0
network 10.1.14.0 0.0.0.255 area 0

R2
interface Loopback0

```
ip address 10.1.4.2 255.255.255.0
!
interface FastEthernet0/0
ip address 10.1.25.2 255.255.255.0
ip ospf cost 10
router ospf 1
network 10.1.4.0 0.0.0.255 area 0
network 10.1.25.0 0.0.0.255 area 0

R4
interface FastEthernet0/0
ip address 10.1.14.4 255.255.255.0
ip ospf cost 10
!
interface FastEthernet0/1
ip address 10.1.46.4 255.255.255.0
ip ospf cost 10
!
interface FastEthernet1/0
ip address 10.1.45.4 255.255.255.0
ip ospf cost 5
!
router ospf 1
area 2 nssa
network 10.1.14.0 0.0.0.255 area 0
network 10.1.45.0 0.0.0.255 area 0
network 10.1.46.0 0.0.0.255 area 2

R5
interface FastEthernet0/0
ip address 10.1.25.5 255.255.255.0
ip ospf cost 10
!
interface FastEthernet0/1
ip address 10.1.56.5 255.255.255.0
!
interface FastEthernet1/0
ip address 10.1.45.5 255.255.255.0
ip ospf cost 15
!
router ospf 1
```

log-adjacency-changes
area 2 nssa
network 10.1.25.0 0.0.0.255 area 0
network 10.1.45.0 0.0.0.255 area 0
network 10.1.56.0 0.0.0.255 area 2

R6
interface FastEthernet0/0
ip address 10.1.46.6 255.255.255.0
ip ospf cost 10
!
interface FastEthernet0/1
ip address 10.1.56.6 255.255.255.0
ip ospf cost 15
!
router ospf 1
area 2 nssa
network 10.1.46.0 0.0.0.255 area 2
network 10.1.56.0 0.0.0.255 area 2
R6#sh ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route

Gateway of last resort is not set
10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O IA 10.1.14.0/24 [110/20] via 10.1.46.4, 00:00:59, FastEthernet0/0
O IA 10.1.4.2/32 [110/26] via 10.1.56.5, 00:01:09, FastEthernet0/1
[110/26] via 10.1.46.4, 00:00:59, FastEthernet0/0
O IA 10.1.5.1/32 [110/21] via 10.1.46.4, 00:00:59, FastEthernet0/0
O IA 10.1.25.0/24 [110/25] via 10.1.56.5, 00:01:09, FastEthernet0/1
[110/25] via 10.1.46.4, 00:00:59, FastEthernet0/0
C 10.1.46.0/24 is directly connected, FastEthernet0/0
O IA 10.1.45.0/24 [110/15] via 10.1.46.4, 00:01:01, FastEthernet0/0
C 10.1.56.0/24 is directly connected, FastEthernet0/1
R6#sh ip ospf nei

Neighbor ID Pri State Dead Time Address Interface
10.1.56.5 1 FULL/DR 00:00:36 10.1.56.5 FastEthernet0/1
10.1.46.4 1 FULL/DR 00:00:39 10.1.46.4 FastEthernet0/0
R6#ping 10.1.4.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.4.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/156/208 ms
R6#traceroute 10.1.4.2
Type escape sequence to abort.
Tracing the route to 10.1.4.2
1 10.1.56.5 124 msec
10.1.46.4 96 msec
10.1.56.5 4 msec
2 10.1.45.5 168 msec
10.1.25.2 128 msec
10.1.45.5 120 msec

Configure OSPF Interface Parameters
Our OSPF implementation allows you to alter certain interface-specific OSPF parameters, as needed. You are not required to alter any of these parameters, but some interface parameters must be consistent across all routers in an attached network. Those parameters are controlled by the ip ospf hello-interval, ip ospf deadinterval, and ip ospf authentication-key commands. Therefore, be sure that if you do configure any of these parameters, the configurations for all routers on your network have compatible values.

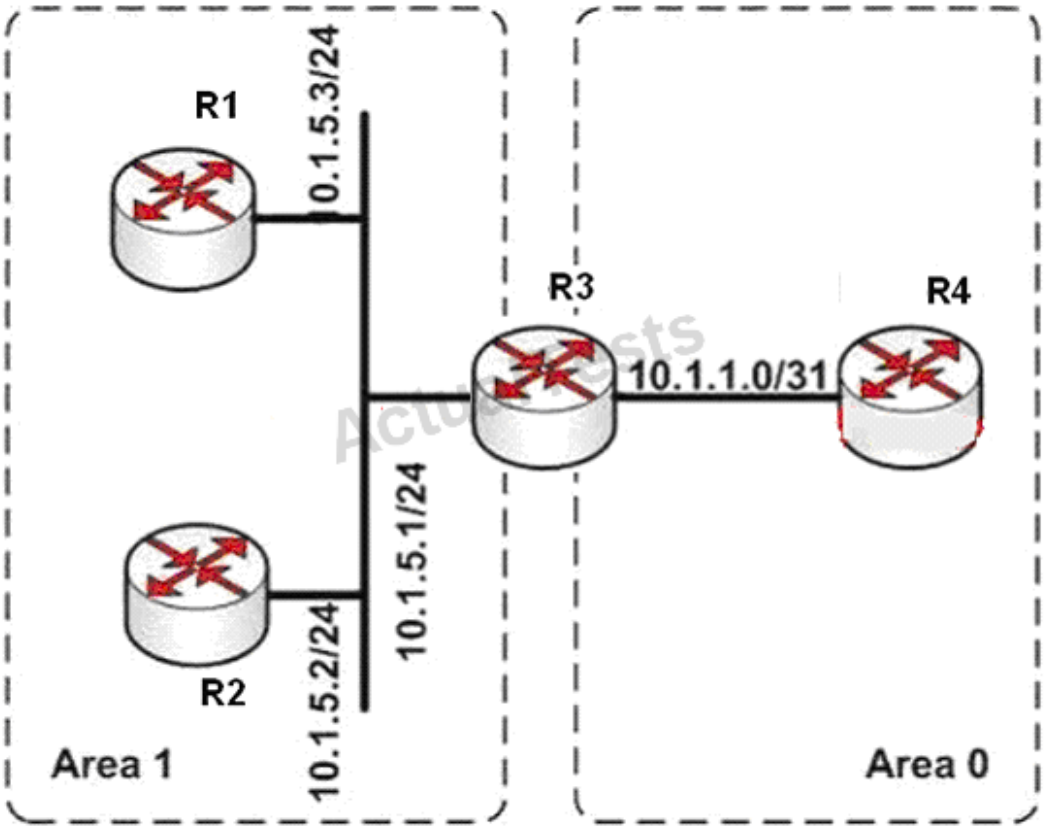
In interface configuration mode, use any of the following commands to specify interface parameters as needed for your network:

Command	Purpose
ip ospf cost cost	Explicitly specify the cost of sending a packet on an OSPF interface.
ip ospf retransmit-interval seconds	Specify the number of seconds between link state advertisement retransmissions for adjacencies belonging to an OSPF interface.
ip ospf transmit-delay seconds	Set the estimated number of seconds it takes to transmit a link state update packet on an OSPF interface.
ip ospf priority number	Set priority to help determine the OSPF designated router for a network.
ip ospf hello-interval seconds	Specify the length of time between the hello packets that the Cisco IOS software sends on an OSPF interface.
ip ospf dead-interval seconds	Set the number of seconds that a device's hello packets must not have been seen before its neighbors declare the OSPF router down.
ip ospf authentication-key key	Assign a password to be used by neighboring OSPF routers on a network segment that is using OSPF's simple password authentication.
ip ospf message-digest-key keyid md5 key	Enable OSPF MD5 authentication.
ip ospf authentication **[message-digest | null]**	Specifies the authentication type for an interface.

Reference
http://www.cisco.com/en/US/docs/ios/12_0/np1/configuration/guide/1cospf.html#wp4937

QUESTION NO: 172

Based on the network provided in the exhibit, in these two areas, all routers are performing OSPF on all interfaces. After examining the OSPF database on R4, do you know which type of LSA will contain 10.1.5.0/24, and which router will have originated it?



A. 10.1.5.0/24 will be in a summary (type 3) LSA originated by R3.
B. 10.1.5.0/24 will be in a router (type 1) LSA generated by R3.
C. 10.1.5.0/24 will be in a network (type 2) LSA originated by R3.
D. 10.1.5.0/24 will not be in any LSA in the OSPF database at R4, because R4 and R3 are in different areas.

Answer: A

Explanation:

R3 is an Area Boarder router and ABR advertise Type 3 Summary LSA's

Type 1 - Router LSA - the router announces its presence and lists the links to other routers or networks in the same area, together with the metrics to them. Type 1 LSAs are flooded across their own area only. The linkstate

ID of the type 1 LSA is the originating router ID.

Type 2 - Network LSA - the designated router (DR) on a broadcast segment (e.g. Ethernet) lists which routers are joined together by the segment. Type 2 LSAs are flooded across their own area only. The link-state ID of the type 2 LSA is the IP interface address of the DR.

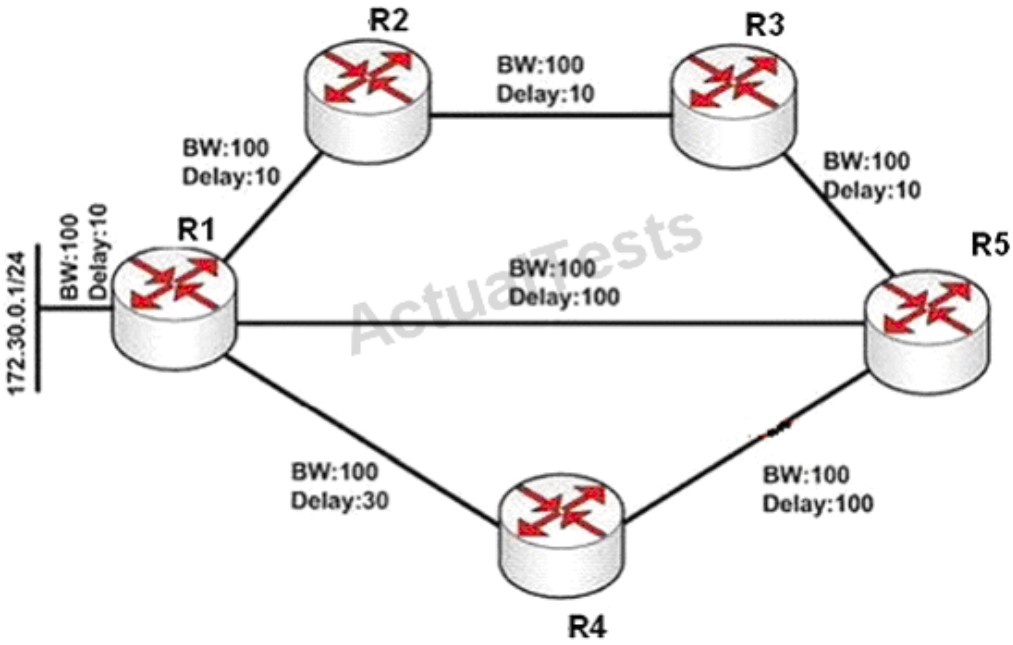
Type 3 - Summary LSA - an Area Border Router (ABR) takes information it has learned on one of its attached areas and it can summarize it (but not by default) before sending it out on other areas it is connected to. This summarization helps provide scalability by removing detailed topology information for other areas, because their routing information is summarized into just an address prefix and metric. The summarization process can also be configured to remove a lot of detailed address prefixes and replace them with a single summary prefix, also helping scalability. The link-

state ID is the destination network number for type 3 LSAs.

Reference
http://en.wikipedia.org/wiki/Link-state_advertisement

QUESTION NO: 173

On the basis of the exhibit provided, assuming that EIGRP is the routing protocol, then at R5, what would be the status of each path to 172.30.1.0/24?



A. the path through R3 would be the successor, the path through R1 would be a feasible successor, and the path through R4 would be neither a successor nor a feasible successor
B. not enough information has been given to figure out what the status of each route would be
C. the path through R3 would be the successor, and the paths through R1 and R4 would be feasible successors
D. the path through R1 would be the successor, the path through R3 would be a feasible successor, and the path through R4 would be neither a successor nor feasible successor

Answer: A

Explanation:

I duplicated this information from within GNS3 and I received the following results

R5#show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route

Gateway of last resort is not set
1.0.0.0/32 is subnetted, 1 subnets
D 1.1.1.1 [90/25735680] via 10.1.35.1, 00:03:53, FastEthernet0/0
2.0.0.0/32 is subnetted, 1 subnets
D 2.2.2.2 [90/25733120] via 10.1.35.1, 00:03:58, FastEthernet0/0
4.0.0.0/32 is subnetted, 1 subnets
D 4.4.4.4 [90/25743360] via 10.1.35.1, 00:03:53, FastEthernet0/0
5.0.0.0/32 is subnetted, 1 subnets
C 5.5.5.5 is directly connected, Loopback0
172.30.0.0/24 is subnetted, 1 subnets
D 172.30.0.0 [90/25610240] via 10.1.35.1, 00:03:53, FastEthernet0/0
10.0.0.0/30 is subnetted, 6 subnets
C 10.1.15.0 is directly connected, FastEthernet1/0
D 10.1.14.0 [90/25615360] via 10.1.35.1, 00:03:55, FastEthernet0/0
D 10.1.12.0 [90/25607680] via 10.1.35.1, 00:04:03, FastEthernet0/0
D 10.1.23.0 [90/25605120] via 10.1.35.1, 00:04:03, FastEthernet0/0
C 10.1.45.0 is directly connected, FastEthernet0/1
C 10.1.35.0 is directly connected, FastEthernet0/0

R5#show ip eigrp top
IP-EIGRP Topology Table for AS(100)/ID(5.5.5.5)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
r - reply Status, s - sia Status

P 1.1.1.1/32, 1 successors, FD is 25735680
via 10.1.35.1 (25735680/25733120), FastEthernet0/0
via 10.1.15.1 (25753600/128256), FastEthernet1/0
P 2.2.2.2/32, 1 successors, FD is 25733120
via 10.1.35.1 (25733120/25730560), FastEthernet0/0
via 10.1.15.1 (25756160/25730560), FastEthernet1/0
P 4.4.4.4/32, 1 successors, FD is 25743360
via 10.1.35.1 (25743360/25740800), FastEthernet0/0
via 10.1.45.1 (25753600/128256), FastEthernet0/1
via 10.1.15.1 (25761280/25735680), FastEthernet1/0
P 5.5.5.5/32, 1 successors, FD is 128256

via Connected, Loopback0
P 10.1.15.0/30, 1 successors, FD is 25625600
via Connected, FastEthernet1/0
P 10.1.14.0/30, 1 successors, FD is 25615360
via 10.1.35.1 (25615360/25612800), FastEthernet0/0
via 10.1.15.1 (25633280/25607680), FastEthernet1/0
via 10.1.45.1 (25633280/25607680), FastEthernet0/1
P 10.1.12.0/30, 1 successors, FD is 25607680
via 10.1.35.1 (25607680/25605120), FastEthernet0/0
via 10.1.15.1 (25628160/25602560), FastEthernet1/0
P 10.1.23.0/30, 1 successors, FD is 25605120
via 10.1.35.1 (25605120/25602560), FastEthernet0/0
P 10.1.45.0/30, 1 successors, FD is 25625600
via Connected, FastEthernet0/1
P 10.1.35.0/30, 1 successors, FD is 25602560
via Connected, FastEthernet0/0
P 172.30.0.0/24, 1 successors, FD is 25610240
via 10.1.35.1 (25610240/25607680), FastEthernet0/0
via 10.1.15.1 (25628160/28160), FastEthernet1/0

Relevant portion of R5 Configuration
!
hostname R5
!
interface Loopback0
ip address 5.5.5.5 255.255.255.255
!
interface FastEthernet0/0
description ##Connection_To_R3##
bandwidth 100
ip address 10.1.35.2 255.255.255.252
delay 10
duplex auto
speed auto
!
interface FastEthernet0/1
description ##Connection_To_R4##
bandwidth 100
ip address 10.1.45.2 255.255.255.252
delay 100
duplex auto

```
speed auto
!
interface FastEthernet1/0
description ##Connection_To_R1##
bandwidth 100
ip address 10.1.15.2 255.255.255.252
delay 100
duplex auto
speed auto
!
router eigrp 100
network 5.5.5.5 0.0.0.0
network 10.1.15.0 0.0.0.3
network 10.1.35.0 0.0.0.3
network 10.1.45.0 0.0.0.3
no auto-summary
```

Successor
A successor for a particular destination is a next hop router that satisfies these two conditions:
It provides the least distance to that destination
It is guaranteed not to be a part of some routing loop

The first condition can be satisfied by comparing metrics from all neighboring routers that advertise that particular destination, increasing the metrics by the cost of the link to that respective neighbor, and selecting the neighbor that yields the least total distance. The second condition can be satisfied by testing a so-called Feasibility Condition for every neighbor advertising that destination. There can be multiple successors for a destination, depending on the actual topology. The successors for a destination are recorded in the topology table and afterwards they are used to populate the routing table as next-hops for that destination.

Feasible Successor
A feasible successor for a particular destination is a next hop router that satisfies this condition:
It is guaranteed not to be a part of some routing loop
This condition is also verified by testing the Feasibility Condition.
Thus, every successor is also a feasible successor. However, in most references about EIGRP the term "feasible successor" is used to denote only those routers which provide a loop-free path but which are not successors (i.e. they do not provide the least distance). From this point of view, for a reachable destination there is always at least one successor, however, there might not be any feasible successors.
A feasible successor provides a working route to the same destination, although with a higher distance. At any time, a router can send a packet to a destination marked "Passive" through any of

its successors or feasible successors without alerting them in the first place, and this packet will be delivered properly. Feasible successors are also recorded in the topology table.
The feasible successor effectively provides a backup route in the case that existing successors die. Also, when performing unequal-cost load-balancing (balancing the network traffic in inverse proportion to the cost of the routes), the feasible successors are used as next hops in the routing table for the load-balanced destination.

By default, the total count of successors and feasible successors for a destination stored in the routing table is limited to four. This limit can be changed in the range from 1 to 6. In more recent versions of Cisco IOS (e.g.
12.4), this range is between 1 and 16.
EIGRP Composite and Vector metrics
Bandwidth
Minimum Bandwidth (in kilobits per second) along the path from router to destination network
Load
Load (number in range 1 to 255; 255 being saturated)
Delay
Total Delay (in 10s of microseconds) along the path from router to destination network
Reliability
Reliability (number in range 1 to 255; 255 being the most reliable)

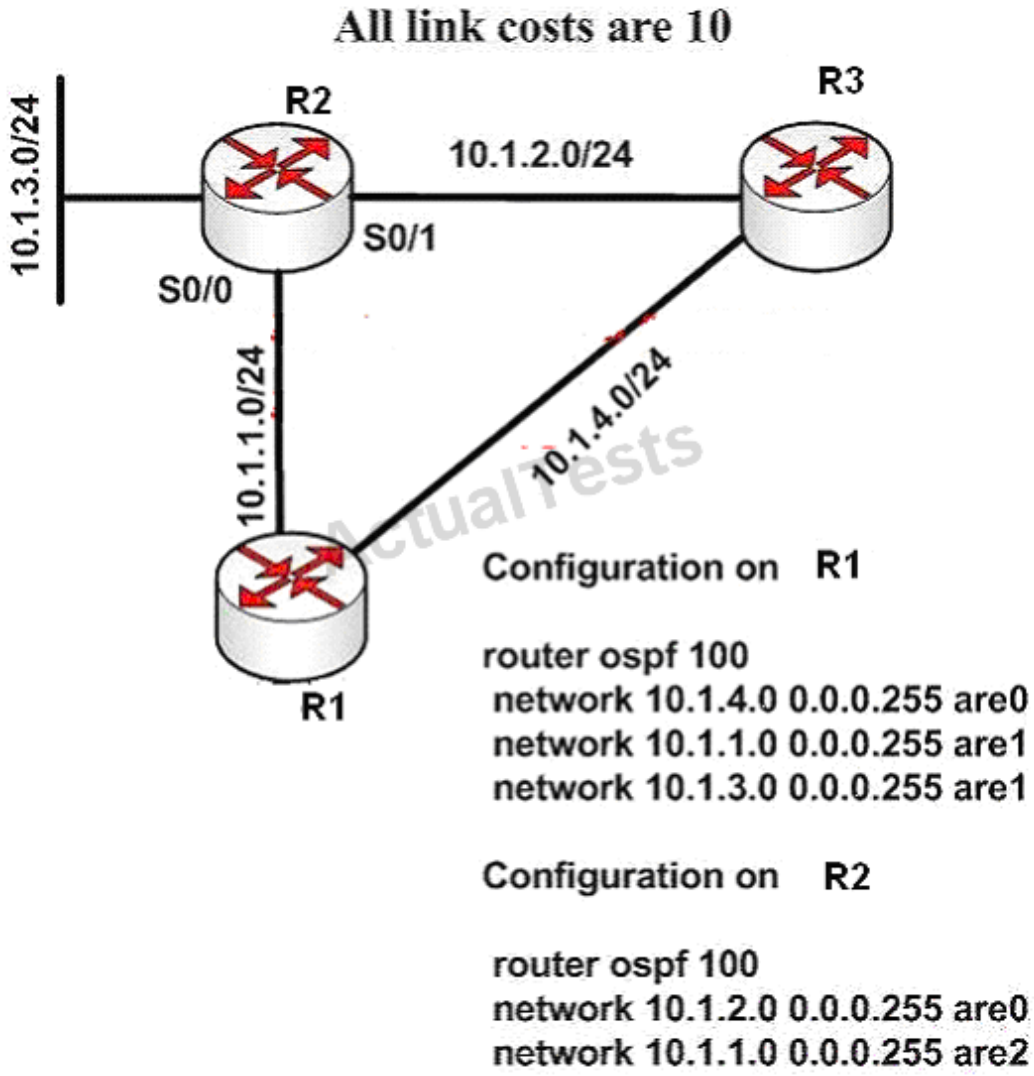
MTU
Minimum path Maximum Transmission Unit (MTU) (never used in the metric calculation)

Hop Count
Number of routers a packet passes through when routing to a remote network, used to limit the EIGRP AS.
The K Values There are five (5) K values used in the Composite metric calculation - K1 through K5. The K values only act as multipliers or modifiers in the composite metric calculation. K1 is not equal to Bandwidth, etc.
By default, only total delay and minimum bandwidth are considered when EIGRP is started on a router, but an administrator can enable or disable all the K values as needed to consider the other Vector metrics.

Reference
http://en.wikipedia.org/wiki/EIGRP#Successor

QUESTION NO: 174

Based on the network provided in the exhibit, how to route the traffic arriving at R1 for 10.1.3.1?



A. through R3, since the path through R3 is through the backbone
B. through R2, because that is the only path available; no neighbor adjacency will be built between R1 and R2
C. through R1, since the path through R1 has the lowest hop count
D. through R1, since the path through R1 has the lowest total metric (10+10=20, versus 10+10+10=30 through R3)

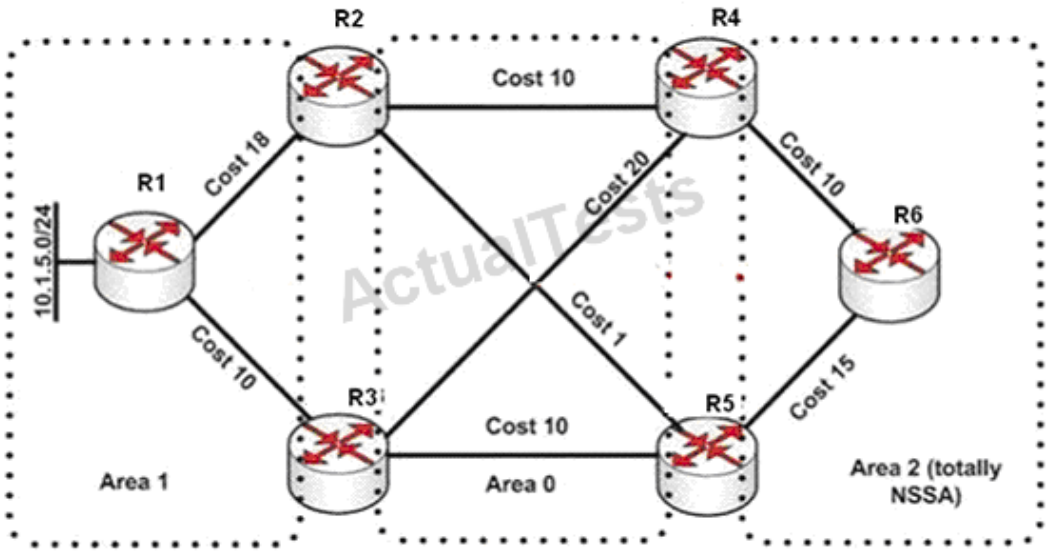
Answer: A

Explanation:

R1 & R2 will not form a OSPF neighbor relationship because they are in two different areas. R2 s0/0(10.1.1.2/24) is in Area 2 and R1 s0/0 is in area 1 (10.1.1.1/24). Therefore traffic will pass from R1 to R3 to R2 when going to the 10.1.3.0/24 network.

QUESTION NO: 175

On the basis of the network provided in the exhibit, R1 is injecting 10.1.5.0/24 by use of a network statement as a network (type 2) LSA. What LSAs will R6 have in its local database for 10.1.5.0/24?



A. R6 will not have any LSAs containing 10.1.5.0/24
B. R6 will have a summary (type 3) LSA containing 10.1.5.0/24, generated by R4 or R5, the Area Border Routers for Area 2, its local area
C. R6 will have the network (type 2) LSA generated by R1 in Area 1 containing 10.1.5.0/24
D. R6 will have a summary (type 3) LSA containing 10.1.5.0/24, generated by either R2 or R3, the Area Border Routers for area 1.

Answer: A

Explanation:

R6 will not have an LSA containing 10.1.5.0/2 because R6 is a total stub area.

QUESTION NO: 176

Which types of prefixes will a router running BGP most likely advertise to an IBGP peer, assuming it is not configured as a route reflector?

A. prefixes received from any other BGP peer and prefixes locally originated via network statements or redistributed
B. all prefixes in its routing table
C. prefixes received from EBGP peers and prefixes locally originated via network statements or redistributed

D. prefixes received from EBGP peers and prefixes received from route reflectors
E. prefixes received from other IBGP peers, prefixes received from EBGP peers, and prefixes redistributed to BGP
F. prefixes received from other IBGP peers and prefixes received from route reflectors

Answer: C

Explanation: Explanation
If your autonomous system will be passing traffic through it from another autonomous system to a third autonomous system, it is very important that your autonomous system be consistent about the routes that it advertises. For example, if your BGP were to advertise a route before all routers in your network had learned about the route through your IGP, your autonomous system could receive traffic that some routers cannot yet route. To prevent this from happening, BGP must wait until the IGP has propagated routing information across your autonomous system. This causes BGP to be synchronized with the IGP. Synchronization is enabled by default

QUESTION NO: 177

Which mechanism can you use to achieve sub-second failover for link failure detection when a switched Ethernet media is used and loss of signal is not supported by the link provider?

A. OSPF standard hellos
B. Cisco Discovery Protocol link detection
C. Bidirectional Forwarding Detection
D. Fast Link Pulse
E. autonegotiation

Answer: C

Explanation:
BFD is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. In addition to fast forwarding path failure detection, BFD provides a consistent failure detection method for network administrators. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning will be easier, and reconvergence time will be consistent and predictable

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

QUESTION NO: 178

You have a router running BGP for the MPLS network and OSPF for the local LAN network at the sales office. A route is being learned from the MPLS network that also exists on the OSPF local network. It is important that the router chooses the local LAN route being learned from the downstream switch running OSPF rather than the upstream BGP neighbor. Also, if the local OSPF route goes away, the BGP route needs to be used. What should be configured to make sure that the router will choose the LAN network as the preferred path?

A. static route needs to be added
B. floating static route needs to be added
C. bgp backdoor command
D. ospf backdoor command

Answer: C

Explanation:

We often run into situations where we have two different routes to a network and we know one is faster than the other. However, a router will only keep one route in the routing table and that will be the one with the shortest administrative distance. Now if two routers (R1 and R2) in an AS are directly connected to each other via a 10meg ethernet running EIGRP (Admin distance 90) but they also have a T1 connection to a service provider running EBGP (Admin distance 20). Now, as human beings, we are smart enough to see that if we are R1 and have to reach a network located behind R2, we should use the EIGRP route via 10meg connection.
Now think of this from a routers perspective, R1 sees two routes to that network, via EIGRP and EBGP. R1 is going to take the slower path via EBGP because EBGPs admin distance is 20. Changing the default administrative distances is not recommended because that may lead to routing loops. So this is where we can use the network backdoor command so router will prefer EIGRP route over EBGP route. BGP treats the network specified by the network backdoor command as a locally assigned network but it doesn't advertise the network in BGP updates. In short, BGP network backdoor command changes the admin distance of that network to 200 thus making it a worst route compared to EIGRP (90). I think this is a great feature that I never utilized before. This actually makes sense and is a better way of doing things than changing the admin distance. Here are a few steps to accomplish this task:

•enable
•configure terminal
•router bgp autonomous-system-number
•neighbor ip-address remote-as autonomous-system-number
•network ip-address backdoor
•end

Example with the backdoor command being used:
R2#sh ip route

Codes: C connected, S static, R RIP, M mobile, B BGP
D EIGRP, EX EIGRP external, O OSPF, IA OSPF inter area
N1 OSPF NSSA external type 1, N2 OSPF NSSA external type 2 E1 OSPF external type 1, E2 OSPF external type 2
i IS-IS, su IS-IS summary, L1 IS-IS level-1, L2 IS-IS level-2 ia IS-IS inter area, * candidate default, U per-user static route o ODR, P periodic downloaded static route
Gateway of last resort is not set
B 1.0.0.0/8 [20/0] via 10.1.12.1, 00:04:42
C 2.0.0.0/8 is directly connected, Loopback0
B 3.0.0.0/8 [20/0] via 10.1.12.1, 00:00:48
10.0.0.0/24 is subnetted, 2 subnets
R 10.1.13.0 [120/1] via 10.1.12.1, 00:00:13, Serial0/0.21 C 10.1.12.0 is directly connected, Serial0/0.21
150.1.0.0/24 is subnetted, 3 subnets
C 150.1.23.0 is directly connected, FastEthernet0/0
D 150.1.3.0 [90/156160] via 150.1.23.3, 00:00:07, FastEthernet0/0 C 150.1.2.0 is directly connected,
Loopback1

Example with link between two routers down:
R2#sh ip route
Codes: C connected, S static, R RIP, M mobile, B BGP
D EIGRP, EX EIGRP external, O OSPF, IA OSPF inter area
N1 OSPF NSSA external type 1, N2 OSPF NSSA external type 2 E1 OSPF external type 1, E2 OSPF external
type 2
i IS-IS, su IS-IS summary, L1 IS-IS level-1, L2 IS-IS level-2 ia IS-IS inter area, * candidate default, U per-user
static route o ODR, P periodic downloaded static route
Gateway of last resort is not set
B 1.0.0.0/8 [20/0] via 10.1.12.1, 00:20:41
C 2.0.0.0/8 is directly connected, Loopback0
B 3.0.0.0/8 [20/0] via 10.1.12.1, 00:16:46
10.0.0.0/24 is subnetted, 2 subnets
R 10.1.13.0 [120/1] via 10.1.12.1, 00:00:04, Serial0/0.21 C 10.1.12.0 is directly connected, Serial0/0.21
150.1.0.0/24 is subnetted, 3 subnets
C 150.1.23.0 is directly connected, FastEthernet0/0
B 150.1.3.0 [200/0] via 10.1.12.1, 00:00:16
C 150.1.2.0 is directly connected, Loopback1

Note: Admin distance of 20 for the eBGP route.

Conditional BGP Route Advertisement:
BGP has a neat feature where you can control what routes to advertise to a certain neighbor. So lets say you are in the middle of two ASs and are passing routes between them. You can use conditional advertisement to say if a network connected to you is down; don't advertise certain network coming in from the other router. It basically revolves around the command advertise-map.

Cisco has some nice summary steps to explain the process:
1. Enable
2. configure terminal
3. router bgp autonomous-system-number
4. neighbor {ip-address | peer-group-name} remote-as autonomous-system-number
5. neighbor ip-address advertise-map map-name {exist-map map-name | non-exist-map map-name}
6. exit
7. route-map map-tag [permit | deny] [sequence-number]
8. Match ip address {access-list-number [access-list-number... | access-list-name...] | access-list-name [access-list-number... | access-list-name] | prefix-list prefix-list-name [prefix-list-name...]}
9. Repeat Steps 7 and 8 for every prefix to be tracked.
10. Exit
11. access-list access-list-number {deny | permit} source [source-wildcard] [log]
12. Repeat Step 11 for every access list to be created.
13. exit

The route map associated with the exist map or nonexist map specifies the prefix that the BGP speaker will track. The route map associated with the advertise map specifies the prefix that will be advertised to the specified neighbor when the condition is met.

QUESTION NO: 179

In BGP routing, what does the rule of synchronization mean?

A. A BGP router can only advertise an EBGP learned route, provided that the route is an IGP route in the routing table.
B. A BGP router can only advertise an IBGP learned route, provided that the route is an IGP route in the routing table.
C. A BGP router can only advertise an IBGP learned route, provided that the route is an IGP route that is not in the routing table.

D. A BGP router can only advertise an EBGP learned route, provided that the route is a metric of 0 in the BGP table.

Answer: B

Explanation: Explanation

When an AS provides transit service to other ASs and if there are non-BGP routers in the AS, transit traffic might be dropped if the intermediate non-BGP routers have not learned routes for that traffic via an IGP. The BGP synchronization rule states that if an AS provides transit service to another AS, BGP should not advertise a route until all of the routers within the AS have learned about the route via an IGP. The topology shown in demonstrates the synchronization rule

QUESTION NO: 180

Router 1 is configured for BGP as dual-homed on the Cisco network. Which three BGP attributes are carried in every BGP update on this router (both IBGP and EBGP)? (Choose three.)

A. origin
B. router-ID
C. AS-path
D. local-preference
E. next-hop

Answer: A,C,E

Explanation: Explanation

There are basically two major types of attribute:
Well Known.
Optional
Well Known:
Well known attributes are must be recognized by each compliant of BGP implementations. Well known attributes are propagated to other neighbors also. Further divided into:
1. Mandatory: It is BGP well known attributes. Mandatory attributes are must be present in all update message passed between BGP peers. It is present in route description. Must be supported and propagated.
2. Discretionary: It is BGP well known attributes. Discretionary attributes may be present on update message.
Must be supported; propagation optional.
Optional:
Optional attributes are recognized by some implementation of BGP & expected that not recognized by everyone. Optional attributes are propagated to their neighbors based on the meanings.

Further divided into:
1. Transitive: Optional transitive attributes don't have to be supported, but must be passed onto peers. Marked as partial if unsupported by neighbor
2. Non Transitive: Optional non-transitive attributes don't have to be supported, and can be ignored. Deleted if unsupported by neighbor
BGP attributes:
1. Weight (Attribute Type Mandatory):
Weight is a Cisco-defined attribute that is local to a router. The weight attribute is not advertised to neighboring routers. If the router learns about more than one route to the same destination, the route with the highest weight is preferred.
2. Local preference (Attribute Type Discretionary):
The local preference attribute is used to prefer an exit point from the local autonomous system. Unlike the weight attribute, the local preference attribute is propagated throughout the local AS. If there are multiple exit points from the AS, the local preference attribute is used to select the exit point for a specific route.
3. AS path (Attribute Type Mandatory):
When a route advertisement passes through an autonomous system, the AS number is added to an ordered list of AS numbers that the route advertisement has traversed.
4. Origin:
The origin attribute indicates how BGP learned about a particular route. The origin attribute can have one of three possible values:
a. IGP The route is interior to the originating AS. This value is set when the network router configuration command is used to inject the route into BGP.
b. EGP -The route is learned via the Exterior Gateway Protocol (EGP).
c. Incomplete
The origin of the route is unknown or is learned some other way. An origin of Incomplete occurs when a route is redistributed into BGP.
5. Multi-exit discriminator (Attribute Type - Non Transitive):
The multi-exit discriminator (MED) or metric attribute is used as a suggestion to an external AS regarding the preferred route into the AS that is advertising the metric.
6. Next-hop (Attribute Type Mandatory):
The EBGP next-hop attribute is the IP address that is used to reach the advertising router. For EBGP peers, the next-hop address is the IP address of the connection between the peers.
7. Community (Attribute Type - Transitive):
The community attribute provides a way of grouping destinations, called communities, to which routing decisions (such as acceptance, preference, and redistribution) can be applied. Route maps are used to set the community attribute. The predefined community attributes are as follows:
a. No-export: Do not advertise this route to EBGP peers.
b. No-advertise: Do not advertise this route to any peer.
c. Internet: Advertise this route to the Internet community; all routers in the network belong to it.
8. Atomic Aggregate (Attribute Type - Discretionary):

Notes that route summarization has been performed.
9. Aggregator (Attribute Type - Transitive):
Identifies the router and AS where summarization was performed.
10. Originator ID (Attribute Type - Non Transitive): Identifies a route reflector.
11. Cluster List (Attribute Type - Non Transitive):
Records the route reflector clusters the route has traversed.

QUESTION NO: 181

In your Cisco EIGRP network, you notice that the neighbor relationship between two of your routers was recently restarted. Which two of these choices could have made this occur? (Choose two.)

A. An update packet with init flag set from a known, already established neighbor relationship was received by one of the routers.
B. The ARP cache was cleared.
C. The counters were cleared.
D. The IP EIGRP neighbor relationship was cleared manually.

Answer: A,D
Explanation:
The following are the most common causes of problems with EIGRP neighbor relationships:
Unidirectional link
Uncommon subnet, primary, and secondary address mismatch
Mismatched masks
K value mismatches
Mismatched AS numbers
Stuck in active
Layer 2 problem
Access list denying multicast packets
Manual change (summary router, metric change, route filter)
According till Ivan Pepelnjak's book "EIGRP Network Design Solutions" the Init flag is set in the initial update packet when to neighbors discover each other and start their initial topology table exchange. There are two basic purposes for the Init flag. First, it's a part of the three way handshake that eigrp uses when building an adjacency.
5. Router B comes up on a wire.
6. Router A receives Router B's hello, and places it in "pending" state. This is a not completely formed adjacency; as long as B is in this state, A won't send any routing information to it.
7. Router A sends an empty unicast update with the Init bit set.

8. Router B receives this update with the Init bit set, and places Router A in the "pending" state.
9. Router B now transmits an empty update with the Init bit set, unicast, to A. This empty update also contains the acknowledgement for Router A's Init update (that this ack is piggybacked is an integral part of the three way handshake process).
10.Router A, on receiving this Init update, places Router B in the "neighbor" state, and sends an acknowledgement for the Init update from Router B.
11.Router B receives this ack, and places A in "neighbor" state. The two routers can now exchange routing information, knowing they have full two way connectivity between them.
The second use of the Init bit is more esoteric. Suppose you have Routers A and B, running along fine, for many hours. Router A reloads, but comes back up before Router B's hold timer has expired. When Router B sees A's hellos, it will assume that A just missed a couple, and everything is fine. But everything isn't fine-A just lost all of its routing information! How can A signal this state, and as B to resynchronize? A can send an empty update, with the Init bit set. This causes Router B to place A in the "pending" state, and wipe out all the information it's learned from A (unless, of course, graceful restart is configured/etc.).

QUESTION NO: 182

Your Cisco network currently runs OSPF and you have a need to policy-route some specific traffic, regardless of what the routing table shows. Which one of these options would enable you to policy-route the traffic?

A. source IP address and the protocol (such as SSL, HTTPS, SSH)
B. the packet Time to Live and the source IP address
C. type of service header and DSCP value
D. destination IP address

Answer: A
Explanation: Explanation
Policy-based routing (PBR) provides a mechanism for expressing and implementing forwarding/routing of data packets based on the policies defined by the network administrators. It provides a more flexible mechanism for routing packets through routers, complementing the existing mechanism provided by routing protocols. Routers forward packets to the destination addresses based on information from static routes or dynamic routing protocols such as Routing Information Protocol (RIP), Open Shortest Path First (OSPF), or Enhanced Interior Gateway Routing Protocol (Enhanced IGRP). Instead of routing by the destination address, policybased routing allows network administrators to determine and implement routing policies to allow or deny paths based on the following:

•Identity of a particular end system
•Application
•Protocol
•Size of packets

QUESTION NO: 183

You use OSPF as your network routing protocol. You use the command show ip route and you see several routes described as 0, 0 IA, 0 E1, and 0 E2. What routes are in your area?

A. OIA
B. OE1
C. O E2
D. 0

Answer: D
Explanation: Explanation
Depending on the point where a network is sourced, there are various types of routes that could be present in an OSPF domain. When there are multiple routes to a particular network in a OSPF domain, the type of the route influences the route that is selected and installed by the router in the routing table. In OSPF, routes that are learned by a router from OSPF sources within the same area are known as intra-area routes. Routes that originate from an OSPF router in a different area are considered as inter-area routes. Certain networks could belong to a domain outside OSPF, which could then be redistributed into the OSPF by an Autonomous System
Boundary Router (ASBR). Such routes are considered external routes. They can be further divided into external type-1 or external type-2 routes, depending on how they are advertised while being redistributing on the ASBR. The difference between these two types is the way in which the metric for the route is calculated.

OSPF-running routers use these criteria to select the best route to be installed in the routing table:
1. Intra-area routes.
2. Inter-area routes.
3. External Type-1 routes.
4. External Type-2 routes.
a. If there are multiple routes to a network with the same route type, the OSPF metric calculated as cost based on the bandwidth is used for selecting the best route. The route with the lowest value for cost is chosen as the best route.
b. If there are multiple routes to a network with the same route type and cost, it chooses all the routes to be installed in the routing table, and the router does equal cost load balancing across

multiple paths.

QUESTION NO: 184

What are the mandatory, well-known BGP attributes?

A. origin, AS-path, next-hop
B. AS-path, origin, MED
C. AS-path, origin, weight
D. AS-path, weight, MED

Answer: A
Explanation:
BGP Path Attributes
Mandatory Well-Known Attributes
Origin: Specifies the router's origin
IGP
EGP
Unknown — Route was redistributed

AS-Path: Sequence of AS numbers through which the route is accessible
Next-Hop: IP address of the next-hop router

Discretionary Well-Known Attributes
Local Preference: Used for consistent routing policy with an AS
Atomic Aggregate: Informs the neighbor AS that the originating router aggregated routes

Nontransitive Attributes
Multiexit Discriminator: Used to discriminate between multiple entry points into an AS

Transitive Attributes
Aggregator: IP address and AS of the router that performed aggregation
Community: Used for route tagging

Reference
CCIE Routing and Switching v4.0 Quick Reference

QUESTION NO: 185

The OSPF RFC 3623 Graceful Restart feature allows you to configure IETF NSF in ultivendor networks. When using OSPF Graceful Restart, which mechanism is used to continue forwarding packets during a switchover?

A. Reverse Path Forwarding
B. Hardware-based forwarding
C. UDP forwarding
D. Layer 2 Forwarding

Answer: B

Explanation: Explanation
Graceful Restart Router Operation
Graceful Restart Initiation
The restarting router becomes aware that it should start the graceful restart process when the network administrator issues the appropriate command or when an RP reloads and forces and Redundancy Facility (RF) switchover. The length of the grace period can be set by the network administrator or calculated by the OSPF software of the restarting router. In order to prevent the LSAs from the restarting router from aging out, the grace period should no t exceed an LSA refresh time of 1800 seconds. In preparation for graceful restart, the restarting router must perform the following action before its software can be reloaded: The restarting router must ensure that its forwarding table is updated and will remain in place during the restart. No OSPF shutdown procedures are performed since neighbor routers must act as if the restarting router is still in service. The OSPF software is reloaded on the router (it undergoes graceful restart).
OSPF Processes during Graceful Restart After the router has reloaded; it must modify its OSPF processes until it reestablishes full adjacencies with all former fully adjacent OSPF neighbors. During graceful restart, the restarting router modifies its OSPF processes in the following ways: The restarting router does not originate LSAs with LS types 1, 5, or 7 so that the other routers in the OSPF domain will use the LSAs that the restarting router had originated prior to reloading. The router does not modify or flush any self-originated LSAs.

The restarting router runs its OSPF routing calculations in order to return any OSPF virtual links to operation.
However, the restarting router does not install OSPF routes into the system's forwarding table, and the router relies on the forwarding entries that it had installed prior to undergoing the graceful restart process. If the restarting router determines that is was the Designated Router on a given segment prior to the graceful restart, it will reelect itself.

Graceful Restart Process Exit
The restarting router exits the graceful restart process when one of the following events occur: The router has reestablished all adjacencies. The graceful restart was successful. The router

receives an LSA that is inconsistent with an LSA from the same router prior to the graceful restart. The inconsistency can be mean either that the router does not support the graceful restart feature or that the router has terminated its helper mode for some reason. The graceful restart was unsuccessful. The grace period has expired. The graceful restart was not successful. Once the restarting router has completed the graceful restart process, it returns to normal OSPF operation, reoriginating LSAs based on the current state of the router and updating its forwarding tables based on current link-state database contents. At this time, it flushes the grace-lsa's that it had originated during the initiation of the graceful restart process.

QUESTION NO: 186

You work as a network engineer for the company, you want to configure two BGP speakers to form an EBGP session across a firewall. On the engineer's network, the firewall always permits TCP sessions that are initiated from the inside network (the network attached to the inside interface of the firewall). What prerequisite is there for enabling BGP to run on this network?

A. EBGP multihop will need to be configured for this to work.
B. This should work with normal BGP peering, with no additional configuration on the BGP speakers or the firewall.
C. The BGP protocol port must be opened on the firewall
D. There is no way to make BGP work across a firewall.

Answer: C
Explanation:
If TCP Port 179 is open for BGP than eBGP multihop must also be enabled
Because BGP uses unicast TCP packets on port 179 to communicate with its peers, you can configure PIX1 and PIX2 to allow unicast traffic on TCP port 179. This way, BGP peering can be established between the routers that are connected through the firewall. Redundancy and the desired routing policies can be achieved through the manipulation of the BGP attributes.

The neighbor ebgp-multihop command enables BGP to override the default one hop eBGP limit because it changes the Time to Live (TTL) of eBGP packets from the default value of 1

Reference
ASA/PIX: BGP through ASA Configuration Example
http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a008009487d.shtml

QUESTION NO: 187

Two directly connected routers, R1 and -R2, are both configured for OSPF graceful restart. R2 is able to switch packets in hardware, but R1 is not. If a network administrator logs on to R2 and performs a system reload, which will be the result?

A. Traffic forwarded from R2 to or through -R1 will continue to be forwarded based on the forwarding table state at the time of the reload.
B. R2 will continue to forward traffic to -R1, but -R1 will drop the traffic because its neighbor adjacency with R2 has failed.
C. R2 will continue forwarding traffic to and through R1, but R1 will drop this traffic because it is not capable of maintaining its forwarding state.
D. All the traffic R2 is forwarding to or through R1 will be dropped while OSPF rebuilds its neighbor adjacency and forwarding tables.

Answer: A

Explanation:

Graceful Restart Router Operation

Graceful Restart Initiation

The restarting router becomes aware that it should start the graceful restart process when the network administrator issues the appropriate command or when an RP reloads and forces and Redundancy Facility (RF) switchover. The length of the grace period can be set by the network administrator or calculated by the OSPF software of the restarting router. In order to prevent the LSAs from the restarting router from aging out, the grace period should not exceed an LSA refresh time of 1800 seconds. In preparation for graceful restart, the restarting router must perform the following action before its software can be reloaded: The restarting router
must ensure that its forwarding table is updated and will remain in place during the restart. No OSPF shutdown procedures are performed since neighbor routers must act as if the restarting router is still in service. The OSPF software is reloaded on the router (it undergoes graceful restart).

OSPF Processes during Graceful Restart

After the router has reloaded; it must modify its OSPF processes until it reestablishes full adjacencies with all former fully adjacent OSPF neighbors. During graceful restart, the restarting router modifies its OSPF processes in the following ways:

The restarting router does not originate LSAs with LSA types 1, 5, or 7 so that the other routers in the OSPF domain will use the LSAs that the restarting router had originated prior to reloading. The router does not modify or flush any self-originated LSAs. The restarting router runs its OSPF routing calculations in order to return any OSPF virtual links to operation. However, the restarting router does not install OSPF routes into the system??s forwarding table, and the router relies on the forwarding entries that it had installed prior to undergoing the graceful restart process. If the restarting router determines that is was the Designated Router on a given
segment prior to the graceful restart, it will reelect itself.

Graceful Restart Process Exit
The restarting router exits the graceful restart process when one of the following events occurs: The router has reestablished all adjacencies. The graceful restart was successful. The router receives an LSA that is inconsistent with an LSA from the same router prior to the graceful restart. The inconsistency can mean either that the router does not support the graceful restart feature or that the router has terminated its helper mode for some reason. The graceful restart was unsuccessful. The grace period has expired.

The graceful restart was not successful.
Once the restarting router has completed the graceful restart process, it returns to normal OSPF operation, reoriginating LSAs based on the current state of the router and updating its forwarding tables based on current link-state database contents. At this time, it flushes the grace-lsa's that it had originated during the initiation of the graceful restart process.

QUESTION NO: 188

Which two steps below should you perform on the hub router while configuring EIGRP routing over DMVPN (mGRE tunnel)? (Choose two.)

A. Set the NHRP hold time to match the EIGRP hold time
B. Add the enable eigrp stub command
C. Add the disable eigrp as-member split-horizon command
D. Add the disable eigrp as-member next-hop-self command

Answer: C,D
Explanation:
To stop EIGRP from assign hub as the next hop for all routes if you don't disable split horizon on the tunnel interface you will only see the routes the hub itself is responsible for from a spoke router. When configuring an EIGRP AS you use the tunnel network id as a network you want to participate on with EIGRP.

QUESTION NO: 189

Area Border Router (ABR) is a router located on the border of one or more OSPF areas that connect those areas to the backbone network. An ABR will inject a default route into which two types of areas? (Choose two.)

A. Area 0
B. NSSA
C. Totally stubby
D. Stub

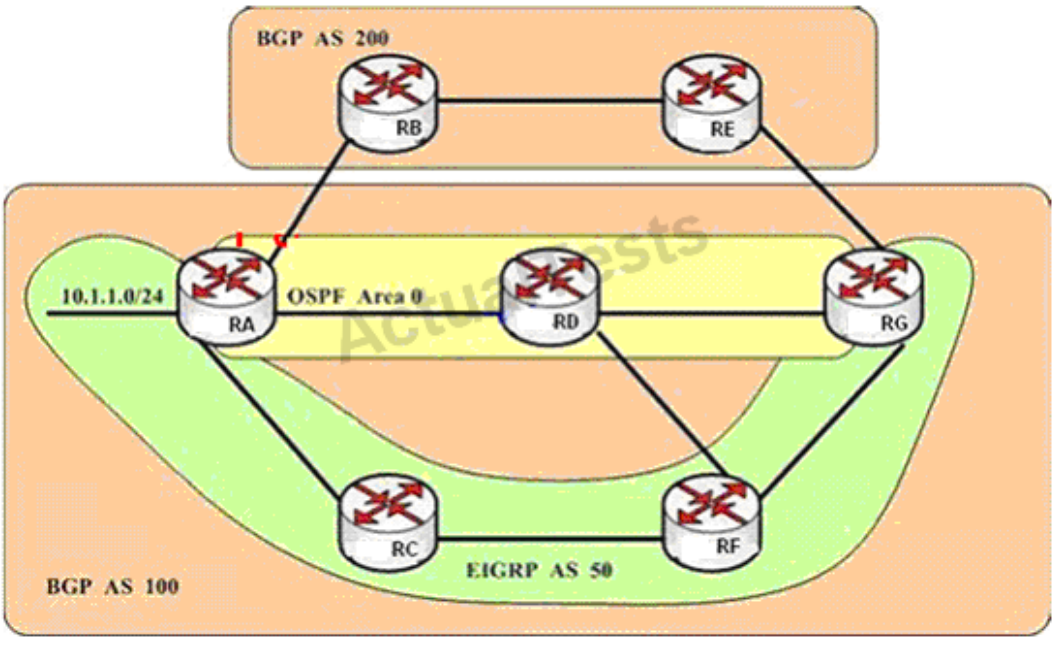
Answer: C,D

Explanation:

Totally Stub Areas: These areas do not allow routes other than intra-area and the default routes to be propagated within the area. The ABR injects a default route into the area and all the routers belonging to this area use the default route to send any traffic outside the area.
Stub Areas: These areas do not accept routes belonging to external autonomous systems (AS); however, these areas have inter-area and intra-area routes. In order to reach the outside networks, the routers in the stub area use a default route which is injected into the area by the Area Border Router (ABR). A stub area is typically configured in situations where the branch office need not know about all the routes to every other office, instead it could use a default route to the central office and get to other places from there. Hence the memory requirements of the leaf node routers is reduced, and so is the size of the OSPF database.

QUESTION NO: 190

Study the exhibit carefully. In this network, if all required configurations are true for routing. Subnet 10.1.1.0/24 is sourced by RA and advertised via BGP, OSPF, and EIGRP. Finally, RG knows this subnet. Which routing protocol and administrative distance can be used by RG to reach subnet 10.1.1.0/24?



A. EIGRP, AD 90
B. EIGRP, AD 170
C. OSPF, AD 110
D. BGP, AD 20

Answer: A

Explanation:

Administrative distance is the feature that routers use in order to select the best path when there are two or more different routes to the same destination from two different routing protocols. Administrative distance defines the reliability of a routing protocol. Each routing protocol is prioritized in order of most to least reliable (believable) with the help of an administrative distance value.

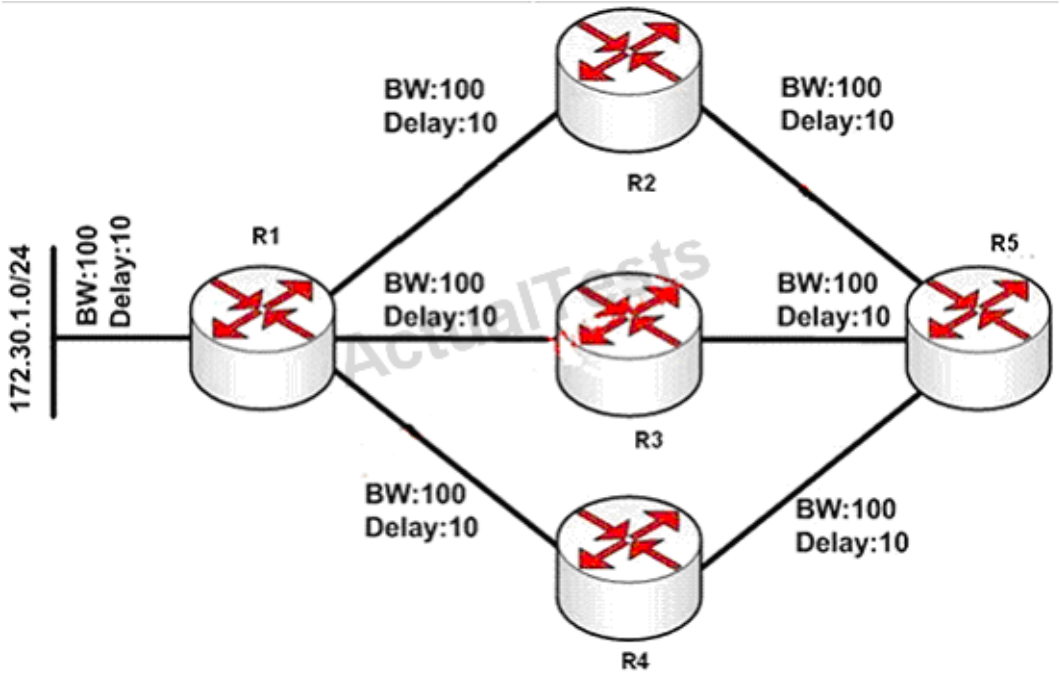
Route Source	Default Distance Values
Connected interface	0
Static route	1
Enhanced Interior Gateway Routing Protocol (EIGRP) summary route	5
External Border Gateway Protocol (BGP)	20
Internal EIGRP	90
IGRP	100
OSPF	110
Intermediate System-to-Intermediate System (IS-IS)	115
Routing Information Protocol (RIP)	120
Exterior Gateway Protocol (EGP)	140
On Demand Routing (ODR)	160
External EIGRP	170
Internal BGP	200
Unknown*	255

ActualTests

Reference
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094195.shtml

QUESTION NO: 191

You are a network technician, study the exhibit carefully. In this EIGRP network, the output of the command show interface for the link between R2 and R5 indicates that the link load varies between 10 and 35. Which K value setting will be used to make sure that this link is not used by EIGRP when the link load reaches 35, but can be used again when the link load drops below 20?



A. Link load is not read in real time, so there is no way to set the K values to make EIGRP choose to use or not use a link based on the link load.
B. There is not enough information in the question to determine the correct answer.
C. Use the K5 setting to include load in EIGRP's metric calculations.
D. Use the K2 setting to include load in EIGRP's metric calculations.

Answer: A

Explanation:

Using the Metrics

When you initially configure EIGRP, remember these two basic rules if you are attempting to influence EIGRP metrics:

The bandwidth should always be set to the real bandwidth of the interface; multipoint serial links and other mismatched media speed situations are the exceptions to this rule.

The delay should always be used to influence EIGRP routing decisions.

Because EIGRP uses the interface bandwidth to determine the rate at which to send packets, it is important that these be set correctly. If it is necessary to influence the path EIGRP chooses, always use delay to do so.

At lower bandwidths, the bandwidth has more influence over the total metric; at higher bandwidths, the delay has more influence over the total metric

Reference
http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094cb7.shtml

QUESTION NO: 192

According to the exhibit provided, what will be the purpose of this route map when applied to traffic passing through a router?



A. take any packet sourced from any address in the 10.2.0.0/16 network or destined to 10.1.14.25 and set the next hop to 10.1.1.1
B. nothing; extended access lists are not allowed in route maps used for policy-based routing
C. take any packet sourced from any address in the 10.2.0.0/16 network and destined to 10.1.14.25 and set the next hop to 10.1.1.1
D. drop any packet sourced from 10.2.0.0/16

Answer: A
Explanation:
In this configuration example, any traffic matching access list 100 will have their next hop set to 10.1.1.1 overriding the normal behavior of the routing table. Access list 100 has two entries, so any traffic matching either will be policy routed.

QUESTION NO: 193

Which reserved 2-byte AS number is seen in the AS path when a BGP router is not capable of understanding 4-byte AS numbers?

A. 00000
B. 12345
C. 23456

D. 65000
E. 99999

Answer: C
Explanation:
When BGP was under development, it was assumed that a 16 bit (2 bytes) number for the Autonomous System (AS) was more than enough (65536 AS Numbers), but the same situation as with IPv4 happened, we are running out of AS numbers “basically” at the same speed as IPv4 addresses. Strategies where put in place to keep the business rolling, like reusing old orphan AS numbers, but something drastically had to be done. The solution was to create a bigger AS number space with 32 bits (4 bytes) that provides a pool of 4,294,967,296 BGP AS numbers.
The problem with this new strategy of a bigger AS space is that Old BGP systems that only support 2 bytes AS number would have problems to deal with this new design. Some kind of backward compatibility implementation was needed. In this post we are going to discuss the details.
The 4 byte AS implementation and compatibility design is built around the following features:
New capability for the 32 bit (4 byte) AS numbers
New optional transitive attributes (AS4_PATH and AS4_AGGREGATOR)
New extended communities
A well known transit AS number “23456 to represent 4 bytes AS numbers to old BGP devices that dont understand the 4 byte as.
Reference: https://blog.initialdraft.com/archives/3985/

Topic 3, Implement IPv6

QUESTION NO: 194

What is the first thing that happens when IPv6 is enabled on an interface on a host?

A. A router solicitation is sent on that interface.
B. There is a duplicate address detection on the host interface.
C. The link local address is assigned on the host interface.
D. A neighbor redirect message is sent on the host interface.

Answer: B
Explanation:
Duplicate address detection (DAD) is used to verify that an IPv6 home address is unique on the LAN before assigning the address to a physical interface (for example, QDIO). z/OS

Communications Server responds to other nodes doing DAD for IP addresses assigned to the interface.

Reference
http://publib.boulder.ibm.com/infocenter/zos/v1r12/index.jsp?topic=%2Fcom.ibm.zos.r12.hale001%2Fipv6d0021002145.htm

QUESTION NO: 195

What is the flooding scope of an OSPFv3 LSA, if the value of the S2 bit is set to 1 and the S1 bit is set to 0?

A. link local
B. area wide
C. AS wide
D. reserved

Answer: C
Explanation:
The Type 1 router LSA is now link local and the Type 2 Network LSA is AS Wide
S2 and S1 indicate the LSA's flooding scope. Table 9-1 shows the possible values of these two bits and the associated flooding scopes.

Table 9-1 S bits in the OSPFv3 LSA Link State Type field and their associated flooding scopes

S2	S1	Flooding Scope
0	0	Link-Local
0	1	Area
1	0	AS (Routing Domain)
1	1	Reserved

ActualTests

LSA Function Code, the last 13 bits of the LS Type field, corresponds to the OSPFv2 Type field. Table 9-2 shows the common LSA types used by OSPFv3 and the values of their corresponding

LS Types. If you decode the hex values, you will see that the default U bit of all of them is 0. The S bits of all LSAs except two indicate area scope. Of the remaining two, AS External LSAs have an AS flooding scope and Link LSAs have a linklocal flooding scope. Most of the OSPFv3 LSAs have functional counterparts in OSPFv2; these OSPFv2 LSAs and their types are also shown in Table 9-2.

Table 9-2 OSPFv3 LSA types and their OSPFv2 counterparts

OSPFv3 LSAs		OSPFv2 LSAs	
LS Type	Name	Type	Name
0x2001	Router LSA	1	Router LSA
0x2002	Network LSA	2	Network LSA
0x2003	Inter-Area Prefix LSA	3	Network Summary LSA
0x2004	Inter-Area Router LSA	4	ASBR Summary LSA
0x4005	AS-External LSA	5	AS-External LSA
0x2006	Group Membership LSA	6	Group Membership LSA
0x2007	Type-7 LSA	7	NSSA External LSA
0x0008	Link LSA		*No Corresponding LSA*
0x2009	Intra-Area Prefix LSA		*No Corresponding LSA*

ActualTests

Reference
http://www.networkworld.com/subnets/cisco/050107-ch9-ospfv3.html?page=1

QUESTION NO: 196

How will EIGRPv6 react if there is an IPv6 subnet mask mismatch between the Global Unicast addresses on a point-to-point link?

A. EIGRPv6 will form a neighbor relationship.
B. EIGRPv6 will not form a neighbor relationship.
C. EIGRPv6 will form a neighbor relationship, but with the log MSG: "EIGRPv6 neighbor not on a common subnet."
D. EIGRPv6 will form a neighbor relationship, but routes learned from that neighbor will not be

installed in the routing table.

Answer: A
Explanation:
http://www.ietf.org/rfc/rfc3587.txt

QUESTION NO: 197

Which two tunneling techniques support IPv6 multicasting? (Choose two.)

A. 6to4
B. 6over4
C. ISATAP
D. 6PE
E. GRE

Answer: B,E
Explanation:
When IPv6 multicast is supported (over a 6to4 tunnel), an IPv6 multicast routing protocol must be used

Restrictions for Implementing IPv6 Multicast
IPv6 multicast for Cisco IOS software uses MLD version 2. This version of MLD is fully backward-compatible with MLD version 1 (described in RFC 2710). Hosts that support only MLD version 1 will interoperate with a router running MLD version 2. Mixed LANs with both MLD version 1 and MLD version 2 hosts are likewise supported.

IPv6 multicast is supported only over IPv4 tunnels in Cisco IOS Release 12.3(2)T, Cisco IOS Release 12.2
(18)S, and Cisco IOS Release 12.0(26)S.
When the bidirectional (bidir) range is used in a network, all routers in that network must be able to understand the bidirectional range in the bootstrap message (BSM).
IPv6 multicast routing is disabled by default when the ipv6 unicast-routing command is configured. On Cisco Catalyst 6500 and Cisco 7600 series routers, the ipv6 multicast-routing also must be enabled in order to use IPv6 unicast routing

Reference
http://www.cisco.com/web/about/ac123/ac147/ac174/ac197/
about_cisco_ipj_archive_article09186a00800c830a.html
http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-multicast.html

https://supportforums.cisco.com/thread/183386

QUESTION NO: 198

Which two OSPF LSA types are new in OSPF version 3? (Choose two.)

A. Link
B. NSSA external
C. Network link
D. Intra-area prefix
E. AS domain

Answer: A,D

Explanation:

New LSA Types

OSPFv3 carries over the seven basic LSA types we're familiar with from OSPFv2. However, the type 1 and 2 LSAs have been re-purposed, as will be discussed in a bit. OSPFv3 also introduces two new LSA types: Link and Intra-area Prefix.

OSPFv3		OSPFv2	
0x2001	Router LSA	1	Router LSA
0x2002	Network LSA	2	Network LSA
0x2003	Inter-area Prefix LSA	3	Network Summary LSA
0x2004	Inter-area Router LSA	4	ASBR Summary LSA
0x4005	AS-External LSA	5	AS-External LSA
0x2006	Group Membership LSA	6	Group Membership LSA
0x2007	Type-7 LSA	7	NSSA External LSA
0x0008	Link LSA		
0x2009	Intra-area Prefix LSA		

Reference
http://packetlife.net/blog/2010/mar/2/ospfv2-versus-ospfv3/

QUESTION NO: 199

In order to maintain security, with which hop count are IPv6 neighbor discovery packets sent?

A. 0
B. 1
C. 255
D. 256

Answer: C
Reference: http://www.txv6tf.org/wp-content/uploads/2010/08/Carter-Tutorial-IPv6-Security-Texas.pdf (see slide 26, note in red)

QUESTION NO: 200

Which command will define a VRF with name 'CCIE' in IPv6?

A. ip vrf CCIE
B. ipv6 vrf CCIE
C. vrf definition CCIE
D. ipv6 vrf definition CCIE

Answer: C
Explanation:

Vrf definition CCIE creates a multiprotocol VRF for both IPv4 and IPv6

QUESTION NO: 201

Which three fields are optional in an OSPFv3 external LSA? (Choose three.)

A. Forwarding Address
B. External Route
C. Reference Link-State ID
D. Option

E. Prefix Options

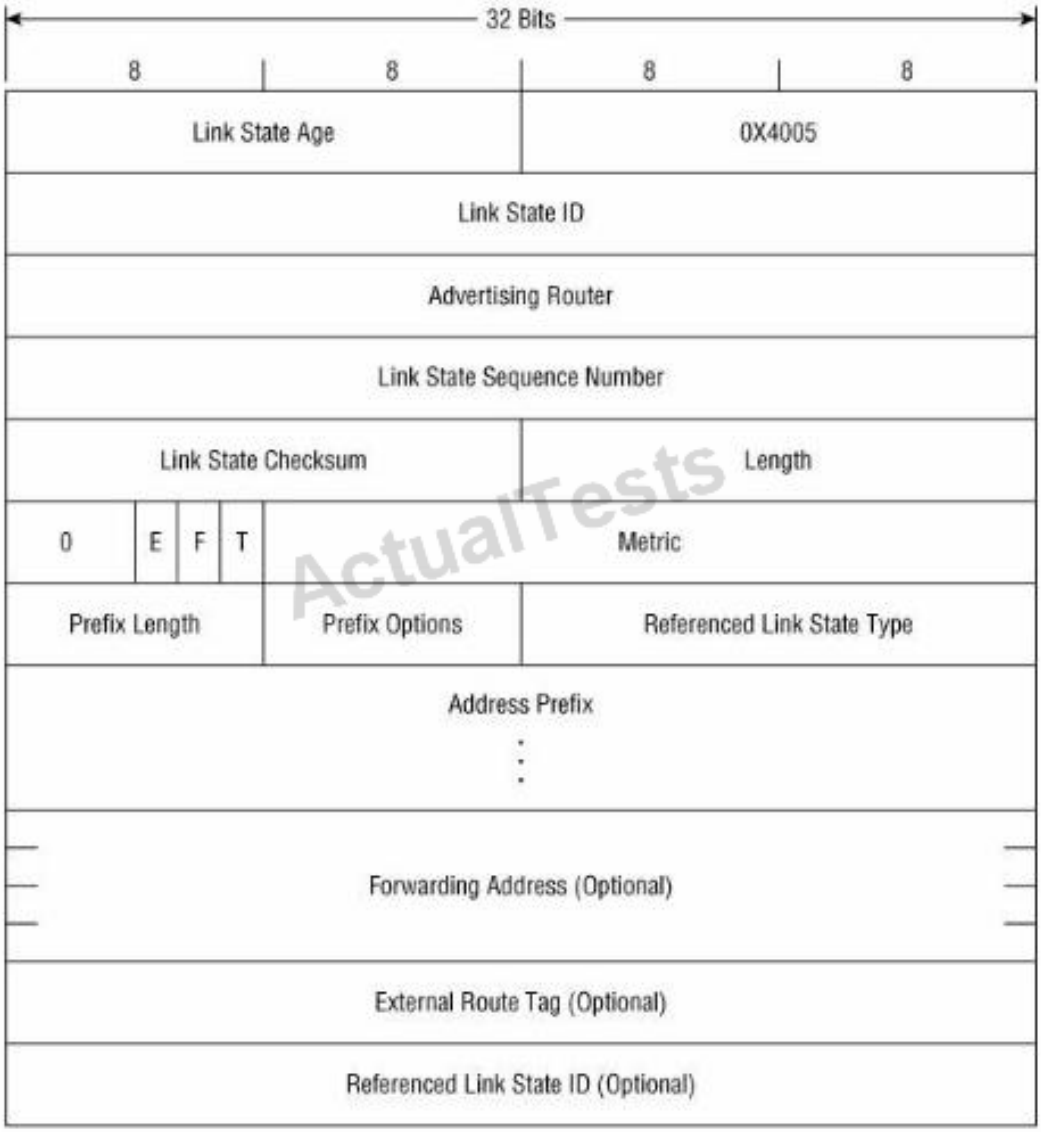
Answer: A,B,C

Explanation:

AS-External LSA

As with OSPFv2, the AS-External LSA advertises prefixes external to the OSPF routing domain; one LSA is required for each external prefix advertised. However, the format of the OSPFv3 As-External LSA (Figure 9-10) is different from its OSPFv2 counterpart.

Figure 9-10. OSPFv3 AS-External LSA



Reference

http://fengnet.com/book/CCIE%20Professional%20Development%20Routing%20TCPIP%20Volume%20I/images/09fig10_alt.jpg

QUESTION NO: 202

On a router, interface S0 is running EIGRPv6, and interface S1 is running OSPFv3. A redistribution command is issued under OSPFv3, redistribute EIGRP 1 metric 20 under ipv6 router ospf 1. What will happen after applying this redistribution command?

A. All routes showing up as D and D EX in the routing table will be redistributed into OSPFv3.
B. All routes showing up as D, D EX, and C in the routing table will be redistributed into OSPFv3.
C. All routes showing up as D and D EX in the routing table and the S0 interface will be redistributed into OSPFv3.
D. All routes showing up as D in the routing table will be redistributed into OSPFv3.
E. All routes showing up as D EX in the routing table will be redistributed into OSPFv3.

Answer: A

Explanation:

D are EIGRP Internal Routes and D EX are EIGRP external routes. Both Internal and External EIGRP routes will be redistributed with the configuration shown above

QUESTION NO: 203

Which address needs to be present on an interface for OSPFv3 to form an adjacency?

A. Global unicast
B. Unique local
C. Link local
D. FF02::5
E. FF02::6

Answer: C

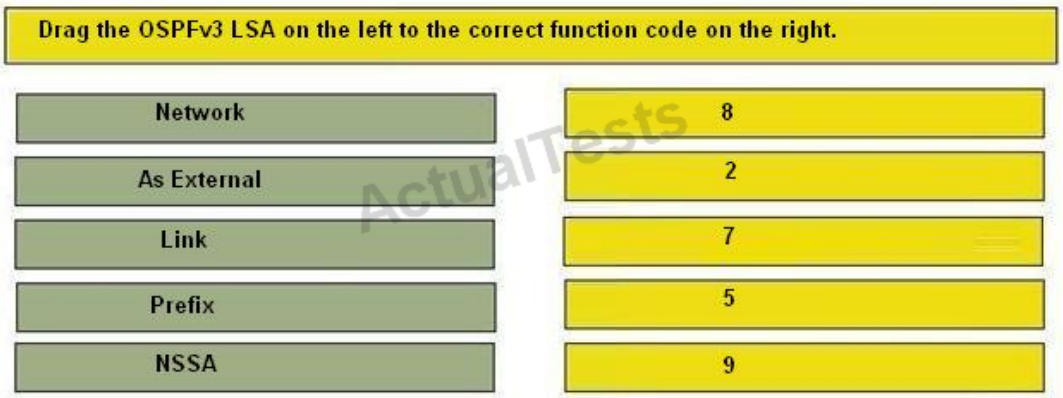
Explanation:

OSPFv3, however, makes use of IPv6's link-local address scope (FE80::/10). All OSPFv3 adjacencies are formed using link-local addresses:

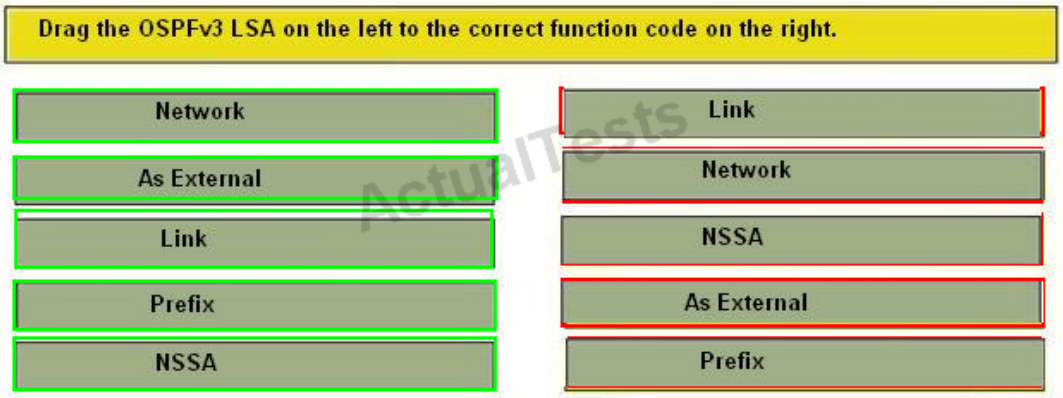
R1# show ipv6 ospf neighbor detail
Neighbor 3.3.3.3
In the area 1 via interface FastEthernet0/0
Neighbor: interface-id 4, link-local address FE80::C003:7DFF:FE07:0
Neighbor priority is 1, State is FULL, 6 state changes

DR is 3.3.3.3 BDR is 2.2.2.2
Options is 0x6670B96D
Dead timer due in 00:00:32
Neighbor is up for 00:16:48
Index 1/1/1, retransmission queue length 0, number of retransmission 0
First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
Last retransmission scan length is 0, maximum is 0
Last retransmission scan time is 0 msec, maximum is 0 msec

QUESTION NO: 204 DRAG DROP



Answer:



Explanation:
Network - 2
AS External – 5
Link – 8
Prefix – 9
NSSA – 7

QUESTION NO: 205

Which mechanism does OSPFv3 use when the router LSA is too big to be sent out?

A. It relies on IPv6 to do the fragmenting.
B. It splits the LS Update packet into smaller packets.
C. It splits the LSA into smaller router LSA packets.
D. It produces an error.

Answer: C
Reference: http://en.wikipedia.org/wiki/User:Wbenton

QUESTION NO: 206

What is the RP for IPv6 multicast address FF76:0130:2001:db8:3811::2323?

A. 2001:db8:3811::1
B. 2001:db8:3811::2323
C. 2001:db8:3811::0130
D. 2001:db8:3811::3
E. 2001:db8:3811::6

Answer: A
Explanation:
Embedded RP defines an address allocation policy in which the address of the RP is encoded in an IPv6 multicast group address. This allows an easy deployment of scalable inter-domain multicast and simplifies the intra-domain multicast configuration as well. IPv6 Multicast group addresses embedded with RP information start with ff70::/12 where the flag value of 7 means embedded RP.

Figure 4. Multicast Address with Embedded RP



There is no need to pre-configure routers with the RP address information. Routers can automatically extract and use the RP information from the IPv6 multicast group address. This allows for a large number of RPs to be deployed anywhere in the Internet. Embedded RP requires no change in protocol operations. It can be considered an automatic replacement for static RP configuration.

The router can learn only one RP address for a multicast group using embedded RP. It cannot support RP redundancy. Proposals are being considered to introduce RP redundancy by mechanisms other than BSR for IPv6 multicast. Embedded RP does not support Bidirectional PIM. Embedded RP allows the application to dictate which router is the RP. There is the possibility that a low-end router could end up becoming the RP for hundreds of high data rate sources if the application defines an erroneous RP address (this can be prevented by disabling Embedded RP learning).
For more information on Embedded RP, see RFC3956.

Reference
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6552/whitepaper_c11-508498.html

QUESTION NO: 207

Refer to the exhibit.

```
Ethernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8BB:CCFF:FE00:6400
  Global unicast address(es):
    2011::1, subnet is 2011::/64 [DUP]
    2011::2, subnet is 2011::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::5
    FF02::6
    FF02::1:FF00:1
    FF02::1:FF00:2
    FF02::1:FF00:6400
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
```

ActualTests

Which statement is correct?

A. One of the IPv6 addresses is already used on another device on the same segment.
B. Multiple addresses on the same subnet and on the same interface are not permitted for IPv6.
C. The MTU is too small for IPv6.
D. Multiple addresses on the same interface are not permitted for IPv6.

Answer: A

Explanation:

As per the exhibit, there is an IPv6 address conflict as another device is using the same IPv6 address on the same segment.

QUESTION NO: 208

Which command can be used to filter a RIPv6 route from getting installed in the routing table?

A. ipv6 router rip ccie
distribute-list..
B. ipv6 router rip ccie
offset-list..
C. interface e0/0
ipv6 rip ccie distribute-list..
D. interface e0/0
ipv6 rip ccie advertise..

Answer: A

Explanation:

Filtering IPv6 RIP Routing Updates

Route filtering using distribute lists provides control over the routes RIP receives and advertises. This control may be exercised globally or per interface.

Filtering is controlled by IPv6 distribute lists. Input distribute lists control route reception, and input filtering is applied to advertisements received from neighbors. Only those routes that pass input filtering will be inserted in the RIP local routing table and become candidates for insertion into the IPv6 routing table.

Output distribute lists control route advertisement; Output filtering is applied to route advertisements sent to neighbors. Only those routes passing output filtering will be advertised.

Global distribute lists (which are distribute lists that do not apply to a specified interface) apply to all interfaces.

If a distribute list specifies an interface, then that distribute list applies only to that interface.

An interface distribute list always takes precedence. For example, for a route received at an interface, with the interface filter set to deny, and the global filter set to permit, the route is blocked, the interface filter is passed, the global filter is blocked, and the route is passed.

IPv6 prefix lists are used to specify certain prefixes or a range of prefixes that must be matched before a permit or deny statement can be applied. Two operand keywords can be used to designate a range of prefix lengths to be matched. A prefix length of less than, or equal to, a value is configured with the le keyword. A prefix length greater than, or equal to, a value is specified using the ge keyword. The ge and le keywords can be used to specify the range of the prefix length to be matched in more detail than the usual ipv6-prefix/prefix-length argument. For a candidate prefix to match against a prefix list entry three conditions can exist:

The candidate prefix must match the specified prefix list and prefix length entry.

The value of the optional le keyword specifies the range of allowed prefix lengths from the prefix-length argument up to, and including, the value of the le keyword.

The value of the optional ge keyword specifies the range of allowed prefix lengths from the value of the ge keyword up to, and including, 128.

Note

The first condition must match before the other conditions take effect.

An exact match is assumed when the ge or le keywords are not specified. If only one keyword operand is specified then the condition for that keyword is applied, and the other condition is not applied. The prefix-length value must be less than the ge value. The ge value must be less than, or equal to, the le value. The le value must be less than or equal to 128.

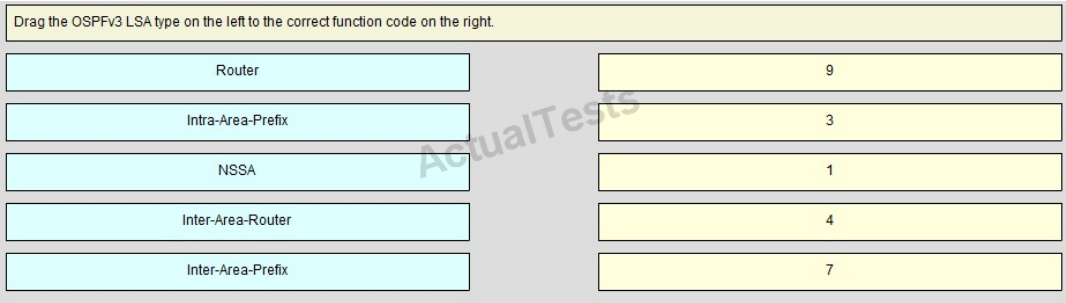
SUMMARY STEPS

1. enable
2. configure terminal

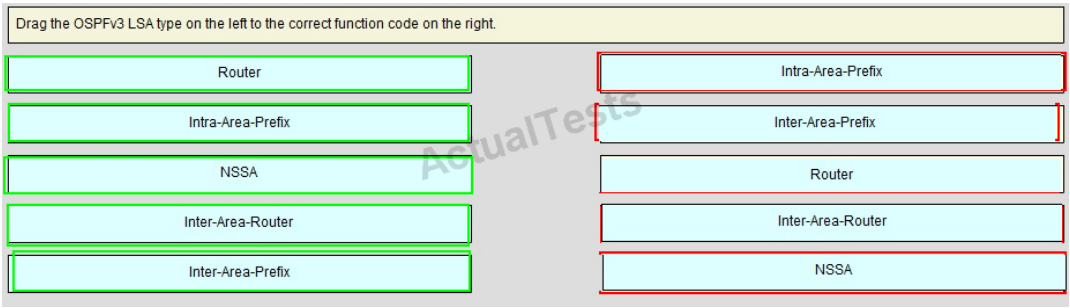
3. ipv6 prefix list prefix-list-name [seq seq-number] {deny ipv6-prefix/prefix-length | description text} [ge gevalue] [le le-value]
4. ipv6 prefix list prefix-list-name [seq seq-number] {permit ipv6-prefix/prefix-length | description text} [ge gevalue] [le le-value]
5. Repeat Steps 3 and 4 as many times as necessary to build the prefix list.
6. ipv6 router rip name
7. distribute-list prefix-list prefix-list-name {in | out} [interface-type interface-number]

Reference
http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-rip.html

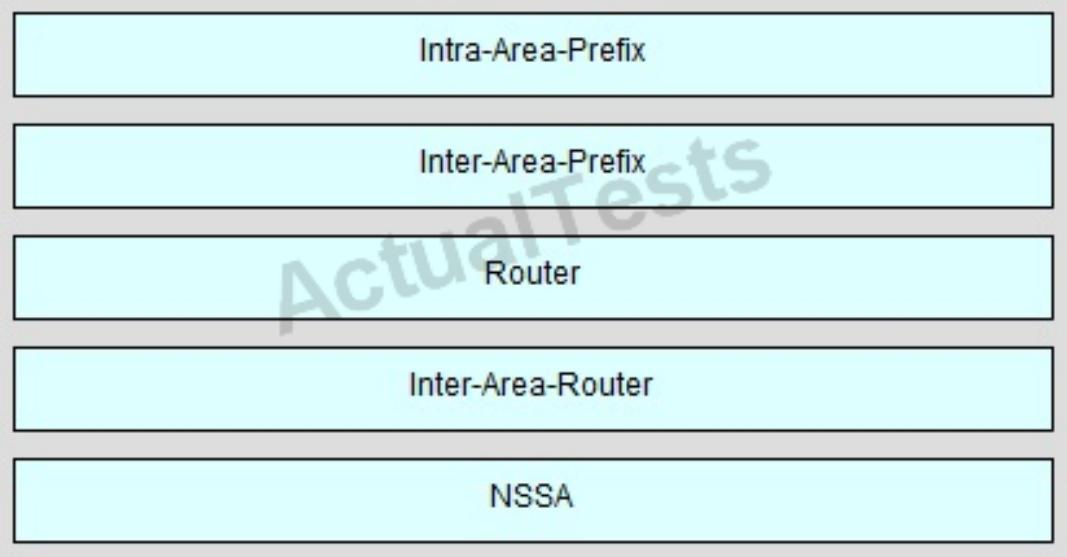
QUESTION NO: 209 DRAG DROP



Answer:



Explanation:



QUESTION NO: 210

You are using IPv6, and would like to configure EIGRPv6. Which three of these correctly describe how you can perform this configuration? (Choose three.)

A. EIGRP for IPv6 is directly configured on the interfaces over which it runs.
B. EIGRP for IPv6 is not configured on the interfaces over which it runs, but if a user uses passive-interface configuration, EIGRP for IPv6 needs to be configured on the interface that is made passive.
C. There is a network statement configuration in EIGRP for IPv6, the same as for IPv4.
D. There is no network statement configuration in EIGRP for IPv6.
E. When a user uses a passive-interface configuration, EIGRP for IPv6 does not need to be configured on the interface that is made passive.
F. When a user uses a non-passive-interface configuration, EIGRP for IPv6 does not need to be configured on the interface that is made passive

Answer: A,D,E
Explanation:

Restrictions for Implementing EIGRP for IPv6:
This section lists ways in which EIGRP for IPv6 differs from EIGRP IPv4 as well as EIGRP for IPv6 restrictions.
•

EIGRP for IPv6 is directly configured on the interfaces over which it runs. This feature allows EIGRP for IPv6 to be configured without the use of a global IPv6 address. There is no network statement in EIGRP for IPv6.

In per-interface configuration at system startup, if EIGRP has been configured on an interface, then the EIGRP protocol may start running before any EIGRP router mode commands have been executed.

•

An EIGRP for IPv6 protocol instance requires a router ID before it can start running.

•

EIGRP for IPv6 has a shutdown feature. The routing process should be in "no shutdown" mode in order to start running.

•

When a user uses passive-interface configuration, EIGRP for IPv6 does not need to be configured on the interface that is made passive.

•

EIGRP for IPv6 provides route filtering using the distribute-list prefix-list command. Use of the routE.map command is not supported for route filtering with a distribute list.

Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00805fc867.html

QUESTION NO: 211

Though many options are supported in EIGRPv6, select two options from the below list that are supported. Choose 2

A. VRF
B. auto-summary
C. per-interface configuration
D. prefix-list support via route-map
E. prefix-list support via distribute-list

Answer: C,E
Explanation:
EIGRPv6 does differ from EIGRPv4 in the following ways:
EIGRPv6 is configured (enabled) directly on Cisco routers interfaces; this means EIGRPv6 can be configured (enabled) on a routers interface, without having to configure (assign) a Global IPv6 address on the interface and without using the network command while the router is in router configuration mode.
Also, when configuring (enabling) EIGRPv6 on a Cisco router, the EIGRP routing process must be configured (assigned) with a router-id (by using the router configuration command router-id); if a router-id is not configured (assigned) the EIGRPv6 routing process will not start.
The EIGRPv6 routing process also uses a shutdown feature; meaning an EIGRPv6 routing

process will not start until the routing process has been placed into no shutdown mode. (by, typing the no shutdown command while the router is in router configuration mode)
Also, on Passive Interfaces; EIGRPv6 is not required to be configured.
Lastly, EIGRPv6 use the router configuration command distribute-list prefix-list to perform route filtering; and when configuring route filtering the route-map command is not supported.
Below is some additional information on EIGRPv6:
IPv6 EIGRP and IPV4 EIGRP are very similar in concept except for the following differences:
IPv6 is configured on interface basis (like OSPFv3 and RIPng) and networks are advertised based on the interface command -> C is correct.

When configured on interface, IPv6 EIGRP is initially placed in "shutdown" state as with OSPFv3, IPv6 EIGRP require a router-id in IPv4 format Passive interfaces can only be configured in the routing process mode. The need for extra memory resources and supported in IOS 12.4(6)T and later.

There is no split horizon in IPv6 because it is possible to get multiple prefixes per interface. Their is no concept of classful routing in IPv6 EIGRP consequently no automatic summary -> B is not correct EIGRPv6 uses the router configuration command "distribute-list prefix-list" to perform route filtering, and when configuring route filtering the "route-map" command is not supported -> E is correct but D is not.
Virtual Routing and Forwarding (VRF) is also supported in EIGRPv6.

QUESTION NO: 212

During the IPv6 address resolution, a node sends a neighbor solicitation message in order to discover which of these?

A. The Layer 2 multicast address of the destination node
B. The solicited node multicast address of the destination node
C. The Layer 2 address of the destination node based on the destination IPv6 address
D. The IPv6 address of the destination node based on the destination Layer 2 address

Answer: C
Explanation:
Nodes send Neighbor Solicitations to request the link-layer address of a target node while also providing their own link-layer address to the target. Neighbor Solicitations are multicast when the node needs to resolve an address and unicast when the node seeks to verify the reachability of a neighbor.

QUESTION NO: 213

Which two of these steps are minimum requirements to configure OSPFv3 under IPv6? (Choose two.

A. Configure a routing process using the command ipv6 router ospf [process-id].
B. Add the network statement for the interfaces on which OSPF will run.
C. Configure OSPF on the interface that it will run on.
D. Use the passive-interface command on the interfaces on which OSPF should not run.
E. Enable routing.

Answer: C,E
Explanation:
The first step to configure OSPFv3 under IPv6 is to enable IPv6 unicast routing:
R1(config)# ipv6 unicast-routing
Also we need to enable the OSPF process:
R1(config)# ipv6 router ospf 1

There are a few changes in configuring OSPFv3 vs OSPF for IPv4. Instead of using the “network” and “area” commands in ospf router configuration mode you now configure OSPFv3 on a per interface basis using the ipv6 ospf area command in interface configuration mode. For example:

R1(config)# interface fa0/0
R1(config-if)# ipv6 ospf 1 area 0
Note: The “network” command does not exist in OSPFv3.

Reference
http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-ospf.html#wp1070061
Note: You will see under the section how to implement ospf for ipv6 it only has 1 REQUIRED thing configure interface and in the comments it says that OSPF IPV6 routing is disabled by default.

QUESTION NO: 214

From which IPv6 address are EIGRPv6 hello messages sourced?

A. Global unicast
B. Link local

C. Site local
D. Unique local
E. FF02::A

Answer: B
Explanation:
Link local address is the address being used in EIGRPv6 as the next hop address in the IPv6 routing table, it is required that the engineer map both the link-local & global addresses to an interface in the EIGRPv6 configuration as well. EIGRPv6-enabled routers will not become adjacent without the link-local mapping, even if the global address mapping exist on the router.

QUESTION NO: 215

In which way can the IPv6 address of 2031:0000:130F:0000:0000:09C0:876A:130B be expressed most efficiently?

A. 2031:0:130F:0:0:09C0:876A:130B
B. 2031::130F::9C0:876A:130B
C. 2031:0:130F::9C0:876A:130B
D. 2031:0:130F:0:0:9C0:876A:130B

Answer: C
Explanation:
IPv6 Addressing Notation
IP addresses change significantly with IPv6. IPv6 addresses are 16 bytes (128 bits) long rather than four bytes (32 bits). This larger size means that IPv6 supports more than 300,000,000,000,000,000,000,000,000,000,000,000,000 possible addresses! As an increasing number of cell phones and other consumer electronics expand their networking capability and require their own addresses, the smaller IPv4 address space will eventually run out and IPv6 become mandatory.

IPv6 addresses are generally written in the following form:
hhhh:hhhh:hhhh:hhhh:hhhh:hhhh:hhhh:hhhh
In this full notation, pairs of IPv6 bytes are separated by a colon and each byte in turns is represented as a pair of hexadecimal numbers, like in the following example:
E3D7:0000:0000:0000:51F4:9BC8:C0A8:6420
As shown above, IPv6 addresses commonly contain many bytes with a zero value. Shorthand notation in
IPv6 removes these values from the text representation (though the bytes are still present in the actual network address) as follows:

E3D7::51F4:9BC8:C0A8:6420
Finally, many IPv6 addresses are extensions of IPv4 addresses. In these cases, the rightmost four bytes of an IPv6 address (the rightmost two byte pairs) may be rewritten in the IPv4 notation.
Converting the above example to mixed notation yields E3D7::51F4:9BC8:192.168.100.32
IPv6 addresses may be written in any of the full, shorthand or mixed notation illustrated above.

Reference
http://compnetworking.about.com/od/tcpiptutorials/a/ipaddrnotation.htm

QUESTION NO: 216

Internet Protocol version 6 (IPv6) is the next-generation Internet Layer protocol for packet-switched internetworks and the Internet. IPv6 router solicitation is:

A. A request made by a node for the IP address of the local router
B. A request made by a node to join a specified multicast group
C. A request made by a node for a DHCP provided IP address
D. A request made by a node for the IP address of the DHCP server

Answer: A
Explanation:
In cases when the host (computer or server) needs to prompt an immediate router advertisement, it sends what is called as a Router Solicitation. Examples of this include commands for re-booting or re-starting a running computer. The system is alerted through router solicitation.
Router solicitation messages belong to the ICMPv6 set of messages, specific to the IPv6 protocol. They are identified by a Next Header value "x'3A and decimal 58.
An IPv6 router solicitation is closely associated to the Neighbor Discovery (ND) function of the IPv6. Under this, the hosts or routers obtain or discover the link-layer addresses for elements that reside on attached links (neighbor) and to cleansed or purge spaces with cached values that are no longer functioning.

QUESTION NO: 217

Which two are IPv6 neighbor discovery packets? (Choose two.)

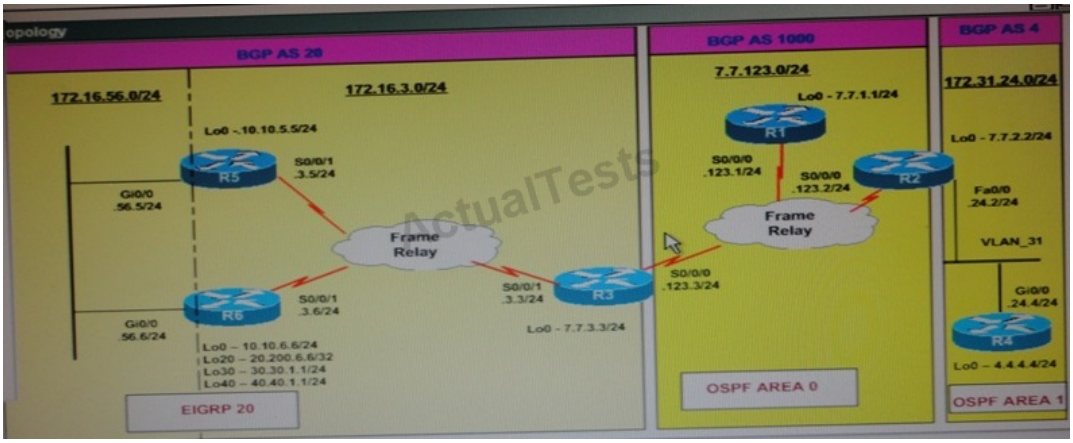
A. Neighbor Solicitation
B. Anycast Solicitation

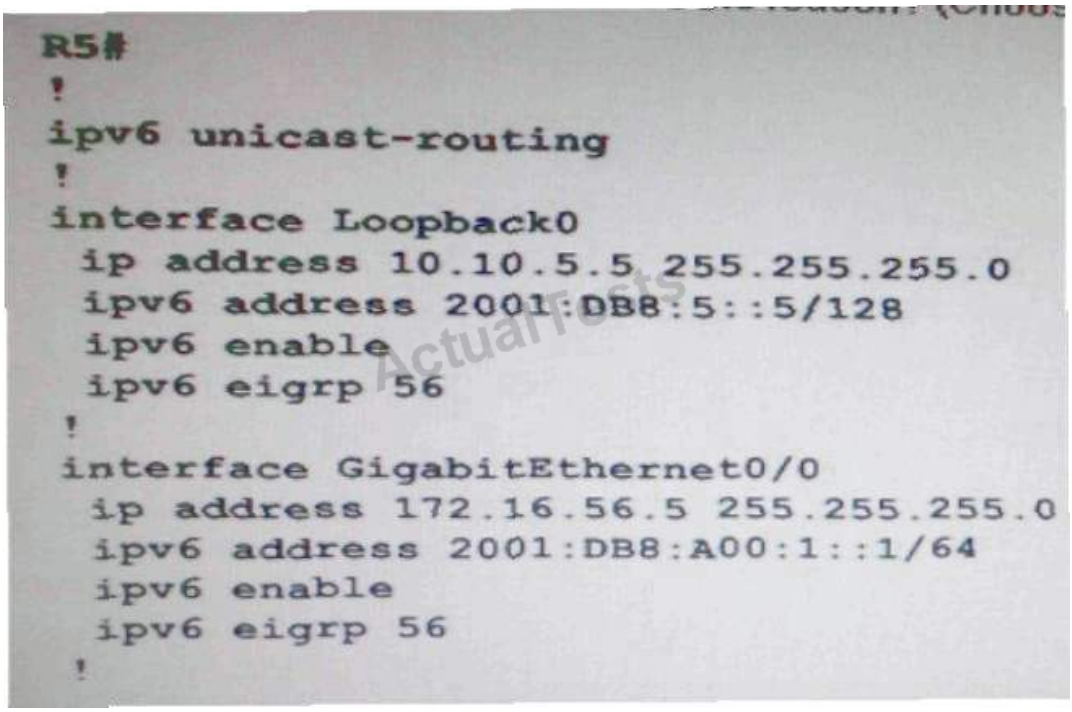
C. Anycast Advertisement
D. Router Advertisement

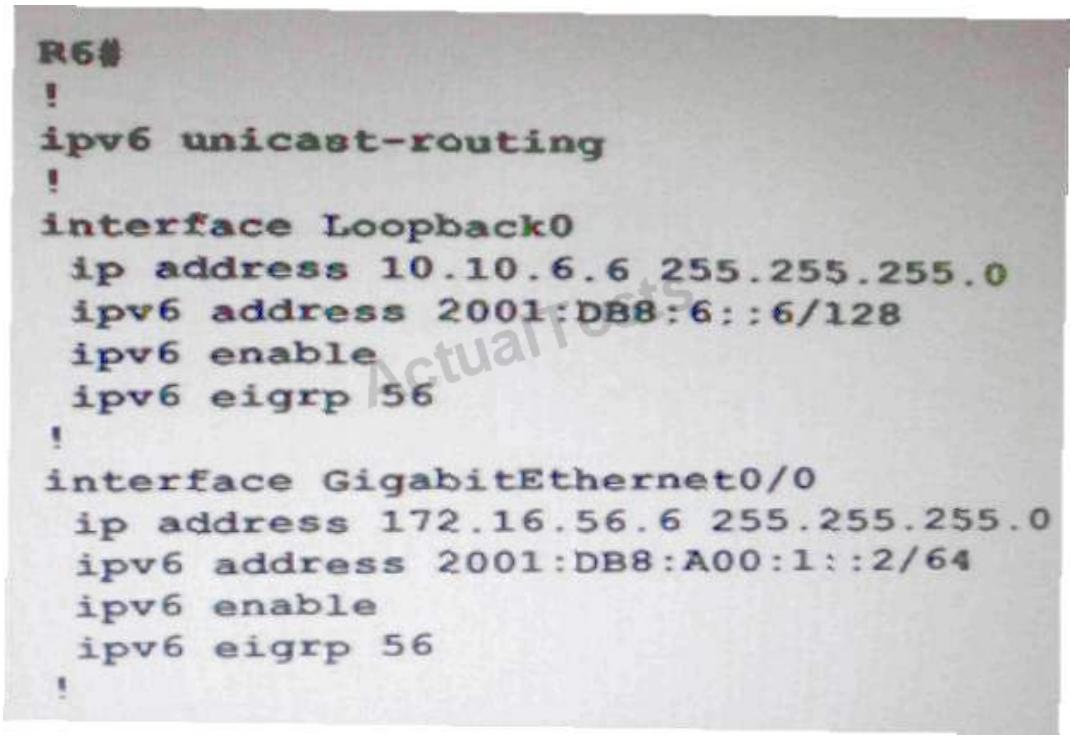
Answer: A,D
Reference: http://packetlife.net/blog/2008/aug/28/ipv6-neighbor-discovery/

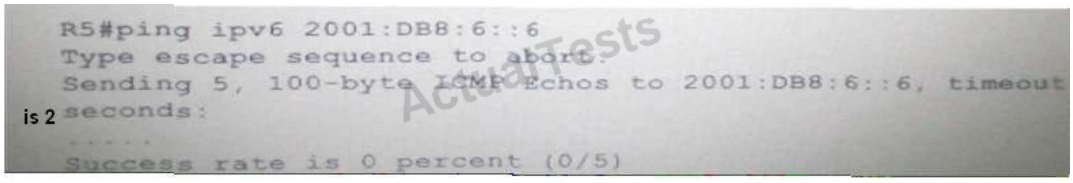
QUESTION NO: 218

Refer to the exhibit.







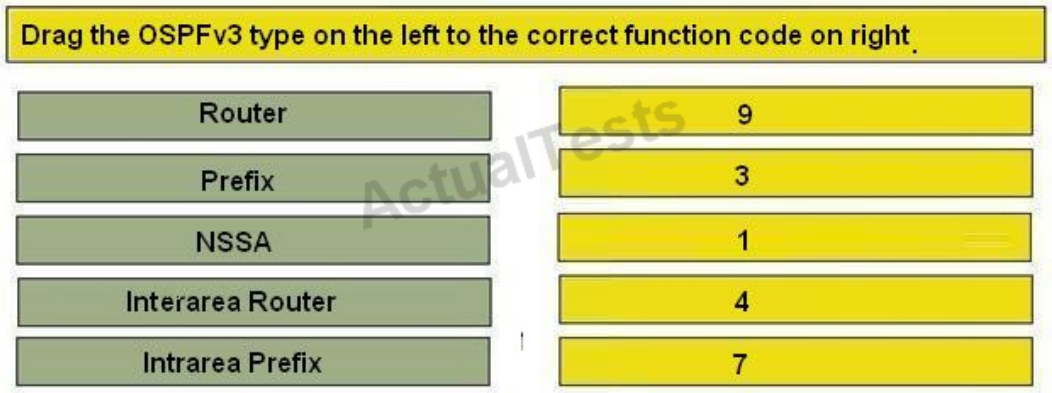


You have just configured R5 and R6 to run EIGRPv6 as shown; the IPv6 ping from R5 to R6-loopback 0 is failing. Which statement could be the reason?

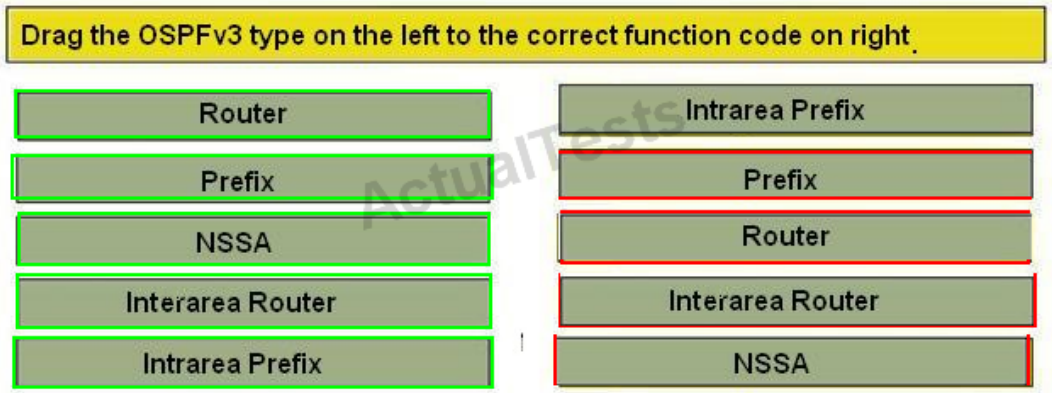
A. The loopback interfaces on R5 and R6 must be configured on an EIGRPv6 As number other than 56.
B. The loopback interfaces on R5 and R6 must be configured to EIGRPv6 As number 56.
C. You need to configure the EIGRPv6 router process on both routers.
D. You need to configure the EIGRPv6 router process in at least one of the routers
E. You should remove the ipv6 eigrp 56 from the loopback interfaces on both routers.

Answer: C
Reference: http://www.slideshare.net/CiscoCanada/ipv6-handsonlabhritter

QUESTION NO: 219 DRAG DROP



Answer:



Explanation:

Router -1

Prefix - 3

NSSA - 7

interarea router - 4

Intraarea Prefix – 9

Reference :

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6553/prod_presentation0900aecd80311e31.pdf

QUESTION NO: 220

What is the correct command to set the router ID for an OSPFv3 process?

A. router-id 1.2.3.4

B. router-id 2011::1

C. router-id Loopback0

D. router-id FF02::5

Answer: A

Explanation: Explanation

Stub Router

ipv6 unicast-routing

ipv6 cef

!

interface serial 0/0

no ip address

ipv6 enable

ipv6 address 2001:ABAB::/64 eui-64

ipv6 ospf 1 area 2

!

ipv6 router ospf 1

router-id 3.3.3.3

area 2 stub

!

Reference

http://www.cisco.com/en/US/tech/tk872/technologies_configuration_example09186a0080b40d8a.shtml

QUESTION NO: 221

Which of these statements best describes the major difference between an IPv4-compatible tunnel and a 6to4 tunnel?

A. An IPv4-compatible tunnel is a static tunnel, but an 6to4 tunnel is a semiautomatic tunnel.
B. The deployment of a IPv4-compatible tunnel requires a special code on the edge routers, but a 6to4 tunnel does not require any special code.
C. An IPv4-compatible tunnel is typically used only between two IPv6 domains, but a 6to4 tunnel is used to connect to connect two or more IPv6 domains.
D. For an IPv4-compatible tunnel, the ISP assigns only IPv4 addresses for each domain, but for a 6to4 tunnel, the ISP assigns only IPv6 addresses for each domain.

Answer: C

Explanation:

Automatic 6to4 Tunnels

An automatic 6to4 tunnel allows isolated IPv6 domains to be connected over an IPv4 network to

remote IPv6 networks. The key difference between automatic 6to4 tunnels and manually configured tunnels is that the tunnel is not point-to-point; it is point-to-multipoint. In automatic 6to4 tunnels, routers are not configured in pairs because they treat the IPv4 infrastructure as a virtual nonbroadcast multi-access (NBMA) link. The IPv4 address embedded in the IPv6 address is used to find the other end of the automatic tunnel.

An automatic 6to4 tunnel may be configured on a border router in an isolated IPv6 network, which creates a tunnel on a per-packet basis to a border router in another IPv6 network over an IPv4 infrastructure. The tunnel destination is determined by the IPv4 address of the border router extracted from the IPv6 address that starts with the prefix 2002::/16, where the format is 2002:border-router-IPv4-address::/48. Following the embedded IPv4 address are 16 bits that can be used to number networks within the site. The border router at each end of a 6to4 tunnel must support both the IPv4 and IPv6 protocol stacks. 6to4 tunnels are configured between border routers or between a border router and a host.

The simplest deployment scenario for 6to4 tunnels is to interconnect multiple IPv6 sites, each of which has at least one connection to a shared IPv4 network. This IPv4 network could be the global Internet or a corporate backbone. The key requirement is that each site have a globally unique IPv4 address; the Cisco IOS software uses this address to construct a globally unique 6to4/48 IPv6 prefix. As with other tunnel mechanisms, appropriate entries in a Domain Name System (DNS) that map between hostnames and IP addresses for both IPv4 and IPv6 allow the applications to choose the required address. Automatic IPv4-Compatible IPv6 Tunnels

Automatic IPv4-compatible tunnels use IPv4-compatible IPv6 addresses. IPv4-compatible IPv6 addresses are IPv6 unicast addresses that have zeros in the high-order 96 bits of the address, and an IPv4 address in the low-order 32 bits. They can be written as 0:0:0:0:0:0:A.B.C.D or ::A.B.C.D, where "A.B.C.D" represents the embedded IPv4 address.

The tunnel destination is automatically determined by the IPv4 address in the low-order 32 bits of IPv4- compatible IPv6 addresses. The host or router at each end of an IPv4-compatible tunnel must support both the IPv4 and IPv6 protocol stacks. IPv4-compatible tunnels can be configured between border- routers or between a border-router and a host. Using IPv4-compatible tunnels is an easy method to create tunnels for IPv6 over IPv4, but the technique does not scale for large networks.

QUESTION NO: 222

Which information is carried in an OSPFv3 intra-area-prefix LSA?

A. IPv6 prefixes
B. link-local addresses
C. solicited node multicast addresses
D. IPv6 prefixes and topology information

Answer: A

Explanation:

The OSPFv3 s new LSA, the Intra-area Prefix LSA (type 9), handles intra-area network information that was previously included in OSPFv2 type 2 LSAs. It is used in order to advertise one or more IPv6 prefixes. The prefixes are associated with router segment, stub network segment or transit network segment. Intra-area prefix LSAs (type 9) & Inter-Area-Prefix-LSA (type 3) carry all IPv6 prefix information, which, in IPv4, is included in router LSAs and network LSAs. Note: An address prefix is represented by three fields: prefix length, prefix options, and address prefix. In OSPFv3, addresses for these LSAs are expressed as prefix, prefix length instead of address, mask. The LSA types defined in OSPF are as follows:

Type 1 - Router LSA - the router announces its presence and lists the links to other routers or networks in the same area, together with the metrics to them. Type 1 LSAs are flooded across their own area only. The linkstate

ID of the type 1 LSA is the originating router ID.

Type 2 - Network LSA - the designated router on a broadcast segment (e.g. Ethernet) lists which routers are joined together by the segment. Type 2 LSAs are flooded across their own area only. The link-state ID of the type 2 LSA is the IP interface address of the DR.

Type 3 - Summary LSA - an Area Border Router (ABR) takes information it has learned on one of its attached areas and it can summarize it (but not by default) before sending it out on other areas it is connected to. This summarization helps provide scalability by removing detailed topology information for other areas, because their routing information is summarized into just an address prefix and metric. The summarization process can also be configured to remove a lot of detailed address prefixes and replace them with a single summary prefix, also helping scalability. The link-state ID is the destination network number for type 3 LSAs. Type 4 - ASBR-Summary LSA - this is needed because Type 5 External LSAs are flooded to all areas and the detailed next-hop information may not be available in those other areas. This is solved by an Area Border Router flooding the information for the router (i.e. the Autonomous System Boundary Router) where the type 5 originated. The link-state ID is the router ID of the described ASBR for type 4 LSAs. Type 5 - External LSA - these LSAs contain information imported into OSPF from other routing processes. They are flooded to all areas (except stub areas). For "External Type 1" LSAs routing decisions are made by adding the OSPF metric to get to the ASBR and the external metric from there on, while for "External Type 2" LSAs only the external metric is used. The link-state ID of the type 5 LSA is the external network number.

Type 6 - Group Membership LSA - this was defined for Multicast extensions to OSPF (MOSPF)[1], a multicast OSPF routing protocol which was not in general use. MOSPF has been deprecated since OSPFv3[2] and is not currently used. It may be reassigned in the future.

Type 7 - Routers in a Not-so-stubby-area (NSSA) do not receive external LSAs from Area Border Routers, but are allowed to send external routing information for redistribution. They use type 7 LSAs to tell the ABRs about these external routes, which the Area Border Router then translates to type 5 external LSAs and floods as normal to the rest of the OSPF network.

Type 8 - A link-local only LSA for OSPFv3. A Type 8 LSA is used to give information about link-

local addresses and a list of IPv6 addresses on the link. In OSPFv2, however, the Type 8 was originally intended to be used as a so-called External-Attributes-LSA for transit autonomous systems where OSPFv2 could replace the internal Border Gateway Protocol (iBGP). In these networks, the BGP destinations would be carried in LSA Type 5 while their BGP attributes would be inserted into LSA Type 8. Most OSPFv2 implementations never supported this feature.
Type 9 - a link-local "opaque" LSA (defined by RFC2370) in OSPFv2 and the Intra-Area-Prefix LSA in OSPFv3. It is the OSPFv3 LSA that contains prefixes for stub and transit networks in the link-state ID.
Type 10 - an area-local "opaque" LSA as defined by RFC2370. Opaque LSAs contain information which should be flooded by other routers even if the router is not able to understand the extended information itself. Typically type 10 LSAs are used for traffic engineering extensions to OSPF, flooding extra information about links beyond just their metric, such as link bandwidth and color.
Type 11 - an AS "opaque" LSA defined by RFC 5250, which is flooded everywhere except stub areas. This is the opaque equivalent of the type 5 external LSA

OSPFv3 LSA Types

LSA Function Code	LS Type	Description
1	0x2001	Router-LSA
2	0x2002	Network-LSA
3	0x2003	Inter-Area-Prefix-LSA Also referred to as Summary-LSA in OSPFv2.
4	0x2004	Inter-Area-Router-LSA Also referred to as ASBR-Summary-LSA in OSPFv2.
5	0x2005	AS-External-LSA Also referred to as External-LSA or AS-External-LSA in OSPFv2.
~~6~~	~~0x2006~~	~~MOSPF-LSA~~ Deprecated in OSPFv3 (may be reassigned) Was referred to as Multicast-OSPF-LSA in OSPFv2.
7	0x2007	NSSA-LSA Also referred to as NSSA-LSA in OSPFv2.
8	0x2008	Link-LSA
9	0x2009	Intra-Area-Prefix-LSA

ActualTests

QUESTION NO: 223

Which IPv6 address would you ping to determine if OSPFv3 is able to send and receive unicast packets across a link?

A. anycast address
B. site-local multicast
C. global address ofthe link
D. unique local address
E. link-local address

Answer: E

Explanation:

A link-local address is an Internet Protocol address that is intended only for communications within the segment of a local network (a link) or a point-to-point connection that a host is connected to. Routers do not forward packets with link-local addresses.

QUESTION NO: 224

You are using IPv6, and would like to configure EIGRPv6. Which three of these correctly describe how you can perform this configuration? (Choose three.)

A. EIGRP for IPv6 is directly configured on the interfaces over which it runs.
B. EIGRP for IPv6 is not configured on the interfaces over which it runs, but if a user uses passive- interface configuration, EIGRP for IPv6 needs to be configured on the interface that is made passive.
C. There is a network statement configuration in EIGRP for IPv6, the same as for IPv4.
D. There is no network statement configuration in EIGRP for IPv6.
E. When a user uses a passive-interface configuration, EIGRP for IPv6 does not need to be configured on the interface that is made passive.
F. When a user uses a non-passive-interface configuration, EIGRP for IPv6 does not need to be configured on the interface that is made passive

Answer: A,D,E

Explanation: This section lists ways in which EIGRP for IPv6 differs from EIGRP IPv4 and lists EIGRP for IPv6 restrictions:

• EIGRP for IPv6 is directly configured on the interfaces over which it runs. This feature allows EIGRP for IPv6 to be configured without the use of a global IPv6 address. There is no network statement in EIGRP for IPv6.
In per-interface configuration at system startup, if EIGRP has been configured on an interface, then the EIGRP protocol may start running before any EIGRP router mode commands have been executed.
• An EIGRP for IPv6 protocol instance requires a router ID before it can start running.
• EIGRP for IPv6 has a shutdown feature. The routing process should be in "no shut" mode in order to start running.
• When a user uses a passive-interface configuration, EIGRP for IPv6 need not be configured on the interface that is made passive.
• EIGRP for IPv6 provides route filtering using the distribute-list prefix-list command. Use of the route-map command is not supported for route filtering with a distribute list.
Reference: http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-eigrp.html#wp1055769

QUESTION NO: 225

Which two IPv6 address pairs can be assigned on a point-to-point link with a /125 mask? (Choose two.)

A. 2001:DB8::1 and 2001:DB8::2
B. 2001:DB8::9 and 2001:DB8::10
C. 2001:DB8::10 and 2001:DB8::17
D. 2001:DB8::1 and 2001:DB8::18

Answer: A,C
Explanation:
2001:DB8::x/125 = 2001:0DB8:0000:0000:0000:0000:0000:00xx

1- The 0's disapear because the 3 zero ommiting rules:
Leading zeros can be omitted
Groups of four zeros can be abbreviated to a zero
Multiple contiguous groups (2-byte groups) can be abbreviated to double colon (::) but only once.
2- IPv6 is made by 128bits, therefore 128-125 = 3
This means that the last 3 bits represent the hosts, just like in IPv4
3- IPv6 is represented in Hex numbers:
Ex: 2001 = 00100001
4- Then you can determine the answer just like in IPv4 (or at least in a very similar way)

2001:DB8::1 and 2001:DB8::2 belong to the /125 mask because of the 3 last bits
2001:DB8::9 and 2001:DB8::10 are obviously not in the same /125 subnet (remember 10 in hex is much different thatn 10 in decimal)
2001:DB8::10 and 2001:DB8::17 fall into the /125 subnet
2001:DB8::1 and 2001:DB8::18 it goes beyond the /125 subnet
5- In case you didn't understand why:
2001:DB8::1 - 2001:DB8::9 = First Subnet's Host
2001:DB8::10 - 2001:DB8::18 = Second Subnet's Hosts

QUESTION NO: 226

Which three causes will prevent a host from getting an IPv6 address with stateless autoconfiguration? (Choose three.)

A. The autoconfig command is missing from the router interface.
B. IPv6 multicast routing is not enabled on the router interface.
C. IPv6 unicast routing is not enabled on the router interface.
D. The router interface is configured with a /63 mask.
E. The router interface is configured with a /65 mask.

Answer: C,D,E
Explanation:
For IPv6 to function IPv6 unicast routing must be enabled on the interface. For stateless autoconfiguration to work properly, the advertised prefix length in router advertisement messages must always be 64 bits.
Reference:
http://www.cisco.com/en/US/docs/security/fwsm/fwsm31/configuration/guide/ipv6_f.html

QUESTION NO: 227

Which two are considered valid router IDs for EIGRPv6? (Choose two.)

A. 0.0.0.1
B. 0.255.0.1
C. 2001:DB8::1
D. 2001:DB8:ABCD:FEEF::1

Answer: A,B

Explanation:

As with OSPFv3, IPv6 EIGRP require a router-id in IPv4 format. Here is an example EIGRPv6 configuration:

ipv6 router eigrp 6
eigrp router-id 6.6.6.6

QUESTION NO: 228

Which two tunneling techniques in IPv6 do not require an IPv4 tunnel destination in the configuration? (Choose two.)

A. 6to4
B. 6over4
C. ISATAP
D. GRE

Answer: A,C

Explanation:

Example: Configuring 6to4 Tunnels

The following example configures a 6to4 tunnel on a border router in an isolated IPv6 network. The IPv4 address is 192.168.99.1, which translates to the IPv6 prefix of 2002:c0a8:6301::/48. The IPv6 prefix is subnetted into 2002:c0a8:6301::/64 for the tunnel interface: 2002:c0a8:6301:1::/64 for the first IPv6 network, and 2002:c0a8:6301:2::/64 for the second IPv6 network. The static route ensures that any other traffic for the IPv6 prefix 2002::/16 is directed to tunnel interface 0 for automatic tunneling.

interface Ethernet0
description IPv4 uplink
ip address 192.168.99.1 255.255.255.0
!
interface Ethernet1
description IPv6 local network 1
ipv6 address 2002:c0a8:6301:1::1/64
!
interface Ethernet2
description IPv6 local network 2

```
ipv6 address 2002:c0a8:6301:2::1/64
!
interface Tunnel0
description IPv6 uplink
no ip address
ipv6 address 2002:c0a8:6301::1/64
tunnel source Ethernet 0
tunnel mode ipv6ip 6to4
!
ipv6 route 2002::/16 tunnel 0
```

Example: Configuring ISATAP Tunnels

The following example shows the tunnel source defined on Ethernet 0 and the tunnel mode command used to configure the ISATAP tunnel. Router advertisements are enabled to allow client autoconfiguration.

```
ipv6 unicast-routing
interface tunnel 1
tunnel source ethernet 0
tunnel mode ipv6ip isatap
ipv6 address 2001:DB8::/64 eui-64
no ipv6 nd ra suppress
exit
```

Reference: http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-tunnel.html#wp1056870

QUESTION NO: 229

What does a unicast prefix-based IPv6 multicast address start with?

A. FF80
B. FF3
C. FF
D. FF7

Answer: B
Explanation:

From RFC 3306, the following are a few examples of the structure of unicast prefix-based multicast addresses.

- Global prefixes - A network with a unicast prefix of 3FFE:FFFF:1::/48 would also have a unicast prefix-based multicast prefix of FF3x:0030:3FFE:FFFF:0001::/96 (where 'x' is any valid scope).

- SSM - All IPv6 SSM multicast addresses will have the format FF3x::/96.

Reference: http://tools.ietf.org/html/rfc3306

QUESTION NO: 230

How does an IPv6 router deal with a packet that is larger than the outgoing interface MTU?

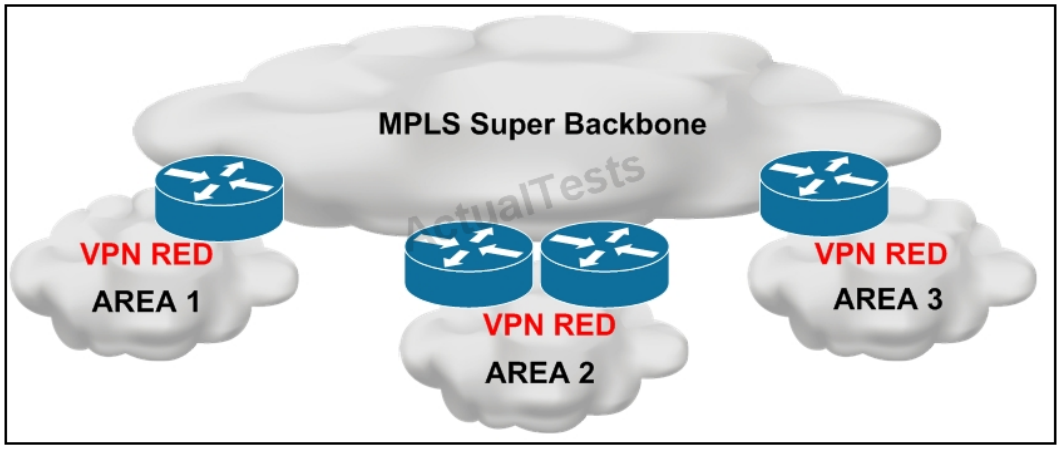
A. It will fragment the packet at Layer 2.
B. It will fragment the packet at Layer 3.
C. It will drop the packet and send an ICMPv6 message "packet too big" back to the source.
D. It will drop the packet.

Answer: C
Explanation:
Routers do not fragment IPv6 packets, as they do for IPv4. When a router receives a packet that is too big for the MTU of the forwarding path, the intermediate mode must drop the packet and send an ICMPv6 “packet too big” back to the source node.
Reference:

http://www.arbornetworks.com/asert/2011/07/ipv6-fragmentation/

Topic 4, Implement MPLS Layer 3 VPNs

QUESTION NO: 231
Refer to the exhibit.



This is an MPLS VPN network with OSPF as the PE-CE routing protocol. Which statement is correct?

A. The routing inside the VPN RED will never work correctly.
B. The routing inside the VPN RED can be enabled by configuring virtual links between the PE routers.
C. The routing inside the VPN RED can be enabled by configuring area 0 inside the VRF on the PE routers.
D. The routing inside the VPN RED will work without any special OSPF configuration.
E. The routing inside the VPN RED will work if the PE routers have a full mesh of sham-links configured for VRF RED.

Answer: D

Explanation:

http://www.cisco.com/en/US/docs/net_mgmt/vpn_solutions_center/1.1/user/guide/VPN_UG1.html

QUESTION NO: 232

Which of the following is true about the MPLS header and its EXP field size?

A. The MPLS header is 2 bytes, and the EXP field is 3 bits long.
B. The MPLS header is 1 byte, and the EXP field is 3 bits long.
C. The MPLS header is 4 bytes, and the EXP field is 3 bits long.
D. The MPLS header is 3 bytes, and the EXP field is 3 bits long.

Answer: C

Reference:
http://books.google.com.pk/books?id=DoOsh4NsCKwC&pg=PA95&lpg=PA95&dq=MPLS+header+is+4+bytes,+and+the+EXP+field+is+3+bits+long&source=bl&ots=lGc3RcpOUs&sig=MwvvUoDkbWfl8elpSuq_wn53i_I&hl=en&sa=X&ei=HTXLUfSzBsap4gT8tYHoBg&ved=0CCoQ6AEwAA#v=one

page&q=MPLS%20header%20is%204%20bytes%2C%20and%20the%20EXP%20field%20is%203%20bits%20long&f=false (page 100)

QUESTION NO: 233

A network is composed of several VRFs. It is required that VRF users VRF_A and VRF_B be able to route to and from VRF_C, which hosts shared services. However, traffic must not be allowed to flow between VRF_A and VRF_B. How can this be accomplished?

A. route redistribution
B. import and export using route descriptors
C. import and export using route targets
D. Cisco MPLS Traffic Engineering

Answer: C
Explanation:
An MPLS VPN implementation is very similar to a dedicated router peer-to-peer model implementation. From a CE router's perspective, only IPv4 updates, as well as data, are forwarded to the PE router. The CE router does not need any specific configuration to enable it to be a part of a MPLS VPN domain. The only requirement on the CE router is a routing protocol (or a static/default route) that enables the router to exchange IPv4 routing information with the connected PE router. In the MPLS VPN implementation, the PE router performs multiple functions. The PE router must first be capable of isolating customer traffic if more than one customer is connected to the PE router. Each customer, therefore, is assigned an independent routing table similar to a dedicated PE router in the initial peer-to-peer discussion. Routing across the SP backbone is performed using a routing process in the global routing table. P routers provide label switching between provider edge routers and are unaware of VPN routes. CE routers in the customer network are not aware of the P routers and, thus, the internal topology of the SP network is transparent to the customer
The P routers are only responsible for label switching of packets. They do not carry VPN routes and do not participate in MPLS VPN routing. The PE routers exchange IPv4 routes with connected CE routers using individual routing protocol contexts. To enable scaling the network to large number of customer VPNs, multiprotocol BGP is configured between PE routers to carry customer routes.
Customer isolation is achieved on the PE router by the use of virtual routing tables or instances, also called virtual routing and forwarding tables/instances (VRFs). In essence, it is similar to maintaining multiple dedicated routers for customers connecting into the provider network. The function of a VRF is similar to a global routing table, except that it contains all routes pertaining to a specific VPN versus the global routing table. The VRF also contains a VRF-specific CEF forwarding table analogous to the global CEF table and defines the connectivity requirements and

protocols for each customer site on a single PE router. The VRF defines routing protocol contexts that are part of a specific VPN as well as the interfaces on the local PE router that are part of a specific VPN and, hence, use the VRF. The interface that is part of the VRF must support CEF switching. The number of interfaces that can be bound to a VRF is only limited by the number of interfaces on the router, and a single interface (logical or physical) can be associated with only one VRF. The VRF contains an IP routing table analogous to the global IP routing table, a CEF table, list of interfaces that are part of the VRF, and a set of rules defining routing protocol exchange with attached CE routers (routing protocol contexts). In addition, the VRF also contains VPN identifiers as well as VPN membership information (RD and RT are covered in the next section).

Route targets (RTs) are additional identifiers used in the MPLS VPN domain in the deployment of MPLS VPN that identify the VPN membership of the routes learned from that particular site. RTs are implemented by the use of extended BGP communities in which the higher order 16 bits of the BGP extended community (64 total bits) are encoded with a value corresponding to the VPN membership of the specific site. When a VPN route learned from a CE router is injected into VPNv4 BGP, a list of VPN route target extended community attributes is associated with it. The export route target is used in identification of VPN membership and is associated to each VRF. This export route target is appended to a customer prefix when it is converted to a VPNv4 prefix by the PE router and propagated in MP-BGP updates. The import route target is associated with each VRF and identifies the VPNv4 routes to be imported into the VRF for the specific customer. The format of a RT is the same as an RD value.

QUESTION NO: 234

A request arrived on your MPLS-vpn-bgp group. Due to a security breach, your customer is experiencing DoS attacks coming from specific subnets (200.0.10.0/24, 200.0.12.0/24). You have checked all MPLS-EBGP routes being advertised to BHK from other VPN sites and found four subnets listed: 200.0.10.0/24, 200.0.11.0/24, 200.0.12.0/24, 200.0.13.0/24. You immediately apply an outbound ACL filter using the appropriate MPLS-EBGP tool: access-list 1 deny 200.0.10.0 255.255.254.0 access-list 1 permit any What happens when you apply this ACL on the MPLS-EBGP connection to BHK?

A. It blocks all routes.
B. B. It blocks the routes 200.0.12.0/24, 200.0.10.0/24 only.
C. It blocks the routes 200.0.12.0/24, 200.0.13.0/24 only.
D. It blocks the routes 200.0.10.0/24, 200.0.13.0/24 only.
E. Nothing happens, no routes are blocked.

Answer: B

Explanation:

Remember, for the wild card mask, 1s are I DON'T CARE, and 0s are I CARE.

In the access-list we put an 0.0.0.0 255.255.254.255 network; of course 255 means "1111 1111. This means we don't care about any of the bits in the first, second & 4th octets.

In fact, the number 0 (in 0.0.0.0) is just smallest numbers we can throw there and it is easy to type but we can use any number, it wouldn't matter, since I DON'T CARE about them except the third octet as the wild card mask is not all "255.

Now let's extract the 0 in the third octet in binary form (so easy, right?)

0 = 0000 0000

With the 254 in the wildcard mask, we only care about the last bit of the third octet because 254 is "1111 1110.

That means, if the third octet is in the form of xxxx xxx0 then it will match my access-list (x can be 0 or 1 because I DON'T CARE).

Now let's write the third octet of 4 above subnets in binary form:

10 = 0000.1010

11 = 0000.1011

12 = 0000.1100

13 = 0000.1101

So, only 10 & 12 satisfy my access list -> I will only block the routes to 200.0.12.0/24, 200.0.10.0/24 -> B is correct.

Here is a simple configuration example explaining the question above. Connect to Routers R1 and BHK via FastEthernet 0/0

Router R1

```
interface Loopback0
ip address 200.0.10.1 255.255.255.0
!
interface Loopback1
ip address 200.0.11.1 255.255.255.0
!
interface Loopback2
ip address 200.0.12.1 255.255.255.0
!
interface Loopback3
ip address 200.0.13.1 255.255.255.0
!
interface FastEthernet0/0
ip address 10.0.1.2 255.255.255.252
router bgp 65500
```

no synchronization
bgp log-neighbor-changes
network 10.0.1.0 mask 255.255.255.252
network 200.0.10.0
network 200.0.11.0
network 200.0.12.0
network 200.0.13.0
neighbor 10.0.1.1 remote-as 65525
no auto-summary
Router BHK
router bgp 65525
no synchronization
bgp log-neighbor-changes
network 10.0.1.0 mask 255.255.255.252
neighbor 10.0.1.2 remote-as 65500
neighbor 10.0.1.2 route-map 1 in
distribute-list list in
no auto-summary
access-list 1 deny 0.0.0.0 255.255.254.255
access-list 1 permit any

Note: You may need to clear the BGP process on Router BHK after applying the route-map

QUESTION NO: 235

Multi Protocol Label Switching (MPLS) is a data-carrying mechanism that belongs to the family of packet-switched networks. For an MPLS label, if the stack bit is set to1, which option is true?

A. The stack bit will only be used when LDP is the label distribution protocol
B. The label is the last entry in the label stack.
C. The stack bit is for Cisco implementations exclusively and will only be used when TDP is the label distribution protocol.
D. The stack bit is reserved for future use.

Answer: B

Explanation: Explanation



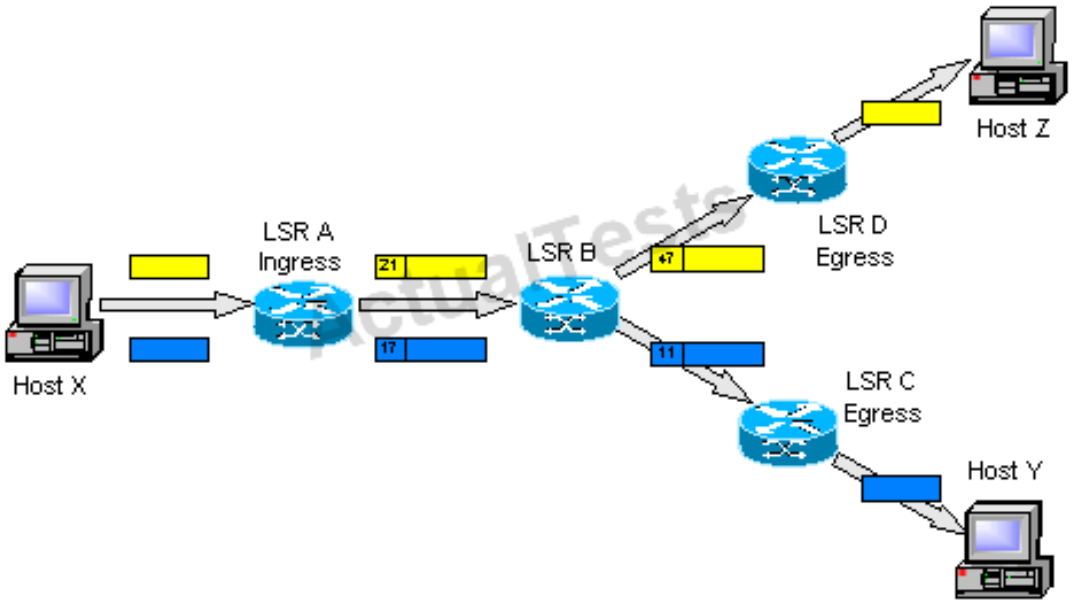
MPLS Header Packet Format
LABEL: 20 bits
EXP: Experimental, 3bits - are reserved for experimental use
S: Bottom of stack, 1 bit
TTL: Time to Live, 8bits - same as IP TTL
The bottom-of-stack bit, or "stack bit", is just used to indicate it is the bottom of the label stack because it is possible (and common) to have more than one label attached to a packet. The bottommost label in a stack has the S bit set to 1, other labels have the S bit set to 0. Sometimes it is useful to know where the bottom of the label stack is and the S bit is the tool to find it.

QUESTION NO: 236

Which two fundamental modifications, related to traffic forwarding, does MPLS introduce? (Choose two.)

A. IP lookup is performed on every hop within the MPLS core.
B. IP destination routing is reduced to label lookup within the MPLS network.
C. For unicast routing, labels are assigned to FECs (in other words, IP prefixes).
D. For multicast routing, labels are assigned to IP multicast groups.

Answer: B,C
Explanation:
MPLS works by tagging packets with an identifier (a label) to distinguish the LSPs. When a packet is received, the router uses this label (and sometimes also the link over which it was received) to identify the LSP. It then looks up the LSP in its own forwarding table to determine the best link over which to forward the packet, and the label to use on this next hop.
A different label is used for each hop, and it is chosen by the router or switch performing the forwarding operation. This allows the use of very fast and simple forwarding engines, as the router can select the label to minimize processing.
Ingress routers at the edge of the MPLS network use the packet's destination address to determine which LSP to use. Inside the network, the MPLS routers use only the LSP labels to forward the packet to the egress router.



In the diagram above, LSR (Label Switched Router) A uses the destination IP address on each packet to select the LSP, which determines the next hop and initial label for each packet (21 and 17). When LSR B receives the packets, it uses these labels to identify the LSPs, from which it determines the next hops (LSRs D and C) and labels (47 and 11). The egress routers (LSRs D and C) strip off the final label and route the packet out of the network.
As MPLS uses only the label to forward packets, it is protocol-independent, hence the term "Multi-Protocol" in MPLS. Packet forwarding has been defined for all types of layer-2 link technologies, with a different label encoding used in each case.

QUESTION NO: 237

For which routes does LDP advertise a label binding?

A. all routes in the routing table
B. only the IGP and BGP routes in the routing table
C. only the BGP routes in the routing table
D. only the IGP routes in the routing table

Answer: D
Explanation:
LDP can only do bindings for IGP learned routes. If the route is learned from BGP, BGP has to do the label binding. For CCIE R&S you don't need to worry about using BGP for label distribution, as this is used for Inter-AS MPLS L3VPN scenarios. If you change your setup so the routes are learned from IGP instead of BGP they will have labels.

QUESTION NO: 238

Which command can be used on a PE router to connect to a CE router (11.1.1.4) in VRF red?

A. telnet 11.1.1.4 /vrf-source red
B. telnet 11.1.1.4 source /vrf red
C. telnet 11.1.1.4 /source vrf red
D. telnet 11.1.1.4 /vrf red
E. telnet 11.1.1.4 vrf red

Answer: D
Explanation:
Telnetting can be done through the VRF using the Management Ethernet interface. In the following example, the router telnets to 172.17.1.1 through the Management Ethernet interface VRF:

Router# telnet 172.17.1.1 /vrf Mgmt-intf

Reference
http://www.cisco.com/en/US/docs/routers/asr1000/configuration/guide/chassis/Management_Ethernet.html

QUESTION NO: 239

What are the "bound IP addresses" used for in LDP?

A. For each such address and mask /32, one label binding is created.
B. As the TCP endpoint, IP addresses are used by the LDP session.
C. These addresses are used to find the label binding to put in the LFIB, by looking up these IP addresses in the routing table.
D. These addresses are used for penultimate hop popping (PHP) when forwarding packets to the next router directly.

Answer: C
Explanation:
In case of liberal label retention Label Information Base (LIB) maintains remote binding information through down streaming or through upcoming hop. The label binding is utilized in Label forwarding

information base (LFIB) but no other labels are kept which are not used for forwarding packets. The cause for storing remote binding in LFIB is subject to topological change and implementation of dynamic routing due to downlink of router. Conservative label retention mode configure on an LSR does not contain all remote bindings except an associated upcoming hop in its LIB. However LLR will help in rapid routing topological change while CLR utilizes memory efficiently.

QUESTION NO: 240

What does the OSPF command capability vrf-lite achieve?

A. It enables provider edge (PE) specific checks on a router when the OSPF process is associated with the VRF.
B. It disables provider edge (PE) specific checks on a router when the OSPF process is associated with the VRF.
C. It enables the exchange of the "VRF-Lite" capability when the OSPF adjacency is formed.
D. It disables the MPLS processing on the OSPF learned routes inside the VRF.

Answer: B
Explanation:
Capability vrf-lite
To suppress the Provider Edge (PE) specific checks on a router when the OSPF process is associated with the VRF, use the capability vrf-lite command in router configuration mode. To restore the checks, use the no form of this command.
capability vrf-lite no capability vrf-lite
Syntax Description
This command has no arguments or keywords.
Defaults Disabled. PE specific checks are performed if the process is associated with VRF
command modes

Prerequisites
CEF must be running on the network.
SUMMARY STEPS
1. enable
2. show ip ospf [process-id]
3. configure terminal
4. router ospf process-id [vrf vpn-name]
5. capability vrf-lite

DETAILED STEPS

	Command or Action	Purpose

Step 1	`enable`
	
	Example:
	`Router> enable`
		Enables higher privilege levels, such as privileged EXEC mode.
		• Enter your password if prompted.

Step 2	`show ip ospf [process-id]`
	
	Example:
	`Router> show ip ospf 1`
		Displays the status of the router. If the display indicates that the router is connected to the VPN backbone, you can use the **capability vrf-lite** command to decouple the PE router from the VPN backbone.

Step 3	`configure terminal`
	
	Example:
	`Router# configure terminal`
		Enters global configuration mode.

Step 4	`router ospf process-id [vrf vpn-name]`
	
	Example:
	`Router(config)# router ospf 1 vrf grc`
		Enables OSPF routing and enters router configuration mode.
		• The *process-id* argument identifies the OSPF process.
		• Use the **vrf** keyword and *vpn-name* argument to identify a VPN.

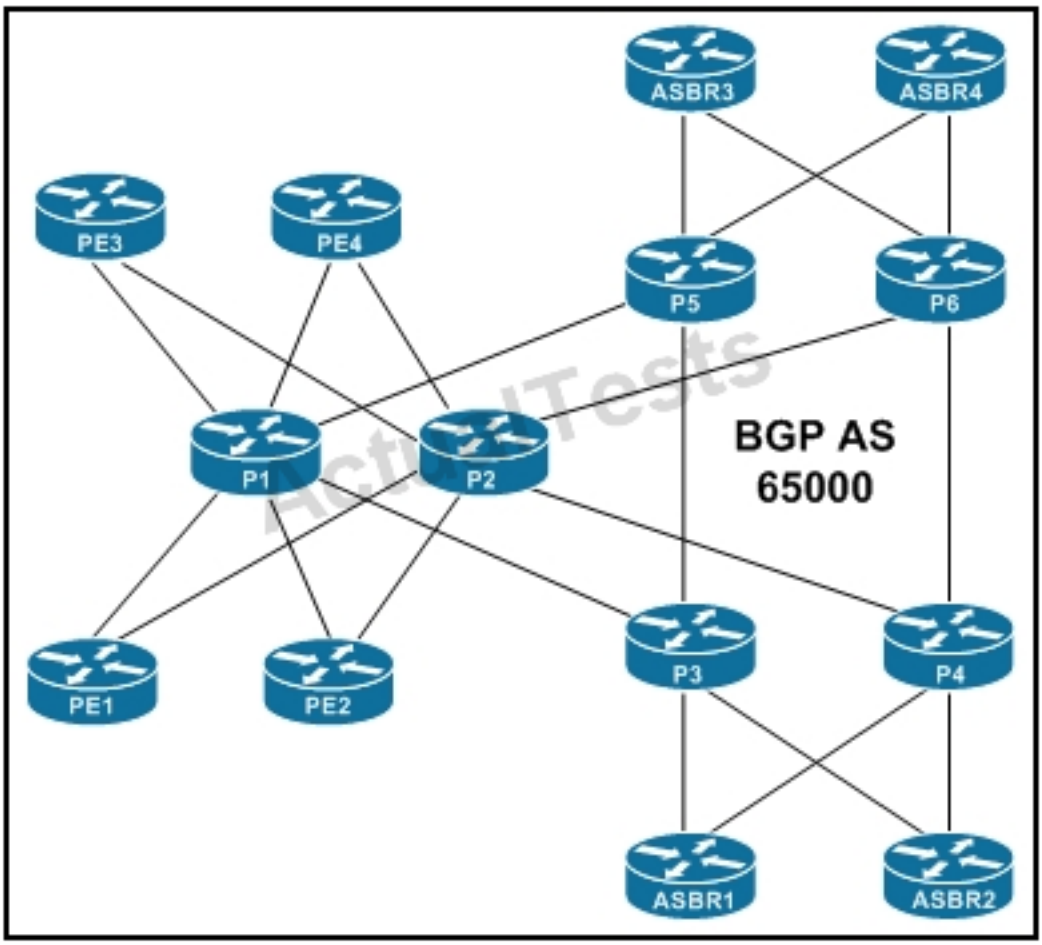
Step 5	`capability vrf-lite`
	
	Example:
	`Router(config)# capability vrf-lite`
		Applies the multi-VRF capability to the OSPF process.

ActualTests

Reference
http://www.cisco.com/en/US/docs/ios/12_0st/12_0st21/feature/guide/ospfvrfl.html

QUESTION NO: 241

Refer to the exhibit.



AS65000 has core network P1 – P6. The eBGP peers to another AS are through ASBR1–2 and ASBR3–4. All business and residential customers POPs are currently connected to the P1–P2 core.

AS65000 has decided to enable MPLS L3VPN services between all ASBR and PE routers. Which routers are the best positioned as VPNv4 RR?

A. ASBR1, ASBR2 and ASBR3, ASBR4, since they have eBGP peers to another AS
B. P1, since it is in the center of the MPLS network
C. a new pair of routers (RR1, RR2) dedicated as VPNv4 RR connected to P1, P2 since they are off the path
D. in MPLS L3VPN, there is no requirement to have VPNv4 RR

Answer: C

Explanation: The new pair of router dedicated as VPNv4RR connected to P1 and P2 are best as VPNv4 RR because they are off the path.

QUESTION NO: 242

What is the Cisco IOS command to turn on explicit null forwarding by LDP?

A. ldp explicit-null
B. mpls forwarding explicit-null
C. mpls ldp advertise-labels explicit null
D. mpls ldp explicit-null

Answer: D

Explanation:

When you issue the mpls ldp explicit-null command, Explicit Null is advertised in place of Implicit Null for directly connected prefixes.

SUMMARY STEPS

1. enable
2. configure terminal
3. mpls ip
4. mpls label protocol {ldp | tdp | both}
5. interface type number
6. mpls ip
7. exit
8. mpls ldp explicit-null [for prefix-acl | to peer-acl | for prefix-acl to peer-acl]
9. exit
10. show mpls forwarding-table [network {mask | length} | labels label [- label] | interface interface | next-hop address | lsp-tunnel [tunnel-id]] [vrf vpn-name] [detail]

Reference

http://www.cisco.com/en/US/docs/ios/12_4t/12_4t2/ftldp41.html

QUESTION NO: 243

Refer to the exhibit.

```
PE1#sh run vrf 999

ip vrf 999
 rd 99.99.99.99:99
 route-target export 99:99
 route-target import 99:99
!
```

ActualTests

Which statement is correct about this configuration?

A. This is not a legal configuration.
B. The RD is a type 0 RD.
C. The RD is a type 1 RD.
D. The RD is a type 2 RD.
E. This is a special RD that is used only for management VRFs.

Answer: C
Explanation:
The route distinguisher (RD) is an 8-octet value consisting of 2 major fields, the Type Field (2 octets) and Value Field (6 octets). The type field determines how the value field should be interpreted. The three Type values, as defined in the internet draft, are:

Type 0:
Type Field (2 octets)
Administrator subfield (2 octets)
Assigned number subfield (4 octets)
The administrator field must contain an AS number (using private AS numbers is discouraged).
The Assigned field contains a number assigned by the service provider.

Type 1:
Type Field (2 octets)
Administrator subfield (4 octets)
Assigned number subfield (2 octets)
The administrator field must contain an IP address (using private IP address space is discouraged). The
Assigned field contains a number assigned by the service provider.

Type 2:
Type Field (2 octets)
Administrator subfield (4 octets)

Assigned number subfield (2 octets)
The administrator field must contain a 4-octet AS number (using private AS numbers is discouraged). The
Assigned field contains a number assigned by the service provider.
Reference
http://en.wikipedia.org/wiki/Route_distinguisher

QUESTION NO: 244

Which command should be used on a PE router to connect to a CE router (11.1.1.4) in VRF red?

A. telnet 11.1.1.4 /vrf-source red
B. telnet 11.1.1.4 source /vrf red
C. telnet 11.1.1.4 /source vrf red
D. telnet 11.1.1.4 /vrf red
E. telnet 11.1.1.4 vrf red

Answer: D
Explanation:

You would need to telnet the ce router in vrf red. The correct syntax is telent ipaddress/vrf red.

QUESTION NO: 245

Which of these statements about penultimate hop popping are true? (Choose three)

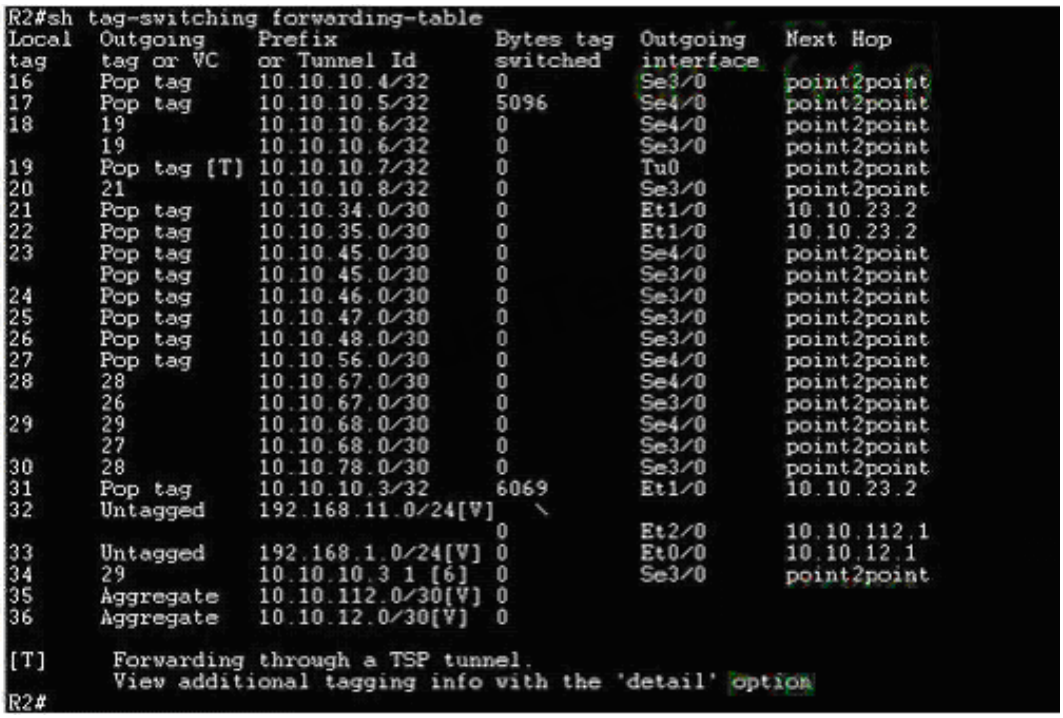
A. It is used only for directly connected subnets or aggregate routes
B. It can only be used with LDP.
C. It is only used when two or more labels are stacked.
D. It enables the Edge LSR to request a label pop operation from Its upstream neighbors
E. It is requested through TDP using a special label value that is also called the implicit-null value.
F. It is requested through LDP using a special label value that is also called the implicit-null value.

Answer: D,E,F
Explanation:
In order to implement penultimate hop popping, the edge LSR requests a label pop operation from its upstream neighbor via LDP or TDP using a special implicit-null label. This label has a value of 3

for LDP and 1 for TDP.

QUESTION NO: 246

Refer to the exhibit. According to the output of the command show tag-switching forwarding-table, which four of these statements are true? (Choose four.



```
R2#sh tag-switching forwarding-table
Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
16     Pop tag     10.10.10.4/32    0          Se3/0       point2point
17     Pop tag     10.10.10.5/32    5096       Se4/0       point2point
18     19          10.10.10.6/32    0          Se4/0       point2point
       19          10.10.10.6/32    0          Se3/0       point2point
19     Pop tag [T] 10.10.10.7/32    0          Tu0         point2point
20     21          10.10.10.8/32    0          Se3/0       point2point
21     Pop tag     10.10.34.0/30    0          Et1/0       10.10.23.2
22     Pop tag     10.10.35.0/30    0          Et1/0       10.10.23.2
23     Pop tag     10.10.45.0/30    0          Se4/0       point2point
       Pop tag     10.10.45.0/30    0          Se3/0       point2point
24     Pop tag     10.10.46.0/30    0          Se3/0       point2point
25     Pop tag     10.10.47.0/30    0          Se3/0       point2point
26     Pop tag     10.10.48.0/30    0          Se3/0       point2point
27     Pop tag     10.10.56.0/30    0          Se4/0       point2point
28     28          10.10.67.0/30    0          Se4/0       point2point
       26          10.10.67.0/30    0          Se3/0       point2point
29     29          10.10.68.0/30    0          Se4/0       point2point
       27          10.10.68.0/30    0          Se3/0       point2point
30     28          10.10.78.0/30    0          Se3/0       point2point
31     Pop tag     10.10.10.3/32    6069       Et1/0       10.10.23.2
32     Untagged    192.168.11.0/24[V]   \
                                    0          Et2/0       10.10.112.1
33     Untagged    192.168.1.0/24[V] 0         Et0/0       10.10.12.1
34     29          10.10.10.3 1 [6]  0         Se3/0       point2point
35     Aggregate   10.10.112.0/30[V] 0
36     Aggregate   10.10.12.0/30[V]  0

[T]     Forwarding through a TSP tunnel.
        View additional tagging info with the 'detail' option
R2#
```

A. Packets to the IP address 10.10.10.5/32 will be tagged with "17" toward the next hop.
B. Label "19" will be advertised to MPLS neighbors so that they can use this label to reach the IP address 10.10.10.6/32.
C. IP address 10.10.10.4/32 is directly connected to the neighbor router on serial 3/0.
D. Packets arriving with label "17" will be forwarded without any label toward serial 4/0.
E. Packets arriving with label "20" will be forwarded with label "21" after label-swapping.
F. Label "20" is advertised to MPLS neighbors so that they can use this information to reach the prefix 10.10.10.8/32.

Answer: C,D,E,F

Explanation:

Label stacking is the encapsulation of an MPLS packet inside another MPLS packet – that is, adding an MPLS header “on top of” (hence stacking) an existing MPLS header. The result of stacking is the ability to tunnel one MPLS LSP inside another LSP.

The primary advantage of LDP is that is scales well. It signals LSPs hop-by-hop, and so routers along the path do not have to maintain state for each LSP. Therefore LDP is useful in edge

applications such as VPNs where hundreds or thousands of LSPs are originated and terminated. But LDP has no traffic engineering capabilities; it just follows the IGP shortest path to find LSP end-points.

A central concept to MPLS is the Forwarding Equivalence Class (FEC), and it's something many people new to the technology struggle to understand. So in this post I'd like to discuss FECs and their role in MPLS.

An FEC is a set of packets that a single router:

(1) Forwards to the same next hop;

(2) Out the same interface; and

(3) With the same treatment (such as queuing).

FECs are nothing new. Every router performing generic IP forwarding determines the next hop to which the packet is to be forwarded, the interface out which the packet is sent to get to that next hop, and how to queue the packet for that interface. But we don't often hear those very basic procedures presented as "determining what FEC a packet belongs to."

QUESTION NO: 247

Which three are needed to run VRF lite for IPv6 on a router? (Choose three.)

A. VRF definition for IPv6
B. MP BGP for IPv6
C. LDP
D. VRF-enabled routing protocol
E. VRF-enabled interface

Answer: A,D,E

Explanation:

The word VRF stands for Virtual Routing and Forwarding, this feature is used to create multiple instances of the routing table on the same routing device. VRFs are usually used in conjunction with MPLS VPN to separate the traffic of multiple MPLS VPN customers. VRF Lite feature is part of Cisco's network virtualization portfolio.

VRF Lite means VRF without the need to run MPLS in the network. VRF Lite allows the network administrator to create multiple routing instances on the same routing device within the enterprise. VRF Lite can be useful when you need to isolate traffic between two networks sharing the same routing platform or if you have multiple networks with overlapping addresses sharing the same physical network. Multiple instances of routing protocols can be used for different VRFs on the same device to exchange routes dynamically with a direct connected device

VRF Lite Configuration:

R2 is connected via Ethernet to R5. Two VRFs (VRF-LITE-A & B) are configured to demonstrate L3 traffic isolation. I am using static routes for this example but dynamic routing protocols can be used. R2 Configuration:

```
ip vrf VRF-LITE-A
rd 100:1
!
ip vrf VRF-LITE-B
rd 100:2
!-- Assign interfaces to VRF
interface FastEthernet0/1.25
encapsulation dot1Q 25
ip vrf forwarding VRF-LITE-A
ip address 25.25.25.2 255.255.255.0
!
interface FastEthernet0/1.52
encapsulation dot1Q 52
ip vrf forwarding VRF-LITE-B
ip address 52.52.52.2 255.255.255.0
interface Loopback20
ip vrf forwarding VRF-LITE-A
ip address 20.20.20.20 255.255.255.255
!
interface Loopback22
ip vrf forwarding VRF-LITE-B
ip address 22.22.22.22 255.255.255.255
ip route vrf VRF-LITE-A 50.50.50.50 255.255.255.255 25.25.25.5
ip route vrf VRF-LITE-B 55.55.55.55 255.255.255.255 52.52.52.5
```

R5 Configuration:

```
ip vrf VRF-LITE-A
rd 100:1
!
ip vrf VRF-LITE-B
rd 100:2
interface Loopback50
ip vrf forwarding VRF-LITE-A
ip address 50.50.50.50 255.255.255.255
!
interface Loopback55
ip vrf forwarding VRF-LITE-B
```

ip address 55.55.55.55 255.255.255.255
!
interface FastEthernet0/1.25
encapsulation dot1Q 25
ip vrf forwarding VRF-LITE-A
ip address 25.25.25.5 255.255.255.0
!
interface FastEthernet0/1.52
encapsulation dot1Q 52
ip vrf forwarding VRF-LITE-B
ip address 52.52.52.5 255.255.255.0
ip route vrf VRF-LITE-A 20.20.20.20 255.255.255.255 25.25.25.2
ip route vrf VRF-LITE-B 22.22.22.22 255.255.255.255 52.52.52.2
Operation Verification: The following tests were taken from R2 only, the same can be done on R5 for verification.

R2#sh ip route vrf VRF-LITE-A
Routing Table: VRF-LITE-A
!-- output omitted----------
Gateway of last resort is not set
50.0.0.0/32 is subnetted, 1 subnets
S 50.50.50.50 [1/0] via 25.25.25.5
20.0.0.0/32 is subnetted, 1 subnets
C 20.20.20.20 is directly connected, Loopback20
25.0.0.0/24 is subnetted, 1 subnets
C 25.25.25.0 is directly connected, FastEthernet0/1.25
R2#sh ip route vrf VRF-LITE-B
Routing Table: VRF-LITE-B
!--output omitted----------
Gateway of last resort is not set
55.0.0.0/32 is subnetted, 1 subnets
S 55.55.55.55 [1/0] via 52.52.52.5
52.0.0.0/24 is subnetted, 1 subnets
C 52.52.52.0 is directly connected, FastEthernet0/1.52
22.0.0.0/32 is subnetted, 1 subnets
C 22.22.22.22 is directly connected, Loopback22
R2#ping 50.50.50.50
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 50.50.50.50, timeout is 2 seconds:

Success rate is 0 percent (0/5)

R2#ping vrf VRF-LITE-A 50.50.50.50
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 50.50.50.50, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/143/396 ms
R2#ping 55.55.55.55
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 55.55.55.55, timeout is 2 seconds:

Success rate is 0 percent (0/5)
R2#ping vrf VRF-LITE-B 55.55.55.55
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 55.55.55.55, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 20/133/340 ms
Reference
http://www.networkers-online.com/blog/2009/02/vrf-lite/

QUESTION NO: 248

Which statement correctly describes the disabling of IP TTL propagation in an MPLS network?

A. The TTL field from the IP packet is copied into the TTL field of the MPLS label header at the ingress edge LSR.
B. TTL propagation cannot be disabled in an MPLS domain.
C. TTL propagation is only disabled on the ingress edge LSR,
D. The TTL field of the MPLS label header is set to 255.
E. The TTL field of the IP packet is set to 0.

Answer: D
Explanation: Explanation
Time-to-Live (TTL) is an 8-bit field in the MPLS label header which has the same function in loop detection of the IP TTL field. Recall that the TTL value is an integer from 0 to 255 that is decremented by one every time the packet transits a router. If the TTL value of an IP packet becomes zero, the router discards the IP packet, and an ICMP message stating that the TTL expired in transit is sent to the source IP address of the IP packet. This mechanism prevents an IP packet from being routed continuously in case of a routing loop.

By default, the TTL propagation is enabled so a user can use traceroute command to view all of

the hops in the network.
We can disable MPLS TTL propagation with the no mpls ip propagate-ttl command under global configuration.

When entering a label-switched path (LSP), the edge router will use a fixed TTL value (255) for the first label.

This increases the security of your MPLS network by hiding provider network from customers.

QUESTION NO: 249

Which of these tables is used by an LSR to perform a forwarding lookup for a packet destined to an address within an RFC 4364 VPN?

A. CEF
B. FIB
C. LFIB
D. IGP

Answer: C
Explanation:
Notice: The term Label Switch Router (LSR) refers to any router that has awareness of MPLS labels
Label Forwarding Information Base (LFIB) is responsible for forwarding incoming packets based on label as it holds necessary label information, as well as the outgoing interface and next-hop information

QUESTION NO: 250

Which two of these parameters are used to determine a forwarding equivalence class? (Choose two.

A. IP prefix
B. Layer 2 circuit
C. RSVP request from CE for bandwidth reservation
D. BGP MED value

Answer: A,B

Explanation:

A Forwarding Equivalence Class (FEC) is a class of packets that should be forwarded in the same manner (i.e. over the same path). A FEC is not a packet, nor is it a label. A FEC is a logical entity created by the router to represent a class (category) of packets. When a packet arrives at the ingress router of an MPLS domain, the router parses the packet's headers, and checks to see if the packet matches a known FEC (class). Once the matching FEC is determined, the path and outgoing label assigned to that FEC are used to forward the packet.

FECs are typically created based on the IP destinations known to the router, so for each different destination a router might create a different FEC, or if a router is doing aggregation, it might represent multiple destinations with a single FEC (for example, if those destinations are reachable through the same immediate next hop anyway). The MPLS framework, however, allows for the creation of FECs using advanced criteria like source and destination address pairs, destination address and TOS, etc.

Topic 5, Implement IP Multicast

QUESTION NO: 251

Which two commands are required to enable multicast on a router, knowing that the receivers only support IGMPv2? (Choose two.)

A. ip pim rp-address
B. ip pim ssm
C. ip pim sparse-mode
D. ip pim passive

Answer: A,C

Explanation:

Sparse mode logic (pull mode) is the opposite of Dense mode logic (push mode), in Dense mode it is supposed that in every network there is someone who is requesting the multicast traffic so PIM-DM routers begin by flooding the multicast traffic out of all their interfaces except those from where a prune message is received to eliminate the “leaf” from the multicasting tree (SPT), the Source-Based Tree (S, G); as opposed to Sparse mode that send the traffic only if someone explicitly requested it.

Not like Dense mode, which build a separated source-based tree (S, G) between the source and the requester of the traffic, Sparse mode mechanism is based on a fixed point in the network

named Rendez-Vous point.

All sources will have to register with the RP to which they send their traffic and thereby build a source-based tree (S, G) between them and the RP (not with the final multicast receiver like in PIM-DM) and all PIM-SM routers, “whatever” multicast traffic they are requesting, have to register with the RP and build a shared-tree (*. G)

Reference
http://tools.ietf.org/html/rfc2236
http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a00800b0871.shtml
http://www.cisco.com/en/US/tech/tk828/technologies_tech_note09186a0080094821.shtml#sparsemode

QUESTION NO: 252

Which type of domains is interconnected using Multicast Source Discovery Protocol?

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

Answer: A
Explanation: Multicast Source Discovery Protocol (MSDP) is a Protocol Independent Multicast (PIM) family multicast routing protocol defined by Experimental RFC 3618. MSDP interconnects multiple IPv4 PIM Sparse-Mode (PIM-SM) domains which enables PIM-SM to have Rendezvous Point (RP) redundancy and inter-domain multicasting.

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

QUESTION NO: 253

Which two multicast address ranges are assigned as source-specific multicast destination addresses and are reserved for use by source-specific applications and protocols? (Choose two.)

A. 232.0.0.0/8
B. 239.0.0.0/8
C. 232.0.0.0/4
D. FF3x::/32
E. FF2x::/32
F. FF3x::/16

Answer: A,D
Explanation: Source-specific multicast (SSM) is a method of delivering multicast packets in which the only packets that are delivered to a receiver are those originating from a specific source address requested by the receiver. By so limiting the source, SSM reduces demands on the network and improves security.

SSM requires that the receiver specify the source address and explicitly excludes the use of the (*, G) join for all multicast groups in RFC 3376, which is possible only in IPv4's IGMPv3 and IPv6's MLDv2.
Source-specific multicast is best understood in contrast to any-source multicast (ASM). In the

ASM service model a receiver expresses interest in traffic to a multicast address. The multicast network must
1. discover all multicast sources sending to that address, and
2. route data from all sources to all interested receivers.

This behavior is particularly well suited to groupware applications where
1. all participants in the group want to be aware of all other participants, and
2. the list of participants is not known in advance.
The source discovery burden on the network can become significant when the number of sources is large.

In the SSM service model, in addition to the receiver expressing interest in traffic to a multicast address, the receiver expresses interest in receiving traffic from only one specific source sending to that multicast address.
This relieves the network of discovering many multicast sources and reduces the amount of multicast routing information that the network must maintain.
SSM requires support in last-hop routers and in the receiver's operating system. SSM support is not required in other network components, including routers and even the sending host. Interest in multicast traffic from a specific source is conveyed from hosts to routers using IGMPv3 as specified in RFC 4607.
SSM destination addresses must be in the ranges 232.0.0.0/8 for IPv4 or FF3x::/96 for IPv6.

Reference
http://en.wikipedia.org/wiki/Source-specific_multicast

QUESTION NO: 254

Refer to the exhibit.

```
r15#show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, 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

(*, 239.1.1.1), 00:00:29/stopped, RP 172.255.255.14, flags: SJCL
  Incoming interface: Ethernet0/0, RPF nbr 192.168.2.2
  Outgoing interface list:
    Loopback15, Forward/Sparse, 00:00:29/00:02:30
```

What does the incoming interface of the above (*, G) entry indicate?

A. the interface closest to the source, according to the unicast routing table
B. the interface where an IGMP join has been received
C. the interface with the highest IP address
D. the last interface to hear a PIM (*, G) join
E. the interface closest to the RP, according to the unicast routing table

Answer: E
Explanation:
Source Trees
A source tree is the simplest form of distribution tree. The source host of the multicast traffic is located at the root of the tree, and the receivers are located at the ends of the branches. Multicast traffic travels from the source host down the tree toward the receivers. The forwarding decision on which interface a multicast packet should be transmitted out is based on the multicast forwarding table. This table consists of a series of multicast state entries that are cached in the router. State entries for a source tree use the notation (S, G) pronounced S comma G. The letters represents the IP address of the source, and G represents the group address.

Shared Trees
Shared trees differ from source trees in that the root of the tree is a common point somewhere in the network.

This common point is referred to as the rendezvous point (RP). The RP is the point at which receivers join to learn of active sources. Multicast sources must transmit their traffic to the RP. When receivers join a multicast group on a shared tree, the root of the tree is always the RP, and multicast traffic is transmitted from the RP down toward the receivers. Therefore, the RP acts as a go-between for the sources and receivers. An RP can be the root for all multicast groups in the network, or different ranges of multicast groups can be associated with different RPs.

Multicast forwarding entries for a shared tree use the notation (*, G), which is pronounced star comma G. This is because all sources for a particular group share the same tree. (The multicast groups go to the same RP.)

Therefore, the * or wildcard represents all sources.
Additional Information from Microsoft
Multicast traffic from source 162.10.4.1 (for example) uses the RPT, meaning the source sends it to the RP rather than to the multicast group (the router would denote this by having a (*, G) entry rather than a (S, G) entry). Before sending this traffic, Router 1 checks its unicast routing table to see if packets from the RP are arriving on the correct interface. In this case they are, because they arrive on interface I1, and the packets are forwarded.

Reference
http://technet.microsoft.com/en-us/library/bb742462.aspx

QUESTION NO: 255

Apart from interdomain multicast routing, what else is MSDP used for?

A. Source Specific Multicast and IGMPv2
B. Announcing multicast sources to BGP speakers
C. Anycast RP
D. Intradomain multicast routing

Answer: C
Explanation:
Reference
http://www.cisco.com/en/US/docs/ios/12_2/ip/configuration/guide/1cfmsdp_ps1835_TSD_Products_Configuration_Guide_Chapter.html

QUESTION NO: 256

Which IGMPv2 message contains a non-zero "Max Response Time"?

A. Membership Query
B. Membership Report
C. Membership Delay
D. Backward Compatible IGMPv1 Report Message

Answer: A

Explanation:
The Max Response Time field is used only in Membership Query messages. It specifies the maximum allowed time before sending a responding report in units of 1/10 second. In all other messages, it is set to zero by the sender and ignored by receivers.

QUESTION NO: 257

Refer to the exhibit.

```
Internet Protocol, Src Addr: 192.168.xxx.xxx (192.168.xxx.xxx), Dst Addr: 239.255.255.250 (239.255.255.250)
   Version: 4
   Header length: 24 bytes
   Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
      0000 00.. = Differentiated Services Codepoint: Default (0x00)
      .... ..0. = ECN-Capable Transport (ECT): 0
      .... ...0 = ECN-CE: 0
   Total Length: 32
   Identification: 0x58dd (22749)
   Flags: 0x00
      0... = Reserved bit: Not set
      .0.. = Don't fragment: Not set
      ..0. = More fragments: Not set
   Fragment offset: 0
   Time to live: 1
   Protocol: IGMP (0x02)
   Header checksum: 0x10f4 (correct)
   Source: 192.168.xxx.xxx (192.168.xxx.xxx)
   Destination: 239.255.255.250 (239.255.255.250)
   Options: (4 bytes)
      Router Alert: Every router examines packet
```

ActualTests

Which IGMPv2 message is displayed in the output of the packet capture?

A. General Query
B. Membership Report
C. Membership Query
D. Membership Delay

Answer: B

Explanation:

Example traffic

```
(Certain addresses munged)

No.     Time        Source                Destination           Protocol Info
     38 18.419525   192.168.xxx.xxx        239.255.255.250       IGMP     V2
Membership Report

Frame 38 (46 bytes on wire, 46 bytes captured)
    Arrival Time: May 29, 2005 16:45:58.272406000
    Time delta from previous packet: 0.481171000 seconds
    Time since reference or first frame: 18.419525000 seconds
    Frame Number: 38
    Packet Length: 46 bytes
    Capture Length: 46 bytes
    Protocols in frame: eth:ip:igmp
Ethernet II, Src: xx:xx:xx:xx:xx:xx, Dst: 01:00:5e:7f:ff:fa
    Destination: 01:00:5e:7f:ff:fa (01:00:5e:7f:ff:fa)
    Source: xx:xx:xx:xx:xx:00 (ABCD:xx:00)
    Type: IP (0x0800)
Internet Protocol, Src Addr: 192.168.xxx.xxx (192.168.xxx.xxx), Dst Addr:
239.255.255.250 (239.255.255.250)
    Version: 4
    Header length: 24 bytes
    Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
        0000 00.. = Differentiated Services Codepoint: Default (0x00)
        .... ..0. = ECN-Capable Transport (ECT): 0
        .... ...0 = ECN-CE: 0
    Total Length: 32
    Identification: 0x58dd (22749)
    Flags: 0x00
        0... = Reserved bit: Not set
        .0.. = Don't fragment: Not set
        ..0. = More fragments: Not set
    Fragment offset: 0
    Time to live: 1
    Protocol: IGMP (0x02)
    Header checksum: 0x10f4 (correct)
    Source: 192.168.xxx.xxx (192.168.xxx.xxx)
    Destination: 239.255.255.250 (239.255.255.250)
    Options: (4 bytes)
        Router Alert: Every router examines packet
Internet Group Management Protocol
    IGMP Version: 2
    Type: Membership Report (0x16)
    Max Response Time: 0,0 sec (0x00)
    Header checksum: 0xfa04 (correct)
    Multicast Address: 239.255.255.250 (239.255.255.250)

0000  01 00 5e 7f ff fa xx xx xx xx xx xx 08 00 46 00   ..^.....Z#.s..F.
0010  00 20 58 dd 00 00 01 02 10 f4 c0 a8 xx xx ef ff   . X..........d..
0020  ff fa 94 04 00 00 16 00 fa 04 ef ff ff fa         ..............
```

ActualTests

Reference

http://wiki.wireshark.org/IGMP

QUESTION NO: 258

What is Phantom RP used for?

A. it is used for load balancing in bidirectional PIM

B. it is used for redundancy in bidirectional PIM

C. it is used for redundancy in PIM-SM
D. it is used for load balancing in PIM-SM

Answer: B
Explanation:
Phantom RP
In Bidirectional PIM (Bidir-PIM), the RP does not have an actual protocol function. The RP acts as a routing vector in which all the traffic converges. The RP can be configured as an address that is not assigned to any particular device called a Phantom RP. This means that the RP address does not need to reside on a physical router interface, but can just be an address in a subnet. The RP can also be a physical router, but it is not necessary.

Reference
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6552/whitepaper_c11-508498.html

QUESTION NO: 259

What does the command ip pim send-rp-discovery scope 16 achieve?

A. it enables PIM BSR
B. it makes the router an Auto-RP mapping agent
C. it advertises the router as an RP with Auto-RP
D. it makes the router an Auto-RP client

Answer: B
Explanation:
If no ip pim rp-announce-filter commands are configured, a router enabled to be an RP mapping agent (using the ip pim send-rp-discovery command) will accept all RP announcements for all groups from all C-RPs.
Configure one or more ip pim rp-announce-filter commands on RP mapping agents to filter unwanted RP messages.

Command	Description
access-list (IP standard)	Defines a standard IP access list.
ip pim send-rp-discovery	Configures the router to be an RP mapping agent.

ip pim send-rp-announce
To use Auto-RP to configure groups for which the router will act as a rendezvous point (RP), use

the ip pim send-rp-announce command in global configuration mode. To unconfigure this router as an RP, use the no form of this command.
ip pim [vrf vrf-name] send-rp-announce {interface-type interface-number | ip-address} scopettl-value [group-list access-list] [interval seconds] [bidir] no ip pim [vrf vrf-name] send-rp-announce {interface-type interface-number | ip-address}

Syntax Description

vrf	(Optional) Supports the multicast Virtual Private Network (VPN) routing and forwarding (VRF) instance.
vrf-name	(Optional) Name assigned to the VRF.
interface-type interface-number	Interface type and number that is used to define the RP address. No space is required between the values.
ip-address	IP address of the RP for the group. The IP address must be a directly connected address. If the command is configured with this argument, the RP-announce message will be sourced by the interface to which this IP address is connected (that is, the source address in the IP header of the RP-announce message is the IP address of that interface).
scope *ttl-value*	Specifies the time-to-live (TTL) value that limits the number of Auto-RP announcements.
group-list *access-list*	(Optional) Specifies the standard IP access list number or name that defines the group prefixes that are advertised in association with the RP address. The access list name cannot contain a space or quotation mark, and must begin with an alphabetic character to avoid confusion with numbered access lists.
interval *seconds*	(Optional) Specifies the interval between RP announcements in seconds. The total holdtime of the RP announcements is automatically set to three times the value of the interval. The default interval is 60 seconds.
bidir	(Optional) Indicates that the multicast groups specified by the *access-list* argument are to operate in bidirectional mode. If the command is configured without this keyword, the groups specified will operate in Protocol Independent Multicast sparse mode (PIM-SM).

Reference
http://www.cisco.com/en/US/docs/ios/ipmulti/command/reference/imc_04.html

QUESTION NO: 260

Which command can be used to check the assignment of RPs to multicast groups?

A. show ip pim rendez-vous
B. show ip rpf rp mapping
C. show ip pim rp mapping
D. show ip pim rp info
E. show ip pim

Answer: C

Explanation:

Explanation show ip pim rp mapping

Use this command to check the RP assignment by multicast group range, and to verify that the source of RP learning (static or auto-RP) and the mapping are correct. If you find an error, check the local router configuration or auto-RP configuration.

R1# show ip pim rp mapping
PIM Group-to-RP Mappings
Group(s) 224.0.1.40/32
RP 192.168.7.2 (?), v1
Info source: local, via Auto-RP
Uptime: 2d00h, expires: never
Group(s): 224.0.0.0/4, Static
RP: 192.168.7.2 (?)

Reference
http://www.cisco.com/en/US/tech/tk828/technologies_tech_note09186a0080093f21.shtml#showippimrpmapping

QUESTION NO: 261

Which three message types are valid PIMv2 message types? (Choose three.)

A. Register
B. Register-Stop
C. Join/Prune
D. Reject
E. Register-Prune
F. Register-Join

Answer: A,B,C

Explanation:

PIMv2 message types:
hello: 0
Register (used in PIM-SM only): 1
Register-Stop (used in PIM-SM only): 2
Join/Prune: 3
Bootstrap (used in PIM-SM only): 4

Assert:5
Graft (used in PIM-DM only): 6
Graft-Ack (used in PIM-DM only): 7
Candidate-RP-Advertisement (used in PIM-SM only): 8

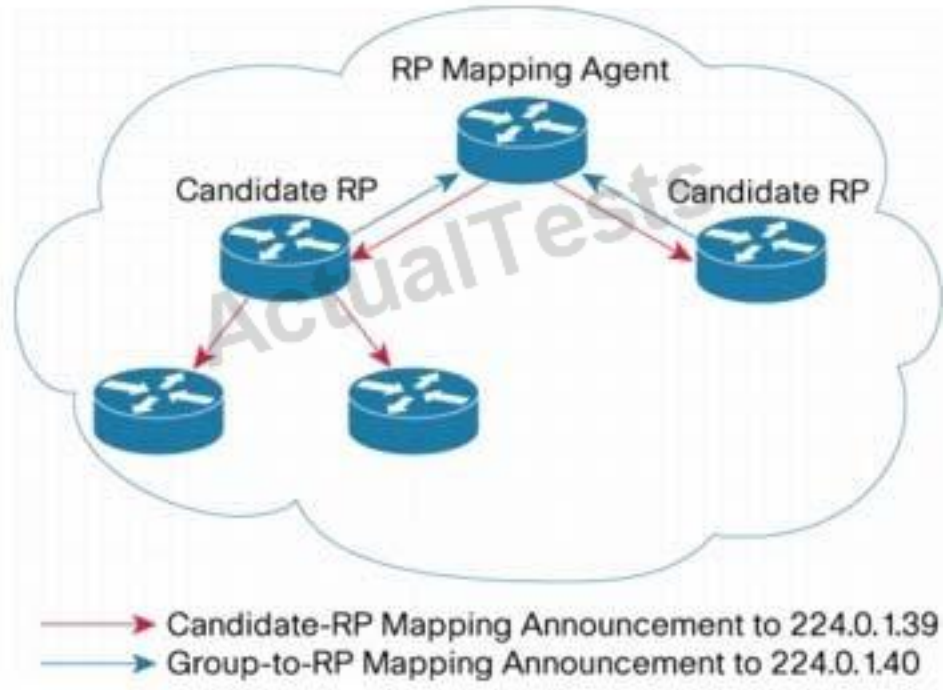
QUESTION NO: 262

Which multicast address is reserved for cisco-rp-discovery?

A. 224.0.1.37
B. 224.0.1.38
C. 224.0.1.39
D. 224.0.1.40

Answer: D
Explanation:
Auto-RP is a mechanism to automate distribution of RP information in a multicast network. The Auto-RP mechanism operates using two basic components, the candidate RPs and the RP mapping agents.
Candidate RPs advertise their willingness to be an RP via "RP-announcement" messages. These messages are periodically sent to a reserved well-known group 224.0.1.39 (CISCO-RP-ANNOUNCE).
RP mapping agents join group 224.0.1.39 and map the RPs to the associated groups. The RP mapping agents advertise the authoritative RP-mappings to another well-known group address 224.0.1.40 (CISCORP-DISCOVERY). All PIM routers join 224.0.1.40 and store the RP-mappings in their private cache.
Figure 2 shows the Auto-RP mechanism where the RP mapping agent periodically multicasts the RP information that it receives to the Cisco-RP-Discovery group.

Figure 2. Auto-RP Mechanism



Reference
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6552/whitepaper_c11-508498.html

QUESTION NO: 263

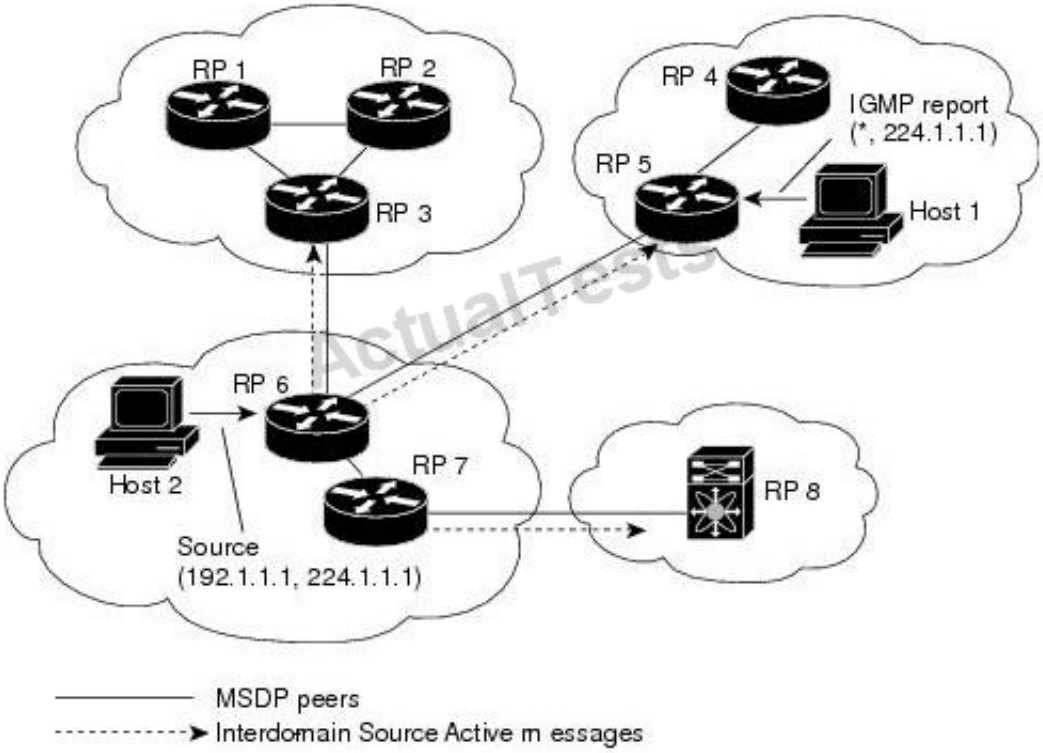
Which message type is constructed and sent from an MSDP router to its MSDP peers when it receives a PIM register message (knowing that the MSDP router is also configured as an RP for the PIM domain)?

A. Source-Active Message
B. PIM Join
C. PIM Hello
D. MSDP Register

Answer: A
Explanation:
Information About MSDP
You can use the Multicast Source Discovery Protocol (MSDP) to exchange multicast source information between multiple BGP-enabled Protocol Independent Multicast (PIM) sparse-mode domains.
When a receiver for a group matches the group transmitted by a source in another domain, the

rendezvous point (RP) sends PIM join messages in the direction of the source to build a shortest path tree. The designated router (DR) sends packets on the source-tree within the source domain, which may travel through the RP in the source domain and along the branches of the source-tree to other domains. In domains where there are receivers, RPs in those domains can be on the source-tree. The peering relationship is conducted over a TCP connection.
Figure 6-1 shows four PIM domains. The connected RPs (routers) are called MSDP peers because each RP maintains its own set of multicast sources. Source host 1 sends the multicast data to group 224.1.1.1. On RP 6, the MSDP process learns about the source through PIM register messages and generates Source-Active (SA) messages to its MSDP peers that contain information about the sources in its domain. When RP 3 and RP 5 receive the SA messages, they forward them to their MSDP peers. When RP 5 receives the request from host 2 for the multicast data on group 224.1.1.1, it builds a shortest path tree to the source by sending a PIM join message in the direction of host 1 at 192.1.1.1.

Figure 6-1 MSDP Peering Between RPs in Different PIM Domains



When you configure MSDP peering between each RP, you create a full mesh. Full MSDP meshing is typically done within an autonomous system, as shown between RPs 1, 2, and 3, but not across autonomous systems.

You use BGP to do loop suppression and MSDP peer-RPF to suppress looping SA messages. For more information about mesh groups, see the "MSDP Mesh Groups" section.

Note
You do not need to configure MSDP in order to use Anycast-RP (a set of RPs that can perform load balancing and failover) within a PIM domain. For more information, see the "Configuring a PIM Anycast-RP Set" section.
For detailed information about MSDP, see RFC 3618.

Reference
http://www.cisco.com/en/US/docs/switches/datacenter/sw/5_x/nx-os/multicast/configuration/guide/msdp.html

QUESTION NO: 264

Which two are differences between IGMPv2 and IGMPv3 reports? (Choose two.)

A. IGMPv3 has the ability to include or exclude source lists.
B. All IGMPv3 hosts send reports to destination address 224.0.0.22.
C. Only IGMPv2 reports may contain multiple group state records.
D. All IGMPv3 hosts send reports to destination address 224.0.0.23.
E. IGMPv2 does not support the Leave Group message.

Answer: A,B
Explanation:
Version 3 Reports are sent with an IP destination address of 224.0.0.22, to which all IGMPv3-capable multicast routers listen. A system that is operating in version 1 or version 2 compatibility modes sends version 1 or version 2 Reports to the multicast group specified in the Group Address field of the Report.
There are a number of different types of Group Records that may be included in a Report message:

A "Current-State Record“ (in response to a Query)
MODE_IS_INCLUDE INCLUDE()
MODE_IS_EXCLUDE EXCLUDE()

A "Filter-Mode-Change Record“ (when the filter mode change)
CHANGE_TO_INCLUDE_MODE TO_IN()
CHANGE_TO_EXCLUDE_MODE TO_EX()

A "Source-List-Change Record“ (when the source list change)
ALLOW_NEW_SOURCES ALLOW()

BLOCK_OLD_SOURCES BLOCK()

Reference
http://alor.antifork.org/talks/IGMP-v3.ppt

QUESTION NO: 265

Which two statements about SA caching are true? (Choose two.)

A. Caching allows pacing of MSDP messages.
B. Caching reduces join latency.
C. Caching should not be done by an MSDP speaker.
D. Caching is used to update the BGP MDT address family.

Answer: A,B
Explanation:
A MSDP speaker MUST cache SA messages. Caching allows pacing of MSDP messages as well as reducing join latency for new receivers of a group G at an originating RP which has existing MSDP (S, G) state. In addition, caching greatly aids in diagnosis and debugging of various problems.

Reference
http://www.ietf.org/rfc/rfc3618.txt

QUESTION NO: 266

In PIM-SM what control plane signaling must a multicast source perform before it begins to send multicast traffic to a group?

A. The source must send a PIM Register message to the rendezvous point (RP).
B. The source must first join the multicast group using IGMP before sending.
C. The source must perform a Request to Send (RTS) and Clear to Send (CTS) handshake with the PIM designated router (DR).
D. No control plane signaling needs to be performed; the source can simply begin sending on the local subnet.

Answer: D

Explanation:
The most common type of multicast issue is the RPF Failure. RPF checks are used both at the control and data plane of multicast routing. Control plane involves PIM signaling some PIM messages are subject to RPF checks. For example, PIM (*,G) Joins are sent toward the shortest path to RP. Next, the BSR/RP address in the BSR messages is subject to RPF check as well. Notice that this logic does not apply to PIM Register messages the unicast register packet may arrive on any interface. However, RPF check is performed on the encapsulated multicast source to construct the SPT toward the multicast source.

Data plane RPF checks are performed every time a multicast data packet is received for forwarding. The source IP address in the packet should be reachable via the receiving interface, or the packet is going to be dropped. Theoretically, with PIM Sparse-Mode RPF checks at the control plane level should preclude and eliminate the data-plane RPF failures, but data-plane RPF failures are common during the moments of IGP reconvergence and on multipoint non-broadcast interfaces.

PIM Dense Mode is different from SM in the sense that data-plane operations preclude control-plane signaling. One typical irresolvable RPF problem with PIM Dense mode is known as split-horizon forwarding, where packet received on one interface, should be forwarded back out of the same interface in the hub-and-spoke topology. The same problem may occur with PIM Sparse mode, but this type of signaling allows for treating the NBMA interface as a collection of point-to-point links by the virtue of PIM NBMA mode.

QUESTION NO: 267

Which of these statements about PIM join messages in classic PIM-SM is correct?

A. PIM join messages are sent every 60 seconds to refresh the upstream router's mroute state for the multicast tree.
B. Routers send a PIM join acknowledgement in response to each PIM join message received from a downstream router.
C. PIM join messages are only sent when the multicast distribution tree is first being established.
D. PIM join messages are sent every three minutes to refresh the upstream router's mroute state for the multicast tree.

Answer: A
Explanation: Explanation
PIM Sparse Mode uses an explicit request approach, where a router has to ask for the multicast feed with a PIM Join message. PIM Sparse Mode is indicated when you need more precise

control, especially when you have large volumes of IP multicast traffic compared to your bandwidth. PIM Sparse Mode scales rather well, because packets only go where they are needed, and because it creates state in routers only as needed. There can be different RP's for different multicast groups, which is one way to spread the load. There is usually one RP per multicast group. Redundancy of RP's is an advanced topic, and requires a little deeper expertise.

One way to do this is with the MSDP protocol (possible later article in the series).
PIM Join message is sent towards a Source (or for PIM-SM, possibly towards an RP), based on unicast routing. The Join message says in effect "we need a copy of the multicasts over here". It connects the sender of the Join and intervening routers to any existing multicast tree, all the way back to the target of the Join if necessary. A Prune message says in effect "we no longer need this over here". A router receiving a Prune sees whether it has any other interfaces requiring the multicast flow, and if not, sends its own Prune message.

One advanced technique is to arrange a separate and perhaps different copy of the unicast routing information just for multicast purposes. This allows "steering" of the Join messages. Multiprotocol BGP, MBGP, for multicast, is one way to do this

PIM-SM SPT-Switchover is controlled by the global configuration command:

ip pim spt-threshold {*kbps* | **infinity**} [**group-list** *access-list-number*]

kbps	Traffic rate in kilobits per second.
infinity	The specified groups will use the shared-tree.
group-list *access-list-number*	Optional. Determines which groups to apply the threshold.

ActualTests

By default, a PIM-SM router sends periodic Join/Prune messages every 60 seconds. To alter this interval, use the global configuration command

ip pim message-interval *seconds*

seconds	Value in the range 1 to 65535

All PIM-SM-enabled routers should be configured with the same message interval time. A router will be pruned from a group if a Join message is not received in the message interval. The default value is three minutes.

Reference
http://ciscoarticles.com/Cisco-Multicast-Routing-and-Switching/PIM-SM-Version-2-RP-

Selection.html

QUESTION NO: 268

The ip pim autorp listener command is used to do which of these?

A. enable a Cisco router to "passively" listen to Auto-RP packets without the router actively sending or forwarding any of the packets
B. allow Auto-RP packets in groups 224.0.1.39 and 224.0.1.40 to be flooded in dense mode out interfaces configured with the ip pim sparse-mode command
C. enable the use of Auto-RP on a router
D. configure the router as an Auto-RP mapping agent

Answer: B
Explanation: Explanation
The IP Pim autorp listener allows the Group 224.0.0.39 & 224.0.0.40 to be dense flooded. As the RP announces 224.0.0.39 to the mapping agent and the mapping agent announces 224.0.0.40 to all routers part of the group. where it can be applied:
It can be applied when dense mode is not configured, for example if you have Sparse-mode Multicast Network, and you need not to statically define your RP or use Autorp.

QUESTION NO: 269

In order to configure two routers as anycast RPs, which of these requirements, af a minimum, must be satisfied?

A. Multicast Source Discovery Protocol mesh-groups must be configured between the two anycast RPs.
B. The RPs must be within the same IGP domain.
C. Multicast Source Discovery Protocol must be configured between the two anycast RPs.
D. The two anycast RPs must be IBGP peers.

Answer: C
Explanation:
Explanation
Multicast Source Discovery Protocol (MSDP) is a mechanism to connect multiple PIM sparse-mode (SM) domains. MSDP allows multicast sources for a group to be known to all rendezvous

point(s) (RPs) in different domains. Each PIM-SM domain uses its own RPs and need not depend on RPs in other domains. An RP runs MSDP over TCP to discover multicast sources in other domains. An RP in a PIM-SM domain has an MSDP peering relationship with MSDP-enabled routers in another domain. The peering relationship occurs over a TCP connection, where primarily a list of sources sending to multicast groups is exchanged. The TCP connections between RPs are achieved by the underlying routing system. The receiving RP uses the source lists to establish a source path. The purpose of this topology is to have domains discover multicast sources in other domains. If the multicast sources are of interest to a domain that has receivers, multicast data is delivered over the normal, source-tree building mechanism in PIM-SM. MSDP is also used to announce sources sending to a group. These announcements must originate at the domain's RP.

MSDP depends heavily on (M)BGP for interdomain operation. It is recommended that you run MSDP in RPs in your domain that are RPs for sources sending to global groups to be announced to the internet.

Each MSDP peer receives and forwards the SA message away from the originating RP to achieve "peer- RPF flooding." The concept of peer-RPF flooding is with respect to forwarding SA messages. The router examines the BGP or MBGP routing table to determine which peer is the next hop toward the originating RP of the SA message. Such a peer is called an "RPF peer" (Reverse-Path Forwarding peer). The router forwards the message to all MSDP peers other than the RPF peer.

If the MSDP peer receives the same SA message from a non-RPF peer toward the originating RP, it drops the message. Otherwise, it forwards the message on to all its MSDP peers.
When an RP for a domain receives an SA message from an MSDP peer, it determines if it has any group members interested in the group the SA message describes. If the (*,G) entry exists with a nonempty outgoing interface list, the domain is interested in the group, and the RP triggers an (S,G) join toward the source.

QUESTION NO: 270

Which two of these statements correctly describe classic PIM-SM? (Choose two.)

A. The lOS default is for a last-hop router to trigger a switch to the shortest path tree as soon as a new source is detected on the shared tree.
B. The lOS default is for every one of the routers on the shared tree to trigger a switch to the shortest path tree as soon as a new source is detected on the shared tree.
C. The default behavior of switching to the shortest path tree as soon as a new source is detected

on the shared tree can be disabled by setting the value in the ip pirn spt-threshold command to "infinity.
D. The default behavior of switching to the shortest path tree as soon as a new source is detected on the shared tree can be disabled by setting the value in the ip pirn spt-threshold command to "zero."

Answer: A,C
Explanation: Explanation
They are checking you for syntax ip pim spt-threshold command to “infinity” is the right answer.
same source as above:
IP pim spt-threshold [vrf vrf-name] spt-threshold {kbps | infinity} [group-list access-list]
To configure when a Protocol Independent Multicast (PIM) leaf router should join the shortest path source tree for the specified group infinity Causes all sources for the specified group to use the shared tree.

Reference
http://www.cisco.com/en/US/docs/ios/12_2/ipmulti/command/reference/1rfmult2.html#wp1020111

QUESTION NO: 271

Which statement is true of a source that wants to transmit multicast traffic to group 239.1.1.1?

A. Before sending traffic, it must first join multicast group 239.1.1.1 by sending an IGMPv2 membership report to the default router on the local subnet.
B. It must send an IGMPv2 Request to Send packet and then wait for an IGMPv2 Clear to Send packet from the IGMPv2 querier router on the local subnet
C. It may begin transmitting multicast traffic to the group only when there is no other host transmitting to the group on the local subnet.
D. It may transmit multicast traffic to the group at any time.

Answer: D
Explanation:
IP multicast is a method of sending Internet Protocol (IP) datagrams to a group of interested receivers in a single transmission. It is often employed for streaming media applications on the Internet and private networks.
The method is the IP-specific version of the general concept of multicast networking. It uses specially reserved multicast address blocks in IPv4 and IPv6. In IPv6, IP multicast addressing replaces broadcast addressing as implemented in IPv4.

Key concepts in IP multicast include an IP multicast group address,[3] a multicast distribution tree

and receiver driven tree creation.[4]
An IP multicast group address is used by sources and the receivers to send and receive multicast messages. Sources use the group address as the IP destination address in their data packets. Receivers use this group address to inform the network that they are interested in receiving packets sent to that group. For example, if some content is associated with group 239.1.1.1, the source will send data packets destined to 239.1.1.1. Receivers for that content will inform the network that they are interested in receiving data packets sent to the group 239.1.1.1. The receiver joins 239.1.1.1. The protocol typically used by receivers to join a group is called the Internet Group Management Protocol (IGMP).
With routing protocols based on shared trees, once the receivers join a particular IP multicast group, a multicast distribution tree is constructed for that group. The protocol most widely used for this is Protocol Independent Multicast (PIM). It sets up multicast distribution trees such that data packets from senders to a multicast group reach all receivers which have joined the group. For example, all data packets sent to the group 239.1.1.1 are received by receivers who joined 239.1.1.1. There are variations of PIM implementations: Sparse Mode (SM), Dense Mode (DM), Source Specific Mode (SSM) and Bidirectional Mode (Bidir, or Sparse-Dense Mode, SDM). Of these, PIM-SM is the most widely deployed as of 2006; SSM and Bidir are simpler and scalable variations developed more recently and are gaining in popularity.

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

QUESTION NO: 272

Which three statements are true about Source Specific Multicast? (Choose three.)

A. Is best suited for applications that are in the one-to-many category.
B. SSM uses shortest path trees only.
C. The use of SSM is recommended when there are many sources and it is desirable to keep the amount of mroute state in the routers in the network to a
D. There are no RPs to worry about

Answer: A,B,D
Explanation:
The Source Specific Multicast feature is an extension of IP multicast where datagram traffic is forwarded to receivers from only those multicast sources to which the receivers have explicitly joined. For multicast groups configured for SSM, only source-specific multicast distribution trees (no shared trees) are created.

The current IP multicast infrastructure in the Internet and many enterprise intranets is based on the PIM- SM protocol and Multicast Source Discovery Protocol (MSDP). These protocols have proven to be reliable, extensive, and efficient. However, they are bound to the complexity and functionality limitations of the Internet Standard Multicast (ISM) service model. For example, with ISM, the network must maintain knowledge about which hosts in the network are actively sending multicast traffic. With SSM, this information is provided by receivers through the source address(es) relayed to the last hop routers by IGMP v3lite or URD. SSM is an incremental response to the issues associated with ISM and is intended to coexist in the network with the protocols developed for ISM. In general, SSM provides a more advantageous IP multicast service for applications that utilize SSM.
ISM service is described in RFC 1112. This service consists of the delivery of IP datagrams from any source to a group of receivers called the multicast host group. The datagram traffic for the multicast host group consists of datagrams with an arbitrary IP unicast source address S and the multicast group address G as the IP destination address. Systems will receive this traffic by becoming members of the host group. Membership to a host group simply requires signalling the host group through IGMP Version 1, 2, or 3. In SSM, delivery of datagrams is based on (S, G) channels. Traffic for one (S, G) channel consists of datagrams with an IP unicast source address S and the multicast group address G as the IP destination address. Systems will receive this traffic by becoming members of the (S, G) channel. In both SSM and ISM, no signalling is required to become a source. However, in SSM, receivers must subscribe or unsubscribe to (S, G) channels to receive or not receive traffic from specific sources. In other words, receivers can receive traffic only from (S, G) channels that they are subscribed to, whereas in ISM, receivers need not know the IP addresses of sources from which they receive their traffic. The proposed standard approach for channel subscription signalling utilizes IGMP INCLUDE mode membership reports, which are only supported in Version 3 of IGMP.

SSM can coexist with the ISM service by applying the SSM delivery model to a configured subset of the IP multicast group address range. The Internet Assigned Numbers Authority (IANA) has reserved the address range 232.0.0.0 through 232.255.255.255 for SSM applications and protocols. Cisco IOS software allows SSM configuration for an arbitrary subset of the IP multicast address range 224.0.0.0 through 239.255.255.255.

When an SSM range is defined, existing IP multicast receiver applications will not receive any traffic when they try to use addresses in the SSM range (unless the application is modified to use explicit (S, G) channel subscription or is SSM enabled through URD).

QUESTION NO: 273

What is the Internet Protocol Number for all PIM control Messages?

A. 15
B. 17
C. 25
D. 103

Answer: D
Reference: http://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml (search '103')

QUESTION NO: 274

Which value is used in the PIM TYPE field to indicate a Join/Prune message?

A. 0
B. 1
C. 2
D. 3

Answer: D
Reference: http://tools.ietf.org/html/draft-venaas-pim-hierarchicaljoinattr-00

QUESTION NO: 275

Which two multicast addresses are reserved for use by ALL-PIM-ROUTERS? (Choose two.)

A. 224.0.0.13
B. 224.0.0.17
C. 224.0.0.39
D. ff02::d
E. ff02::39
F. ff02::13

Answer: A,D
Reference: http://www.routeralley.com/ra/docs/multicast.pdf (see page 2)

QUESTION NO: 276

Which value is used in the PIM TYPE field to indicate a Register message?

A. 0
B. 1
C. 2
D. 3

Answer: B
Reference: http://technet.microsoft.com/en-us/library/bb742462.aspx (see Dr election)

QUESTION NO: 277

What does the beginning of a multicast address look like, if it is used for embedded RP?

A. FF7
B. FF
C. FF3
D. Embedded RP does not use any special IPv6 address.

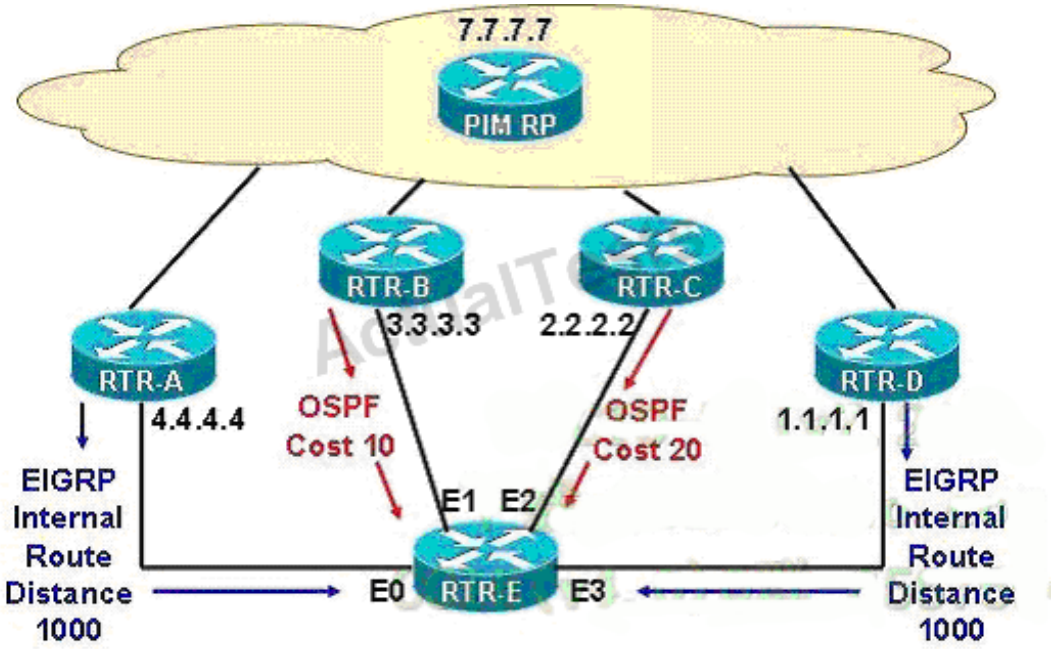
Answer: A
Explanation:
Embedded RP defines an address allocation policy in which the address of the RP is encoded in an IPv6 multicast group address. This allows an easy deployment of scalable inter-domain multicast and simplifies the intra-domain multicast configuration as well. IPv6 Multicast group addresses embedded with RP information start with ff70::/12 where the flag value of 7 means embedded RP.

Reference: http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6552/whitepaper_c11-508498.html

QUESTION NO: 278

Refer to the exhibit. Router E learned about the PIM RP (designated as 7.7.7.7) from four different sources. Routers A and D advertised the 7.0.0.0 network via EIGRP. Routers B and C advertised the 7.0.0.0 network via OSPF. Considering that all four Ethernet interfaces on router E could potentially lead back to the PIM-RP, when router E receives the first multicast packet down the

shared tree, which incoming interface will be used to successfully pass the RPF check?



A. E0
B. E1
C. E2
D. E3
E. None of these interfaces will be used to successfully pass the RPF check.
F. All of these interfaces would successfully pass the RPF check.

Answer: A

Explanation:

PIM will use the neighbor with the highest IP. E0 and E3 is taken in to account as EIGRP has lower AD the OSPF, if both E0 and E3 cost are the same next there is rule that “When faced with multiple equal cost paths to a source, IP multicast chooses the interface that has a Protocol Independent Multicast (PIM) neighbor with the highest IP address as the incoming interface and then sends prunes to PIM neighbors on the other links.”

QUESTION NO: 279

Refer to the exhibit.

```
P1#show ip mfib 239.1.1.1 10.1.100.3
(10.1.100.3,239.1.1.1) Flags:
   SW Forwarding: 74326/0/81/0, Other: 0/0/0
   Ethernet0/1 Flags: A
   Ethernet1/0 Flags: F NS
     Pkts: 23/0
```

Which statement is correct?

A. Ethernet0/1 is in the outgoing interface list, and Ethernet1/0 is in the incoming interface list.
B. Ethernet1/0 is in the outgoing interface list, and Ethernet0/1 is in the incoming interface list.
C. This is not a valid MFIB entry.
D. This MFIB entry is currently not used to forward multicast traffic.
E. All multicast packets for this MFIB entry are punted.

Answer: B
Reference:
http://www.cisco.com/en/US/tech/tk828/technologies_tech_note09186a0080094b55.shtml (see 'diagnose the problem')

QUESTION NO: 280

Refer to the exhibit.

```
ip pim ssm range ssm-groups
!
ip access-list standard ssm-groups
 permit 232.0.0.0 0.255.255.255
```

Which statement is correct?

A. Group 232.1.1.1 will be in SSM mode, but without the access-list ssm-groups it would also be in SSM mode.
B. Group 232.1.1.1 will not be in SSM mode yet; some further configuration is needed.
C. Group 232.1.1.1 can never be in SSM mode, regardless of any configuration.
D. Group 232.1.1.1 will be in SSM mode because of the access-list ssm-groups.

E. Group 232.1.1.1 is not in SSM mode yet; a different access-list would be needed to achieve this.

Answer: D
Reference: http://networklessons.com/multicast/source-specific-multicast-ssm/

QUESTION NO: 281

In which two scenarios would MSDP be used? (Choose two.)

A. Interdomain multicast
B. Anycast RP
C. Distributing the RP address to other routers
D. Implementing PIM-SSM

Answer: A,B
Reference: http://tools.ietf.org/html/rfc4611 (see introduction)

QUESTION NO: 282

Refer to the exhibit. From the MAC addresses shown in the command output, to which two ports is the multicast stream 225.230.57.199 being forwarded on this switch? (Choose two.)

Switch#show mac-address-table multicast

vlan	mac address	type	ports
2	0100.5ee6.39c7	igmp	Gi3/7,Fa6/28,Fa7/20
2	0100.5e00.0123	igmp	Fa5/7,Fa6/28,Fa7/20
2	0100.5e66.39c7	igmp	Gi3/4,Gi3/7,Fa4/10,Fa4/14 Fa7/31,Fa7/40
3	0100.5e00.017f	igmp	Gi3/7,Fa6/21
3	0100.5e50.4453	igmp	Gi3/7,Fa4/2,Fa4/3,Fa4/14,Fa4/38,Fa5/3

ActualTests

A. Fa6/28
B. Fa7/20
C. Gi3/7
D. Fa4/2
E. Fa4/14

F. Fa4/38
G. Fa6/28
H. Fa5/7

Answer: C,E

Explanation:

IP 225.230.57.199 will become MAC 0100.5ee6.39c7, so the interfaces G3/7, F6/28 and F7/20 will receive the traffic.

Mac address 0100.5e66.39c7 is accessible via Gi3/4, Gi3/7, Fa4/10, Fa4/14, Fa7/31, and Fa7/40. So options C and E i.e., Gi3/7 & Fa4/14 are the correct answers.

First thing is to convert the IP address to binary so: 225.230.57.199 in binary is 11100001 11100110 00111001

11000111. Then put the first 6 hex characters in front of the binary address and remove the first 4 bits (which are always 1110) this is the IEEE OUI for layer 2 multicast addressing 0100.5e | 0001 11100110 00111001

11000111. Then we "always" change the next 5 bits to a binary 0 (which leaves us with 24 bits for the conversion to hex) 0100.5e | 01100110 00111001 11000111

This leaves you with 24 bits to convert into hex from binary. Thus 01100110 00111001 11000111 = 66.39.c7 and the question asks which interfaces listed are forwarding this group...

All of the following interfaces below are forwarding this:

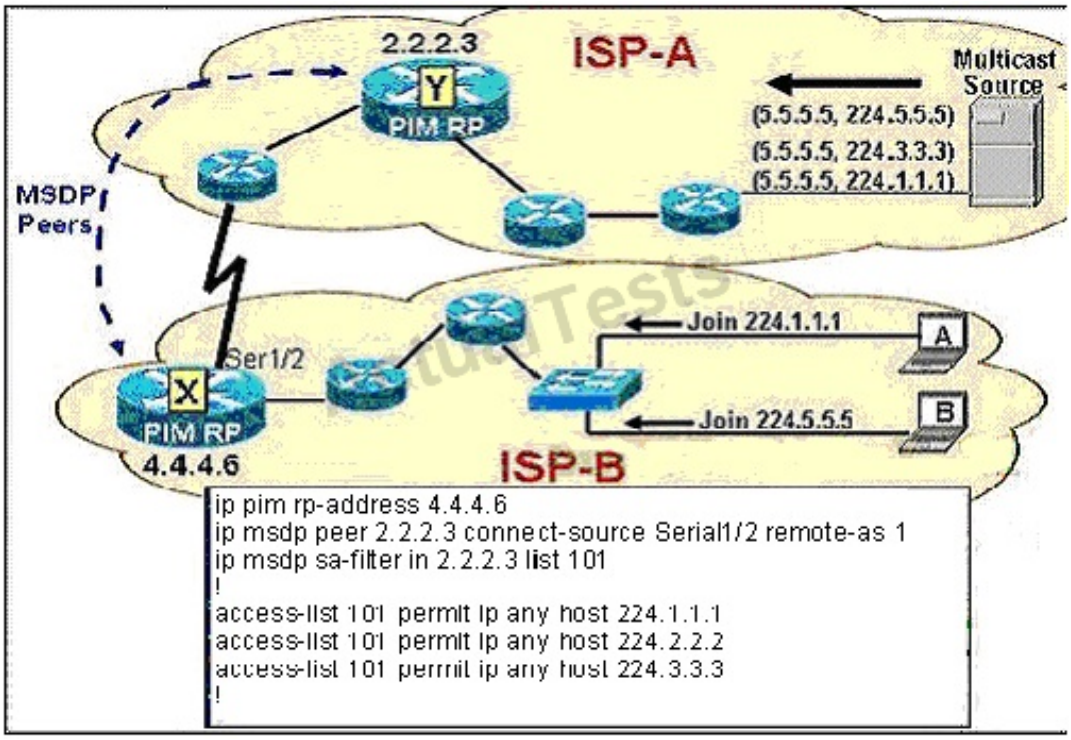
Gi3/4, Gi3/7, Fa4/10, Fa4/14, Fa7/31, Fa7/40

But only 2 are listed in the multiple choice:

Gi3/7 & Fa4/14

QUESTION NO: 283

Refer to the exhibit.



Two ISPs have decided to use MSDP and configured routers X and Y (both are PIM RPs) as MSDP peers. In the domain of ISP B, PC A has sent an IGMP membership report for the group 224.1.1.1 and PC B has sent an IGMP membership report for the group 224.5.5.5. Assuming that the MSDP peering relationship between routers X and Y is functional, and given the partial configuration output shown from router X, which two of these statements are true? Choose two.)

A. Router X will contain an entry for 224.1.1.1 in its SA cache and will also have an installed (S, G) entry for this in its mroute table.
B. Router X will not contain an entry for 224.1.1.1 in its SA cache but will have an installed (*, G) entry for this in its mroute table.
C. Router X will not contain an entry for 224.5.5.5 in its SA cache but will have an installed (S, G) entry for this in its mroute table.
D. Router X will not contain an entry for 224.5.5.5 in its SA cache but will have an installed (*, G) entry for this in its mroute table.
E. Router X will have no entries for 224.5.5.5 in neither its SA cache nor in its mroute table. ^F. Router F. X will have no entries for 224.1.1.1 in neither its SA cache nor in its mroute table.

Answer: A,D

Explanation:

Configuring an MSDP Peer

You enable MSDP by configuring an MSDP peer to the local router.

Note

The router you specify by Domain Naming System (DNS) name or IP address as an MSDP peer is probably a Border Gateway Protocol (BGP) neighbor. If it is not, see the section "Configuring a Default MSDP Peer" later in this document.

To configure an MSDP peer, use the following commands in global configuration mode as needed. The second command is optional.

Command	Purpose
Router(config)# **ip msdp peer** {*peer-name* | *peer-address*} [**connect-source** *type number*] [**remote-as** *as-number*]	Enables MSDP and configures an MSDP peer as specified by the DNS name or IP address. If you specify the **connect-source** keyword, the primary address of the specified local interface *type* and *number* values are used as the source IP address for the TCP connection. The **connect-source** keyword is recommended, especially for MSDP peers on a border that peer with a router inside the remote domain.
Router(config)# **ip msdp description** {*peer-name* | *peer-address*} *text*	Configures a description for a specified peer to make it easier to identify in a configuration or in **show** command output.

Filtering SA Request Messages

By default, only routers that are caching SA information can respond to SA request messages. By default, such a router honors all SA request messages from its MSDP peers. That is, it will supply the IP addresses of the sources that are active.

However, you can configure the router to ignore all SA request messages from an MSDP peer. Or, you can honor only those SA request messages from a peer for groups described by a standard access list. If the access list passes, SA request messages will be accepted. All other such messages from the peer for other groups will be ignored.

To configure one of these options, use either of the following commands in global configuration mode:

Command	Purpose
Router(config)# **ip msdp filter-sa-request** {*peer-address* | *peer-name*}	Filters all SA request messages from the specified MSDP peer.
Router(config)# **ip msdp filter-sa-request** {*peer-address* | *peer-name*} **list** *access-list*	Filters SA request messages from the specified MSDP peer for groups that pass the standard access list. The access list describes a multicast group address.

Reference

http://www.cisco.com/en/US/docs/ios/12_2/ip/configuration/guide/1cfmsdp_ps1835_TSD_Products_Configuration_Guide_Chapter.html#wp1001055

QUESTION NO: 284

What is the default behavior of PIM-SSM If there are multiple equal-cost paths to the multicast

source?

A. It will send the join only to the neighbor with the highest IP address.
B. It will send the join only to the neighbor with the lowest IP address.
C. It will send the join in a round-robin fashion across all neighbors.
D. It will send the join to all neighbors.

Answer: A
Explanation:
By default, for Protocol Independent Multicast sparse mode (PIM-SM), Source Specific Multicast (PIMSSM), bidirectional PIM (bidir-PIM), and PIM dense mode (PIM-DM) groups, if multiple equal-cost paths are available, Reverse Path Forwarding (RPF) for IPv4 multicast traffic is based on the PIM neighbor with the highest IP address. This method is referred to as the highest PIM neighbor behavior. This behavior is in accordance with RFC 2362 for PIM-SM, but also applies to PIM-SSM, PIM-DM, and bidir-PIM.

QUESTION NO: 285

Which statement is incorrect in reference to IPv6 multicast?

A. IPv6 multicast uses Multicast Listener Discovery.
B. The first 8 bits of an IPv6 multicast address are always FF (1111 1111).
C. IPv6 multicast requires MSDP.
D. PIM dense mode is not part of IPv6 multicast.

Answer: C
Explanation:

QUESTION NO: 286

Which IPv6 multicast address is reserved for use by all PIM routers?

A. ff02::39
B. ff02::13
C. ff02::d
D. ff02::17

Answer: C

Explanation:
If the IPv6 Destination Address field is the multicast address ALL-PIM-ROUTERS, the IPv6 form of the address (ff02::d) is used. These IPv6 PIM control messages are of course not transmitted natively over the service provider's network, but rather are encapsulated in GRE/IPv4.

QUESTION NO: 287

Which RFC number contains the specification for Protocol Independent Multicast sparse mode?

A. 1011
B. 1918
C. 2044
D. 4601

Answer: D
erence: http://www.ietf.org/rfc/rfc4601.txt

QUESTION NO: 288

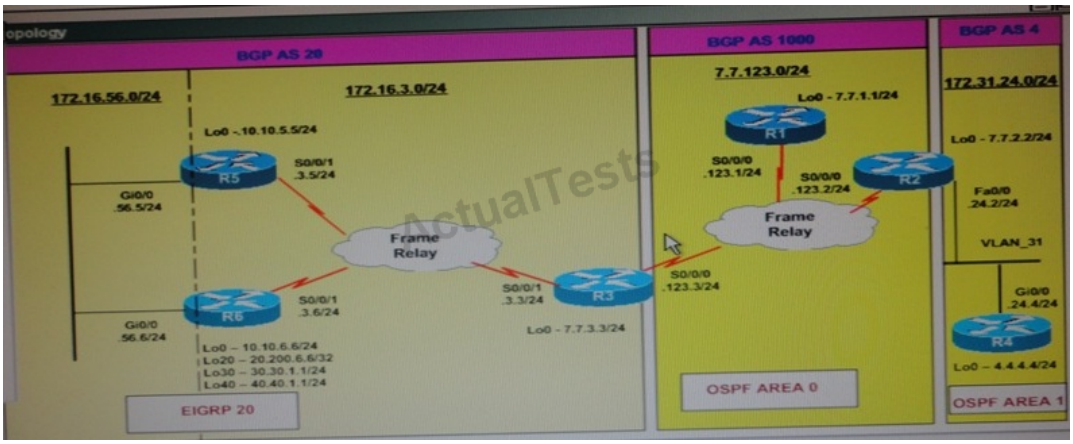
Which two are differences between IGMPv2 and IGMPv3 reports? (Choose two.)

A. IGMPv3 adds the ability to include or exclude source lists.
B. All IGMPv2 hosts send reports to destination address 224.0.0.22.
C. Only IGMPv3 reports may contain multiple group state records.
D. All IGMPv3 hosts send reports to destination address 224.0.0.23.
E. IGMPv2 does not support the Leave Group message.

Answer: A,C
Reference: http://www.juniper.net/techpubs/software/erx/junose700/swconfig-routing-vol1/html/ip-multicast-config14.html (see including and excluding traffic)

QUESTION NO: 289

Refer to the exhibit.



R6#sh ipv6 mroute

<…>

(*,FF04::10), 01:15:32/never, RP 2001:DB8:5::5, flags: SPC

Incoming interface: GigabitEthernet0/0

RPF nbr: FE80::216:47FF:FEBB:FF0

(*,FF04::30), 00:00:07/never, RP 2001:DB8:5::5, flags: SPC

Incoming Interface: GigabitEthernet0/0

RPF nbr: FE80::216:47FF:FEBB:FF0

We have IPv6 multicast configured between R5 and R6, which three statements are true based on the partial command output shown? (Choose three)

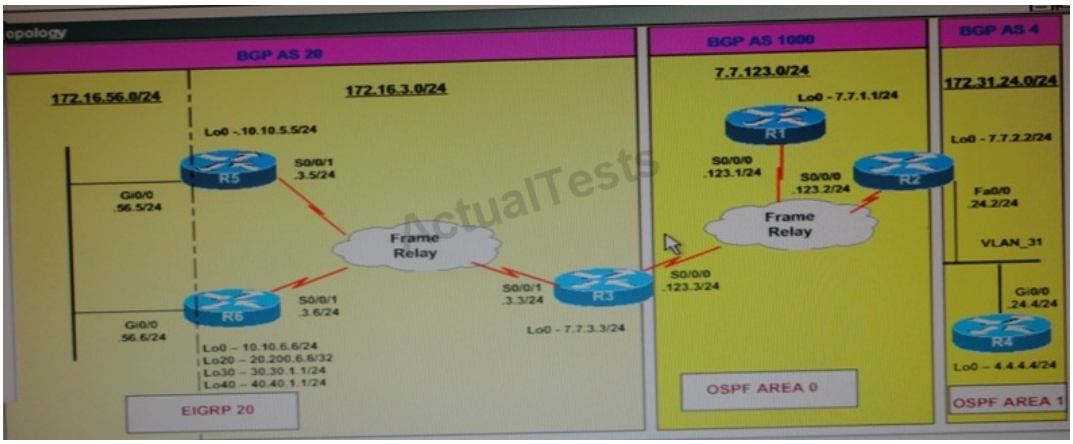
A. I R6 has joined one multicast group, and it expires in 46 seconds.
B. The rendezvous point address is 2001:DB8:5::5.
C. The multicast group address is FE80::216.
D. R6 has joined two multicast groups, and it expires in 7 seconds.
E. The multicast entry is operating in sparse mode.
F. The multicast groups are FF04::10 and FF04::30.

Answer: B,E,F
Reference:
http://www.cisco.com/en/US/docs/ios/ipv6/command/reference/ipv6_14.html#wp2378397

QUESTION NO: 290

Refer to the exhibit.

```
R6#show ipv6 mroute
<...>
(*, FF04::10), 00:00:46/never, RP 2001:DB8:5::5, flags: SPC
  Incoming interface: GigabitEthernet0/0
  RPF nbr: FE80::216:47FF:FEBB:FF0
```



We have IPv6 multicast configured between R5 and R6. Which two statements are true based on the partial command output shown? (Choose two.)

A. R6 has joined the multicast group, and it expires in 46 seconds
B. The rendezvous point address is FE80::21 6:47FF:FEBB:FFO.
C. The multicast group address is FF04::1 0.
D. The multicast entry is operating in dense mode.
E. The multicast route has been pruned.

Answer: C,E
: http://www.slideshare.net/CiscoCanada/ipv6-handsonlabhritter

QUESTION NO: 291

How is RPF used in multicast routing?

A. to prevent multicast packets from looping
B. to prevent PIM packets from looping
C. to instruct PIM where to send a (*, G) or (S, G) join message
D. to prevent multicast packets from looping and to instruct PIM where to send a (*, G) or (S, G) join message

Answer: D
Explanation:

QUESTION NO: 292

Refer to the exhibit.

```
Router#sh ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/32 is subnetted, 2 subnets
C       1.1.1.1 is directly connected, Loopback0
O       1.1.1.2 [110/11] via 3.3.3.2, 00:00:06, Ethernet1/0
                [110/11] via 2.2.2.2, 00:00:06, Ethernet0/0
     2.0.0.0/24 is subnetted, 1 subnets
C       2.2.2.0 is directly connected, Ethernet0/0
     3.0.0.0/24 is subnetted, 1 subnets
C       3.3.3.0 is directly connected, Ethernet1/0

Router#sh ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
2.2.2.2           Ethernet0/0              00:21:13/00:01:43 v2    1 / DR S
3.3.3.2           Ethernet1/0              00:01:04/00:01:39 v2    1 / DR S
```

Which interface(s) will show ip rpf 1.1.1.2 indicate as RPF interface(s)?

A. Ethernet 1/0
B. Ethernet 0/0
C. Both Ethernet 0/0 and Ethernet 1/0
D. RPF will fail

Answer: A

Explanation:

When troubleshooting multicast routing, the primary concern is the source address. Multicast has a concept of Reverse Path Forwarding check (RPF check). When a multicast packet arrives on an interface, the RPF process checks to ensure that this incoming interface is the outgoing interface used by unicast routing to reach the source of the multicast packet. This RPF check process prevents loops. Multicast routing does not forward a packet unless the source of the packet passes a reverse path forwarding (RPF) check. Once a packet passes this RPF check, multicast routing forwards the packet based only upon the destination address.

Reference
http://www.cisco.com/en/US/tech/tk828/technologies_tech_note09186a0080094b55.shtml

QUESTION NO: 293

Which command is used to enable SSM with the range 232.0.0.0 - 255.0.0.0?

A. ip pim ssm
B. ip pim ssm 232.0.0.0 255.0.0.0
C. ip pim ssm range 50
D. access-list 50 permit 232.0.0.0 15.255.255.255
E. ip pim enable default

Answer: A
Reference: http://books.google.com.pk/books?id=2Nhv8mdg-o4C&pg=PA202&lpg=PA202&dq=ip+pim+ssm+enable+SSM+with+the+range+232.0.0.0+-+255.0.0.0&source=bl&ots=ycFzpApLQk&sig=HUz9vxpttyXErIthl14DkZwIo8w&hl=en&sa=X&ei=ZTLLUamUCozY4QTXxoHYBQ&ved=0CEUQ6AEwBQ#v=onepage&q=ip%20pim%20ssm%20enable%20SSM%20with%20the%20range%20232.0.0.0%20-%20255.0.0.0&f=false

QUESTION NO: 294

Which command is used to enable SSM with the range 232.0.0.0/8?

A. ip pim ssm default
B. ip pim ssm 232.0.0.0 255.0.0.0
C. ip pim ssm range 50 access-list 50 permit 232.0.0.0 15.255.255.255
D. ip pim enable default

Answer: A
Reference: http://packetlife.net/blog/2010/jul/27/source-specific-multicast-pim-ssm/ (see PIM-SSM configuration)

QUESTION NO: 295

Which value is used in the PIM TYPE field to indicate a register-stop message?

A. 0
B. 1
C. 2
D. 3

Answer: C
Reference: http://www.iana.org/assignments/pim-parameters/pim-parameters.xhtml (see PIM message types)

QUESTION NO: 296

Multicast is being deployed in the network, and only ip pim sparse-dense mode has been configured on all interfaces in the network to support a new video streaming application. No other multicast configuration was applied anywhere in the network. Since enabling multicast, the network monitoring tools show periodic spikes in link utilization throughout the network, even in areas where the video application is not being used. What could be a possible cause?

A. PIM sparse mode is being used.
B. PIM dense mode is being used.
C. The BGP multicast address family has not been configured.
D. IGMP version 3 is being used.
E. IP PIM neighbor filters have not been applied.

Answer: B
Explanation:
PIM DM builds source-based multicast distribution trees. In dense mode, a PIM DM router or multilayer switch assumes that all other routers or multilayer switches forward multicast packets for a group. If a PIM DM device receives a multicast packet and has no directly connected members or PIM neighbors present, a prune message is sent back to the source to stop unwanted multicast traffic. Subsequent multicast packets are not flooded to this router or switch on this pruned branch because branches without receivers are pruned from the distribution tree, leaving only branches that contain receivers.

QUESTION NO: 297

Choose two commands that are required to enable multicast on a router, when it is known that the receivers use a specific functionality of IGMPv3. (Choose two.)

A. ip pim rp-address
B. ip pim ssm
C. ip pim sparse-mode
D. ip pim passive

Answer: B,C

Explanation: Explanation

Source specific multicast only works with IGMPv3

Reference

http://www.cisco.com/en/US/docs/ios/ipmulti/command/reference/imc_04.html#wp1043638

QUESTION NO: 298

Which multicast address is reserved for cisco-rp-announce?

A. 224.0.331.37
B. 224.0.1.38
C. 224.0.1.39
D. 224.0.1.40

Answer: C

Explanation:

Auto-RP is a mechanism to automate distribution of RP information in a multicast network. The Auto-RP mechanism operates using two basic components, the candidate RPs and the RP mapping agents.

• Candidate RPs advertize their willingness to be an RP via "RP-announcement" messages. These messages are periodically sent to a reserved well-known group 224.0.1.39 (CISCO-RP-ANNOUNCE).

• RP mapping agents join group 224.0.1.39 and map the RPs to the associated groups. The RP mapping agents advertise the authoritative RP-mappings to another well-known group address 224.0.1.40 (CISCO-RP-DISCOVERY). All PIM routers join 224.0.1.40 and store the RP-mappings in their private cache.

Reference: http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6552/whitepaper_c11-508498.html

QUESTION NO: 299

What is the value in seconds of the "Max Response Time" in an IGMPv2 report message?

A. 0
B. 1

C. 2
D. 3

Answer: A
Explanation:
GMPv2 messages**IGMPv2 packet structure**
+
Bits 0–7
8–15
16–31
0
Type
Max Resp Time
Checksum
32
Group Address
Where:
Type
Indicates the message type as follows: Membership Query (0x11), Membership Report (IGMPv1: 0x12, IGMPv2: 0x16, IGMPv3: 0x22), Leave Group (0x17)

Max Resp Time
Specifies the time limit for the corresponding report. The field has a resolution of 100 milliseconds, the value is taken directly. **This field is meaningful only in Membership Query (0x11); in other messages it is set to 0 and ignored by the receiver.**
Reference: http://en.wikipedia.org/wiki/Internet_Group_Management_Protocol

Topic 6, Implement Network Security

QUESTION NO: 300

Which three statements are true about TACACS+? (Choose three.)

A. It is a Cisco proprietary protocol.
B. It runs on TCP port 59. .
C. Authentication and authorization are done at different stages.
D. TACACS+ encrypts the entire body of the packet, but leaves a standard TACACS+ header.

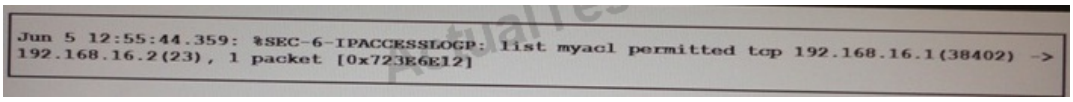
E. It is an industry standard protocol.
F. TACACS+ encrypts both the entire body of the packet and the TACACS+ header.

Answer: A,C,D
Explanation:
TACACS+ utilizes TCP port 49. It consists of three separate protocols, which can be implemented on separate servers.[1]
TACACS+ offers multiprotocol support, such as IP and AppleTalk. Normal operation fully encrypts the body of the packet for more secure communications. It is a Cisco proprietary enhancement to the original TACACS protocol.

Reference
http://en.wikipedia.org/wiki/TACACS%2B

QUESTION NO: 301

Refer to the exhibit.



Which two statements are correct? (Choose two.)

A. The hexadecimal value of the number of packets that hit the access list is 0x723E6E12.
B. The access list has logging enabled.
C. The packet was discarded.
D. The command ip access-list logging hash-generation is enabled.
E. The Telnet connection is successfully set up.

Answer: B,D
Reference: http://www.packetu.com/2009/07/31/ccie-security-quick-tips/

QUESTION NO: 302

Refer to the exhibit.

```
router#show control-plane host open-ports
Active internet connections (servers and established)
Prot                Local Address              Foreign Address      Service    State
 tcp                        *:22                             *:0  SSH-Server   LISTEN
 tcp                        *:23                             *:0      Telnet   LISTEN
 tcp                        *:80                             *:0   HTTP CORE   LISTEN
!
class-map type port-filter match-all CMAP_PF
 match  closed-ports
!
policy-map type port-filter PMAP_PF
 class CMAP_PF
   drop
!
control-plane host
 service-policy type port-filter input PMAP_PF
!
```

ActualTests

Which statement is correct?

A. This configuration is not valid.
B. Control Plane Policing is configured; however you cannot determine on which interface it is configured.
C. NTP is not configured on the router.
D. Telnet traffic will be dropped.

Answer: C

Explanation: Display the current connection table for the router. For example, for IOS routers, type "show control-plane host open-ports" Press "Enter." The router will display a table with one line for each currently open connection. The fourth column from the left is labeled "Service;" the entries under that column correspond to the protocols currently in use. The current TCP traffic corresponds exactly to the table entries that have "TCP" under the fourth column.

QUESTION NO: 303

What does Cisco recommend when you are enabling Cisco IOS IPS?

A. Do not enable all the signatures at the same time.
B. Do not enable the ICMP signature.
C. Disable the Zone-Based Policy Firewall because it is not compatible with Cisco IOS IPS.
D. Disable CEF because it is not compatible with Cisco IOS IPS. .

Answer: A

Explanation:

Router memory and resource constraints prevent a router from loading all Cisco IOS IPS signatures. Thus, it is recommended that you load only a selected set of signatures that are defined by the categories. Because the categories are applied in a “top-down” order, you should

first retire all signatures, followed by “unretiring” specific categories. Retiring signatures enables the router to load information for all signatures, but the router does not build the parallel scanning data structure.

QUESTION NO: 304

Refer to the exhibit.

```
R1-2503# show ip ospf interface serial0
Serial0 is up, line protocol is up
  Internet Address 192.16.64.1/24, Area 0
  Process ID 10, Router ID  172.16.10.36 , 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 00:00:05
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 70.70.70.70
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
    Youngest key id is 1
```

ActualTests

Which statement is correct?

A. OSPF peers are using Type 1 authentication
B. OSPF peers are using Type 2 authentication
C. Authentication is used, but there is a password mismatch
D. The OSPF peer IP address is 172.16.10.36

Answer: B
Explanation:
These are the three different types of authentication supported by OSPF.
Null Authentication—This is also called Type 0 and it means no authentication information is included in the packet header. It is the default.

Plain Text Authentication—This is also called Type 1 and it uses simple clear-text passwords.
MD5 Authentication—This is also called Type 2 and it uses MD5 cryptographic passwords.
Authentication does not need to be set. However, if it is set, all peer routers on the same segment must have the same password and authentication method. The examples in this document demonstrate configurations for both plain text and MD5 authentication.

Reference

http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080094069.shtml

QUESTION NO: 305

Which two statements are true about Unicast Reverse Path Forwarding Loose Mode? (Choose two.)

A. It is used in multihome network scenarios.
B. It can be used with BGP to mitigate DoS and DDoS.
C. It does not need to have CEF enabled.
D. It is enabled via the interface level command ip verify unicast reverse-path.
E. It cannot be used with "classification" access lists.

Answer: A,B
Explanation: The Unicast Reverse Path Forwarding Loose Mode feature creates a new option for Unicast Reverse Path Forwarding (Unicast RPF), providing a scalable anti-spoofing mechanism suitable for use in multihome network scenarios. This mechanism is especially relevant for Internet Service Providers (ISPs), specifically on routers that have multiple links to multiple ISPs. In addition, Unicast RPF (strict or loose mode), when used in conjunction with a Border Gateway Protocol (BGP) "trigger, " provides an excellent quick reaction mechanism that allows network traffic to be dropped on the basis of either the source or destination IP address, giving network administrators an efficient tool for mitigating denial of service (DoS) and distributed denial of service (DDoS) attacks.

Reference
http://www.cisco.com/en/US/docs/ios/12_2t/12_2t13/feature/guide/ft_urpf.html

QUESTION NO: 306

Refer to the exhibit.

```
Router# show run
aaa authentication login default group tacacs+ none
```

What would be the security risk when you are using the above configuration?

A. The locally configured users would override the TACACS+ security policy.
B. It would be impossible to log in to the router if the TACACS+ server is down.
C. The default login policy would override the TACACS+ configuration.
D. If the TACACS+ server failed, no authentication would be required.

Answer: D
Explanation:
You could use the aaa authentication login default tacacs+ enable command to specify that if your TACACS+ server fails to respond, you can log in to the access server by using your enable password. If you do not have an enable password set on the router, you will not be able to log in to it until you have a functioning TACACS+ UNIX daemon or Windows NT or Windows 2000 server process configured with usernames and passwords. The enable password in this case is a last-resort authentication method. You also can specify none as the last-resort method, which means that no authentication is required if all other methods failed.

QUESTION NO: 307

Which three protocols should be explicitly managed by using a CoPP policy on an Internet border router? (Choose three.)

A. SMTP
B. ICMP
C. BGP
D. SSH
E. RTP
F. BitTorrent
G. VTP

Answer: B,C,D
Explanation: Control Plane Policing (CoPP) is a Cisco IOS-wide feature designed to allow users to manage the flow of traffic handled by the route processor of their network devices. CoPP is designed to prevent unnecessary traffic from overwhelming the route processor that, if left unabated, could affect system performance. Route processor resource exhaustion, in this case, refers to all resources associated with the punt path and route processor(s) such as Cisco IOS process memory and buffers, and ingress packet queues.

Reference
http://www.cisco.com/web/about/security/intelligence/coppwp_gs.html#3

QUESTION NO: 308

What is true about IP Source Guard with port security?

A. Binding should be manually configured.
B. It is not supported if IEEE 802.1x port-based authentication is enabled
C. The DHCP server must support option 82, or the client is not assigned an IP address.
D. It filters based on source IP address only.

Answer: C
Reference:
http://www.cisco.com/en/US/docs/switches/lan/catalyst3750x_3560x/software/release/12.2_55_se/configuration/guide/swdhcp82.html (see enabling IP source guard, see the table – step 3)

QUESTION NO: 309

After applying a new ACL on a device, its CPU utilization rose significantly and many messages starting with "%SEC-6-IPACCESSLOG" appeared on the Syslog server.

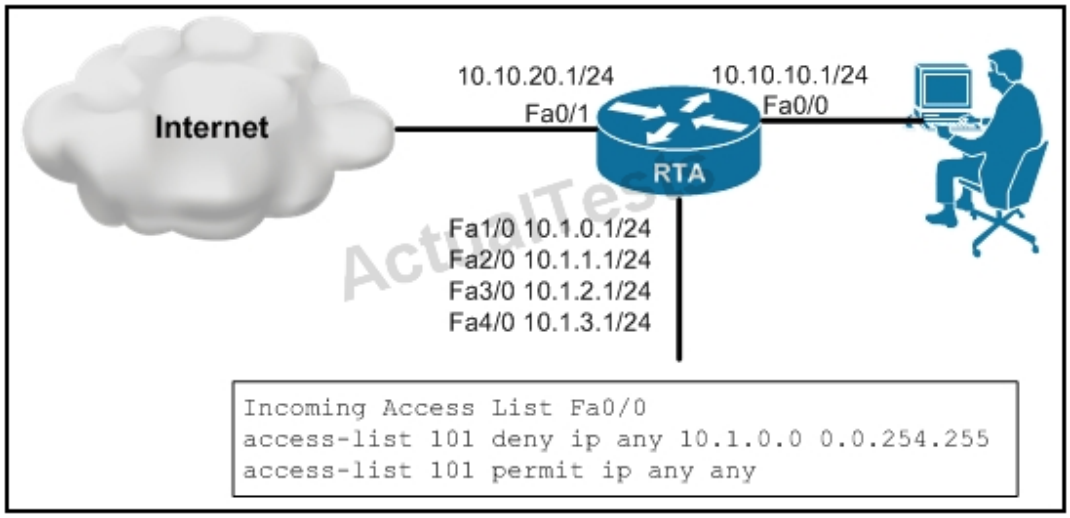
What can be done to resolve this situation?

A. Increase memory allocation for ACLs.
B. Remove all entries from the ACL and use a single permit ip any any statement.
C. Remove the log keyword from each ACL entry.
D. Reboot the device after the ACL has been applied.

Answer: C
Reference: http://www.networksa.org/?p=347

QUESTION NO: 310

Refer to the exhibit.



Clients that are connected to Fa0/0 of RTA are only allowed to connect to the Internet and networks, but not the networks on Fa1/0, Fa2/0, Fa3/0 and Fa4/0. To achieve this, you have configured an ACL on RTA and applied it on the incoming direction of interface Fa0/0.

After you apply this ACL, you learn that some of these networks are still accessible for clients that are connected to the 10.10.10.0/24 network. What is the correct ACL configuration to solve this issue?

A. access-list 101 deny ip any 10.1.0.0 0.0.1.255
access-list 101 permit ip any any
B. access-list 101 permit ip any 10.1.0.0 0.0.1.255
access-list 101 deny ip any any
C. access-list 101 deny ip any 10.1.0.0 0.0.252.255
access-list 101 permit ip any any
D. access-list 101 deny ip any 10.1.0.0 0.0.3.255
access-list 101 permit ip any any

Answer: D

Explanation:

Access-lists use a wild card mask which is incorrectly configured in the above example

Reference
http://www.cisco.com/en/US/products/sw/secursw/ps1018/products_tech_note09186a00800a5b9a.shtml

QUESTION NO: 311

What needs to be enabled for Unicast RPF?

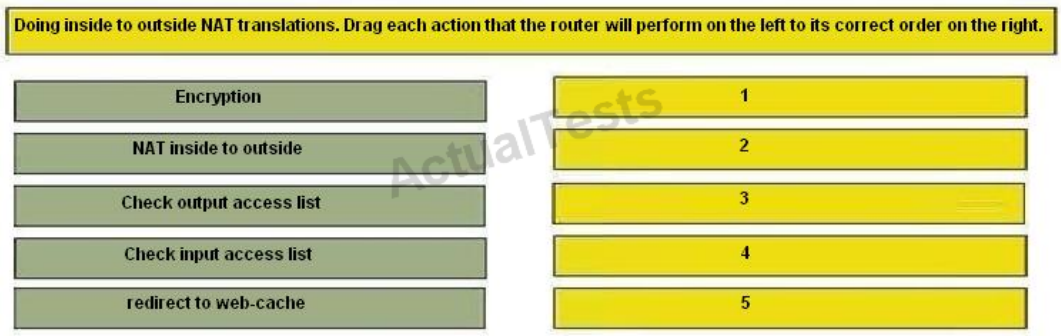
A. BGP
B. OSPF
C. CEF
D. RIP

Answer: C

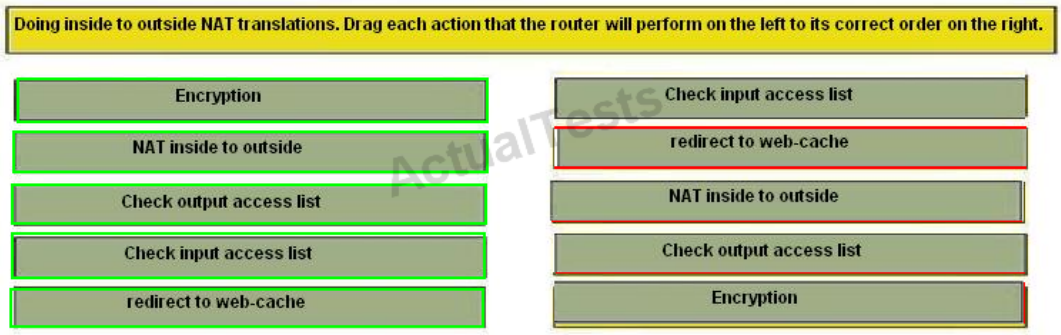
Explanation:

Cisco Express Forwarding switching must be enabled for Unicast RPF to function. This command has been enabled by default as of IOS version 12.2. If it is not enabled, administrators can enable it with the following global configuration command: ip cef

QUESTION NO: 312 DRAG DROP



Answer:



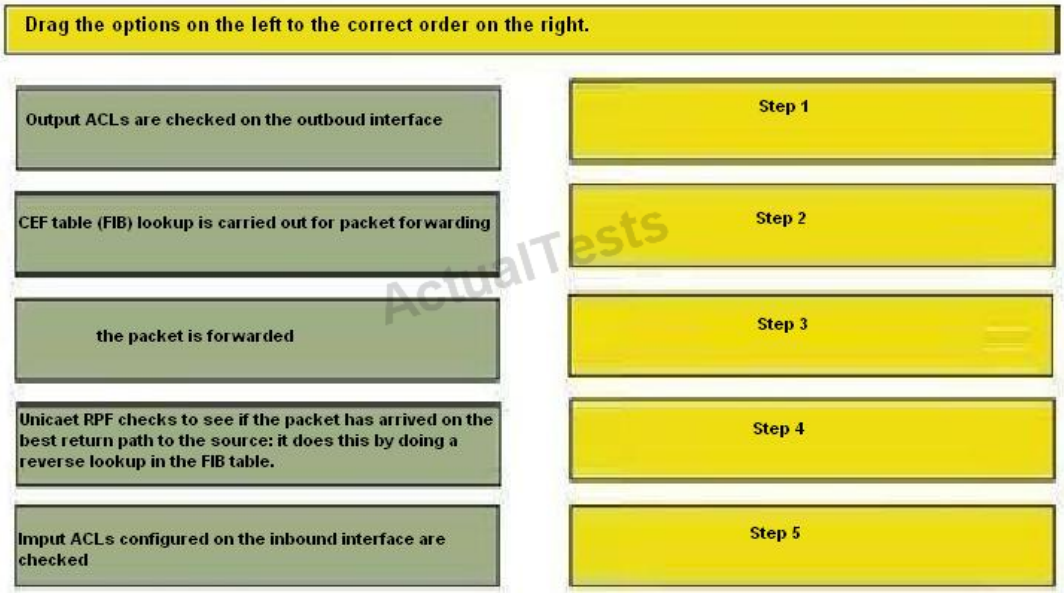
Explanation:

Encryption – 5
NAT inside to outside – 3
Check output access list – 4
Check input access list – 1
Redirect to web-cache – 2

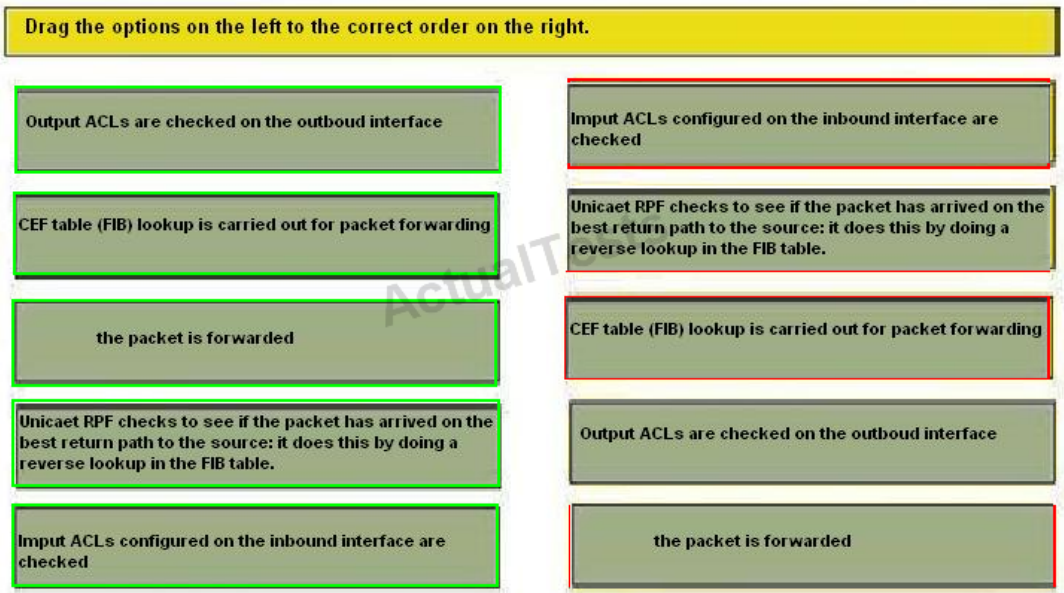
Reference:

http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080133ddd.shtml

QUESTION NO: 313 DRAG DROP



Answer:



Explanation:

Output ACLs are checked on the outbound interface – Step 4

CEF table (FIB) lookup is carried out for packet forwarding – Step 3

The packet is forwarded – Step 5

Unicast RPF checks to see if the packet has arrived on the best return path to the source, it does

this by doing a reverse lookup in the FIB table – Step 2
Input ACLs configured on the inbound interface are checked – Step 1

QUESTION NO: 314

What is a requirement to enable Cisco IOS IPS with 5.x signature?

A. disable Zone-Based Firewall as the two features are not compatible
B. disable Cisco Express Forwarding as the two features are not compatible
C. generate a certificate and export on Cisco.com to receive a signature update
D. import the public RSA key from the Cisco IPS team that allows the router to verify that a signature update (which was signed by this key) comes from Cisco

Answer: D
Reference: http://www.cisco.com/en/US/docs/ios-xml/ios/sec_data_ios_ips/configuration/12-4t/sec-ips5-sig-fs-ue.html

QUESTION NO: 315

What is the minimum key size to enable SSH v2?

A. 512 bits
B. 768 bits
C. 1024 bits
D. 2048 bits

Answer: B
Explanation:
According to Cisco it is important to generate a key-pair with at least 768 as bit size when you enable SSH v2

Reference
http://www.cisco.com/en/US/tech/tk583/tk617/technologies_tech_note09186a00800949e2.shtml

QUESTION NO: 316

Refer to the exhibit.

```
R2#sh ip ospf int fa 0/0
FastEthernet0/0 is up, line protocol is up
  Internet Address 4.4.42.2/24, Area 142
  Process ID 4, Router ID 4.4.2.2, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 4.4.10.10, Interface address 4.4.42.10
  Backup Designated router (ID) 4.4.2.2, Interface address 4.4.42.2
  Flush timer for old DR LSA due in 00:00:27
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:09
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 3/3, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 11
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 4.4.10.10  (Designated Router)
  Suppress hello for 0 neighbor(s)
  Simple password authentication enabled
```

Which of the following options is correct?

A. OSPF peers are using Type 1 authentication.
B. OSPF peers are using Type 2 authentication.
C. Authentication is used, but there is a password mismatch.
D. The OSPF peer IP address is 172.16.10.36.

Answer: A

Explanation:

This document shows sample configurations for Open Shortest Path First (OSPF) authentication which allows the flexibility to authenticate OSPF neighbors. You can enable authentication in OSPF in order to exchange routing update information in a secure manner. OSPF authentication can either be none (or null), simple, or MD5. The authentication method "none" means that no authentication is used for OSPF and it is the default method. With simple authentication, the password goes in clear-text over the network. With MD5 authentication, the password does not pass over the network. MD5 is a message-digest algorithm specified in RFC 1321. MD5 is considered the most secure OSPF authentication mode. When you configure authentication, you must configure an entire area with the same type of authentication. Starting with Cisco IOS® Software Release 12.0(8), authentication is supported on a per-interface basis. This is also mentioned in RFC 2328, Appendix D. This feature is added in Cisco bug ID CSCdk33792 (registered customers only).

These are the three different types of authentication supported by OSPF.
Null Authentication—This is also called Type 0 and it means no authentication information is

included in the packet header. It is the default.

Plain Text Authentication—This is also called Type 1 and it uses simple clear-text passwords.
MD5 Authentication—This is also called Type 2 and it uses MD5 cryptographic passwords

Reference
http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080094069.shtml.com/

QUESTION NO: 317

Which protocol and port have been assigned by IANA for RADIUS authentication?

A. UDP/1812
B. UDP/1813
C. TCP/1812
D. TCP/1813

Answer: A
Explanation:
Technical description for port 1812:
The RADIUS (Remote Authentication Dial-In User Service) protocol running on the system port 1812 is related to its authentication module. This service is primarily an element of a networking protocol which allows for a deployment of centralized accounting, authorization and access procedures. This protocol allows for the management of network resources for the efficient usage of network services.

When a user or a device attempts to connect to a network service, the authentication process is normally executed. The protocol using the computer port 1812 determines the appropriate user privileges for the entered credentials. A corresponding record of the network access is recorded into the accounting server for the implementation of the Triple A process.

The protocol supported by the network port 1812 is normally deployed by ISPs (Internet Service Providers) due to its ubiquitous support nature.

This service also supports the implementation of VPNs (Virtual Private Networks) and wireless networking environments.

UDP port numbers

RADIUS has been officially assigned UDP ports 1812 for RADIUS Authentication and 1813 for RADIUS Accounting by the Internet Assigned Numbers Authority (IANA). However, prior to IANA allocation of ports 1812 and 1813, ports 1645 and 1646 (authentication and accounting, respectively) were used unofficially and became the default ports assigned by many RADIUS Client/Server implementations of the time. The tradition of using 1645 and 1646 for backwards compatibility continues to this day. For this reason many RADIUS Server implementations monitor both sets of UDP ports for RADIUS requests. Microsoft RADIUS servers default to 1812 and 1813. Cisco RADIUS servers listen on RADIUS ports UDP 1645 and UDP 1812 for authentication; on ports 1646 and 1813 for accounting and can be configured with non-standard ports. Juniper Networks' RADIUS servers listen on both unofficial and official ports 1645, 1812, 1646 and 1813 by default but can be configured with arbitrary ports.

Reference
http://en.wikipedia.org/wiki/RADIUS
http://www.pc-library.com/ports/tcp-udp-port/1812/

QUESTION NO: 318

What is also called Type 0 authentication in OSPF on Cisco Routers?

A. MD5
B. There is no Type 0 authentication
C. SHA1
D. Null

Answer: D
Explanation:
These are the three different types of authentication supported by OSPF.
Null Authentication—This is also called Type 0 and it means no authentication information is included in the packet header. It is the default.

Plain Text Authentication—This is also called Type 1 and it uses simple clear-text passwords.
MD5 Authentication—This is also called Type 2 and it uses MD5 cryptographic passwords.

Reference
http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080094069.shtml

QUESTION NO: 319

What is true about Unicast RPF in strict mode?

A. It works well with a multihomed environment.
B. It will inspect IP packets that are encapsulated in tunnels, such as GRE, LT2P, or PPTP.
C. uRPF is performed within the CEF switching path.
D. There might be a problem with DHCP as Unicast RPF is blocking packets with a 0.0.0.0 source address.

Answer: C
Explanation:
Understanding uRPF – Unicast Reverse Path Forwarding
Spoofed packets are a big problem with on the Internet, they are commonly used in DNS amplification attacks, and TCP SYN floods. Unfortunately there is no simple way to totally fix all spoofed packets on the Internet but if service providers implement ingress filtering on their network, it effectively stops such attacks with spoofed source addresses coming from their patch. The process is actually standardised Best Practice in BCP 38 “Network Ingress Filtering” which all service providers should implement if they have Internet facing services for good karma.
There are a number of ways of implementing ingress filtering, one of the technically simplest is to create ACLs of your customers global address ranges and only allow packets sourced from those ranges to leave your network. Configuration wise Unicast Reverse Path Forwarding (uRPF) is in my opinion the simplest way of managing this and it has a couple of extra features.

uRPF checks incoming unicast packets and validates that a return path exists, there is not much point in forwarding a packet if it doesnt know how to return it right?
There are 2 methods of implementation of uRPF strict and loose. Strict mode is where the source of the packet is reachable via the interface that it came from, this is nice for extra security on the edge of your network but not so good if you have multiple edges towards the Internet eg you peer at multiple IXPs where you might expect asymmetric routing. In such cases loose mode is used which checks that a return route exists in the routing table.

Configuration
The configuration is super simple, after CEF has been enabled just go to the interface you wish to check inbound traffic and use the following command, with the “rx” option for strict mode or “any” for loose mode.

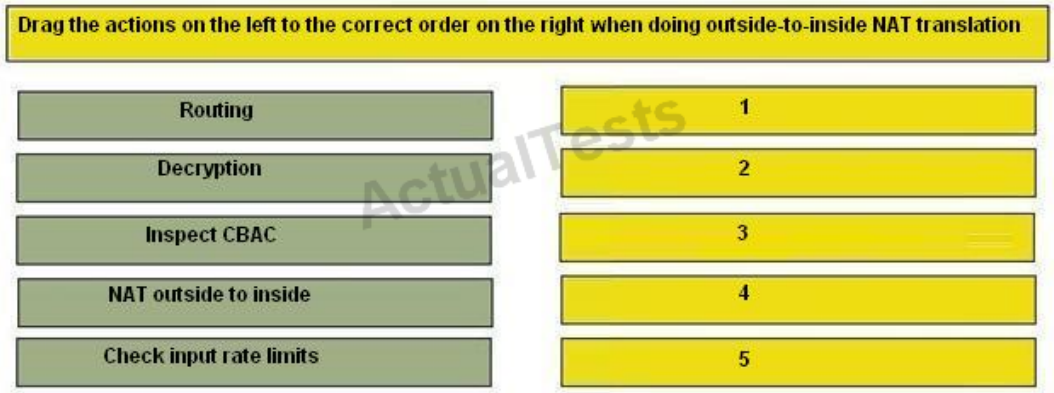
Router(config-if)#ip verify unicast source reachable-via ?

Any Source is reachable via any interface

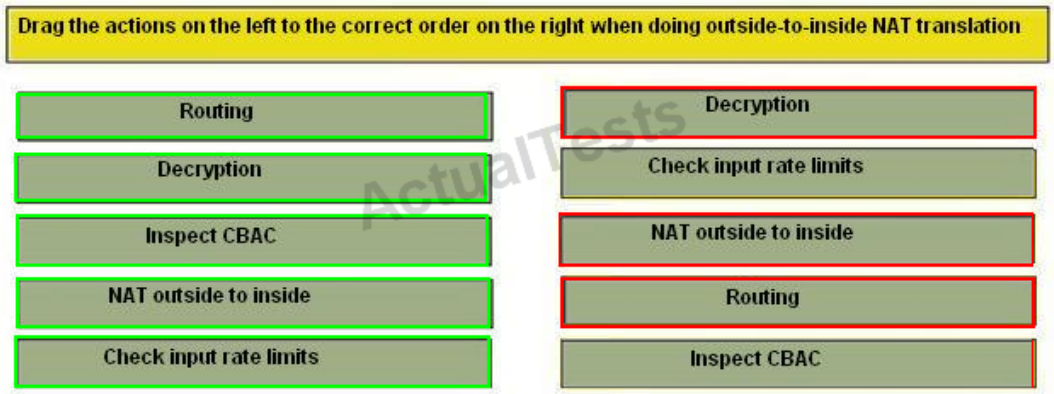
rx Source is reachable via interface on which packet was received
Verification
Obviously you can check the running config to see if its configured but if your a fan of using other show commands its visible under the sh cef interface and sh ip interface as shown below;

Router#sh cef interface fastEthernet 0/0 | i RPF
IP unicast RPF check is enabled
Router# sh ip int fa0/0 | i verify
IP verify source reachable-via RX

QUESTION NO: 320 DRAG DROP



Answer:



Explanation:
1.Decryption
2.Check input rate limits
3.NAT outside to inside
4.Routing

5.Inspect CBAC

Reference: NAT Order of Operation

http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080133ddd.shtml

QUESTION NO: 321

What is the purpose of an explicit "deny any" statement at the end of an ACL?

A. none, since it is implicit
B. to enable Cisco lOS IPS to work properly; however, it is the deny all traffic entry that is actually required
C. to enable Cisco lOS Firewall to work properly; however, it is the deny all traffic entry that is actually required
D. to allow the log option to be used to log any matches
E. to prevent sync flood attacks
F. to prevent half-opened TCP connections

Answer: D
Explanation:
As we know, there is always a “deny all” line at the end of each access-list to drop all other traffic that doesn’t match any “permit” lines. You can enter your own explicit deny with the “log” keyword to see what are actually blocked, like this:

Router(config)# access-list 1 permit 192.168.30.0 0.0.0.255
Router(config)# access-list 1 deny any log

Note: The log keyword can be used to provide additional detail about source and destinations for a given protocol. Although this keyword provides valuable insight into the details of ACL hits, excessive hits to an ACL entry that uses the log keyword increase CPU utilization. The performance impact associated with logging varies by platform. Also, using the log keyword disables Cisco Express Forwarding (CEF) switching for packets that match the access-list statement. Those packets are fast switched instead.

QUESTION NO: 322

Which of these is mandatory when configuring Cisco IOS Firewall?

A. Cisco IOS IPS enabled on the untrusted interface
B. NBAR enabled to perform protocol discovery and deep packet inspection
C. a route map to define the trusted outgoing traffic
D. a route map to define the application inspection rules
E. an inbound extended ACL applied to the untrusted interface

Answer: E

Explanation:

After the ACL is defined, it must be applied to the interface (inbound or outbound). In early software releases, out was the default when a keyword out or in was not specified. The direction must be specified in later software releases.

QUESTION NO: 323

Which feature is used to translate several internal addresses to only one or a few external addresses (also referred to as "overload")?

A. Network Address Translation
B. Address Translation Table
C. Overload Address Method
D. Port Address Translation

Answer: D

Explanation:

The PAT feature, a subset of NAT functionality, can be used to translate several internal addresses into only one or a few external addresses. PAT uses unique source port numbers on the private global IP address to distinguish between translations. Because the port number is encoded in 16 bits, the total number could theoretically be as high as 65,536 per IP address. PAT will attempt to preserve the original source port number. If this number is already allocated then PAT will attempt to find the first available port number starting from the beginning of the appropriate port group 0–511, 512–1023, or 1024–655351. If there is still no port number available from the appropriate group and more than one IP address is configured, PAT will move to the next IP address in the pool and try to allocate the original source port number again. This continues until it runs out of available ports and IP addresses.

QUESTION NO: 324

Which feature would prevent guest users from gaining network access by unplugging an IP phone

and connecting a laptop computer?

A. IPSec VPN
B. SSL VPN
C. port security
D. port security with statically configured MAC addresses
E. private VLANs

Answer: D
Explanation:
Port Security with Dynamically Learned and Static MAC Addresses
You can use port security with dynamically learned and static MAC addresses to restrict a port's ingress traffic by limiting the MAC addresses that are allowed to send traffic into the port. When you assign secure MAC addresses to a secure port, the port does not forward ingress traffic that has source addresses outside the group of defined addresses. If you limit the number of secure MAC addresses to one and assign a single secure MAC address, the device attached to that port has the full bandwidth of the port.

A security violation occurs in either of these situations:

When the maximum number of secure MAC addresses is reached on a secure port and the source MAC address of the ingress traffic is different from any of the identified secure MAC addresses, port security applies the configured violation mode.
If traffic with a secure MAC address that is configured or learned on one secure port attempts to access another secure port in the same VLAN, applies the configured violation mode.
Note After a secure MAC address is configured or learned on one secure port, the sequence of events that occurs when port security detects that secure MAC address on a different port in the same VLAN is known as a MAC move violation.
See the "Configuring the Port Security Violation Mode on a Port" section for more information about the violation modes.
After you have set the maximum number of secure MAC addresses on a port, port security includes the secure addresses in the address table in one of these ways:
You can statically configure all secure MAC addresses by using the switchport port-security mac-address

mac_address interface configuration command.
You can allow the port to dynamically configure secure MAC addresses with the MAC addresses of connected devices.
You can statically configure a number of addresses and allow the rest to be dynamically configured.
If the port has a link-down condition, all dynamically learned addresses are removed.
Following bootup, a reload, or a link-down condition, port security does not populate the address

table with dynamically learned MAC addresses until the port receives ingress traffic.
A security violation occurs if the maximum number of secure MAC addresses have been added to the address table and the port receives traffic from a MAC address that is not in the address table. You can configure the port for one of three violation modes: protect, restrict, or shutdown. See the "Configuring Port Security" section.
To ensure that an attached device has the full bandwidth of the port, set the maximum number of addresses to one and configure the MAC address of the attached device.

Reference
http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/port_sec.html#wp1061587

QUESTION NO: 325

For a router connected to two ISPs for redundancy, using IPSLA and static routing, how would you configure uRPF on the uplink interface?

A. ip verify unicast source reachable-via any
B. ip verify unicast reverse-path
C. ip verify unicast reverse-path loose
D. ip verify unicast reverse-path strict

Answer: A
Explanation:
Unicast RPF Examples
Cisco IOS Devices
An important consideration for deployment is that Cisco Express Forwarding switching must be enabled for Unicast RPF to function. This command has been enabled by default as of IOS version 12.2. If it is not enabled, administrators can enable it with the following global configuration command: ip cef Unicast RPF is enabled on a per-interface basis. The ip verify unicast source reachable-via rx command enables Unicast RPF in strict mode. To enable loose mode, administrators can use the any option to enforce the requirement that the source IP address for a packet must appear in the routing table. The allow-default option may be used with either therx or any option to include IP addresses not specifically contained in the routing table. The allow-self-ping option should not be used because it could create a denial of service condition. An access list such as the one that follows may also be configured to specifically permit or deny a list of addresses through Unicast RPF:

interface FastEthernet 0/0

ip verify unicast source reachable-via {rx | any} [allow-default]
[allow-self-ping] [list]
Addresses that should never appear on a network can be dropped by entering a route to a null interface. The following command will cause all traffic received from the 10.0.0.0/8 network to be dropped even if Unicast RPF is enabled in loose mode with the allow-default option: ip route 10.0.0.0 255.0.0.0 Null0

Reference
http://www.cisco.com/web/about/security/intelligence/unicast-rpf.html

QUESTION NO: 326

Refer to the exhibit.

```
*May 1 22:12:13.243: %SEC-6-IPACCESSLOGP: list ACL-IPv4-E0/0-IN permitted tcp
192.168.1.3(1024) -> 192.168.2.1(22), 1 packet
*May 1 22:17:16.647: %SEC-6-IPACCESSLOGP: list ACL-IPv4-E0/0-IN permitted tcp
192.168.1.3(1024) -> 192.168.2.1(22), 9 packets
*May 3 19:05:38.183: %IPV6-6-ACCESSLOGP: list ACL-IPv6-E0/0-IN/10 permitted tcp
2001:DB8::3(1027) -> 2001:DB8:1000::1(22), 1 packet
*May 3 19:11:32.619: %IPV6-6-ACCESSLOGP: list ACL-IPv6-E0/0-IN/10 permitted tcp
2001:DB8::3(1027) -> 2001:DB8:1000::1(22), 9 packets
```

Which two statements are correct? (Choose two.)

A. The packet is received on the private interface.
B. Zone-Based Firewall is configured on the router.
C. Logging is enabled for the access list.
D. Private is an identification of that ACE.
E. The Telnet connection is successfully set up.

Answer: C,D
Explanation:
The log and log-input Access Control Entry Options
The log and log-input options apply to an individual ACE and cause packets that match the ACE to be logged.
The log-input option enables logging of the ingress interface and source MAC address in addition to the packet's source and destination IP addresses and ports.
The first packet logged via the log or log-input options will generate a syslog message. There are two scenarios in which subsequent log messages will not be sent immediately. If the log-enabled ACE matches another packet with identical characteristics to the packet that generated a log message, the number of packets matched is incremented and then reported at five-minute

intervals. Similarly, if any log-enabled ACE in any ACL on any interface matches a packet within one second of the initial log message, the match or matches are counted for five minutes and then reported. These periodic updates will contain the number of packets matched since the previous message.

Reference
http://www.cisco.com/web/about/security/intelligence/acl-logging.html#2
http://www.cisco.com/en/US/docs/app_ntwk_services/data_center_app_services/ace_appliances/vA1_7_/ configuration/system/message/guide/config.html

http://www.cisco.com/en/US/docs/interfaces_modules/services_modules/ace/v3.00_A2/configuration/system/message/guide/config.html

QUESTION NO: 327

Which of these statements accurately identifies how Unicast Reverse Path Forwarding can be employed to prevent the use of malformed or forged IP sources addresses?

A. It is applied only on the input interface of a router.
B. It is applied only on the output interface of a router.
C. It can be configured either on the input or output interface of a router.
D. It cannot be configured on a router interface.
E. It is configured under any routing protocol process.

Answer: A
Explanation:
Unicast Reverse Path Forwarding:
Is a small security feature, when configured on an interface, the router checks the incoming packet's source address with its routing table. If the incoming packet's source is reachable via the same interface it was received, the packet is allowed. URPF provides protection again spoofed packets with unverifiable source.
http://www.cciecandidate.com/?p=494
Unicast RPF can be used in any "single-homed" environment where there is essentially only one access point out of the network; that is, one upstream connection. Networks having one access point offer the best example of symmetric routing, which means that the interface where a packet enters the network is also the best return path to the source of the IP packet. Unicast RPF is best used at the network perimeter for Internet, intranet, or extranet environments, or in ISP environments for customer network terminations.

Feature Overview
The Unicast RPF feature helps to mitigate problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network by discarding IP packets that lack a verifiable IP source address.
For example, a number of common types of denial-of-service (DoS) attacks, including Smurf and Tribe Flood Network (TFN), can take advantage of forged or rapidly changing source IP addresses to allow attackers to thwart efforts to locate or filter the attacks. For Internet service providers (ISPs) that provide public access, Unicast RPF deflects such attacks by forwarding only packets that have source addresses that are valid and consistent with the IP routing table. This action protects the network of the ISP, its customer, and the rest of the Internet.

How It Works
When Unicast RPF is enabled on an interface, the router examines all packets received as input on that interface to make sure that the source address and source interface appear in the routing table and match the interface on which the packet was received. This "look backwards" ability is available only when Cisco express forwarding (CEF) is enabled on the router, because the lookup relies on the presence of the Forwarding Information Base (FIB). CEF generates the FIB as part of its operation.

Note Unicast RPF is an input function and is applied only on the input interface of a router at the upstream end of a connection.

Reference
http://www.cisco.com/en/US/docs/ios/11_1/feature/guide/uni_rpf.html

QUESTION NO: 328

Unicast Reverse Path Forwarding can perform all of these actions except which one?

A. examine all packets received to make sure that the source addresses and source interfaces appear in the routing table and match the interfaces where the packets were received
B. check to see if any packet received at a router interface arrives on the best return path
C. combine with a configured ACL
D. log its events, if you specify the logging options for the ACL entries used by the unicast rpf command
E. inspect IP packets encapsulated in tunnels, such as GRE

Answer: E
Explanation:

For RPF to function, CEF must be enabled on the router. This is because the router uses the Forwarding Information Base (FIB) of CEF to perform the lookup process, which is built from the router's routing table. In other words, RPF does not really look at the router's routing table; instead, it uses the CEF FIB to determine spoofing.

Also, RPF cannot detect all spoofed packets. For the network in this example, the perimeter router cannot determine spoofing from packets received on the external E1 interface if they match the default route statement. Therefore, the more routes your router has in its CEF FIB table, the more likely the router will be capable of detecting spoofing attacks. In addition, RPF cannot detect any spoofed packets that are encapsulated, such as packets encapsulated in GRE, IPSec, L2TP, and other packets.

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. Unicast RPF works in one of three different modes: strict mode, loose mode, or VRF mode. Note that not all network devices support all three modes of operation. Unicast RPF in VRF mode will not be covered in this document.

When administrators use Unicast RPF in strict mode, the packet must be received on the interface that the router would use to forward the return packet. Unicast RPF configured in strict mode may drop legitimate traffic that is received on an interface that was not the router's choice for sending return traffic. Dropping this legitimate traffic could occur when asymmetric routing paths are present in the network.

When administrators use Unicast RPF in loose mode, the source address must appear in the routing table. Administrators can change this behavior using the allow-default option, which allows the use of the default route in the source verification process. Additionally, a packet that contains a source address for which the return route points to the Null 0 interface will be dropped. An access list may also be specified that permits or denies certain source addresses in Unicast RPF loose mode.

Care must be taken to ensure that the appropriate Unicast RPF mode (loose or strict) is configured during the deployment of this feature because it can drop legitimate traffic. Although asymmetric traffic flows may be of concern when deploying this feature, Unicast RPF loose mode is a scalable option for networks that contain asymmetric routing paths.

Unicast RPF in an Enterprise Network In many enterprise environments, it is necessary to use a combination of strict mode and loose mode Unicast RPF. The choice of the Unicast RPF mode that will be used will depend on the design of the network segment connected to the interface on which Unicast RPF is deployed.

Administrators should use Unicast RPF in strict mode on network interfaces for which all packets received on an interface are guaranteed to originate from the subnet assigned to the interface. A subnet composed of end stations or network resources fulfills this requirement. Such a design would be in place for an access layer network or a branch office where there is only one path into and out of the branch network. No other traffic originating from the subnet is allowed and no other

routes are available past the subnet. Unicast RPF loose mode can be used on an uplink network interface that has a default route associated with it.

Reference
http://www.cisco.com/web/about/security/intelligence/unicast-rpf.html

QUESTION NO: 329

Which two of these elements need to be configured prior to enabling SSH? (Choose two.)

A. hostname
B. loopback address
C. default gateway
D. domain name
E. SSH peer address

Answer: A,D
Explanation:
To enable Secure Shell (SSHv2) version 2 (disable version 1) on a Cisco router an IOS with 3des encryption is required. When there is no SSH version configured, version 1 and 2 will be supported both.

Follow the next steps to enable SSH:
1. Configure the hostname command.
2. Configure the DNS domain.
3. Generate RSA key to be used.
4. Enable SSH transport support for the virtual type terminal (vty)
Example SSH version 2 configuration:
hostname ssh-router
aaa new-model
username cisco password cisco
ip domain-name routers.local
!
Specifies which RSA keypair to use for SSH usage.
ip ssh rsa keypair-name sshkeys
!
Enables the SSH server for local and remote authentication on the router.
!

For SSH Version 2, the modulus size must be at least 768 bits. crypto key generate rsa usage-keys label
sshkeys modulus 768 !
Configures SSH control variables on your router.
ip ssh timeout 120
!
configure SSH version 2 (will disable SSH version 1)
ip ssh version 2
!
disable Telnet and enable SSH
line vty 0 4
transport input SSH
Commands to verify SSH configuration:
show ssh
show ip ssh
debug ip ssh

QUESTION NO: 330

You are the network administrator of a small Layer 2 network with 50 users. Lately, users have been complaining that the network is very slow. While troubleshooting, you notice that the CAM table of your switch is full, although it supports up to 12,000 MAC addresses. How should you solve this issue and prevent it from happening in the future?

A. Upgrade the switches.
B. Configure BPDU guard.
C. Configure VLAN access lists.
D. Configure port security.
E. Configure Dynamic ARP inspection.

Answer: D
Reference: http://www.slideshare.net/techdude/router-and-switch-presentation (slide 14)

QUESTION NO: 331

Refer to the exhibit.

```
class-map type inspect match-all c1
 match access-group 101
 match protocol http
policy-map type inspect p1
 class type inspect c1
 drop
```

What is true about the configuration in this exhibit?

A. It is an invalid configuration because it includes both an application layer match and a Layer 3 ACL.
B. It will create a class map that matches the content of ACL 101 and the HTTP protocol, and will then create an inspection policy that will drop packets at the class map.
C. It will create a class map that matches the content of ACL 101 and the HTTP protocol, and will then create an inspection policy that will allow packets at the class map.
D. It will create a class map that matches the content of ACL 101 or the HTTP protocol (depending on the zone of the interface), and will then create an inspection policy that will drop packets at the class map.
E. It will create a class map that matches the content of ACL 101 or the HTTP protocol (depending on the zone of the interface), and will then create an inspection policy that will allow packets at the class map.
F. It is an invalid configuration because the class map and policy map names must match.

Answer: B

Explanation: Explanation
Technically the syntax is incorrect as the application that is being inspected should be listed after the keyword type. However, this is not listed as one of the options. The correct configuration should be as follows:

class-map type inspect http match-all el
match access-group 101
policy-map type inspect http pl
class type inspect el drop

When multiple match criteria exist in the traffic class, you can identify evaluation instructions using the match any or match-all keywords. If you specify match-any as the evaluation instruction, the traffic being evaluated must match one of the specified criteria, typically match commands of the same type. If you specify match-all as the evaluation instruction, the traffic being evaluated must match all the specified criteria, typically match commands of different types.

Identifying Traffic in an Inspection Class Map
This type of class map allows you to match criteria that is specific to an application. For example, for DNS traffic, you can match the domain name in a DNS query.
Note Not all applications support inspection class maps. See the CLI help for a list of supported applications. A class map groups multiple traffic matches (in a match-all class map), or lets you match any of a list of matches (in a match-any class map). The difference between creating a class map and defining the traffic match directly in the inspection policy map is that the class map lets you group multiple match commands, and you can reuse class maps. For the traffic that you identify in this class map, you can specify actions such as dropping, resetting, and/or logging the connection in the inspection policy map. If you want to perform different actions on different types of traffic, you should identify the traffic directly in the policy map.

To define an inspection class map, perform the following steps:
Step 1 (Optional) If you want to match based on a regular expression, see the "Creating a Regular Expression" section and the "Creating a Regular Expression Class Map" section.
Step 2 Create a class map by entering the following command:
hostname(config)# class-map type inspect application [match-all | match-any]
class_map_name hostname(config-cmap)#

Reference
http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.8/vfw/command/reference/vfr38cm.html

QUESTION NO: 332

Why would a rogue host that is running a DHCP server on a campus LAN network present a security risk?

A. It may allocate IP addresses from an unknown subnet to the users.
B. All multicast traffic can be sniffed by using the DHCP multicast capabilities.
C. The CPU utilization of the first hop router can be overloaded by exploiting DHCP relay open ports.
D. A potential man-in-the-middle attack can be used against the clients.

Answer: D
Explanation:
A rogue DHCP server is typically used in conjunction with a network attacker who launches man-in-the-middle (MitM) attacks. MitM is an attack technique in which the attacker exploits normal protocol processing behavior to reroute normal traffic flow between two endpoints. A hacker will broadcast DHCP requests with spoofed MAC addresses, thereby exhausting the address space of

the legitimate DHCP server. Once the addresses are exhausted, the rogue DHCP server provides DHCP responses to users' DHCP requests. These responses would include DNS servers and a default gateway, which would be used to launch a MitM attack.

QUESTION NO: 333

Which is the result of enabling IP Source Guard on an untrusted switch port that does not have DHCP snooping enabled?

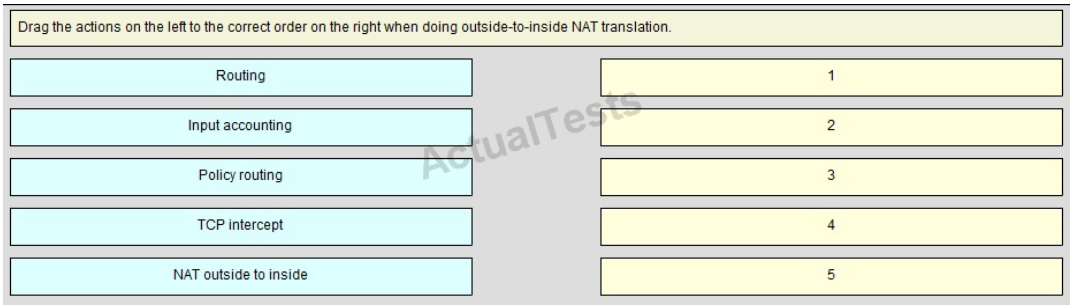
A. DHCP requests will be switched in the software, which may result in lengthy response times.
B. The switch will run out of ACL hardware resources.
C. All DHCP requests will pass through the switch untested.
D. The DHCP server reply will be dropped and the client will not be able to obtain an IP address.

Answer: D
Explanation:
DHCP snooping is a feature that provides network security by filtering untrusted DHCP messages and by building and maintaining a DHCP snooping binding database. DHCP snooping acts like a firewall between untrusted hosts and DHCP servers. DHCP snooping allows all DHCP messages on trusted ports, but it lters DHCP messages on untrusted ports.

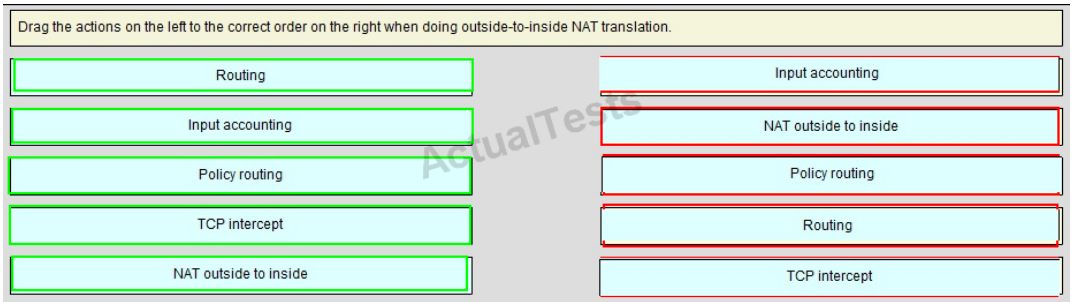
Cisco switches can use DHCP snooping feature to mitigate this type of attack. When DHCP snooping is enabled, switch ports are classified as trusted or untrusted. Trusted ports are allowed to send all types of DHCP messages while untrusted ports can send only DHCP requests. If a DHCP reply is seen on an untrusted port, the port is shut down.
By default, if you enable IP source guard without any DHCP snooping bindings on the port, a default port access-list (PACL) that denies all IP traffic expect the DHCP Request (DHCP Discover) is installed on the port. Therefore the DHCP Server can hear the DHCP Request from the Client but its reply is filtered by the switch and the client can’t obtain an IP address -> D is correct.
Some useful information about DHCP snooping & IP Source Guard:
When enabled along with DHCP snooping, IP Source Guard checks both the source IP and source MAC addresses against the DHCP snooping binding database (or a static IP source entry). If the entries do not match, the frame is ltered. For example, assume that the**show ip dhcp snooping binding** command displays the following binding table entry:
MacAddress
IpAddress
LeaseSec
Type

VLAN
Interface
01:25:4A:5E:6D:25
10.0.0.20
6943
dhcp-snooping
2
FastEthernet0/1
If the switch receives an IP packet with an IP address of 10.0.0.20, IP Source Guard forwards the packet only if the MAC address of the packet is 01:25:4A:5E:6D:25.

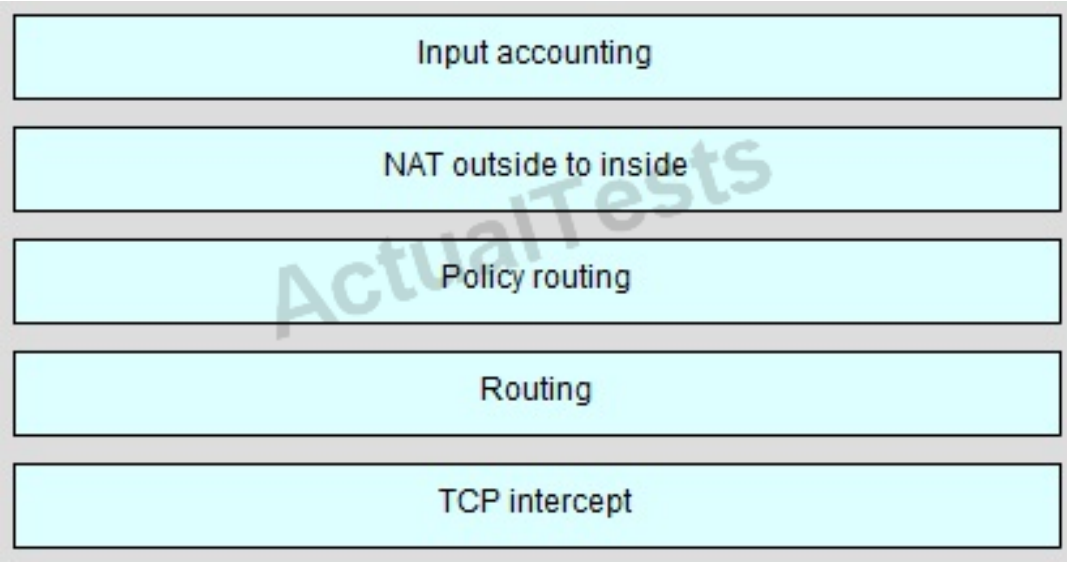
QUESTION NO: 334 DRAG DROP



Answer:



Explanation:



QUESTION NO: 335

If you have overlapping IP address between two different networks or routing domains, which two commands are needed to globally configure NAT to get this to work?

A. ip nat outside source static udp x.x.x.x y.y.y.y and ip nat inside source udp x.x.x.x y.y.y.y
B. ip nat outside source static x.x.x.x y.y.y.y and ip nat inside source static x.x.x.x y.y.y.y
C. ip nat outside source static tcp x.x.x.x y.y.y.y and ip nat outside source tcp x.x.x.x y.y.y.y
D. ip nat outside source list 1 interface x and ip nat inside source list 1 interface x

Answer: B
Explanation:
IP nat outside source list
Translates the source of the IP packets that are traveling outside to inside. Translates the destination of the IP packets that are traveling inside to outside
IP nat inside source list Translates the source of IP packets that are traveling inside to outside. Translates the destination of the IP packets that are traveling outside to inside

QUESTION NO: 336

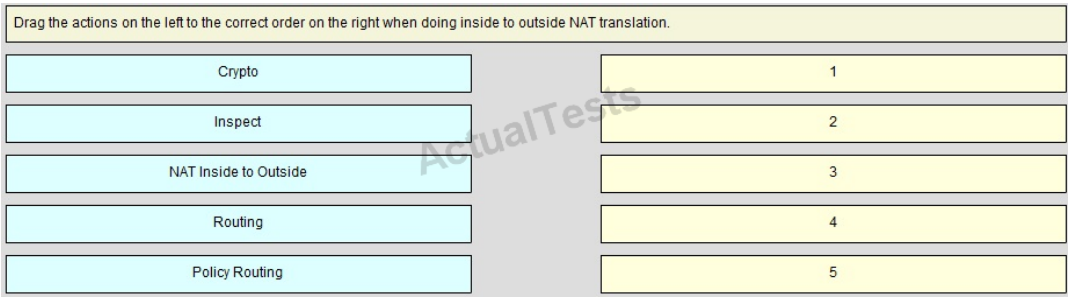
Which IOS security feature is configured by the ip inspect inspection-name {in | out} command?

A. IPsec site-to-site VPN
B. Cisco AutoSecure
C. Cisco IOS Firewall
D. IPS

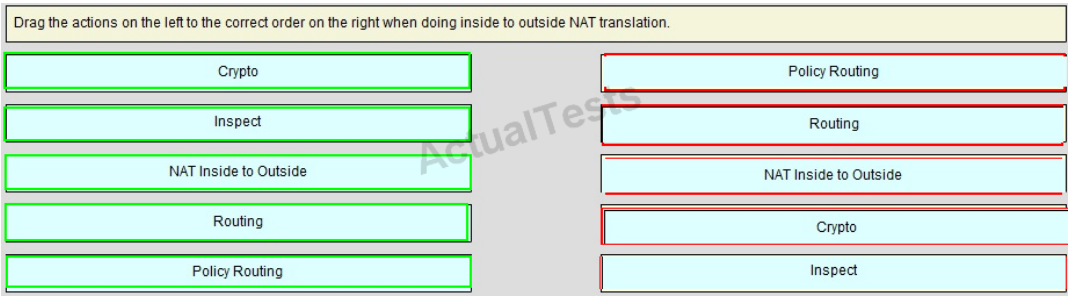
Answer: C
Explanation:
CBAC is a function of the Cisco IOS feature set. CBAC is configured using the "ip inspect" command. The ip inspect inspection-name {in | out} command is used to apply the inspection rule to an interface. The keyword in is used for inbound traffic when the CBAC is applied on the internal (trusted, or secure) interface. The keyword out is used for outbound traffic when the CBAC is applied on the external, unsecured interface

Reference
http://www.ciscopress.com/articles/article.asp?p=26533

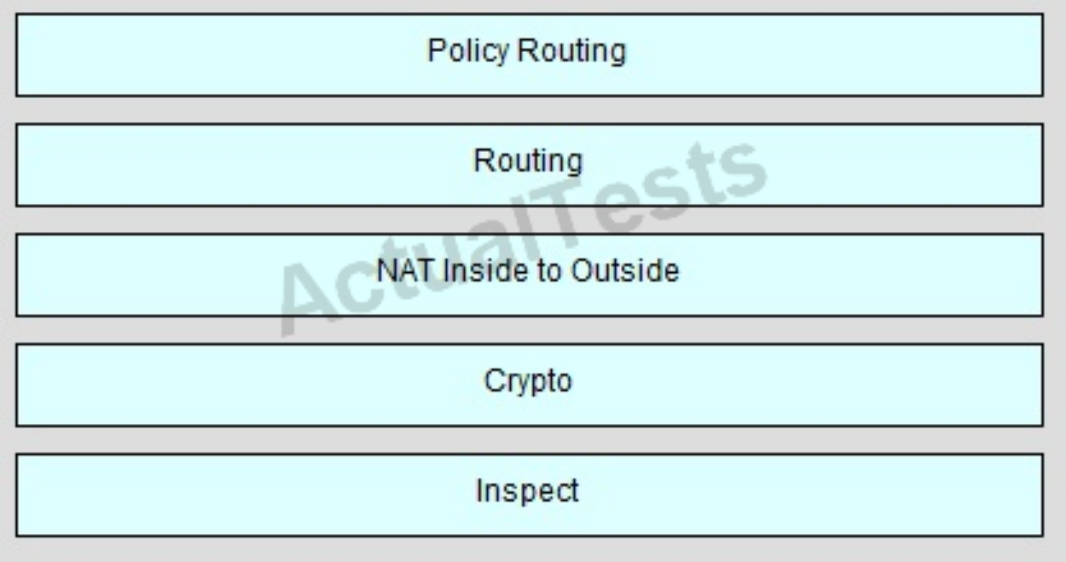
QUESTION NO: 337 DRAG DROP



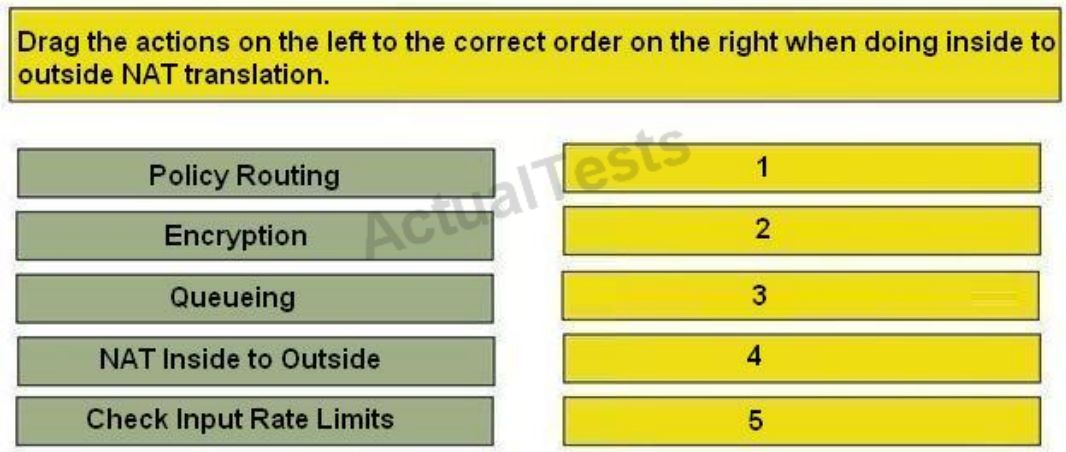
Answer:



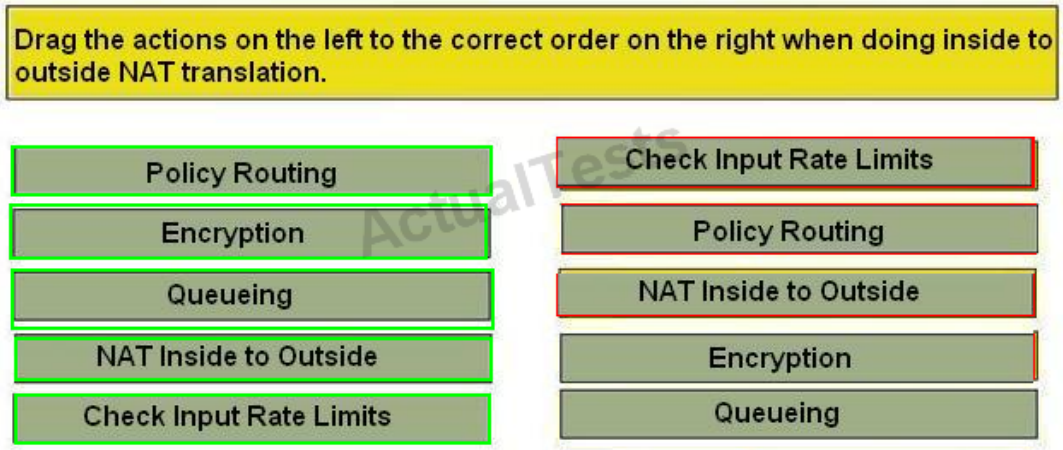
Explanation:



QUESTION NO: 338 DRAG DROP



Answer:



Explanation:

Routing - 2Encryption - 4Queueing - 5NAT Inside to Outside - 3Check Input Rate Limits - 1
Reference: NAT Order of Operationhttp://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080133ddd.shtml

QUESTION NO: 339

Which two statements are true about the Inside Global address in NAT? (Choose two.)

A. the IP address of an inside host as it appears to the outside network
B. the IP address of an outside host as it appears to the inside network
C. if the enterprise is connected to the global Internet, this address can be allocated from a globally unique address space
D. if the enterprise is connected to the global Internet, this address can be allocated from the space defined by RFC 1918

Answer: A,C

Explanation: "global" is what you see on the Internet, "local" is what you see in your company "inside"/"outside" is where the box physically resides

Obviously, you can only use globally unique addresses for global addresses.

QUESTION NO: 340

Which three configuration items are required to enable SSH on a router? (Choose three.)

A. a domain name
B. an RSA key
C. a hostname
D. a self-signed certificate
E. a RADIUS server
F. a username and password

Answer: A,B,C
Reference: http://www.cisco.com/en/US/docs/ios/12_2/security/configuration/guide/scfssh.html (prerequisites to configuring SSH)

QUESTION NO: 341

Which three protocols should be explicitly managed by using Control Plane Policing on an Internet border router? (Choose three.)

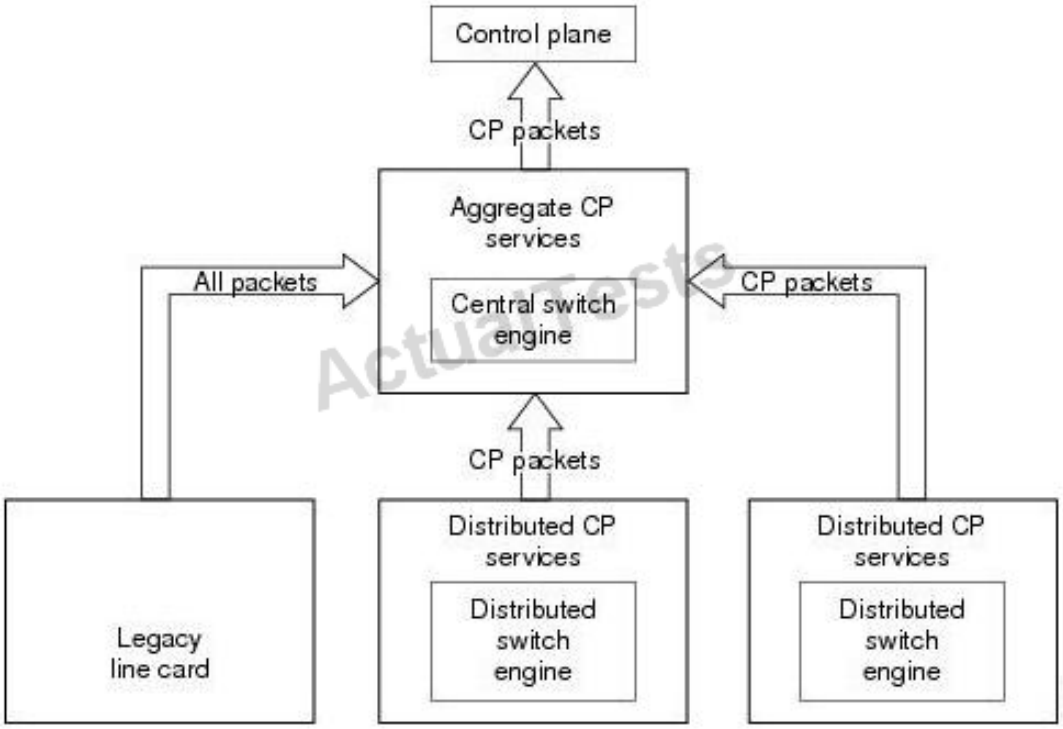
A. LDAP
B. ICMP
C. RTP
D. BGP
E. SSH
F. RDP

Answer: B,D,E
Explanation:
Control Plane Security and Packet QoS Overview
To protect the CP on a router from DoS attacks and to provide packet QoS, the Control Plane Policing feature treats the CP as a separate entity with its own ingress (input) and egress (output) ports, which are like ports on a router and switch. Because the Control Plane Policing feature treats the CP as a separate entity, a set of rules can be established and associated with the ingress and egress port of the CP.

These rules are applied only after the packet has been determined to have the CP as its destination or when a packet exits from the CP. Thereafter, you can configure a service policy to prevent unwanted packets from progressing after a specified rate limit has been reached; for example, a system administrator can limit all TCP/SYN packets that are destined for the CP to a maximum rate of 1 megabit per second. Input CP services are executed after router input port services and a routing decision on the input path have been made. As shown in Figure 2, CP security and packet QoS are applied on:

An aggregate level by the central switch engine and applied to all CP packets received from all

line cards on the router (see Aggregate Control Plane Services)
A distributed level by the distributed switch engine of a line card and applied to all CP packets received from all interfaces on the line card (see Distributed Control Plane Services)
Figure 2 Input Control Plane Services: Aggregate and Distributed Services



The following types of Layer 3 packets are forwarded to the control plane and processed by aggregate and distributed control plane policing:
Routing protocol control packets
Packets destined for the local IP address of the router
Packets from management protocols (such as Simple Network Management Protocol [SNMP], Telnet, and secure shell [SSH])

Note Ensure that Layer 3 control packets have priority over other packet types that are destined for the control plane.

Reference
http://www.cisco.com/en/US/docs/ios/12_3t/12_3t4/feature/guide/gtrtlimt.html

QUESTION NO: 342

Refer to the exhibit.

```
!
class-map type inspect match-all CMAP_ICMP
 match protocol icmp
!
policy-map type inspect PMAP_A_B
 class type inspect CMAP_ICMP
  inspect
 class class-default
  drop
!
zone security INSIDE
zone security OUTSIDE
zone-pair security A_B source INSIDE destination OUTSIDE
 service-policy type inspect PMAP_A_B
!
interface FastEthernet0/0
 ip address 10.48.67.125 255.255.254.0
 zone-member security OUTSIDE
 duplex auto
 speed auto
 !
!
interface FastEthernet0/1
 ip address 192.168.101.1 255.255.255.0
 zone-member security INSIDE
 duplex auto
 speed auto
 !
!
```

ActualTests

What is true about traffic from the INSIDE zone to the OUTSIDE zone?

A. All icmp echo requests will be inspected.
B. All IP traffic will be dropped.
C. All icmp echo requests will be passed, but the icmp echo reply to the echo request from the OUTSIDE zone will be dropped. /..l
D. All IP traffic will be inspected.

Answer: A

Explanation:

When the traffic moves from inside to outside zone, ICMP echo requests will be inspected because the inspection is set using policy-map-type command is used.

QUESTION NO: 343

Refer to the exhibit.

```
class-map match-all telnet-allowed
  match intput-interface fa0/0
  match access-list telnet
class-map telnet-deny
  match access-list telnet

Policy-map telnet
  class telnet-allowed
    bandwidth 64
  class telnet-deny
    drop

control-plane
  service-policy input telnet
```

Which option is correct?

A. This configuration is not valid.
B. The user can use Telnet to any interface on the device as long as the input interface for Telnet is FastEthernet 0/0.
C. Telnet access is not allowed on the router.
D. The user can use Telnet from the router only if the Telnet is sourced from FastEthernet 0/0.

Answer: B

Explanation:

As in the exhibit, the access list deny any input from the interface for telnet if the fastethernet is not 0/0

QUESTION NO: 344

If a certificate authority trustpoint is not configured when enabling HTTPS and the remote HTTPS server requires client authentication, connections to the secure HTTP client will fail. Which command must be enabled for correct operation?

A. ip http client secure-ciphersuite 3des-ede-cbc-sha
B. ip https max-connections 10
C. ip http timeout-policy idle 30 life_120 requests 100
D. ip http client secure-trustpoint trustpoint-name

Answer: D

Explanation:

IP http client secure-trustpoint

To specify the remote certificate authority (CA) trustpoint that should be used if certification is needed for the secure HTTP client, use the ip http client secure-trustpoint command in global configuration mode. To remove a client trustpoint from the configuration, use the no form of this command.

IP http client secure-trustpoint trustpoint-name

Topic 7, Implement Network Services

QUESTION NO: 345

Refer to the exhibit.

```
Router#show standby
Ethernet0/0 - Group 35 (version 2)
  State is Standby
    6 state changes, last state change 00:01:22
  Virtual IP address is 10.1.1.1
  Active virtual MAC address is 0000.0c9f.f023
    Local virtual MAC address is 0000.0c9f.f023 (v2 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.688 secs
  Preemption disabled
  Active router is 10.1.1.101, priority 100 (expires in 9.440 sec)
    MAC address is aabb.cc00.6500
  Standby router is local
  Priority 150 (configured 150)
  Group name is "hsrp-Et0/0-25" (default)
```

Which option best describes how the virtual MAC address is composed?

A. based on a randomly generated number
B. based on the burned-in MAC address of the router
C. based on a number manually configured by the administrator
D. based on the configured standby group number

Answer: C

Explanation: Explanation

http://www.thebryantadvantage.com/CCNPCertificationBCMSNExamHSRPMACAddress.htm

QUESTION NO: 346

Refer to the exhibit.

```
R100#show glbp
Ethernet0/0 - Group 150
  State is Active
  Virtual IP address is 10.1.1.150
  Preemption enabled, min delay 0 sec
  Active is local
  Standby is 10.1.1.101, priority 100 (expires in 8.128 sec)
  Priority 150 (configured)
  Weighting 50 (configured 50), thresholds: lower 1, upper 50
  Load balancing: round-robin
  Group members:
    aabb.cc00.6400 (10.1.1.100) local
    aabb.cc00.6500 (10.1.1.101)
  There are 2 forwarders (1 active)
  Forwarder 1
    State is Active
      1 state change, last state change 00:07:56
    Redirection enabled
    Preemption enabled, min delay 30 sec
    Active is local, weighting 10
  Forwarder 2
    State is Listen
    Redirection enabled, 599.360 sec remaining (maximum 600 sec)
    Time to live: 14399.360 sec (maximum 14400 sec)
    Active is 10.1.1.101 (primary), weighting 30 (expires in 10.432 sec)
```

ActualTests

How will traffic be split between the routers, assuming that there are many hosts on this subnet?

A. All traffic will be sent to the primary router (10.1.1.100).
B. Traffic will be split equally between the two routers (10.1.1.100 and 10.1.1.101).
C. Traffic will be split 25% (10.1.1.101) / 75% (10.1.1.100) between the two routers.
D. Traffic will be split 75% (10.1.1.101) / 25% (10.1.1.100) between the two routers.

Answer: D

Explanation: Explanation
In addition to being able to set priorities on different gateway routers, GLBP allows a weighting parameter to be set. Based on this weighting (compared to others in the same virtual router group), ARP requests will be answered with MAC addresses pointing to different routers. Thus, load balancing is not based on traffic load, but rather on the number of hosts that will use each gateway router. By default GLBP load balances in roundrobin fashion.
Load Balancing Modes

There will be three types of load balancing methods that can be configured:

Weighted
Host dependant
Round robin
Round Robin Load Balancing Algorithm

Each Virtual Forwarder MAC address takes turns being included in address resolution replies for the virtual IP address. Round robin load balancing is recommended for situations where there are a small number of end hosts.
If no load-balance algorithm is specified then GLBP will operate in a similar fashion to HSRP, i.e. the AVG will only respond to ARP requests with its own VF MAC address, and all traffic will therefore be directed to the AVG. No load balancing is defined using the following configuration statement:
no glbp <glbp-group> load-balancing
The load balancing method will be set to default (round-robin) if any load balancing statement is omitted.

Load Sharing
GLBP weighting has the ability to place a weight on each device when calculating the amount of load sharing that will occur through MAC assignment. Each GLBP router in the group will advertise its weighting and assignment. The AVG will act based on that value.
The only reason you would use this is if you have a larger circuit on the primary router than on the backup router. So the higher weight of 160 will take twice as much traffic as the lower weight of 80. If the weights are

120 and 80 respectively, the primary router will take a quarter more traffic than the backup.

```
track 1 interface Serial0/0/0.1 ip routing  ! Track 1 is watching Serial0/0/0.1 So
it can be weighted.
!
interface GigabitEthernet0/0
 ip address x.x.x.x 255.255.255.0           ! Make sure this IP is in the same
network as your Virtual Gateway IP.
 glbp 1 ip x.x.x.x                          ! You Virtual Gateway
 glbp 1 weighting 100 lower 91 upper 99     ! Threshold of 91-99 If it fails the
weight will drop below 91 and state will be standby
 glbp 1 load-balancing weighted             ! Will Load Balance according to
weight.
 glbp 1 weighting 160                       ! Higher the weight the Higher
Priority
 glbp 1 weighting track 1                   ! The weight is based on the Track 1
interface watching Serial 0/0/0.101
 no ip redirect                             ! Does not redirect traffic if the
router has to Resend to sender.
 speed 100
 duplex full
 track 1 interface Serial0/0.101 ip routing ! Track 1 is watching Serial0/0/0.101
So it can be weighted.
!
interface GigabitEthernet0/0
 ip address x.x.x.x 255.255.255.0           ! Make sure this IP is in the same
network as your Virtual Gateway IP.
 glbp 1 ip x.x.x.x                          ! You Virtual Gateway
 glbp 1 weighting 100 lower 91 upper 99     ! Range of 91-99 If it fails the
weight will drop below 91 and state will be standby
 glbp 1 load-balancing weighted             ! Will Load Balance according to
weight.
 glbp 1 weighting 80                        ! Weight is lower on this router,
traffic will be directed through the other router
 glbp 1 weighting track 1                   ! The weight is based on the Track 1
interface watching Serial 0/0/0.101
 no ip redirect                             ! Does not redirect traffic if the
router has to Resend to sender.
 speed 100
 duplex full
```

ActualTests

Reference
http://en.wikipedia.org/wiki/Gateway_Load_Balancing_Protocol
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6550/prod_presentation0900aecd801790a3_ps6600_Products_Presentation.html
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6554/ps6600/product_data_sheet0900aecd803a546c.html
http://www.cisco.com/en/US/docs/ios/12_2t/12_2t15/feature/guide/ft_glbp.html#wp1027129

QUESTION NO: 347

For the following protocols, which one provides a mechanism to transparently intercept and redirect CIFS traffic from a client to a local Cisco Wide Area Application engine?

A. Virtual Router Redundancy Protocol (VRRP)
B. File Transport Protocol (FTP)
C. Hot Standby Routing Protocol (HSRP)
D. Web Cache Communication Protocol (WCCP)

Answer: D

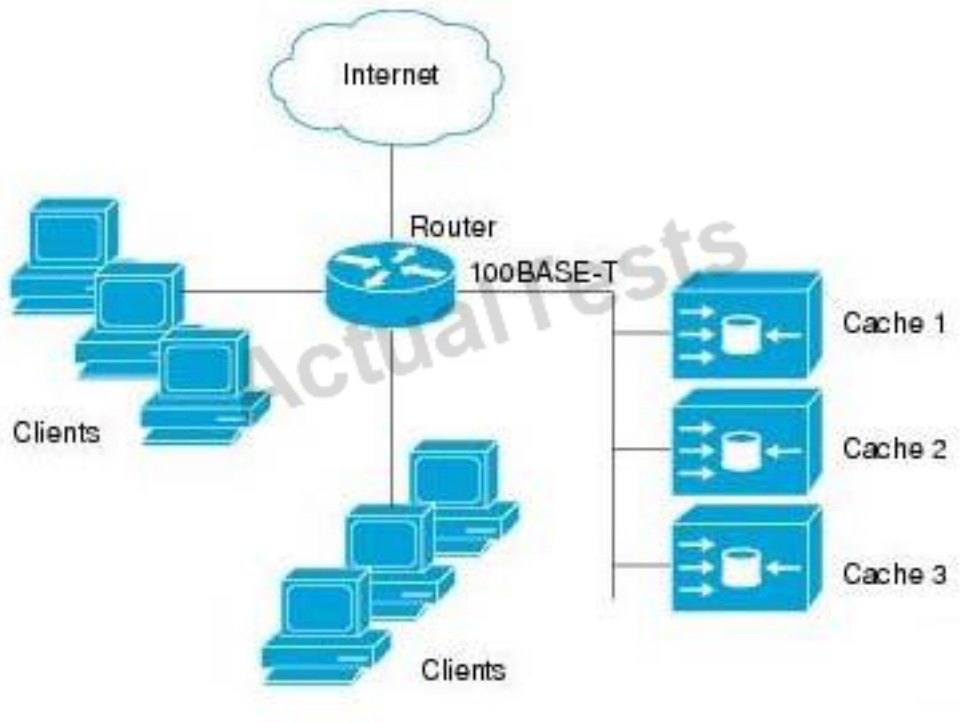
Explanation: Explanation

Web Cache Control Protocol

Description

The Web Cache Control Protocol (WCCP) feature allows you to use a Cisco Cache Engine to handle web traffic, thus reducing transmission costs and downloading time. This traffic includes user requests to view pages and graphics on World Wide Web servers, whether internal or external to your network, and the replies to those requests. shows a sample WCCP network configuration.

Figure 1 Sample Cisco Cache Engine Network Configuration



When a user (client) requests a page from a web server (located in the Internet, in this case), the router sends the request to a Cisco Cache Engine (Cache 1, Cache 2, or Cache 3). If the cache engine has a copy of the requested page in storage, the engine sends the user that page. Otherwise, the engine gets the requested page and the objects on that page from the web server, stores a copy of the page and its objects (caches them), and forwards the page and objects to the user.

WCCP transparently redirects HTTP requests from the intended server to a Cisco Cache Engine. End users do not know that the page came from the cache engine rather than the originally requested web server. The Using the Cisco Cache Engine publication contains detailed information about the Cisco Cache Engine and discusses alternative network configurations.

Reference

http://www.cisco.com/en/US/docs/ios/11_2/feature/guide/wccp.html#wp2040

QUESTION NO: 348

What is a characteristic of Network Time Protocol?

A. NTP updates are sent in the timezone they are collected in, and the syslog server will adjust based on the input time.
B. NTP authentication verifies the source, not the recipient.
C. NTP authentication requires that the recipient has multiple strata clocks to ensure accuracy.
D. Secure NTP can be configured to use SHA-1 hashing, since NTP is very insecure.
E. A stratum 0 clock should be configured at the core of every network, so it can connect to an accurate time source.

Answer: B
Reference: http://www.ine.com/resources/01700369.htm

QUESTION NO: 349

In GLBP, which router will answer on client ARP requests?

A. all active AVF routers as the first response is used by the client
B. the AVG router, replying with a different AVF MAC address each time
C. a random AVF router, based on a GLBP seed hash key
D. only the AVG router that received the ARP request first

Answer: B
Explanation:
LBP Active Virtual Gateway
Members of a GLBP group elect one gateway to be the active virtual gateway (AVG) for that group. Other group members provide backup for the AVG in the event that the AVG becomes unavailable. The AVG assigns a virtual MAC address to each member of the GLBP group. Each gateway assumes responsibility for forwarding packets sent to the virtual MAC address assigned to it by the AVG. These gateways are known as active virtual forwarders (AVFs) for their virtual MAC address.

The AVG is responsible for answering Address Resolution Protocol (ARP) requests for the virtual

IP address.

Load sharing is achieved by the AVG replying to the ARP requests with different virtual MAC addresses.

Reference
http://www.cisco.com/en/US/docs/ios/12_2t/12_2t15/feature/guide/ft_glbp.html#wp1039649

QUESTION NO: 350

What s the default stratum clock on a Cisco router, when you see the key word "master" configured on the NTP line?

A. 1
B. 2
C. 4
D. 6
E. 7

Answer: E
Explanation: Explanation
NTP master
The "ntp master" is used to configure the device as a master clock when external time synchronization is not possible; for example, the router is not connected to the Internet. If the network has ntp master configured and it cannot reach any clock with a lower stratum number, the system claims to be synchronized at the configured stratum number, and other systems synchronize to it via NTP. By default, the master clock function is disabled.
When enabled, the default stratum is 7.
In the world of NTP, stratum levels define the distance from the reference clock. A reference clock is a stratum-0 device that is assumed to be accurate and has little or no delay associated with it (typically an atomic clock).
A server that is directly connected to a stratum-0 device is called a stratum-1 server, a server that is directly connected to a stratum-1 is called a stratum-2 server and so on.

Reference
http://www.cisco.com/en/US/products/hw/switches/ps1893/
products_command_reference_chapter09186a008007dec6.html

QUESTION NO: 351

Which protocol should be used in order to configure first hop redundancy between a Cisco router and a router from another company?

A. HSRP
B. VRRP
C. GLBP
D. IRDP

Answer: B
Reference: http://www.cisco.com/en/US/docs/ios-xml/ios/ipapp_fhrp/configuration/xe-3s/asr1000/fhp-vrrs.html

QUESTION NO: 352

Which two of these statements about WCCP version 2 are false? (Choose two.)

A. It allows for the redirection of traffic other than HTTP, including a variety of UDP and TCP traffic.
B. Only one router can redirect content requests.
C. Multiple routers can redirect content requests.
D. It works only with IP networks.
E. The Cache Engine defines one central "home router" and stores it in its memory.

Answer: B,E
Explanation:
WCCP transparently redirects Hypertext Transfer Protocol (HTTP) requests going to the intended server to a Cache Engine. End users do not know that the page came from the Cache Engine rather than the originally requested web server.
WCCP Version 2 now contains the following new features:

- Multiple router support
- Improved security
- Faster throughput

Load distributing applications capability

•

Client IP addressing transparency

Multirouter Support:

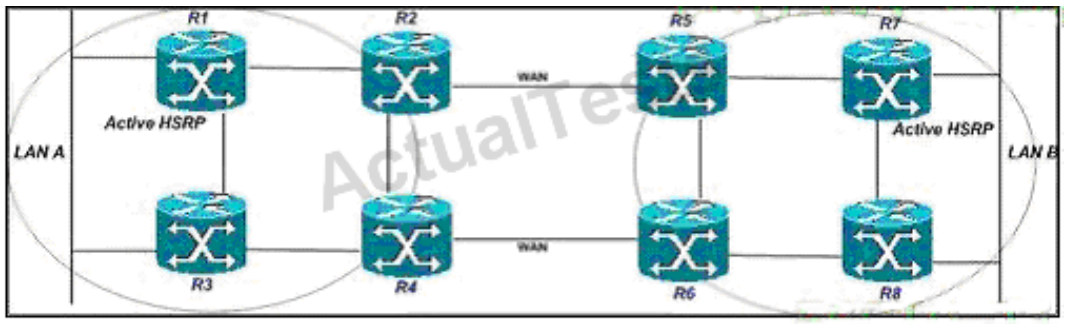
WCCP Version 2 enables a series of Cache Engines, called a *Cache Engine cluster*, to connect to multiple routers. This feature provides redundancy and a more distributed architecture for instances when a Cache Engine needs to connect to a large number of interfaces. This strategy also has the benefit of keeping all the Cache Engines in a single cluster, avoiding unnecessary duplication of web pages across several clusters.

Reference:

http://www.cisco.com/en/US/products/sw/conntsw/ps547/products_user_guide_chapter09186a008009f1ae.html

QUESTION NO: 353

Refer to the exhibit. There are two sites connected across WAN links. All intersite and intrasite links always have the same routing metric. The network administrator sees only the top routers and links being used by hosts at both LAN A and LAN B. What would be two suggestions to load-balance the traffic across both WAN links? Choose two.



A. Make HSRP track interfaces between the edge and core routers.
B. Replace HSRP with GLBP.
C. Add crossed intrasite links: R1-R4, R2-R3, R5-R8, and R6-R7.
D. Make R3 and R8 have lower HSRP priority than R1 and R7.
E. Replace HSRP with VRRP.

Answer: B,C

Explanation:

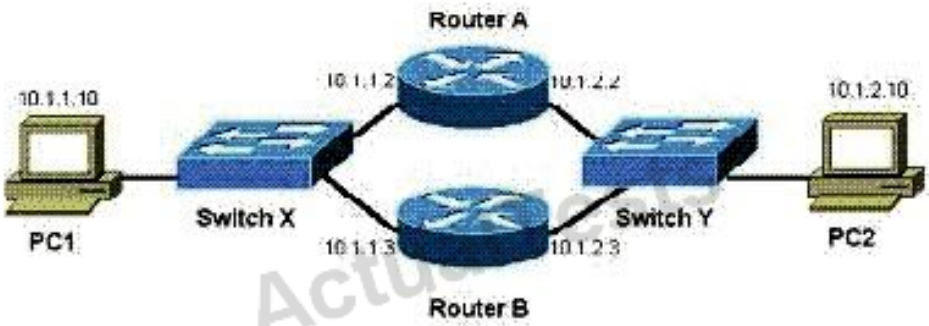
The administrator sees only the top routers (R1,R2,R5 & R7) and links being used by hosts at both LAN A and LAN B because R1 & R7 are currently active HSRP routers (notice that all the data will need to go through these routers). Next, all intersite and intrasite links have the same routing metric so these active routers will send packets to R2 or R5, not R3, R4, R6 or R8 because of the

lower metric of the top routers. For example, hosts in LAN A want to send data to hosts in LAN B, they will send data to R1 -> R2 -> R5 -> R7, which has lower metric than the path R1 -> R3 -> R4 -> R6 -> R5 (or R8) -> R7.
To make the network better, we should add crossed intrasite links so that R1 & R7 can send data to both R2/R4 & R5/R6 as they have the same routing metric now -> C is correct.
Cisco Gateway Load Balancing Protocol (GLBP) differs from Cisco Hot Standby Redundancy Protocol (HSRP) and IETF RFC 3768 Virtual Router Redundancy Protocol (VRRP) in that it has the ability to load balance over multiple gateways. Like HSRP and VRRP an election occurs, but rather than a single active router winning the election, GLBP elects an Active Virtual Gateway (AVG) to assign virtual MAC addresses to each of the other GLBP routers and to assign each network host to one of the GLBP routers -> B is correct.

Note: The routers that receive this MAC address assignment are known as Active Virtual Forwarders (AVF).

QUESTION NO: 354

Refer to the exhibit. What type of issue does this error log indicate if the IP address in the error log is located off of the Router A WAN?




```
Oct 12 13:15:41: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
                 on Vlan25, sourced by 0000.0c07.ac19
Oct 13 16:25:41: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
                 on Vlan25, sourced by 0000.0c07.ac19
Oct 15 22:31:03: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
                 on Vlan25, sourced by 0000.0c07.ac19
Oct 15 22:41:01: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
                 on Vlan25, sourced by 0000.0c07.ac19
```

A. HSRP standby configuration error
B. HSRP burned-in address error
C. HSRP secondary address configuration error
D. this is not an HSRP problem, but rather an STP error or router or switch configuration issue

Answer: D

Explanation:

As you can see in the exhibit, that this is an STP error rather than HSRP problem. It can be router or switch configuration issue.

QUESTION NO: 355

Refer to the exhibit.

```
Router#show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Et0/0       25   150   Active  local           10.1.1.101      10.1.1.1
Et0/0       100  100 P Active  local           10.1.1.101      10.1.1.2
```

ActualTests

If the router that produced the output in the exhibit fails, which virtual IP address will become active when the router recovers?

A. 10.1.1.1
B. 10.1.1.2
C. 10.1.1.101
D. All groups

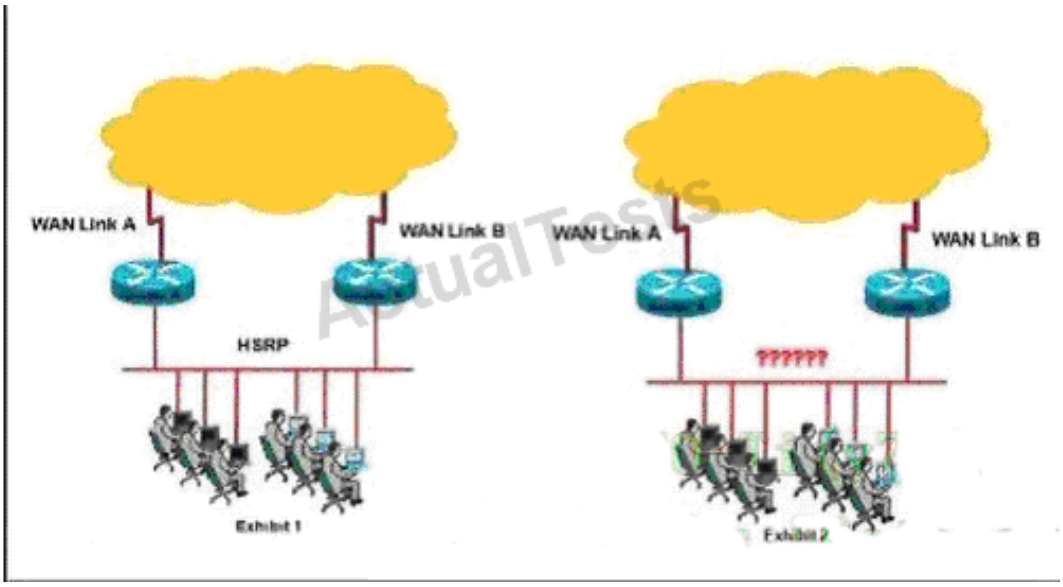
Answer: B

Explanation:

10.1.1.2 will become active when router recovers because preempt is configured on 10.1.1.2

QUESTION NO: 356

Refer to exhibits 1 and 2. In exhibit 1, all users on the LAN segment use router A as the active HSRP router. Router B is the standby router for the HSRP. In exhibit 2, the network management team reported that there is no utilization on the WAN link B. To solve this problem, you decide to change the logical topology of your LAN, but you are not sure about what changes must be made. You must manage HSRP or change it to another protocol in order to provide the most scalable design, automatic redundancy, and load balancing. Which one of these actions would be the best choice?



A. Use MHSRP, with three users using router A as the default gateway and three users using router B as the default gateway.
B. Keep HSRP and activate PBR to redirect half of the traffic to the other WAN link.
C. Use the backup interface on the WAN link B to provide load balancing for all users.
D. Use GLBP instead, because it provides you with up to three MAC addresses for the same default gateway virtual IP address.
E. Use GLBP instead, because it provides you with up to four MAC addresses for the same default gateway virtual IP address.

Answer: E

Explanation:

In HSRP and VRRP, only the primary router is used to forward traffic, others routers must wait for the primary one down before they are used. Also, the bandwidth of the standby (and other) routers are not used and wasted. With GLBP, up to four gateways can be used simultaneously. There is still one virtual IP address in a group, but GLBP can automatically select which router in the group to forward traffic by sending the virtual MAC address of a selected router to that host.

QUESTION NO: 357

Refer to the exhibit.

```
Router#show vrrp
Ethernet0/0 - Group 30
  State is Backup
  Virtual IP address is 10.1.1.30
  Virtual MAC address is 0000.5e00.011e
  Advertisement interval is 1.000 sec
  Preemption enabled
  Priority is 45  (cfgd 100)
    Track object 1 state Down decrement 55
    Track object 2 state Up decrement 55
  Master Router is 10.1.1.101, priority is 50
  Master Advertisement interval is 1.000 sec
  Master Down interval is 3.414 sec (expires in 2.982 sec)

Router#show run | inc track
track 1 interface Serial1/0 line-protocol
track 2 interface Serial1/1 line-protocol
```

ActualTests

Which action would make the router the active VRRP router?

A. Recover interface Serial 1/0.
B. Increase priority in the configuration to 100.
C. Change the interface tracking priority to 100.
D. Recover interface Serial 1/1.

Answer: A

Explanation: Explanation
As VRRP Group 30 is configured with preemption all that is required is that the VRRP Priority be higher than that of the current active VRRP router and the current master router priority is 50.

Topic 8, Implement Quality of Service (QoS)

QUESTION NO: 358

A branch router is configured with an egress QoS policy that was designed for a total number of 10 concurrent VOIP calls.

Due to expansion, 15 VOIP calls are now running over the link, but after the 14th call was established, all calls were affected and the voice quality was dramatically degraded.

Assuming that there is enough bandwidth on the link for all of this traffic, which part of the QoS configuration should be updated due to the new traffic profile?

A. Increase the shaping rate for the priority queue.
B. Remove the policer applied on the priority queue.
C. Remove the shaper applied on the priority queue.
D. Increase the policing rate for the priority queue.

Answer: D
Explanation:
The question works on the premise that there was no congestion on the link upto the 13th call. When you please the 14th call there is congestion on the link.

When there is NO congestion the priority command is allowed to take as much bandwidth as required. When there is congestion on the link the Priority command has to only use the configured bandwidth.

Adding the 14th call caused congestion, which in turn made the priority command restrict the calls to the configured value of 10 hence affect the quality of all calls.

QUESTION NO: 359

Voice quality is bad due to high delay and jitter on a link. Which two actions will improve the quality of voice calls? (Choose two.)

A. Increase the queue size of the voice class.
B. Guarantee bandwidth during congestion to the voice class with a bandwidth command.
C. Increase the tx-ring of the egress interface.
D. Implement LLQ for the voice class.
E. Decrease the rx-ring of the egress interface.
F. Decrease the queue size of the voice class.

Answer: D,F
Reference:
http://www.cisco.com/en/US/docs/ios/voice/monitor/configuration/guide/vt_qos_voip_tbsh.pdf

QUESTION NO: 360

Refer to the exhibit.

```
access-list 101 permit udp any any range 16384 32767
!
Class-map match-any  VOICE
Match access-group  101
Match ip dscp ef
```

On what will the config class-map VOICE match?

A. only on UDP traffic between port ranges 16384 and 32767
B. only on DSCP EF traffic
C. on UDP traffic between port ranges 16384 and 32767, and on DSCP EF traffic
D. only on EF traffic that is UDP and within the UDP range of 16384 and 32767

Answer: C
Reference:
http://www.cisco.com/en/US/docs/ios/solutions_docs/qos_solutions/QoSVoIP/QoSVoIP.html

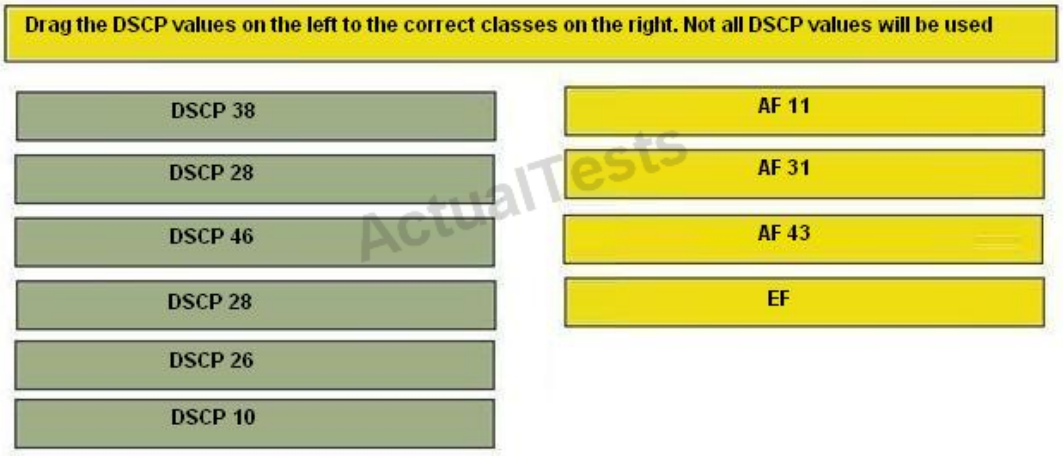
QUESTION NO: 361

Which two statements are true about bandwidth guarantee? (Choose two)

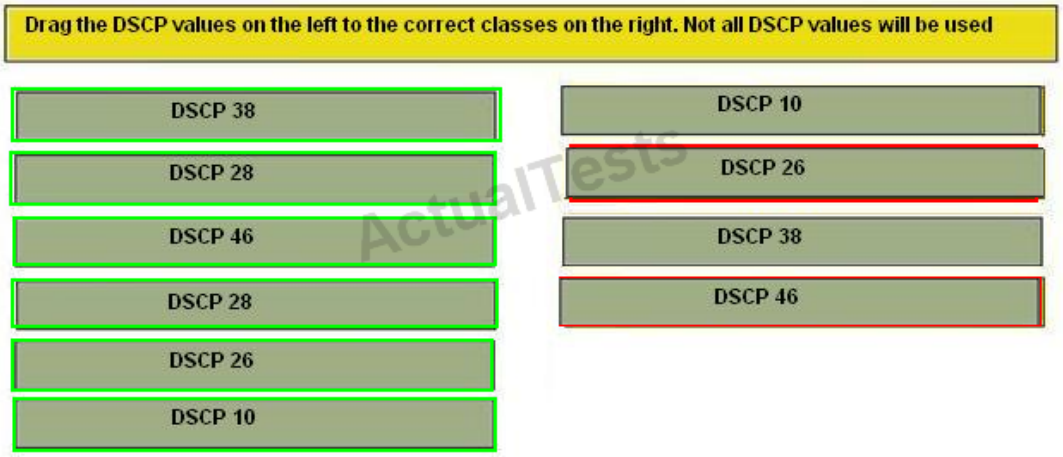
A. When congestion isn't present, the bandwidth command doesn't allow exceeding the allocated rate.
B. When congestion is present, the bandwidth command allows exceeding the allocated rate
C. When congestion is present, the bandwidth command doesn't allow exceeding the allocated rate
D. When congestion isn't present, the bandwidth command allows exceeding the allocated rate

Answer: C,D
Explanation: Reference
http://www.cisco.com/en/US/tech/tk543/tk757/technologies_tech_note09186a0080103eae.shtml

QUESTION NO: 362 DRAG DROP



Answer:



Explanation:

DSCP 38 – AF43

DSCP 28 –

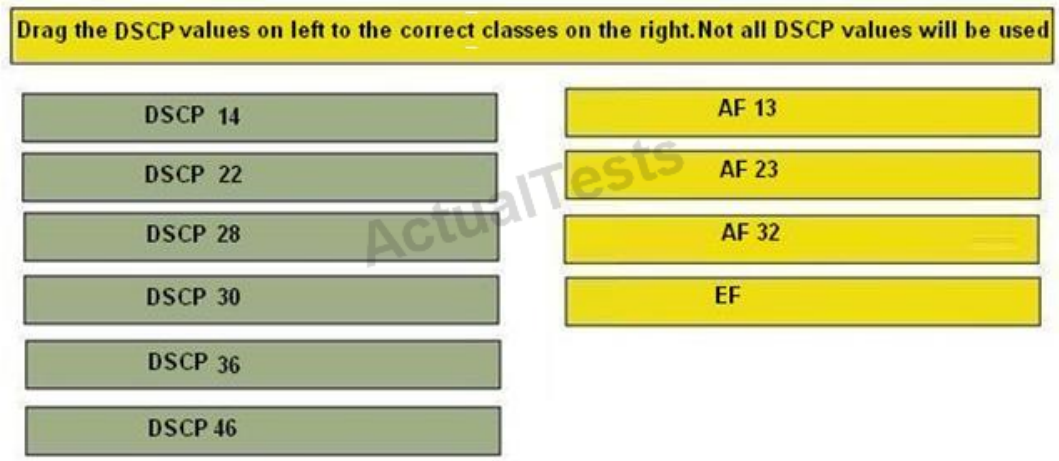
DSCP 46 – EF

DSCP 18 –

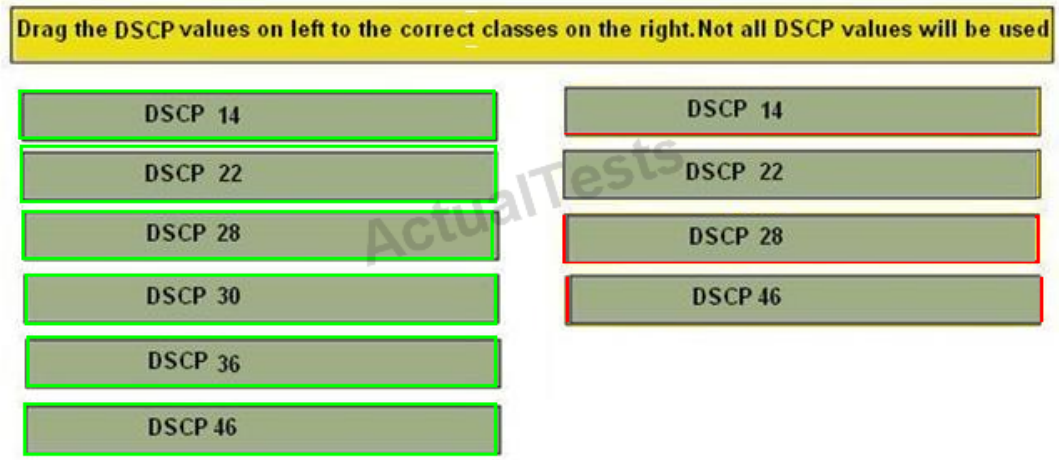
DSCP 26 – AF 31

DSCP 10 – AF11

QUESTION NO: 363 DRAG DROP



Answer:



Explanation:

AF13 – DSCP 14

AF23 – DSCP 22

AF 32 – DSCP 28

EF – DSCP 46

Reference:
http://www.cisco.com/en/US/tech/tk543/tk757/technologies_tech_note09186a00800949f2.shtml

QUESTION NO: 364

Refer to the exhibit. Based on this configuration, what type of marker is achieved?

```
policy-map QoSPolicer
 class PolicedTraffic
  police cir 512000 pir 1024000
    conform-action transmit
    exceed-action set-dscp-transmit cs1
    violate-action drop
```

ActualTests

A. Single-rate, two-color marker V
B. Three-rate, two-color marker
C. Two-rate, three-color marker
D. Single-rate, three-color marker

Answer: C
Explanation:
Networks police traffic by limiting the input or output transmission rate of a class of traffic based on userdefined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or class of service (CoS).
The Two-Rate Policer performs the following functions:
Limits the input or output transmission rate of a class of traffic based on user-defined criteria.
Marks packets by setting the IP precedence value, IP differentiated services code point (DSCP) value, Multiprotocol Label Switching (MPLS) experimental value, Quality of Service (QoS) group, ATM Cell Loss Priority (CLP) bit, and the Frame Relay Discard Eligibility (DE) bit.
With the Two-Rate Policer, you can enforce traffic policing according to two separate rates—committed information rate (CIR) and peak information rate (PIR). You can specify the use of these two rates, along with their corresponding values, by using two keywords, cir and pir, of the police command. For more information about the police command, see the "Command Reference" section of this document. The Two-Rate Policer manages the maximum rate of traffic through a token bucket algorithm. The token bucket algorithm can use the user-configured values to determine the maximum rate of traffic allowed on an interface at a given moment in time. The token bucket algorithm is affected by all traffic entering or leaving the interface (depending on the location of the interface on which the Two-Rate Policer is configured) and is useful in managing network bandwidth in cases where several large packets are sent in the same traffic stream.

Configuration Tasks
See the following sections for configuration tasks for the Two-Rate Policer feature. Each task in the list is identified as either required or optional.
Configuring the Two-Rate Policer (required)
Verifying the Two-Rate Policer Configuration (optional)
Configuring the Two-Rate Policer
The Two-Rate Policer is configured in the service policy. To configure the Two-Rate Policer, use

the following command in policy-map class configuration mode:

Command	Purpose
`Router(config-pmap-c)# police {cir cir} [bc conform-burst] {pir pir} [be peak-burst]`	Specifies that both the CIR and the PIR are to be used for two-rate traffic policing. The **bc** and **be** keywords and their associated arguments (*conform-burst* and *peak-burst*, respectively) are optional.

Although not required for configuring the Two-Rate Policer, the command syntax of the police command also allows you to specify the action to be taken on a packet when you enable an optional action argument. The resulting action corresponding to the keyword choices are listed in Table 1.

Table 1 police Command Action Keywords

Keyword	Resulting Action
drop	Drops the packet.
set-clp-transmit	Sets the ATM CLP bit from 0 to 1 on the ATM cell and sends the packet with the ATM CLP bit set to 1.
set-dscp-transmit *new-dscp*	Sets the IP DSCP value and sends the packet with the new IP DSCP value setting.
set-frde-transmit	Sets the Frame Relay DE bit from 0 to 1 on the Frame Relay frame and sends the packet with the DE bit set to 1.
set-mpls-exp-transmit	Sets the MPLS experimental bits from 0 to 7 and sends the packet with the new MPLS experimental bit value setting.
set-prec-transmit *new-prec*	Sets the IP precedence and sends the packet with the new IP precedence value setting.
set-qos-transmit *new-qos*	Sets the QoS group value and sends the packet with the new QoS group value setting.
transmit	Sends the packet with no alteration.

Related Documents
Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.2
Cisco IOS Quality of Service Solutions Command Reference, Release 12.2
RFC 2698, A Two Rate Three Color Marker

The two-rate policer is often configured on interfaces at the edge of a network to limit the rate of traffic entering or leaving the network. In addition to rate-limiting traffic, the policer's three-color marker can mark packets according to whether the packet conforms (green), exceeds (yellow), or

violates (red) a specified rate. You decide the actions you want the router to take for conforming, exceeding, and violating traffic. For example, you can configure conforming packets to be sent, exceeding packets to be sent with a decreased priority, and violating packets to be dropped. In most common configurations, traffic that conforms is sent and traffic that exceeds is sent with decreased priority or is dropped. You can change these actions according to your network needs. With packet marking, you can partition your network into multiple priority levels or classes of service (CoS). For example, you can configure the two-rate three-color marker to do the following: Assign packets to a QoS group, which the router then uses to determine how to prioritize packets within the router. Set the IP precedence level, IP DSCP value, or the MPLS experimental value of packets entering the network. Networking devices within your network can then use this setting to determine how to treat the traffic. For example, a weighted random early detection (WRED) drop policy can use the IP precedence value to determine the drop probability of a packet.

Set the ATM cell loss priority (CLP) bit in ATM cells. The ATM CLP bit is used to prioritize packets in ATM networks and is set to either 0 or 1. During congestion, the router discards cells with a CLP bit setting of 1 before it discards cells with a CLP bit setting of 0.
The three-color marker distinguishes between the nonconforming traffic that occasionally bursts a certain number of bytes more than the CIR and violating traffic that continually violates the PIR allowance. Applications can utilize the three-color marker to provide three service levels: guaranteed, best effort, and deny. The threecolor marker is useful in marking packets in a packet stream with different, decreasing levels of assurances (either absolute or relative). For example, a service might discard all red packets because they exceed both the committed and excess burst sizes, forward yellow packets as best effort, and forward green packets with a low drop probability.

Reference
http://www.cisco.com/en/US/docs/ios/12_2t/12_2t4/feature/guide/ft2RTplc.html#wp1023708
http://www.cisco.com/en/US/docs/routers/10000/10008/configuration/guides/qos/10qpolce.html#wp1041731

QUESTION NO: 365

On the basis of the definitions of different services in various RFCs, traffic with Expedited

Forwarding per-hop behavior should be marked as which of these?

A. IP ToS of 0xEF
B. IP experimental ECN
C. DSCP decimal 5
D. Binary value of 101110

Answer: D

Explanation:

The assured forwarding (AF) model is used to provide priority values to different data applications. 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. DSCP 46 is backward compatible with an IP Precedence value of 5 as seen in the following binary pattern:

101110 = DSCP 46

The EF marking of 46 does NOT follow the drop preference rules of the assured forwarding model. Please do

NOT think that the 11 means high drop preference. The EF model is used for voice over IP media traffic (RTP) by default in most vendors phones. Cisco IP Phones mark signaling packets (SCCP or SIP) to CS3 (24), while media (RTP) is marked to EF (DSCP 46) by default. All EF traffic is normally mapped to the priority queue (PQ) on Cisco switches and routers. The priority queue guarantees three critical services:

Packet Loss

Delay

Jitter (delay variation)

The three most significant bits of 101 are only considered if IP Precedence was being used. The binary digits of 4 2 1 are used to factor the 101 binary pattern when only three digits are under consideration. The DSCP binary pattern of 101110 (46) uses six digits or binary values-32 16 8 4 2 1.

It is good to know how to convert a DSCP decimal value to an entire ToS octet (byte) values as well. The ToS byte uses all eight bits, while the DSCP is only using the leading six digits. The EF pattern discussed above will become 10111000 when considering the entire octet. Notice the two least significant zeros that were added to the 101110 binary pattern. Many network management utilities will only allow administrators to configure or display the entire ToS byte. A ping V from a Microsoft operating system requires setting the entire ToS byte. An extended ping from a Cisco router will also allow administrators to see the entire ToS byte. Sniffer Pro LAN and Wire Shark sniffers show the entire ToS field as well. IP accounting shows the entire ToS byte, while Netflow shows the ToS byte in hexadecimal format. The ToS byte value for EF is as follows:

128 64 32 16 8 4 2 1

1 0 1 1 1 0 0 0

A DSCP value of 46 results in a ToS byte value of 184. Although you can mark a ping with a ToS value of 184, the ICMP (ping) traffic will probably not be mapped to the proper application class. In the next blog, we will learn QoS models for using markings for different application classes.

QUESTION NO: 366

Refer to the exhibit.

```
police 8000 1000 1000 conform-action transmit exceed-action drop
```

Considering the following policer, which statement is valid?

A. The Tc interval equals 125 ms
B. Traffic exceeding 8 kb/s is systematically dropped
C. The policer allows an excess burst of 1 kb
D. The policer allows traffic to peak to 16 kb/s for the duration of a Tc interval if no traffic passed over the previous interval
E. Traffic exceeding 8 kb/s is never dropped

Answer: D

Explanation:

As per the exhibit, you can see the policer allows traffic to peak to 16 kb/s for the duration of a Tc interval if no traffic passed over the previous interval

QUESTION NO: 367

Refer to the Exhibit.

```
!
class-map match-all CLASS1
 match ip dscp af41
!
```

The displayed QoS configuration has been configured on a router.

IPv6 is being implemented on the router, and it is required to convert the QoS policy to support both IPv4 and IPv6 on the same class.

Which alternative configuration would allow matching DSCP AF41 for both IPv4 and IPv6 on the same class map?

A. Class-map match-all CLASS1

Match dscp af41
B. Class-map match-all CLASS1
Match ip dscp af41
Match ipv6 dscp af41
C. Class-map match-any CLASS1
Match ip dscp af41
Match ipv6 dscp af41
D. Class-map match-any CLASS1
Match qos-group af41

Answer: A

Explanation: Explanation

"match dscp" matches both IPv4 and IPv6 traffic while "match ip dscp" matches only IPv4 traffic

QUESTION NO: 368

Which QoS mechanism can help classify different kinds of Citrix traffic (print job versus real-time terminal control)?

A. qos-group
B. DSCP
C. LFI
D. NBAR

Answer: D

Explanation:

NBAR can recognize different types of Citrix applications; CB marking can use NBAR to classify based on these application types.

QUESTION NO: 369

Refer to the exhibit.

```
R4#sh run | sec VOICE
class-map match-all VOICE
 match access-group name VOICE-ACL
ip access-list extended VOICE-ACL
 permit udp any any range 16384 32787 dscp ef
```

ActualTests

On what will the class-map VOICE match?

A. only UDP traffic between port ranges 16384 and 32787
B. only DSCP EF traffic
C. both UDP traffic between port ranges 16384 and 32787 and DSCP EF traffic
D. only EF traffic that is UDP and is within the UDP range 16384-32787

Answer: D
Explanation:

See the last line of the command. Only EF traffic that is UDP is permitted. The UDP should be within the range 16384 32787.

QUESTION NO: 370

Which two statements are true about bandwidth guarantee? (Choose two.)

A. When congestion is present, the priority command doesn't allow exceeding the allocated rate
B. When congestion isn't present, the priority command doesn't allow exceeding the allocated rate
C. When congestion is present, the priority command allows exceeding the allocated rate
D. When congestion isn't present, the priority command allows exceeding the allocated rate

Answer: A,D
Explanation:
During congestion conditions, the traffic class is guaranteed bandwidth equal to the specified rate. (Recall that bandwidth guarantees are only an issue when an interface is congested.) In other words, the priority command provides a minimum bandwidth guarantee.

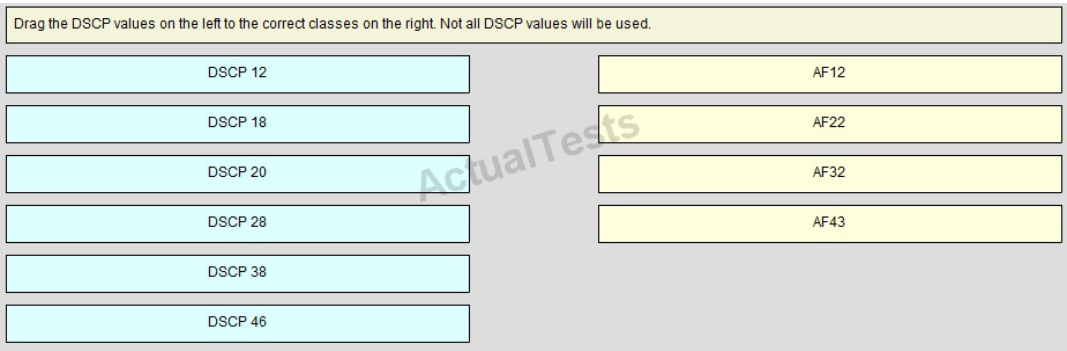
In addition, the priority command implements a maximum bandwidth guarantee. Internally, the priority queue uses a token bucket that measures the offered load and ensures that the traffic stream conforms to the configured rate. Only traffic that conforms to the token bucket is guaranteed low latency. Any excess traffic is sent if the link is not congested or is dropped if the link is congested.

"The purpose of the built-in policer is to ensure that the other queues are serviced by the queueing scheduler. In the original Cisco priority queueing feature, which uses the priority-group and priority-list commands, the scheduler always serviced the highest priority queue first. In extreme cases, the lower priority queues rarely were serviced and effectively were starved of bandwidth.

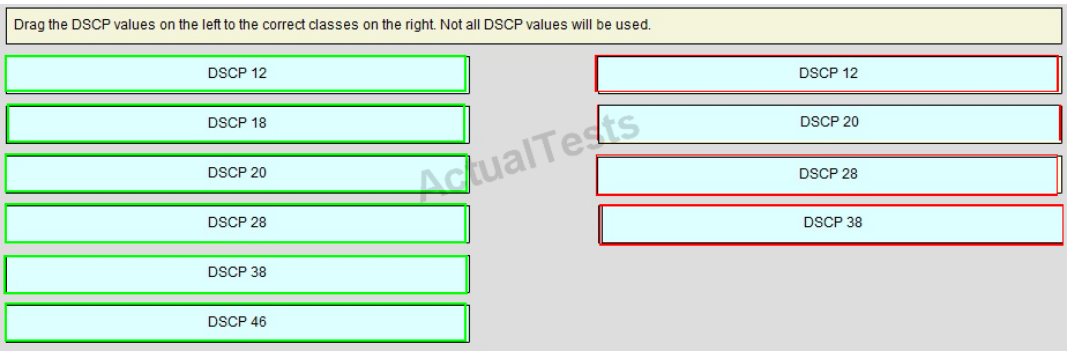
The real benefit of the priority command—and its major difference from the bandwidth

command—is how it provides a strict de-queueing priority to provide a bound on latency. Here is how the Cisco IOS Configuration Guide describes this benefit: "A strict priority queue (PQ) allows delay-sensitive data such as voice to be de-queued and sent before packets in other queues are de-queued....."

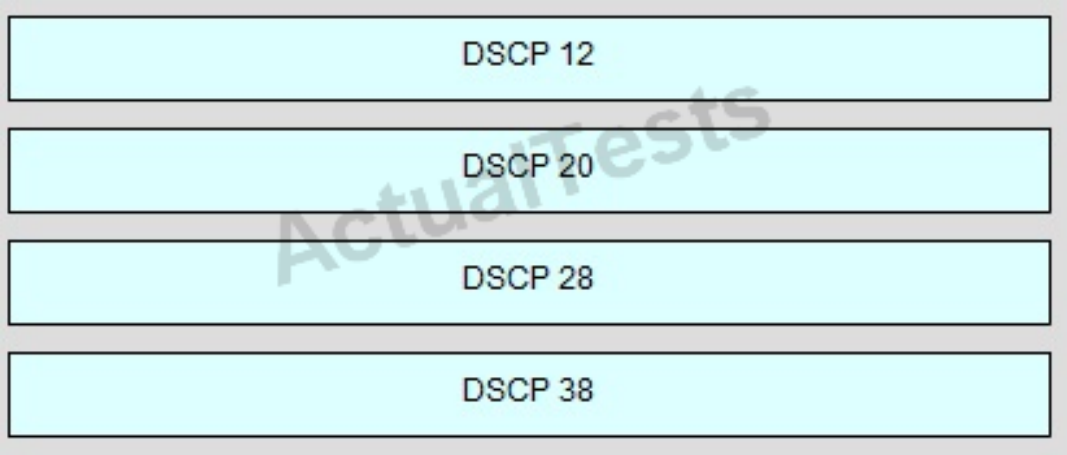
QUESTION NO: 371 DRAG DROP



Answer:



Explanation:



QUESTION NO: 372

You are the network administrator of an enterprise company that just deployed a global IP telephony environment. In order to guarantee good voice quality, you asked your provider to implement QoS on the CE routers of your MPLS network. On your LAN, you have also deployed QoS.

Users, however, keep complaining about bad voice quality, and the provider does not see matches on the DSCP values that you have asked him to match upon. What is most likely the problem?

A. The phones are not sending traffic with the correct DSCP value.
B. The Cisco CallManager is not marking the traffic correctly.
C. This is most likely a bug on the CE routers.
D. Your LAN QoS is incorrectly configured.

Answer: D
Explanation:
The most likely problem is that the LAN QoS is not configured correctly. When the DSCP values do not match, it means LAN QoS has a problem.

QUESTION NO: 373

Which three protocols or applications should be placed in a class that is configured with WRED? (Choose three.)

A. HTTP
B. RTP
C. streaming video
D. BitTorrent
E. POP3

Answer: A,D,E
Reference:
http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/qcfwred_ps1835_TSD_Products_Configuration_Guide_Chapter.html

WRED is useful with adaptive traffic such as TCP/IP. With TCP, dropped packets indicate congestion, so the packet source will reduce its transmission rate. With other protocols, packet sources may not respond or may resend dropped packets at the same rate. Thus, dropping packets does not decrease congestion.

QUESTION NO: 374

What is the command to configure RSVP to reserve up to one-tenth of a Gigabit link, but only allow each individual flow to use 1 MB/s?

A. ip rsvp bandwidth 100000 1000
B. ip rsvp bandwidth 100000 1
C. ip rsvp bandwidth 10 1
D. ip rsvp bandwidth 1000000 1000

Answer: A
Explanation:
This command enables the traffic-engineering tunnels on the interface. It configures the interface to send and receive RSVP signaling to establish traffic-engineering tunnels across this interface; both sides of the link need to have this configuration enabled.

Define the bandwidth allocation on the interfaces:
ip rsvp bandwidth interface-kbps single-flow-kbps [sub-pool kbps]

This command enables RSVP reservations for traffic-engineering tunnels.
interface-kbps is the amount of bandwidth (in kbps) on the interface that is available for reservation, and it is referred to as global pool.
single-flow-kbps is the maximum amount of bandwidth (in kbps) allowed for a single flow. This parameter is ignored for traffic-engineering tunnel reservations.
[sub-pool kbps] is the amount of bandwidth (in kbps) from the global pool available for reservations in a subpool.

ip rsvp bandwidth
To enable RSVP for IP on an interface, use the ip rsvp bandwidth interface configuration command. To disable

RSVP, use the no form of the command.
ip rsvp bandwidth [interface-kbps] [single-flow-kbps]
no ip rsvp bandwidth [interface-kbps] [single-flow-kbps]
Syntax Description interface-kbps

(Optional) Amount of bandwidth (in kbps) on interface to be reserved. The range is 1 to 10, 000, 000.

single-flow-kbps (Optional) Amount of bandwidth (in kbps) allocated to a single flow. The range is 1 to 10, 000, 000.

QUESTION NO: 375

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

A. RED randomly drops packets before the queue becomes full.
B. RED is always useful, without dependency on flow.
C. RED increases the drop rate as the average queue size increases.
D. RED has a per-flow intelligence.

Answer: A,C
Explanation:
RED aims to control the average queue size by indicating to the end hosts when they should temporarily slow down transmission of packets.
RED takes advantage of the congestion control mechanism of TCP. By randomly dropping packets prior to periods of high congestion, RED tells the packet source to decrease its transmission rate. Assuming the packet source is using TCP, it will decrease its transmission rate until all the packets reach their destination, indicating that the congestion is cleared. You can use RED as a way to cause TCP to slow down transmission of packets.

TCP not only pauses, but it also restarts quickly and adapts its transmission rate to the rate that the network can support.

RED distributes losses in time and maintains normally low queue depth while absorbing spikes. When enabled on an interface, RED begins dropping packets when congestion occurs at a rate you select during configuration.

Packet Drop Probability
The packet drop probability is based on the minimum threshold, maximum threshold, and mark probability denominator.

When the average queue depth is above the minimum threshold, RED starts dropping packets. The rate of packet drop increases linearly as the average queue size increases until the average queue size reaches the maximum threshold.

The mark probability denominator is the fraction of packets dropped when the average queue depth is at the maximum threshold. For example, if the denominator is 512, one out of every 512 packets is dropped when the average queue is at the maximum threshold.

Reference
http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/qcfconav.html

QUESTION NO: 376

Which two statements best describe CBWFQ? (Choose two.)

A. The CBWFQ scheduler provides a guaranteed minimum amount of bandwidth to each class.
B. CBWFQ services each class queue using a strict priority scheduler.
C. The class-default queue only supports WFQ.
D. Inside a class queue, processing is always FIFO, except for the class-default queue.

Answer: A,D
Explanation:
Class-based weighted fair queuing (CBWFQ) extends the standard WFQ functionality to provide support for user-defined traffic classes. For CBWFQ, you define traffic classes based on match criteria including protocols, access control lists (ACLs), and input interfaces. Packets satisfying the match criteria for a class constitute the traffic for that class. A queue is reserved for each class, and traffic belonging to a class is directed to the queue for that class.
Once a class has been defined according to its match criteria, you can assign it characteristics. To characterize a class, you assign it bandwidth, weight, and maximum packet limit. The bandwidth assigned to a class is the guaranteed bandwidth delivered to the class during congestion. Within each queue processing is FIFO, except for the class-default queue.

QUESTION NO: 377

Refer to the exhibit.

```
Class-map VIDEO
Match ip precedence 5
Class-map FTP
Match ip precedence 1

policy-map TEST
 class VIDEO
  priority 200
 class FTP
  bandwidth 1000
class class-default
 random-detect
```

ActualTests

Which two statements are correct, when the QoS configuration is applied in an outbound direction on a 10-Mb/s interface? (Choose two.)

A. When reaching 10 Mb/s of input rate, the video class will be policed to 200 kb/s.
B. The class FTP is allowed to reach more than 1 Mb/s in the event of congestion.
C. IP precedence 1 traffic is affected by a drop probability.
D. Video traffic above 200 kb/s is allowed to pass when the total interface output rate does not reach 10 Mb/s.
E. Video traffic above 200 kb/s is allowed to pass when congestion is present.

Answer: B,D
Reference: http://www.cisco.com/en/US/docs/security/asa/asa82/configuration/guide/intrface.html

QUESTION NO: 378

Refer to the exhibit. Voice traffic is marked "precedence 5." How much bandwidth is allocated for voice traffic during periods of congestion?

```
!
class-map match-all Signal
  match ip precedence 3
class-map match-any System
  match access-group name Security
  match ip precedence 6
  match ip precedence 7
class-map match-all Bearer
  match ip precedence 5
!
!
policy-map ProviderOut
  class Bearer
   priority 48
  class Signal
   bandwidth 15
  class System
   bandwidth 15
  class class-default
   fair-queue
   random-detect
   shape average 512000
!
interface Ethernet0/1
 description Provider Interface
 ip address dhcp client-id Ethernet0/1
 ip access-group 111 in
 ip nat outside
 full-duplex
 no cdp enable
 service-policy output ProviderOut
!
```

ActualTests

A. a minimum of 48 kb/s
B. a maximum of 48 kb/s
C. a minimum of 48% of the available bandwidth
D. a maximum of 48% of the available bandwidth

Answer: B

Explanation:

Class-Based Shaping Configuration Task List

To configure Class-Based Shaping, perform the tasks described in the following sections. The task in the first section is required; the tasks in the remaining sections are optional.

• Configuring Class-Based Shaping (Required)
• Configuring CBWFQ Inside Generic Traffic Shaping (Optional)
• Verifying the Configuration of Policy Maps and Their Classes (Optional)

Configuring Class-Based Shaping

To configure Class-Based Shaping, use the first two commands in global configuration mode to specify the name of the policy map and the name of the class map. To specify average or peak rate, use the remaining commands in class-map configuration mode:

	Command	Purpose
Step 1	Router(config)# **policy-map** *policy-map*	Specifies the name of the policy map to be created or modified.
Step 2	Router(config)# **class-map** *class-map-name*	Specifies the name of the class map to be created.
Step 3	Router(config-pmap-c)# **shape** {**average** | **peak**} *cir* [*bc*] [*be*]	Specifies average or peak rate shaping.
Step 4	Router(config-pmap-c)# **shape max-buffers***number-of-buffers*	(Optional) Specifies the maximum number of buffers allowed on shaping queues.

Configuring CBWFQ

Inside Generic Traffic Shaping

To configure class-based weighted fair queueing (CBWFQ) inside GTS, use the first two commands in global configuration mode to specify the name of the policy map and the name of the class map. To specify average or peak rate and to attach the service policy to the class, use the remaining commands in class-map configuration mode:

	Command	Purpose
Step 1	Router(config)# **policy-map** *policy-map*	Specifies the name of the policy map to be created or modified.
Step 2	Router(config)# **class-map** *class-map-name*	Specifies the name of the class map to be created.
Step 3	Router(config-pmap-c)# **shape** {**average** | **peak**} *cir* [*bc*] [*be*]	Specifies average or peak rate shaping.
Step 4	Router(config-pmap-c)# **service-policy** *policy-map*	Attaches the service policy with CBWFQ to the class.

Verifying the Configuration of Policy Maps and Their Classes to display the contents of a specific policy map, a specific class from a specific policy map, or all policy maps configured on an interface, use the following commands in EXEC mode, as needed:

Command	Purpose
Router# **show policy** *policy-map*	Displays the configuration of all classes comprising the specified policy map.
Router# **show policy** *policy-map* **class** *class-name*	Displays the configuration of the specified class of the specified policy map.
Router# **show policy interface** *interface-name*	Displays the configuration of all classes configured for all policy maps on the specified interface.

http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/qcfcbshp.html

The bandwidth and priority commands both define actions that can be applied within a modular quality of service command-line interface (MQC) policy-map, which you apply to an interface, subinterface or virtual circuit (VC) via the service-policy command. Specifically, these commands provide a bandwidth guarantee to the packets which match the criteria of a traffic class. However, the two commands have important functional differences in those guarantees
Summary of Differences between bandwidth and priority commands this table lists the functional differences between the bandwidth and priority commands:

Function	bandwidth Command	priority Command
Minimum bandwidth guarantee	Yes	Yes
Maximum bandwidth guarantee	No	Yes
Built-in policer	No	Yes
Provides low latency	No	Yes

ActualTests

In addition, the bandwidth and priority commands are designed to meet different quality of service (QoS) policy objectives. This table lists those differing objectives:

Application	bandwidth Command	priority Command
Bandwidth management for WAN links	Yes	Somewhat
Manage delay and variations in delay (jitter)	No	Yes
Improve application response time	No	Yes

ActualTests

Reference
http://www.cisco.com/en/US/tech/tk543/tk757/technologies_tech_note09186a0080103eae.shtml#configuringthepriritycommand

QUESTION NO: 379

Refer to the exhibit. Which of these is applied to the Bearer class?

```
!
class-map match-all Signal
  match ip precedence 3
class-map match-any System
  match access-group name Security
  match ip precedence 6
  match ip precedence 7
class-map match-all Bearer
  match ip precedence 5
!
!
policy-map ProviderOut
  class Bearer
   priority 48
  class Signal
   bandwidth 15
  class System
   bandwidth 15
  class class-default
   fair-queue
   random-detect
   shape average 512000
!
interface Ethernet0/1
 description Provider Interface
 ip address dhcp client-id Ethernet0/1
 ip access-group 111 in
 ip nat outside
 full-duplex
 no cdp enable
 service-policy output ProviderOut
!
```

ActualTests

A. wred
B. traffic shaping
C. packet marking
D. packet classification
E. FIFO queuing within the class

Answer: E

Explanation:

Within a class queue, processing is always FIFO, except for the class-default queue.

CBWFQ supports 64 queues, with a maximum and default queue length varying depending on the model of router and the amount of memory installed. All 64 queues can be configured, but one class
queue, called class-default, is automatically configured.

If the explicitly configured classification does not match a packet, IOS places the packet into the class-default class.

Currently, CBWFQ can use either FIFO or WFQ inside the class-default queue. With Flow-Based WFQ in the class-default queue, when CBWFQ decides to take one or more packets from the queue, it takes the packet with the best sequence number (SN) — just like WFQ normally does.

Cisco 7500 series routers support either FIFO or WFQ inside each and every CBWFQ queue, whereas other platforms only support both FIFO and WFQ inside CBWFQ's class-default queue.

If the default class is allocated a bandwidth, WFQ cannot be enabled for the traffic within the default class. This is true for all platforms except the 7500(and soon the 7200).

Currently (except for the Cisco 7500 router platform) all traffic classes except for the default traffic class support only FIFO queuing within the class.

On all platforms, the default traffic class can support either FIFO or WFQ within the class. But if the default traffic class is allocated a minimum bandwidth as shown in the figure, WFQ will not be supported in the default traffic class. The only current exception is for the Cisco 7500 series platforms. In this case, the default traffic class will support only FIFO queuing.

QUESTION NO: 380

Refer to the exhibit. What is the overall type of queuing being used on the outgoing data for interface Ethernet0/1?

```
!
class-map match-all Signal
  match ip precedence 3
class-map match-any System
  match access-group name Security
  match ip precedence 6
  match ip precedence 7
class-map match-all Bearer
  match ip precedence 5
!
!
policy-map ProviderOut
  class Bearer
   priority 48
  class Signal
   bandwidth 15
  class System
   bandwidth 15
  class class-default
   fair-queue
   random-detect
   shape average 512000
!
interface Ethernet0/1
 description Provider Interface
 ip address dhcp client-id Ethernet0/1
 ip access-group 111 in
 ip nat outside
 full-duplex
 no cdp enable
 service-policy output ProviderOut
!
```

ActualTests

A. LLQ
B. FIFO
C. CBWFQ
D. priority queuing
E. weighted fair queuing

Answer: A

Explanation: Explanation

The above exhibit is an example of Class-Based Weighted Fair Queueing (CBWFQ).

After the weight for a packet is assigned, the packet is enqueued in the appropriate class queue. CBWFQ uses the weights assigned to the queued packets to ensure that the class queue is serviced fairly.

Configuring a class policy—thus, configuring CBWFQ—entails these three processes:
Defining traffic classes to specify the classification policy (class maps). This process determines how many types of packets are to be differentiated from one another.
Associating policies—that is, class characteristics—with each traffic class (policy maps). This process entails configuration of policies to be applied to packets belonging to one of the classes previously defined through a class map. For this process, you configure a policy map that specifies the policy for each traffic class.

Attaching policies to interfaces (service policies). This process requires that you associate an existing policy map, or service policy, with an interface to apply the particular set of policies for the map to that interface.

Reference
http://www.cisco.com/en/US/docs/ios/12_0t/12_0t5/feature/guide/cbwfq.html#wp17641

Low Latency Queuing (LLQ) is a feature developed by Cisco to bring strict priority queuing (PQ) to Class-Based Weighted Fair Queuing (CBWFQ). LLQ allows delay-sensitive data (such as voice) to be given preferential treatment over other traffic by letting the data to be dequeued and sent first
Low Latency Queueing Configuration Task List

To configure LLQ, perform the tasks described in the following sections. The task in the first section is required; the tasks in the remaining sections are optional.
• Configuring LLQ (Required)
• Configuring the Bandwidth Limiting Factor (Optional)
• Verifying LLQ (Optional)
• Monitoring and Maintaining LLQ (Optional)
See the end of this chapter for the section "LLQ Configuration Examples."

Configuring LLQ
To give priority to a class within a policy map, use the following command in policy-map class configuration mode:
Configuring the Bandwidth Limiting Factor
To change the maximum reserved bandwidth allocated for CBWFQ, LLQ, and IP RTP Priority, use the following command in interface configuration mode:

Verifying LLQ
To display the contents of the priority queue, such as queue depth and the first packet queued, use the following command in EXEC mode:
The priority queue is the queue whose conversation ID is equal to the number of dynamic queues plus 8. The packets in the priority queue have a weight of 0.

Monitoring and Maintaining LLQ
To tune your RTP bandwidth or decrease RTP traffic if the priority queue is experiencing drops, use the following commands in EXEC mode, as needed:

LLQ
The Low Latency Queueing feature brings strict priority queueing to Class-Based Weighted Fair Queueing (CBWFQ). Strict priority queueing allows delay-sensitive data such as voice to be dequeued and sent first (before packets in other queues are dequeued), giving delay-sensitive data preferential treatment over other traffic.
Without Low Latency Queueing, CBWFQ provides weighted fair queueing based on defined classes with no strict priority queue available for real-time traffic. CBWFQ allows you to define traffic classes and then assign characteristics to that class. For example, you can designate the minimum bandwidth delivered to the class during congestion.
For CBWFQ, the weight for a packet belonging to a specific class is derived from the bandwidth you assigned to the class when you configured it. Therefore, the bandwidth assigned to the packets of a class determines the order in which packets are sent. All packets are serviced fairly based on weight; no class of packets may be granted strict priority. This scheme poses problems for voice traffic that is largely intolerant of delay, especially variation in delay. For voice traffic, variations in delay introduce irregularities of transmission manifesting as jitter in the heard conversation.

The Low Latency Queueing feature provides strict priority queueing for CBWFQ, reducing jitter in voice conversations. Configured by the priority command, Low Latency Queueing enables use of a single, strict priority queue within CBWFQ at the class level, allowing you to direct traffic belonging to a class to the CBWFQ strict priority queue. To enqueue class traffic to the strict priority queue, you configure the priority command for the class after you specify the named class within a policy map. (Classes to which the priority command is applied are considered priority classes.) Within a policy map, you can give one or more classes priority status.
When multiple classes within a single policy map are configured as priority classes, all traffic from these classes is enqueued to the same, single, strict priority queue. One of the ways in which the strict priority queueing used within CBWFQ differs from its use outside CBWFQ is
in the parameters it takes. Outside CBWFQ, by using the ip rtp priority command, you specify the range of UDP ports whose voice traffic flows are to be given priority service. Using the priority command, because you can configure the priority status for a class within CBWFQ, you are no longer limited to a UDP port number to stipulate priority flows. Instead, all of the valid match criteria used to specify traffic for a class now applies to priority traffic. These methods of specifying traffic for a class include matching on access lists, protocols, and input interfaces. Moreover, within an access list you can specify that traffic matches are allowed based on the
IP Differentiated Services Code Point (DSCP) value that is set using the first six bits of the Type of Service (ToS) byte in the IP header. Although it is possible to enqueue various types of real-time

traffic to the strict priority queue, we strongly recommend that you direct only voice traffic to it. This recommendation is made because voice traffic is wellbehaved, whereas other types of real-time traffic are not. Moreover, voice traffic requires that delay be
nonvariable in order to avoid jitter. Real-time traffic such as video could introduce variation in delay, thereby thwarting the steadiness of delay required for successful voice traffic transmission. Configuration Tasks See the following sections for configuration tasks for the Low Latency Queueing feature. Each task in the list indicates if the task is optional or required.

Configuring Low Latency Queueing (Required)
Verifying Low Latency Queueing (Optional)
Configuring Low Latency Queueing

Command	Purpose
`Router(config-pmap-c)# priority bandwidth`	Reserves a strict priority queue for CBWFQ traffic.

To give priority to a class within a policy map, use the following command in policy-map class configuration mode:

Verifying Low Latency Queueing
To see the contents of the priority queue (such as queue depth and the first packet queued), use the following command in EXEC mode:

Command	Purpose
`Router# show queue interface-type interface-number`	Lists fair queueing configuration and statistics for a particular interface.

The priority queue is the queue whose conversation ID is equal to the number of dynamic queues plus 8. The packets in the priority queue have a weight of 0.

Reference
http://www.cisco.com/en/US/docs/ios/12_0t/12_0t7/feature/guide/pqcbwfq.html#wp5329
http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/
qcfwfq_ps1835_TSD_Products_Configuration_Guide_Chapter.html#wp1001719

QUESTION NO: 381

Which two of these are differences between traffic policing and traffic shaping? (Choose two.)

A. with traffic shaping, a router stores excess traffic in packet buffers until bandwidth is available again
B. with policing you can tune the buffer usage for traffic exceeding the specified CIR
C. with shaping you can tune the buffer usage for traffic exceeding the specified CIR
D. shaping should only be applied for ingress traffic, policing only for egress
E. policing uses a token bucket algorithm, shaping uses an SPD algorithm

Answer: A,C
Explanation:
Traffic policing propagates bursts. When the traffic rate reaches the configured maximum rate, excess traffic is dropped (or remarked). The result is an output rate that appears as a saw-tooth with crests and troughs. In contrast to policing, traffic shaping retains excess packets in a queue and then schedules the excess for later transmission over increments of time. The result of traffic shaping is a smoothed packet output rate. Shaping implies the existence of a queue and of sufficient memory to buffer delayed packets, while policing does not. Queueing is an outbound concept; packets going out an interface get queued and can be shaped. Only policing can be applied to inbound traffic on an interface. Ensure that you have sufficient memory when enabling shaping. In addition, shaping requires a scheduling function for later transmission of any delayed packets. This scheduling function allows you to organize the shaping queue into different queues. Note: in particular that the term CIR refers to the traffic rate for a VC based on a business contract, and shaping rate refers to the rate configured for a shaper on a router.

QUESTION NO: 382

Which of these is a valid differentiated services PHB?

A. Guaranteed PHB
B. Class-Selector PHB
C. Reserved Forwarding PHB
D. Discard Eligible PHB
E. Priority PHB

Answer: B

Explanation:

Differentiated Services Definition

Differentiated Services is a multiple service model that can satisfy differing QoS requirements. With Differentiated Services, the network tries to deliver a particular kind of service based on the QoS specified by each packet. This specification can occur in different ways, for example, using the 6-bit differentiated services code point (DSCP) setting in IP packets or source and destination addresses. The network uses the QoS specification to classify, mark, shape, and police traffic and to perform intelligent queuing.

Differentiated Services is used for several mission-critical applications and for providing end-to-end QoS.

Typically, Differentiated Services is appropriate for aggregate flows because it performs a relatively coarse level of traffic classification.

DS Field Definition

A replacement header field, called the DS field, is defined by Differentiated Services. The DS field supersedes the existing definitions of the IP version 4 (IPv4) type of service (ToS) octet (RFC 791) and the IPv6 traffic class octet. Six bits of the DS field are used as the DSCP to select the Per-Hop Behavior (PHB) at each interface. A currently unused 2-bit (CU) field is reserved for explicit congestion notification (ECN). The value of the CU bits is ignored by DS-compliant interfaces when determining the PHB to apply to a received packet.

Per-Hop Behaviors

RFC 2475 defines PHB as the externally observable forwarding behavior applied at a DiffServ-compliant node to a DiffServ Behavior Aggregate (BA).

With the ability of the system to mark packets according to DSCP setting, collections of packets with the same DSCP setting that are sent in a particular direction can be grouped into a BA. Packets from multiple sources or applications can belong to the same BA.

In other words, a PHB refers to the packet scheduling, queueing, policing, or shaping behavior of a node on any given packet belonging to a BA, as configured by a service level agreement (SLA) or a policy map.

The following sections describe the four available standard PHBs:

Default PHB

Class-Selector PHB (as defined in RFC 2474)

Assured Forwarding PHB (as defined in RFC 2597)

Expedited Forwarding PHB (as defined in RFC 2598)

For more information about default PHB, see RFC 2474, Definition of the Differentiated Services Field (DS

Field) in the IPv4 and IPv6 Headers.
Class-Selector PHB
To preserve backward-compatibility with any IP precedence scheme currently in use on the network, DiffServ has defined a DSCP value in the form xxx000, where x is either 0 or 1. These DSCP values are called Class-Selector Code Points. (The DSCP value for a packet with default PHB 000000 is also called the Class-Selector Code Point.)
The PHB associated with a Class-Selector Code Point is a Class-Selector PHB. These Class-Selector PHBs retain most of the forwarding behavior as nodes that implement IP Precedence-based classification and forwarding.
For example, packets with a DSCP value of 11000 (the equivalent of the IP Precedence-based value of 110) have preferential forwarding treatment (for scheduling, queueing, and so on), as compared to packets with a DSCP value of 100000 (the equivalent of the IP Precedence-based value of 100). These Class-Selector PHBs ensure that DS-compliant nodes can coexist with IP Precedence-based nodes.

For more information about Class-Selector PHB, see RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers.

Reference
http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/qcfdfsrv_ps1835_TSD_Products_Configuration_Guide_Chapter.html#wp998695

QUESTION NO: 383

An expanding company is deploying leased lines between its main site and two remote sites. The bandwidth of the leased lines is 128kb/s each, terminated on different serial interfaces on the main router. These links are used for combined VOIP and data traffic. The network administrator has implemented a VOIP solution to reduce costs, and has therefore reserved sufficient bandwidth in a low latency queue on each interface for the VOIP traffic. Users now complain about bad voice quality although no drops are observed in the low latency queue. What action will likely fix this problem?

A. mark VOIP traffic with IP precedence 6 and configure only fair-queue' on the links
B. configure the scheduler allocate 3000 1000 command to allow the OoS code to have enough CPU cycles
C. enable class-based traffic shaping on the VOIP traffic class
D. enable Layer 2 fragmentation and interleaving on the links
E. enable Frame Relay on the links and send voice and data on different Frame Relay PVCs

Answer: D

Explanation: Explanation

Link Fragmentation and Interleaving

Link fragmentation and interleaving (LFI) is a Layer 2 technique in which all Layer 2 frames are broken into small, equal-size fragments, and transmitted over the link in an interleaved fashion. When fragmentation and interleaving are in effect, the network device fragments all frames waiting in the queuing system where it prioritizes smaller frames. Then, the network device sends the fragments over the link. Small frames may be scheduled behind larger frames in the WFQ system. LFI fragments all frames, which reduces the queuing delay

of small frames because they are sent almost immediately. Link fragmentation reduces delay and jitter by normalizing packet sizes of larger packets in order to offer more regular transmission opportunities to the voice packets.

The following LFI mechanisms are implemented in Cisco IOS:

Multilink PPP with interleaving is by far the most common and widely used form of LFI.

FRF.11 Annex C LFI is used with Voice over Frame Relay (VoFR).

FRF.12 Frame Relay LFI is used with Frame Relay data connections.

Interleaving for Multilink PPP Configuration Task List

To configure MLP, perform the tasks described in the following sections. The task in the first section is required; the task in the remaining section is optional.

• Configuring MLP Interleaving (Required)

• Displaying Interleaving Statistics (Optional)

• Monitoring PPP and MLP Interfaces (Optional)

Configuring MLP Interleaving

MLP support for interleaving can be configured on virtual templates, dialer interfaces, and ISDN BRI or PRI interfaces. To configure interleaving, perform the following steps:

Step 1 Configure the dialer interface, BRI interface, PRI interface, or virtual interface template, as defined in the relevant Cisco IOS documents, Step 2 Configure MLP and interleaving on the interface or template.

Note Fair queueing, which is enabled by default, must remain enabled on the interface.

To configure MLP and interleaving on a configured and operational interface or virtual interface template, use the following commands in interface configuration mode:

Monitoring PPP and MLP Interfaces To monitor virtual interfaces, use the following command in EXEC mode:

Reference

http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/qcflfi.html#wp1000892

QUESTION NO: 384

You are the network administrator of an enterprise with a main site and multiple remote sites. Your network carries both VOIP and data traffic. You agree with your service provider to classify VOIP and data traffic according to the different service RFCs. How can your data and VOIP traffic be marked?

A. data marked with DSCP AF21, VOIP marked with DSCP EF
B. data marked with DSCP AF51, VOIP marked with DSCP EF
C. data marked with the DE-bit. VOIP marked with the CLP-bit
D. data marked with DSCP EF, VOIP marked with DSCP AF31
E. data marked with IP precedence 5, VOIP marked with DSCP EF

Answer: A

Explanation: Expedited Forwarding
RFC 2598 defines the Expedited Forwarding (EF) PHB: "The EF PHB can be used to build a low loss, low latency, low jitter, assured bandwidth, end-to-end service through DS (Diffserv) domains. Such a service appears to the endpoints like a point-to- point connection or a "virtual leased line." This service has also been described as Premium service." Codepoint 101110 is recommended for the EF PHB, which corresponds to a DSCP value of 46. Vendor-specific mechanisms need to be configured to implement these PHBs. Refer to RFC 2598 for more information about EF PHB.
DSCP was designed to be more granular and more scalable than IP precedence BUT with backward compatibility.
The priority field (or type of service, ToS) was originally 3 bits, giving it the IP prec values 0-7 (0 being the lowest priority, 7 the highest). DSCP has an 8 bit field, of which 6 bits are used for markings (the 6th bit is always 0). This gives it a larger number of values (both for per hop behavior, or PHB, and drop precedence.

More on this later)
The last 2 bits are used for ECN or explicit congestion notification. This is a brand new feature (as of Windows Vista) and is purportedly going to revolutionize internetwork traffic flow.
With the 6 bits allotted to DSCP, the first three (left to right) are used for Major Class, or Per Hop Behavior.

These match up with the old IP Prec values of 0-7. The second 3 bits identify the drop precedence. Higher = more likely to be dropped. This means that a DSCP marking of AF21 (major class 2, drop precedence of 1) will be preferred over AF22 or AF23. An AF3x will beat any AF1x or AF2x–the major class of 3 is higher than the major class, or PHB, of 2.
It is important to note the drop precedence is only used on classes 1-4. (Here is a table from Wikipedia) A marking of 0 indicates ‘best effort’
The notation DSCP xx is the bit notation. Eg AF12 = 001 100 which is a decimal value of 12.
Here’s another one. AF43 = 100 110 (38) EF, which it equal to IP Prec of 5 is assigned a major class of 5 and a drop precedence of 3 (this is odd, I don’t know why this was done) The decimal value for EF is DSCP 46 or 101 110.

QUESTION NO: 385

Refer to the exhibit. When applying this hierarchical policy map on the on the tunnel1 interface, you measure high jitter for traffic going through class 1234. What is the most likely cause of this jitter?

```
class-map match-all 1234
 match ip precedence 5
class-map match-all 5555
 match access-group 105
class-map match-all 5554
 match access-group 104

policy-map tun-shap
 class class-default
  shape average 150000 300000
  service-policy mark

policy-map mark
 class 1234
  priority 64
 class 5555
  set dscp af31
  bandwidth remaining percent 50
  random-detect dscp-based
 class 5554
  set dscp af32
  bandwidth remaining percent 25
  random-detect dscp-based

interface Tunnel1
 ip address 20.2.2.1 255.255.255.252
 ip load-sharing per-packet
 load-interval 30
 qos pre-classify
 tunnel source 4.4.4.1
 tunnel destination 4.4.4.2
 service-policy output tun-shap

access-list 104 permit ip any host 5.5.5.4
access-list 105 permit ip any host 5.5.5.5
```

A. The configuration of a hierarchical policy map on a tunnel interface is not supported.
B. Class 5555 and class 5554 are both taking up 100% of the bandwidth, leaving nothing for class 1234.
C. The burst size for the traffic shaping is wrongly configured to 15000; this would require an interface capable of sending at 150Mb/s.
D. The burst size for the traffic shaping has been wrongly configured; it should be set as low as possible.
E. The burst size for the traffic shaping has been wrongly configured; it should be set as high as possible.

Answer: D

Explanation:

Displaying Interleaving Statistics

To display interleaving statistics, use the following command in EXEC mode:

What Is a Token Bucket?

A token bucket is a formal definition of a rate of transfer. It has three components: a burst size, a mean rate, and a time interval (Tc). Although the mean rate is generally represented as bits per second, any two values may be derived from the third by the relation shown as follows:

mean rate = burst size / time interval

Here are some definitions of these terms:

Mean rate—Also called the committed information rate (CIR), it specifies how much data can be sent or forwarded per unit time on average.

Burst size—Also called the Committed Burst (Bc) size, it specifies in bits (or bytes) per burst how much traffic can be sent within a given unit of time to not create scheduling concerns. (For a shaper, such as GTS, it specifies bits per burst; for a policer, such as CAR, it specifies bytes per burst.)

Time interval—Also called the measurement interval, it specifies the time quantum in seconds per burst.

By definition, over any integral multiple of the interval, the bit rate of the interface will not exceed the mean rate.

The bit rate, however, may be arbitrarily fast within the interval.

A token bucket is used to manage a device that regulates the data in a flow. For example, the regulator might be a traffic policer, such as CAR, or a traffic shaper, such as FRTS or GTS. A token bucket itself has no discard or priority policy. Rather, a token bucket discards tokens and leaves to the flow the problem of managing its transmission queue if the flow overdrives the regulator. (Neither CAR nor FRTS and GTS implement either a true token bucket or true leaky bucket.)

In the token bucket metaphor, tokens are put into the bucket at a certain rate. The bucket itself has a specified capacity. If the bucket fills to capacity, newly arriving tokens are discarded. Each token is permission for the source to send a certain number of bits into the network. To send a packet, the regulator must remove from the bucket a number of tokens equal in representation to the packet size.

If not enough tokens are in the bucket to send a packet, the packet either waits until the bucket has enough tokens (in the case of GTS) or the packet is discarded or marked down (in the case of CAR). If the bucket is already full of tokens, incoming tokens overflow and are not available to future packets. Thus, at any time, the largest burst a source can send into the network is roughly proportional to the size of the bucket.

Note that the token bucket mechanism used for traffic shaping has both a token bucket and a data buffer, orqueue; if it did not have a data buffer, it would be a policer. For traffic shaping, packets that arrive that cannot be sent immediately are delayed in the data buffer.

For traffic shaping, a token bucket permits burstiness but bounds it. It guarantees that the

burstiness is bounded so that the flow will never send faster than the token bucket's capacity, divided by the time interval, plus the established rate at which tokens are placed in the token bucket. See the following formula: (token bucket capacity in bits / time interval in seconds) + established rate in bps = maximum flow speed in bps This method of bounding burstiness also guarantees that the long-term transmission rate will not exceed the established rate at which tokens are placed in the bucket.
Recommended Burst Values Cisco recommends the following values for the normal and extended burst parameters: normal burst = configured rate * (1 byte)/(8 bits) * 1.5 seconds
extended burst = 2 * normal burst

Reference
http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/qcfpolsh.html

QUESTION NO: 386

Refer to the exhibit. When applying this policy map on the tunnel1 interface, you see packet loss for the TCP class starting at around 100000 b/s, instead of the configured 150000 b/s. What is the most likely cause of the discrepancy?



A. The violate-action command should not be configured.
B. The current configuration of the load-interval command on the tunnel interface is preventing proper policing calculations.
C. The burst size is too low.
D. Policing on tunnel interfaces is not supported.
E. The CIR keyword is missing in the policer.

Answer: C
Explanation: Explanation
Configuration Tasks
See the following sections for configuration tasks for the Configuring Burst Size in Low Latency Queueing feature. Each task in the list is identified as optional or required.
Configuring the LLQ Bandwidth (Required)

Configuring the LLQ Burst Size (Required)
Verifying the LLQ Burst Size (Optional)
Configuring the LLQ Bandwidth
To configure the LLQ bandwidth, use the following command in policy-map class configuration mode:

Command	Purpose
Router# priority *bandwidth*	Specifies the maximum amount of bandwidth (in kpbs) for the priority traffic.

Configuring the LLQ Burst Size
To configure the LLQ burst size, use the following command in policy-map class configuration mode:

Command	Purpose
Router# priority *bandwidth burst*	Specifies the burst size in bytes. The range is 32 to 2 million.

Verifying the LLQ Burst Size
To verify the LLQ burst size, use one of the following commands in EXEC mode:

Command	Purpose
Router# show policy-map	Displays the configuration of all classes comprising the specified service policy map or all classes for all existing policy maps.
Router# show policy-map interface	Displays the configuration of classes configured for service polices on the specified interface or permanent virtual circuit (PVC).

Reference
http://www.cisco.com/en/US/docs/ios/12_1t/12_1t3/feature/guide/dtcfgbst.html#wp1015329

QUESTION NO: 387

NBAR supports all of these with the exception of which one?

A. HTTP
B. IP multicast
C. TCP flows with dynamically assigned port numbers
D. non-UDP protocols

Answer: B
Explanation:
Restrictions for Using NBAR
NBAR does not support the following:
More than 24 concurrent URLs, hosts, or Multipurpose Internet Mail Extension (MIME) type matches.
Matching beyond the first 400 bytes in a packet payload in Cisco IOS releases before Cisco IOS Release 12.3 (7)T. In Cisco IOS Release 12.3(7)T, this restriction was removed, and NBAR now supports full payload inspection. The only exception is that NBAR can inspect custom protocol traffic for only 255 bytes into the payload.

Non-IP traffic
Multiprotocol Label Switching (MPLS)-labeled packets - NBAR classifies IP packets only. You can, however, use NBAR to classify IP traffic before the traffic is handed over to MPLS. Use the Modular Quality of Service (QoS) Command-Line Interface (CLI) (MQC) to set the IP differentiated services code point (DSCP) field on the NBAR-classified packets and make MPLS map the DSCP setting to the MPLS experimental (EXP) setting inside the MPLS header.
Multicast and other non-CEF switching modes Fragmented packets Pipelined persistent HTTP requests

URL/host/MIME classification with secure HTTP
Asymmetric flows with stateful protocols
Packets that originate from or that are destined to the router running NBAR
NBAR is not supported on the following logical interfaces:
Fast EtherChannel
Dialer interfaces until Cisco IOS Release 12.2(4) T
Interfaces where tunneling or encryption is used

QUESTION NO: 388

Modified deficit round robin supports which of these functionalities?

A. priority queue
B. weighted fair queues
C. round-robin service of output queues

D. LLQ

Answer: A,C

Explanation:

Modified deficit round robin (MDRR)—MDRR, a traffic class prioritization mechanism used only on GSR platforms, incorporates emission priority as a facet of quality of service. MDRR is similar in function to WFQ on non-GSR platforms.

In MDRR, IP traffic is mapped to different classes of service queues. A group of queues is assigned to each traffic destination. On the transmit side of the platform, a group of queues is defined on a per-interface basis; on the receive side of the platform, a group of queues is defined on a per-destination basis. IP packets are then mapped to these queues, based on their IP precedence value.

These queues are serviced on a round-robin basis, except for a queue that has been defined to run in either of two ways: a) strict priority mode, or b) alternate priority mode.

In strict priority mode, the high priority queue is serviced whenever it is not empty; this ensures the lowest possible delay for high priority traffic. In this mode, however, the possibility exists that other traffic might not be serviced for long periods of time if the high priority queue is consuming most of the available bandwidth.

In alternate priority mode, the traffic queues are serviced in turn, alternating between the high priority queue and the remaining queues.

Reference

http://www.cisco.com/en/US/docs/ios/12_0st/12_0st10/feature/guide/10st_cos.pdf

QUESTION NO: 389

Refer to the exhibit. A network engineer received a sudden request to prioritize voice over his Cisco network and he has decided to leverage the AutoQoS feature. Based on the output shown, which two tasks need to be performed prior to issuing the autoqos voip command in this router? (Choose two.)

```
RouterA#sh run
Building configuration...
Current configuration : 440 bytes
!
version 12.3
!
hostname RouterA
!
no ip cef
!
class-map match-all NonIPVoice
 match not protocol rtp audio
class-map match-all IPVoice
 match protocol rtp audio
!
policy-map VoIP
 class IPVoice
  priority 300
!
interface FastEthernet0/0
 ip address 192.168.1.2 255.255.255.0
 duplex auto
 speed auto
!
interface Serial1/0
 ip address 10.1.1.1 255.255.255.0
 serial restart-delay 0
 no dce-terminal-timing-enable
 service-policy output VoIP
!
ip classless
!
end

RouterA#
```

ActualTests

A. Enable Cisco Express Forwarding.
B. Enable fast switching.
C. Delete all policy maps.
D. Remove service-policy commands from interface serial1/0.
E. Delete all the currently configured class maps.

Answer: A,D

Explanation:

AutoQoS VoIP Default Configuration

Before configuring AutoQoS VoIP, you should refer to the IOS 12.3 QoS Configuration Guide, which lists several considerations and conditions for the right environment for enabling this feature. For QOS exam purposes, repeating the full list here is not helpful; however, considering a

couple of the most common considerations can help. For instance
AutoQoS VoIP requires that CEF be enabled first.
AutoQoS VoIP cannot be used if the interface already has a service-policy command configured.
Because AutoQoS VoIP relies on the bandwidth settings configured in the bandwidth command, the routers should be configured with correct bandwidth settings on each interface before enabling AutoQoS VoIP. (If you change the bandwidth after enabling AutoQoS VoIP, AutoQoS VoIP does not react and does not change the QoS configuration.)
Supports only point-to-point subinterfaces on Frame Relay interfaces.
Supports HDLC, PPP, Frame Relay, and ATM data link protocols.

Reference
http://www.ciscopress.com/articles/article.asp?p=358548&seqNum=6

QUESTION NO: 390

What is an important consideration that should be taken into account when configuring shaped round robin?

A. It enables policing.
B. Strict priority is not supported.
C. WRED must be previously enabled.
D. It enables WRR.

Answer: B
Explanation:
First we need to understand how round robin algorithm works. The round robin uses multiple queues and dispatches one packet from each queue in each round with no prioritization. For example, it dispatches:
Dispatch one packet from Queue 1
Dispatch one packet from Queue 2
Dispatch one packet from Queue 3
Repeat from Queue 1

There are three implementations of Round Robin scheduling on the Catalyst 6500 and they include Weighted Round Robin (WRR), Deficit Weighted Round Robin (DWRR) and Shaped Round Robin (SRR).

The Weighted Round Robin allows prioritization, meaning that it assigns a “weight” to each queue and dispatches packets from each queue proportionally to an assigned weight. For example:

Dispatch 3 packets from Queue 1 (Weight 3)
Dispatch 2 packets from Queue 2 (Weight 2)
Dispatch 1 packet from Queue 1 (Weight 1)
Repeat from Queue 1 (dispatch 3 next packets)

Unlike Priority Queuing, which always empties the first queue before going to the next queue, this kind of queue prevents starvation of other applications such as if a large download is in progress. The Weighted Round Robin can be used with Strict Priority by setting its weight to 0. That means packets in the other queues will not be serviced until queue 4 is emptied. The problem of WRR is the router is allowed to send the entire packet even if the sum of all bytes is more than the threshold and can make other applications starved. The Deficit Round Robin solves problem of WRR by keeping track of the number of “extra” bytes dispatched in each round – the “deficit” and then add the “deficit” to the number of bytes dispatched in the next round. Shaped Round Robin (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. Shaped egress queues reserve a set of port bandwidth and then send evenly spaced packets as per the reservation.

Shared egress queues are also guaranteed a configured share of bandwidth, but do not reserve the bandwidth. That is, in shared mode, if a higher priority queue is empty, instead of the servicer waiting for that reserved bandwidth to expire, the lower priority queue can take the unused bandwidth. Neither shaped SRR nor shared SRR is better than the other. Shared SRR is used to get the maximum efficiency out of a queuing system, because unused time slots can be reused by queues with excess traffic. This is not possible in a standard Weighted Round Robin. Shaped SRR is used to shape a queue or set a hard limit on how much bandwidth a queue can use. When you use shaped SRR, you can shape queues within a port’s overall shaped rate.

Reference
http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps7078/prod_qas0900aecd805bacc7.html

QUESTION NO: 391

Refer to the exhibit. Based on the configuration shown, which queuing mechanism has been configured on interface serial.1/0?

```
RouterC#show policy-map interface serial 1/0
 Serial1/0

  Service-policy output: BW-Assignment

    Class-map: VoIP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: protocol rtp audio
      Queueing
        Output Queue: Conversation 265
        Bandwidth 50 (%)
        Bandwidth 772 (kbps) Max Threshold 64 (packets)
        (pkts matched/bytes matched) 0/0
        (depth/total drops/no-buffer drops) 0/0/0

    Class-map: FTP-Data (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name FTP-Data
      Queueing
        Output Queue: Conversation 266
        Bandwidth 10 (%)
        Bandwidth 154 (kbps) Max Threshold 64 (packets)
        (pkts matched/bytes matched) 0/0
        (depth/total drops/no-buffer drops) 0/0/0

    Class-map: class-default (match-any)
      4 packets, 954 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
RouterC#
```

A. PQ
B. CQ
C. WFQ
D. LLQ
E. CBWFQ

Answer: E

Explanation:

Class-based weighted fair queueing (CBWFQ) extends the standard WFQ functionality to provide support for user-defined traffic classes. For CBWFQ, you define traffic classes based on match criteria including protocols, access control lists (ACLs), and input interfaces. Packets satisfying the match criteria for a class constitute the traffic for that class. A queue is reserved for each class, and traffic belonging to a class is directed to the queue for that class.

Once a class has been defined according to its match criteria, you can assign it characteristics. To characterize a class, you assign it bandwidth, weight, and maximum packet limit. The bandwidth assigned to a class is the guaranteed bandwidth delivered to the class during congestion.

To characterize a class, you also specify the queue limit for that class, which is the maximum number of packets allowed to accumulate in the queue for the class. Packets belonging to a class are subject to the bandwidth and queue limits that characterize the class.
After a queue has reached its configured queue limit, enqueuing of additional packets to the class causes tail drop or packet drop to take effect, depending on how class policy is configured.
Tail drop is used for CBWFQ classes unless you explicitly configure policy for a class to use Weighted Random Early Detection (WRED) to drop packets as a means of avoiding congestion. Note that if you use WRED packet drop instead of tail drop for one or more classes comprising a policy map, you must ensure that WRED is not configured for the interface to which you attach that service policy.

If a default class is configured with the bandwidth policy-map class configuration command, all unclassified traffic is put into a single queue and given treatment according to the configured bandwidth. If a default class is configured with the fair-queue command, all unclassified traffic is flow classified and given best-effort treatment.

If no default class is configured, then by default the traffic that does not match any of the configured classes is flow classified and given best-effort treatment. Once a packet is classified, all of the standard mechanisms that can be used to differentiate service among the classes apply. Flow classification is standard WFQ treatment. That is, packets with the same source IP address, destination IP address, source Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port, or destination TCP or UDP port are classified as belonging to the same flow. WFQ allocates an equal share of bandwidth to each flow. Flow-based WFQ is also called fair queueing because all flows are equally weighted.

For CBWFQ, which extends the standard WFQ fair queueing, the weight specified for the class becomes the weight of each packet that meets the match criteria of the class. Packets that arrive at the output interface are classified according to the match criteria filters you define, then each one is assigned the appropriate weight.

The weight for a packet belonging to a specific class is derived from the bandwidth you assigned to the class when you configured it; in this sense the weight for a class is user-configurable.
After the weight for a packet is assigned, the packet is enqueued in the appropriate class queue. CBWFQ uses the weights assigned to the queued packets to ensure that the class queue is serviced fairly.

Configuring a class policy—thus, configuring CBWFQ—entails these three processes:
• Defining traffic classes to specify the classification policy (class maps).
This process determines how many types of packets are to be differentiated from one another.
• Associating policies—that is, class characteristics—with each traffic class (policy maps).
This process entails configuration of policies to be applied to packets belonging to one of the

classes previously defined through a class map. For this process, you configure a policy map that specifies the policy for each traffic class.
• Attaching policies to interfaces (service policies).
This process requires that you associate an existing policy map, or service policy, with an interface to apply the particular set of policies for the map to that interface

Reference
http://www.cisco.com/en/US/docs/ios/12_0t/12_0t5/feature/guide/cbwfq.html

QUESTION NO: 392

Which two benefits are of applying WRED? (Choose two.)

A. helps to avoid TCP synchronization
B. allows a different drop profile to be manually enabled for each IP precedence or DSCP
C. provides minimal bandwidth guarantees
D. provides bounded low latency

Answer: A,B
Explanation:
WRED and distributed WRED (DWRED)—both of which are the Cisco implementations of RED—combine the capabilities of the RED algorithm with the IP Precedence feature. Within the section on WRED, the following related features are discussed:

–

Flow-based WRED. Flow-based WRED extends WRED to provide greater fairness to all flows on an interface in regard to how packets are dropped.
–

DiffServ Compliant WRED-DiffServ Compliant WRED extends WRED to support Differentiated Services (DiffServ) and Assured Forwarding (AF) Per Hop Behavior (PHB). This feature enables customers to implement AF PHB by coloring packets according to differentiated services code point (DSCP) values and then assigning preferential drop probabilities to those packets.

WRED avoids the globalization problems that occur when tail drop is used as the congestion avoidance mechanism on the router. Global TCP synchronization occurs as waves of congestion crest only to be followed by troughs during which the transmission link is not fully utilized. Global synchronization of TCP hosts, for example, can occur because packets are dropped all at once. Global synchronization manifests when multiple TCP hosts reduce their transmission rates in response to packet dropping, then increase their transmission rates once again when the congestion is reduced.

Reference:
http://www.google.com/search?hl=en&q=WRED%2C+A+different+drop+profile+can+be+manually+enabled+per+IP+Precedence+or+DSCP

QUESTION NO: 393

NBAR is used to provide which QoS function?

A. classification
B. policing
C. CBWFQ bandwidth guarantees
D. shaping

Answer: A
Explanation:
NBAR addresses IP QoS classification requirements by classifying application-level protocols so that QoS policies can be applied to the classified traffic. NBAR addresses the ongoing need to extend the classification engine for the many existing and emerging application protocols by providing an extensible Packet Description Language (PDL). NBAR can determine which protocols and applications are currently running on a network so that an appropriate QoS policy can be created based upon the current traffic mix and application requirements.

Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800c75d1.html#54116

QUESTION NO: 394

How can excess packets be remarked?

A. shaping
B. policing
C. priority
D. all of the above

Answer: B
Reference:

http://www.cisco.com/en/US/tech/tk543/tk545/technologies_tech_note09186a00800a3a25.shtml (see the table - first row)

QUESTION NO: 395

What is the default maximum reservable bandwidth (percentage) by any single flow on an interface after enabling RSVP?

A. 75 percent
B. 60 percent
C. 56 percent
D. 50 percent
E. 25 percent

Answer: A
Explanation:
You must plan carefully to successfully configure and use RSVP on your network. At a minimum, RSVP must reflect your assessment of bandwidth needs on router interfaces. Consider the following questions as you plan for RSVP configuration:
How much bandwidth should RSVP allow per end-user application flow? You must understand the "feeds and speeds" of your applications. By default, the amount reservable by a single flow can be the entire reservable bandwidth. You can, however, limit individual reservations to smaller amounts using the single flow bandwidth parameter. This value may not exceed the interface reservable amount, and no one flow may reserve more than the amount specified.

How much bandwidth is available for RSVP? By default, 75 percent of the bandwidth available on an interface is reservable. If you are using a tunnel interface, RSVP can make a reservation for the tunnel whose bandwidth is the sum of the bandwidths reserved within the tunnel.
How much bandwidth must be excluded from RSVP so that it can fairly provide the timely service required by low-volume data conversations? End-to-end controls for data traffic assumes that all sessions will behave so as to avoid congestion dynamically. Real-time demands do not follow this behavior. Determine the bandwidth to set aside so bursty data traffic will not be deprived as a side effect of the RSVP QOS configuration.

QUESTION NO: 396

Which two protocols can have their headers compressed through MQC? (Choose two.)

A. RTP
B. RTSP
C. HTTP
D. TCP
E. UDP

Answer: A,D
Explanation: Explanation
RTP or TCP IP header compression is a mechanism that compresses the IP header in a data packet before the packet is transmitted. Header compression reduces network overhead and speeds up transmission of RTP and TCP packets.
Cisco IOS software provides a related feature called Express RTP/TCP Header Compression. Before this feature was available, if compression of TCP or RTP headers was enabled, compression was performed in the process-switching path. Compression performed in this manner meant that packets traversing interfaces that had TCP or RTP header compression enabled were queued and passed up the process to be switched. This procedure slowed down transmission of the packet, and therefore some users preferred to fast-switch uncompressed TCP and RTP packets.

Now, if TCP or RTP header compression is enabled, it occurs by default in the fast-switched path or the Cisco Express Forwarding-switched (CEF-switched) path, depending on which switching method is enabled on the interface. Furthermore, the number of TCP and RTP header compression connections was increased.

If neither fast-switching nor CEF-switching is enabled, then if TCP or RTP header compression is enabled, it will occur in the process-switched path as before.

The Express RTP and TCP Header Compression feature has the following benefits:
1. It reduces network overhead.
2. It speeds up transmission of TCP and RTP packets. The faster speed provides a greater benefit on slower links than faster links.

QUESTION NO: 397

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

A. NBAR performs protocol discovery.
B. NBAR is not dependent on CEF.
C. NBAR is used for traffic statistics collection.
D. NBAR performs traffic classification.

Answer: A,D

Explanation:

NBAR Can perform protocol discovery using the sniffing capability of its classification Engine. In cases where is NBAR is not required for defining the QoS Policy, NBAR protocol discovery mode can be used to get information about traffic present on the network and how much bandwidth each traffic type is using.

Command: LABRouter#show ip nbar protocol-discovery stats byte-rate <Interface>

QUESTION NO: 398

Which two types of QoS functionality will be provided by Network-Based Application Recognition? (Choose two.)

A. NBAR provides the ability to configure MCQ; it is a mandatory MCQ component.
B. NBAR provides deep packet inspection and is used for advanced packet classification.
C. NBAR provides per-protocol packet and byte accounting functionality; it is used to track bandwidth utilization for all protocols described in the loaded PDLMs.
D. NBAR provides scheduling in an MQC policy map using an advanced algorithm.

Answer: B,C

Explanation:

Explanation NBAR classes packets that are normally difficult to classify. For instance, some applications use dynamic port numbers. NBAR can look past the UDP and TCP header, and refer to the host name, URL, or MIME type in HTTP requests.

QUESTION NO: 399

You work as a network technician at a famous Company.com, study the exhibit provided. You are implementing this QoS configuration to improve the bandwidth guarantees for traffic towards two servers, one with the IP address 5.5.5.5 and the other with the IP address 5.5.5.4. Even after the configuration is applied, performance does not seem to improve. Which will be the most likely cause of this problem?

```
  random-detect dscp-based
class GOLD
  priority percent 25
  random-detect dscp-based
policy-map mark
  class 5555
  set dscp af31
class GOLD
  set dscp 5

interface Tunnel1
  ip address 20.1.1.1 255.255.255.252
  ip load-sharing per-packet
  load-interval 30
  tunnel source 4.4.4.1
  tunnel destination 4.4.4.2

interface Serial0/0
  ip address 4.4.4.1 255.255.255.252
  ip load-sharing per-packet
  ip nbar protocol-discovery
  load-interval 30
  service-policy output queue

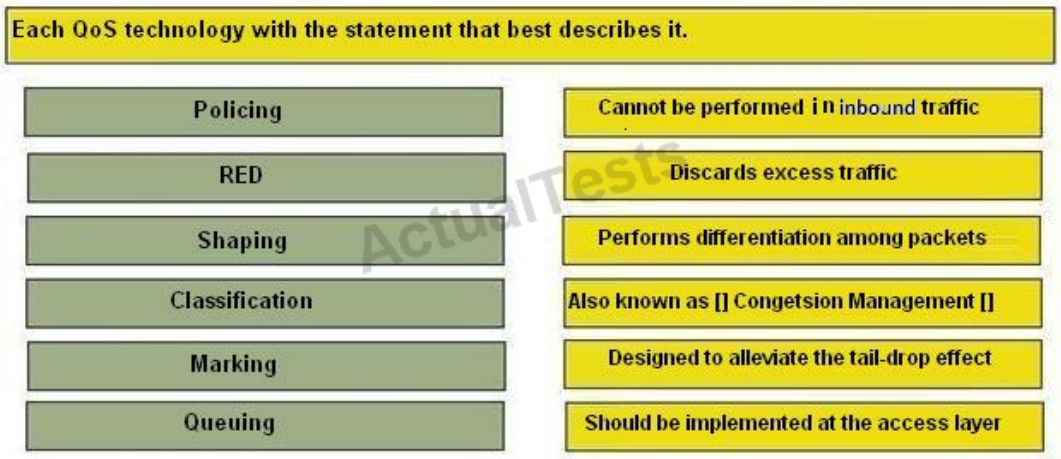
interface Ethernet0/1
  ip address 3.3.3.1 255.255.255.0
  no ip proxy-arp
  load-interval 30
  half-duplex
  no keepalive
  no cdp enable
  service-policy input mark
```

A. The policy map mark has been applied on a half-duplex Ethernet interface; this is not supported.
B. The policy map queue is configured on the wrong interface; it is applied on the serial interface whereas traffic is going over the tunnel interface.
C. The class maps are wrongly configured
D. The ip nbar protocol-discover command cannot be configured together with a service policy output on the serial interface.
E. This is probably a software bug

Answer: C

Explanation:

QUESTION NO: 400 DRAG DROP



Answer:



Explanation:

Policing - Discards excess trafficRED - Designated to alleviate tail drop affectShaping - Can not be performed on inbound trafficClassification - Performs differentiation among packetMarking - should be implemented at the access layerQueuing - Also known as congestion management

QUESTION NO: 401

Which three protocols or applications should be placed in a class that is configured with WRED? (Choose three.)

A. RTP
B. streaming video
C. SMTP

D. SSH
E. BitTorrent

Answer: C,D,E
Explanation: Reference;
http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/qcfwred_ps1835_TSD_Products_Configuration_Guide_Chapter.html

QUESTION NO: 402

Which QoS mechanism will rate-limit traffic, and limit drops by implementing queuing?

A. Shaper
B. Policer
C. WRED
D. Rate-Limit
E. LLQ
F. Fair-Queue

Answer: A
Reference:
http://www.cisco.com/en/US/tech/tk543/tk545/technologies_tech_note09186a00800a3a25.shtml

QUESTION NO: 403

What is the default queuing mechanism on slow serial interfaces?

A. FIFO
B. WFQ
C. CQ
D. LLQ
E. WRR

Answer: B
Explanation:
http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/qcfconmg.html#wp1001019

Given this handling, WFQ ensures satisfactory response time to critical applications, such as

interactive, transaction-based applications, that are intolerant of performance degradation. For serial interfaces at E1 (2.048 Mbps) and below, flow-based WFQ is used by default. When no other queueing strategies are configured, all other interfaces use FIFO by default.

QUESTION NO: 404

Which QoS mechanism will prevent a decrease in TCP performance?

A. Shaper
B. Policer
C. WRED
D. Rate-Limit
E. LLQ
F. Fair-Queue

Answer: C
Reference: http://www.routeralley.com/ra/docs/qos_congestion_avoidance.pdf (page 20

QUESTION NO: 405

Which of the following is true about WRED?

A. WRED cannot be applied to the same interface as CQ, PQ, and WFQ.
B. WRED drops packets from all flows.
C. WRED cannot mark with a probability denominator.
D. WRED cannot be applied to the voice queue.

Answer: A
Explanation: http://www.cisco.com/en/US/docs/ios/qos/configuration/guide/config_wred.html

QUESTION NO: 406

Refer to the Exhibit.

```
Router#show policy-map interface atm 0
 ATM0: VC 8/48 -

  Service-policy output: ATM0-EGRESS
    queue stats for all priority classes:
      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 243078/74868024

    Class-map: voice (match-all)
      243078 packets, 74868024 bytes
      30 second offered rate 12000 bps, drop rate 0 bps
      Match: ip precedence 5
      Priority: 64 kbps, burst bytes 5000, b/w exceed drops: 0

    Class-map: class-default (match-any)
      6504994 packets, 2354790276 bytes
      30 second offered rate 203000 bps, drop rate 3000 bps
      Match: any
      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/8065/0
      (pkts output/bytes output) 6496929/2353490427
      bandwidth remaining 20% (38 kbps)
```

The show command was taken on a router, while a large file was uploading to a server, and a VOIP call was running at the same time:

During the file upload, the remote user on the call complains about poor call quality. After the upload is complete, everything operates properly, and the remote user can hear the local user perfectly. Which QoS mechanism will solve the issue with the VOIP quality?

A. LLQ
B. LFI
C. WRED
D. WFQ

Answer: B
Reference: http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/qcflem.html

QUESTION NO: 407

Which QoS mechanism will allow traffic flows an equal share of the bandwidth?

A. Shaper

B. Policer
C. WRED
D. Rate-Limit
E. LLQ
F. Fair-Queue

Answer: F
Reference:
http://www.cisco.com/en/US/docs/ios/solutions_docs/qos_solutions/QoSVoIP/QoSVoIP.html (see low latency queuing)

QUESTION NO: 408

If shaping is configured with 128 KB/s and a committed burst of 3200 B/s, what would be the value of TC?

A. 25 sec
B. 25 ms
C. 125 sec
D. 125 ms

Answer: B
Explanation:
3200 bits / 128000 bps == 0.025 sec (25 ms)
Reference:
http://www.cisco.com/en/US/tech/tk543/tk545/technologies_tech_note09186a00800a3a25.shtml#comm-anc (search for value of Tc)

QUESTION NO: 409

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

A. Out-of-profile packets are queued.
B. It causes TCP retransmits.
C. Marking and remarking are not supported.
D. It does not respond to BECN and foresight messages.
E. It uses a single- and two-bucket mechanism for metering.

Answer: B,D,E

Reference:
http://books.google.com.pk/books?id=eLa5crPXSOYC&pg=PA642&lpg=PA642&dq=policing+causes+TCP+retransmits&source=bl&ots=KezstlVBgx&sig=HhykMycdtbgEVwP30wqd5yJ3ZCE&hl=en&sa=X&ei=eiTLUeX3PIeL7Aa4k4BA&redir_esc=y#v=onepage&q=policing%20causes%20TCP%20retransmits&f=false

QUESTION NO: 410

Refer to the exhibit. As a network administrator, you have configured a dual-rate, dual-bucket policer in accordance with RFC 2698 on the serial interface of you router, connecting to your provider. The SLA with your provider states that you should only send AF31 (limited to 150 kb/s), AF32 (limited to 50 kb/s)and AF33 (best effort). Your service provider claims you are not conforming to the SLA. Which two things are wrong with this configuration? (Choose two.)



```
class-map match-all af31
 match  dscp af31
class-map match-all af32
 match  dscp af32
class-map match-all af33
 match  dscp af33

policy-map marking
 class af31
  set dscp af31
 class af32
  set dscp af32
 class af33
  set dscp af33

policy-map limit
 class af33
  police cir 150000 bc 50000 pir 200000 be 50000
    conform-action set-dscp-transmit af31
    exceed-action set-dscp-transmit af32
    violate-action set-dscp-transmit default
 class class-default
  bandwidth 300

interface Ethernet0/1
 ip address 3.3.3.1 255.255.255.0
 no ip proxy-arp
 load-interval 30
 half-duplex
 no keepalive
 no cdp enable
 service-policy input marking

interface Serial0/0
 ip address 4.4.4.1 255.255.255.252
 ip load-sharing per-packet
 encapsulation ppp
 load-interval 30
 no dce-terminal-timing-enable
 service-policy output limit
```

A. The configuration of a service policy on half-duplex Ethernet interfaces is not supported.
B. The class class-default sub-command of the policy-map limit command should be set to the DSCP default.
C. The violate action is wrong.
D. This policer configuration is not implementing RFC 2698 dual-bucket, dual-rate.
E. The policer is configured in the wrong class.

Answer: C,E

Explanation: Explanation

All of the policing is supposed to be applied to AF31 & AF32 but instead it is applied to AF33

Feature Overview

Networks police traffic by limiting the input or output transmission rate of a class of traffic based on userdefined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or class of service (CoS).

The Two-Rate Policer performs the following functions:

Limits the input or output transmission rate of a class of traffic based on user-defined criteria.

Marks packets by setting the IP precedence value, IP differentiated services code point (DSCP) value, Multiprotocol Label Switching (MPLS) experimental value, Quality of Service (QoS) group, ATM Cell Loss

Priority (CLP) bit, and the Frame Relay Discard Eligibility (DE) bit.

With the Two-Rate Policer, you can enforce traffic policing according to two separate rates—committed information rate (CIR) and peak information rate (PIR). You can specify the use of these two rates, along with their corresponding values, by using two keywords, cir and pir, of the police command. For more information about the police command, see the "Command Reference" section of this document.

The Two-Rate Policer manages the maximum rate of traffic through a token bucket algorithm. The token bucket algorithm can use the user-configured values to determine the maximum rate of traffic allowed on an interface at a given moment in time. The token bucket algorithm is affected by all traffic entering or leaving the interface (depending on the location of the interface on which the Two-Rate Policer is configured) and is useful in managing network bandwidth in cases where several large packets are sent in the same traffic stream.

The token bucket algorithm provides users with three actions for each packet: a conform action, an exceed action, and an optional violate action. Traffic entering the interface with Two-Rate Policer configured is placed in to one of these categories. Within these three categories, users can decide packet treatments. For instance, packets that conform can be configured to be sent, packets that exceed can be configured to be sent with a decreased priority, and packets that violate can be configured to be dropped.

The Two-Rate Policer is often configured on interfaces at the edge of a network to limit the rate of traffic entering or leaving the network. In the most common configurations, traffic that conforms is sent and traffic that exceeds is sent with a decreased priority or is dropped. Users can change these configuration options to suit their network needs.

Note Additionally, the Two-Rate Policer enables you to implement Differentiated Services (DiffServ) Assured Forwarding (AF) Per-Hop Behavior (PHB) traffic conditioning. For more information about DiffServ, refer to the
"Implementing DiffServ for End-to-End Quality of Service" chapter of the Cisco IOS Quality of Service Solutions
Configuration Guide, Release 12.2.
Police To configure traffic policing, use the police command in policy-map class configuration mode. To remove traffic policing from the configuration, use the no form of this command.
police {cir cir} [bc conform-burst] {pir pir} [be peak-burst] [conform-action action[exceed-action action [violateaction action]]]
no police {cir cir} [bc conform-burst] {pir pir} [be peak-burst] [conform-action action [exceed-action action[violate-action action]]]
Syntax Description

cir	The committed information rate (CIR) at which the first token bucket is updated.
cir	Specifies the CIR value in bits per second.
bc	(Optional) The conform burst size (bc) used by the first token bucket for policing.
conform-burst	(Optional) Specifies the bc value in bytes.
pir	The peak information rate (PIR) at which the second token bucket is updated.
pir	Specifies the PIR value in bits per second.
be	(Optional) The peak burst size (be) used by the second taken bucket for policing.
peak-burst	(Optional) Specifies the be size in bytes.
conform-action	(Optional) Action to take on packets that conform to the CIR and PIR.
exceed-action	(Optional) Action to take on packets that conform to the PIR but not the CIR.
violate-action	(Optional) Action to take on packets exceed the PIR.

ActualTests

action	(Optional) Action to take on packets. Specify one of the following keywords: • **drop**—Drops the packet. • **set-clp-transmit**—Sets the ATM Cell Loss Priority (CLP) bit from 0 to 1 on the ATM cell and sends the packet with the ATM CLP bit set to 1. • **set-dscp-transmit** *new-dscp*—Sets the IP differentiated services code point (DSCP) value and sends the packet with the new IP DSCP value setting. • **set-frde-transmit**—Sets the Frame Relay Discard Eligibility (DE) bit from 0 to 1 on the Frame Relay frame and sends the packet with the DE bit set to 1. • **set-mpls-exp-transmit**—Sets the Multiprotocol Label Switching (MPLS) experimental bits from 0 to 7 and sends the packet with the new MPLS experimental bit value setting. • **set-prec-transmit** *new-prec*—Sets the IP precedence and sends the packet with the new IP precedence value setting. • **set-qos-transmit** *new-qos*—Sets the Quality of Service (QoS) group value and sends the packet with the new QoS group value setting. • **transmit**—Sends the packet with no alteration.

Reference
http://www.cisco.com/en/US/docs/ios/12_2t/12_2t4/feature/guide/ft2RTplc.html#wp1015327

QUESTION NO: 411

Refer to the exhibit. You have noticed that several users in the network are consuming a great deal of bandwidth for the peer-to-peer application Kazaa2. You would like to limit this traffic, and at the same time provide a guaranteed 100 kb/s bandwidth for one of your servers. After applying the configuration in the exhibit, you notice no change in the bandwidth utilization on the serial link; it is still heavily oversubscribing the interface. What is the cause of this problem?

```
no ip cef

class-map match-all kazaa2
 match protocol kazaa2
class-map match-all server
 match access-group 105

policy-map p2p
 class kazaa2
  drop
 class server
  bandwidth 100
 class class-default
  fair-queue

interface Serial0/0
 bandwidth 1234
 ip address 20.1.34.1 255.255.255.252
 ip load-sharing per-packet
 encapsulation ppp
 load-interval 30
 no dce-terminal-timing-enable
 service-policy output p2p
```

ActualTests

A. CEF needs to be enabled for NBAR.
B. In class Kazaa2, you should configure a policer instead of a drop command.
C. The server class should have a priority of 100.
D. The bandwidth parameter on serial 0/0 is wrong.
E. Kazaa2 is not a valid protocol.

Answer: A

Explanation: Explanation
You need to enable Cisco Express Forwarding (CEF) in order to use NBAR.
How do you configure Cisco IOS NBAR?

Keep in mind that in its simplest form NBAR is a traffic identification and marking system. What you do with the marked packets is up to you. For example, you could choose to drop them or choose to give them a higher quality of service.

Configuring and using NBAR to identify and block traffic is actually very easy. Let’s walk through the steps.

Step 1
Make sure that CEF is on using the following command:
Router(config)# ip cef

Step 2
Create a class-map, identifying the traffic you want to block. Here’s an example that would stop any HTTP or MIME e-mail that contains the Readme.exe program:

Router(config)#class-map match-any bad-traffic
Router(config-cmap)# match protocol http url "*readme.exe*"
Router(config-cmap)# match protocol http mime "*readme.exe*"

I want to stress here that HTTP is just one of the many applications that NBAR can identify. For list of NBAR applications recognized with IOS version 12.3, use the following commands:
Router(config)#class-map match-all nbar
Router(config-cmap)#match pro ?

Step 3
Create a policy to mark the traffic. Here's an example:
Router(config)# policy-map mark-bad-traffic
Router(config-pmap)# class bad-traffic
Router(config-pmap)# set ip dscp 1

Step 4
Apply the policy to the interface that faces the Internet or the source of the traffic that you want to block. This marks the traffic when it enters the router. Here's an example:
Router(config)# interface serial 0/0
Router(config-if)#service-policy input mark-bad-traffic

Step 5
Create an access control list (ACL) that denies the marked traffic. Here's an example:
Router(config)# access-list 190 deny ip any any dscp 1
Router(config)# access-list 190 permit ip any any

Step 6
Deny the marked traffic as it's about to exit your router by applying the ACL to an interface. Here's an example:
Router(config)# interface GigabitEthernet 0/0
Router(config-if)# ip access-group 190 out
When you've finished applying the configuration, you can check to see if the router marked and dropped any traffic that met this criteria. To do this, use the show access-lists command.
Summary NBAR is a very powerful application-layer firewall that you may already have installed on your Cisco router.

While traditional firewalls can only recognize traffic based on IOS Layers 3 or 4, Cisco's NBAR can go all the way to Layer 7.

Reference
http://www.cisco.com/en/US/products/hw/routers/ps359/products_tech_note09186a00800fc176.sh

tml
http://www.techrepublic.com/blog/networking/what-can-ciscos-network-based-application-recognition-nbar-dofor-you/399
http://www.cisco.com/en/US/products/ps6616/products_ios_protocol_group_home.html

QUESTION NO: 412

All of these are fundamental building blocks of differentiated services Traffic Conditioner Block except which one?

A. dropper
B. classifier
C. marker
D. querier
E. meter
F. shaper

Answer: D
Explanation:
Differentiated Services Components
The following components make up the foundation of a Cisco Differentiated Services implementation:
Traffic conditioning (traffic policing and traffic shaping)Traffic conditioning is performed at the edges of a DiffServ domain. Traffic conditioners perform traffic shaping and policing functions to ensure that traffic entering the DiffServ domain conforms to the rules specified by the Traffic Conditioning Agreement (TCA) and complies with the service provisioning policy of the domain. Traffic conditioning may range from simple code point remarking to complex policing and shaping operations.
Packet classification - Packet classification uses a traffic descriptor (for example, the DSCP) to categorize a packet within a specific group in order to define that packet. After the packet has been defined (that is, classified), the packet is accessible for QoS handling on the network. Using packet classification, you can partition network traffic into multiple priority levels or classes of service. When traffic descriptors are used to classify traffic, the source agrees to adhere to the contracted terms and the network promises a QoS. Traffic policers and traffic shapers use the traffic descriptor of the packet (that is, the classification of the packet) to ensure adherence to that agreement.
Packet marking - Packet marking is related to packet classification. Packet marking allows you to classify a packet based on a specific traffic descriptor (such as the DSCP value). This classification can then be used to apply user-defined differentiated services to the packet and to

associate a packet with a local QoS group.
Associating a packet with a local QoS group allows users to associate a group ID with a packet. The group ID can be used to classify packets into QoS groups based on prefix, autonomous system, and community string. A user can set up to 64 DSCP values and 100 QoS group markings. Congestion management - Congestion management (or scheduling) is achieved through traffic scheduling and traffic queueing. When there is network congestion, a scheduling mechanism such as CBWFQ is used to provide guaranteed bandwidth to the different classes of traffic. Congestion avoidance - Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottlenecks. Congestion avoidance is achieved through packet dropping. Among the more commonly used congestion avoidance mechanisms is WRED. With WRED and Differentiated Services, you have the option of allowing WRED to use the DSCP value when WRED calculates the drop probability of a packet.

QUESTION NO: 413

Which two statements are true about traffic shaping? (Choose two.)

A. Shaping buffers excess packets.
B. It is applied in the incoming direction only.
C. Shaping can remark excess packets.
D. It supports interaction with frame relay congestion indication.

Answer: A,D
Explanation:
In contrast to policing, traffic shaping retains excess packets in a queue and then schedules the excess for later transmission over increments of time.
Traffic shaping can also interact with a Frame Relay network, adapting to indications of Layer 2 congestion in the WAN. For example, if the backward-explicit congestion notification (BECN) bit is received, the router can lower the rate limit to help reduce congestion in the Frame Relay network.
Reference: http://cisco.jjc.edu/cnt208/ch4/4_7_3/index.html

Topic 9, Troubleshoot a Network

QUESTION NO: 414

When you are troubleshooting duplex mismatches, which two errors are typically seen on the full-duplex end? (Choose two.)

A. runts
B. FCS errors
C. interface resets
D. late collisions

Answer: A,B

Explanation:

FCS, or File Check Sequence Errors, are one of the more common errors found in a network. When packets are transmitted and received, each contains a File Check Sequence that allows the receiving device to determine if the packet is complete without having to examine each bit. This is a type of CRC, or Cyclical Redundancy Check. Barring a station powering up or down during a transmission, the most common cause of these errors is noise. Network noise can be caused by cabling being located too close to noise sources such as lights, heavy machinery, etc. If a cabling installation is particularly faulty -- such as pairs being untwisted, improper terminations, field terminated patch cables, etc. -- these errors will occur on your network. Poorly manufactured components or minimally compliant components that are improperly installed can compound this issue. Cabling segments that are too long can also cause these errors.

Cabling issues, as defined above, or MAC layer packet formation issues (possibly hardware related) cause these errors. A faulty LAN driver can also cause this. Replacement of the driver will correct the latter issue. These errors can also be seen in correlation with RUNT packets or packets that are too short. Noise, however, is the most common cause and can generally be corrected by addressing the cabling channel.

QUESTION NO: 415

Which RMON group stores statistics for conversations between sets of two addresses?

A. hostTopN
B. matrix
C. statistics
D. history
E. packet capture
F. host

Answer: B

Explanation:

RMON tables can be created for buffer capture, filter, hosts, and matrix information. The buffer capture table details a list of packets captured off of a channel or a logical data or events stream. The filter table details a list of packet filter entries that screen packets for specified conditions as they travel between interfaces. The hosts table details a list of host entries. The matrix table details a list of traffic matrix entries indexed by source and destination MAC addresses.

QUESTION NO: 416

Which of the following describes the appropriate port assignment and message exchange in a standard TFTP transaction?

A. Server: 10.0.0.1:69 RRQ/WRQ Sent
Client: 10.0.0.2:1888 RRQ/WRQ Received
B. Server: 10.0.0.1:1888 RRQ/WRQ Received Client: 10.0.0.2:69 RRQ/WRQ Received
C. Server: 10.0.0.1:69 RRQ/WRQ Received Client: 10.0.0.2:69 RRQ/WRQ Sent
D. Server: 10.0.0.1:69 RRQ/WRQ Received Client: 10.0.0.2:1888 RRQ/WRQ Sent
E. Server: 10.0.0.1:1888 RRQ/WRQ Sent Client: 10.0.0.2:69 RRQ/WRQ Sent
F. Server: 10.0.0.1:1888 RRQ/WRQ Received Client: 10.0.0.2:69 RRQ/WRQ Sent

Answer: D
Explanation:
TFTP Daemons listen on UDP port 69 but respond from a dynamically allocated high port. Therefore, enabling this port will allow the TFTP service to receive incoming TFTP requests but will not allow the selected server to respond to those requests. Allowing the selected server to respond to inbound TFTP requests cannot be accomplished unless the TFTP server is configured to respond from port 69.

Reference
http://en.wikipedia.org/wiki/Trivial_File_Transfer_Protocol
http://social.technet.microsoft.com/Forums/en-CA/configmgrosd/thread/9b9bd9e2-6b2e-4073-96af-2703ad6a3249

QUESTION NO: 417

What two features in Cisco switches help prevent Layer 2 loops? (Choose two.)

A. Unidirectional Link Detection

B. Hot Standby Router Protocol
C. Virtual Router Redundancy Protocol
D. PortFast
E. root guard
F. loop guard

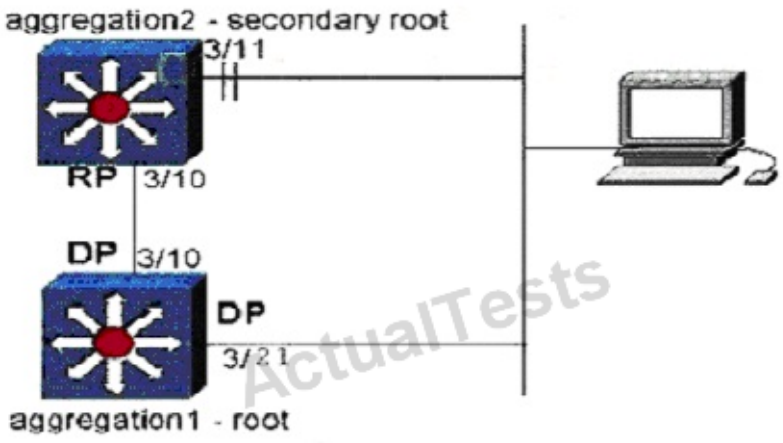
Answer: A,F

Explanation:

The STP loop guard feature provides additional protection against Layer 2 forwarding loops (STP loops). An STP loop is created when an STP blocking port in a redundant topology erroneously transitions to the forwarding state. This usually happens because one of the ports of a physically redundant topology (not necessarily the STP blocking port) no longer receives STP BPDUs. In its operation, STP relies on continuous reception or transmission of BPDUs based on the port role. The designated port transmits BPDUs, and the non-designated port receives BPDUs.

QUESTION NO: 418

Refer to the exhibit. Which switching feature is being tested?




```
aggregation 2 (enable) set spantree portfast 3/11 ena
Warning spantree port fast start should only be enabled on ports connected to a single
host
Connecting hubs, concentrators, switches, bridges,etc to a fast port can cause temporary
spanning tree loops. Use with caution
Spantree port 3/11 fast start enabled
aggregation-2 (enable) set spantree portfast bpdu-filter ena
Spantree portfast bpdu filter enabled on this switch
2001 Feb 06 13 32 14 %SPANTREE-4-LOOPGUARDBLOCK No BPDUs were
received on port 3/21 in VLAN 99 Moved to loop-inconsistent state
```

A. loop guard
B. PortFast
C. root guard
D. BDPU guard

Answer: A

Explanation:

In the exhibit, read the last lines. Loop guard is been tested.

QUESTION NO: 419

When troubleshooting a network, the output of the command show interfaces indicates a large number of runts. What is a runt?

A. the number of packets that are discarded because they exceed the maximum packet size of the medium
B. errors created when the CRC generated by the originating LAN station or far-end device does not match the checksum calculated from the data received.
C. the number of packets that are discarded because they are smaller than the minimum packet size of the medium
D. the number of received packets that were iqnored bv the interface because the interface hardware ran low on internal buffers
E. the number of times that the interface requested another interface within the router to slow down

Answer: C

Explanation:

In networks, a runt is a packet that is too small. For example, the Ethernet protocol requires that each packet be at least 64 bytes long. In Ethernet, which operates on the idea that two parties can attempt to get use of the line at the same time and sometimes do, runts are usually the fragments of packet collisions. Runts can also be the result of bad wiring or electrical interference. Runts are recorded by programs that use the Remote Network Monitoring (RNM) standard information base for network administration. RMON calls them "undersize packets".

A giant is a packet that's oversize.

QUESTION NO: 420

You deployed new fibers in your network to replace copper spans that were too long. While reconnecting the network, you experienced network problems because you reconnected wrong fibers to wrong ports. What could you do to prevent this type of problem in the future, particularly when connecting and reconnecting fiber pairs?

A. Only use fiber in pairs.
B. Configure root guard on your switches.
C. Do not use fiber but use copper.
D. Configure UDLD to prevent one-way link conditions.

Answer: D

Explanation:

UDLD is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. All connected devices must support UDLD for the protocol to successfully identify and disable unidirectional links. When UDLD detects a unidirectional link, it administratively shuts down the affected port and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree topology loops.
nce:
http://www.cisco.com/en/US/docs/switches/lan/catalyst3550/software/release/12.1_19_ea1/configuration/guide/swudld.html#wp1020819

QUESTION NO: 421

While deploying a new switch, you accidentally connect ports 3/12 and 3/18 together, creating a loop. STP detected it and placed port 3/18 in blocking mode. Why did STP not place port 3/12 in blocking mode instead?

A. Port 3/12 was already up and forwarding before the loop was created.
B. Port priority is based on lowest priority and lowest port number.
C. You connected the wire on port 3/18 last.
D. None of the above, it is purely random.

Answer: B

Explanation:

Spanning-tree select the port, by first calculating the cost (which depends on the bandwidth) then the port priority which is based on the lowest priority (in case configured) and the lowest port number.

QUESTION NO: 422

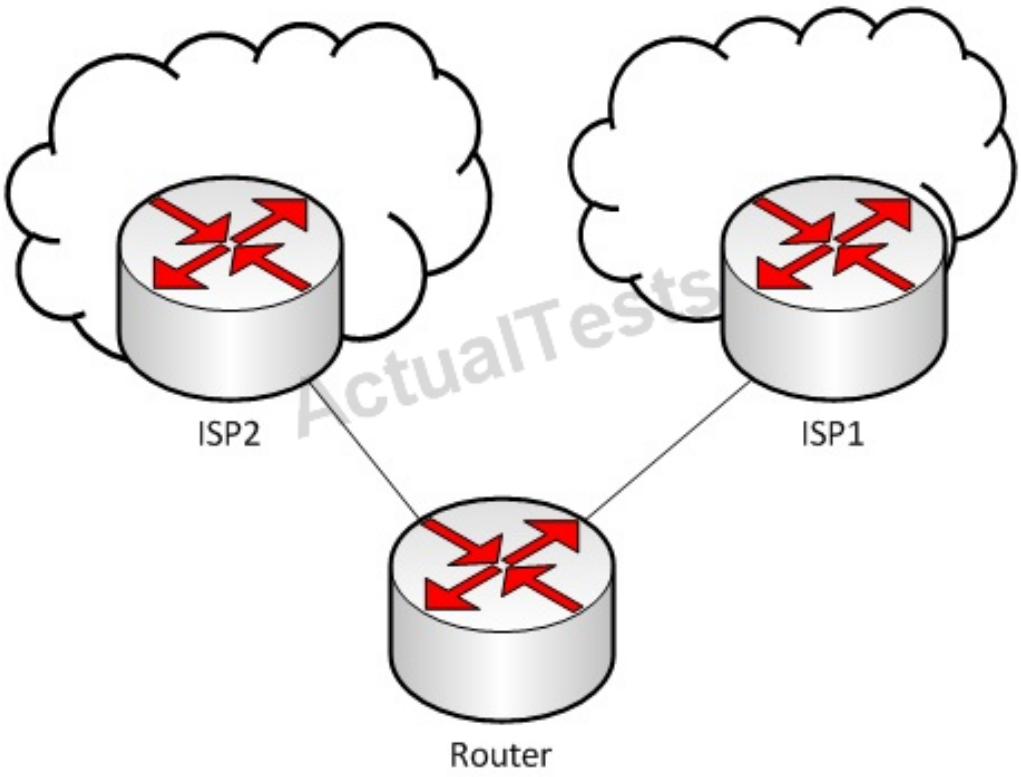
Which value should be used in the configuration register of a router in order to boot into bootstrap?

A. 0x2102
B. 0x2142
C. 0x2124
D. 0x2101

Answer: D
Reference:
http://www.cisco.com/en/US/products/hw/routers/ps133/products_tech_note09186a008022493f.shtml (configreg - see the table, 3rd row)

QUESTION NO: 423

Refer to the exhibit.



An enterprise network has an upstream connection to two different ISPs that are using eBGP and a publicly assigned PI network.

ISP1 is used as the primary provider, while ISP2 is used for backup. Due to packet loss on the link to ISP1, the network engineers changed the default route on R1 to point to ISP2, but they could not establish any connection until they changed the default route to point back at ISP1.

Which two are possible root causes of the issue? (Choose two.)

A. "ip verify unicast source reachable-via rx" is configured on the R1 uplinks.
B. "ip verify unicast source reachable-via any" is configured on the R1 uplinks.
C. "ip verify unicast source reachable-via rx" is configured on the ISP1 link to R1.
D. "ip verify unicast source reachable-via any" is configured on the ISP1 link to R1.
E. "ip verify unicast source reachable-via rx" is configured on the ISP2 link to R1.
F. "ip verify unicast source reachable-via any" is configured on the ISP2 link to R1.

Answer: A,E

Explanation:

ip verify unicast source reachable-via

To enable Unicast Reverse Path Forwarding (Unicast RPF), use the ip verify unicast source reachable-via command in interface configuration mode. To disable Unicast RPF, use the no form of this command.

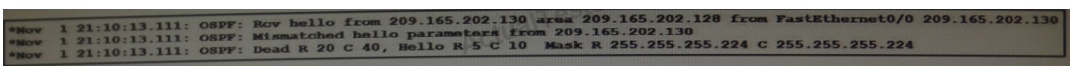
ip verify unicast source reachable-via {rx | any} [allow-default] [allow-self-ping] [list] no ip verify unicast source reachable-via

rx	Examines incoming packets to determine whether the source address is in the Forwarding Information Base (FIB) and permits the packet only if the source is reachable through the interface on which the packet was received (sometimes referred to as strict mode).
any	Examines incoming packets to determine whether the source address is in the FIB and permits the packet if the source is reachable through any interface (sometimes referred to as loose mode).
allow-default	(Optional) Allows the use of the default route for RPF verification.
allow-self-ping	(Optional) Allows a router to ping its own interface or interfaces. ⚠ **Caution** Use caution when enabling the **allow-self-ping** keyword. This keyword opens a denial-of-service (DoS) hole.
list	(Optional) Specifies a numbered access control list (ACL) in the following ranges: • 1 to 99 (IP standard access list) • 100 to 199 (IP extended access list) • 1300 to 2699 (IP standard access list, expanded range)

Reference
http://www.cisco.com/en/US/docs/ios/12_3/security/command/reference/sec_i2g.html#wp1103740

QUESTION NO: 424

Refer to the Exhibit.



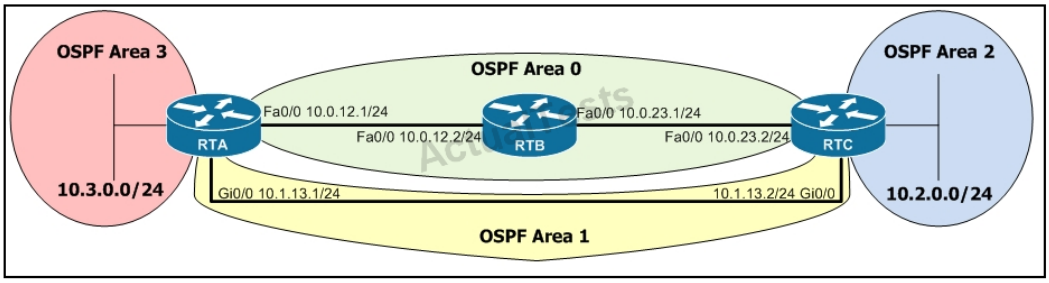
Which two problems does the debug ip ospf event command output from R4 indicate? (Choose Two)

A. mismatched masks between R4 and 209.165.202.130
B. mismatched dead timer between R4 and 209.165.202.130
C. mismatched hello timer between R4 and 209.165.202.130
D. mismatched areas between R4 and 209.165.202.130

Answer: B,C
Reference: http://packetlife.net/blog/2008/may/14/hello-timer-behaviors/

QUESTION NO: 425

Refer to the exhibit.



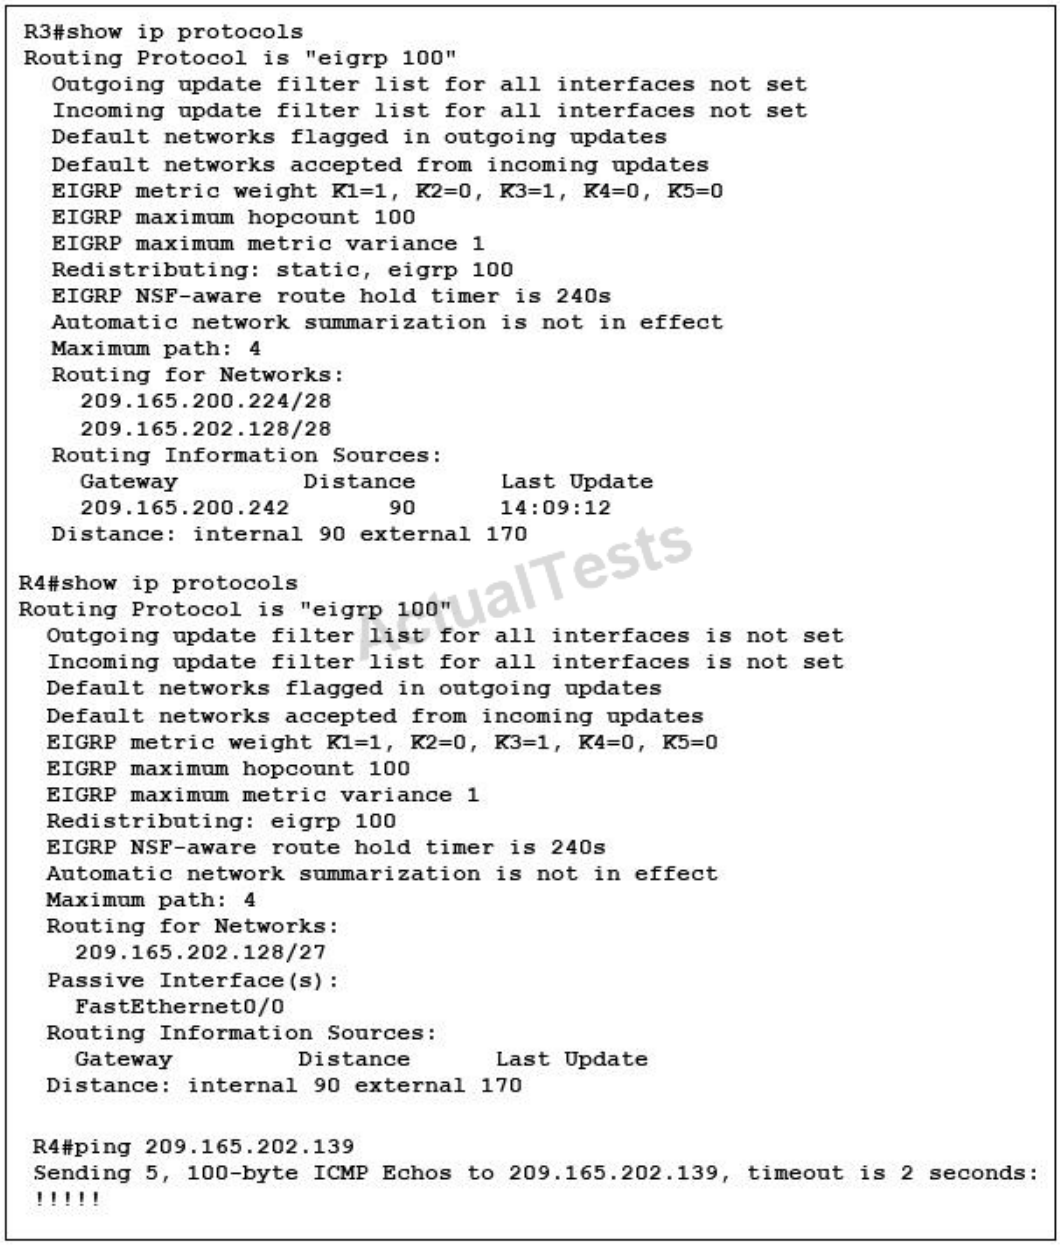
A new link has been provisioned between routers RTA and RTC. After all routing is configured and connectivity is established, you are facing an application performance issue between network 10.2.0.0/24 and 10.3.0.0/24. When you check the routing table of router RTC, you find out that packets toward network 10.3.3.0/24 are sent out via the slower path out of the Fa0/0 interface. What would you do to solve this issue?

A. change Gi0/0 on RTA to area 3 and Gi0/0 on RTC to area 2
B. change Gi0/0 on RTA to area 0 and Gi0/0 on RTC to area 2
C. change Gi0/0 on RTA to area 3 and Gi0/0 on RTC to area 0
D. change Gi0/0 on RTA to area 0 and Gi0/0 on RTC to area 0
E. change Gi0/0 on RTA to area 1 and Gi0/0 on RTC to area 1

Answer: D
Explanation:
To change the slower path, change Gi0/0 on RTA to area 0 and Gi0/0 on RTC to area 0. This way the path will be changed to the faster one.

QUESTION NO: 426

Refer to the exhibit.



```
R3#show ip protocols
Routing Protocol is "eigrp 100"
  Outgoing update filter list for all interfaces not set
  Incoming update filter list for all interfaces not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: static, eigrp 100
  EIGRP NSF-aware route hold timer is 240s
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    209.165.200.224/28
    209.165.202.128/28
  Routing Information Sources:
    Gateway         Distance      Last Update
    209.165.200.242       90      14:09:12
  Distance: internal 90 external 170

R4#show ip protocols
Routing Protocol is "eigrp 100"
  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=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 100
  EIGRP NSF-aware route hold timer is 240s
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    209.165.202.128/27
  Passive Interface(s):
    FastEthernet0/0
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: internal 90 external 170

 R4#ping 209.165.202.139
 Sending 5, 100-byte ICMP Echos to 209.165.202.139, timeout is 2 seconds:
 !!!!!
```

R4 is unable to establish an EIGRP adjacency with R3, the only other router on the Fa0/0 LAN segment, although it is able to ping R3. An EIGRP debug on R4 does not provide any clues. What might be the cause of the problem?

A. The passive interface has disabled the transmission of EIGRP multicast hello packets.
B. EIGRP on R4 should be routing to network 209.165.202.128/28.
C. The designated router/backup designated router (DR/BDR) selection requires that at least three

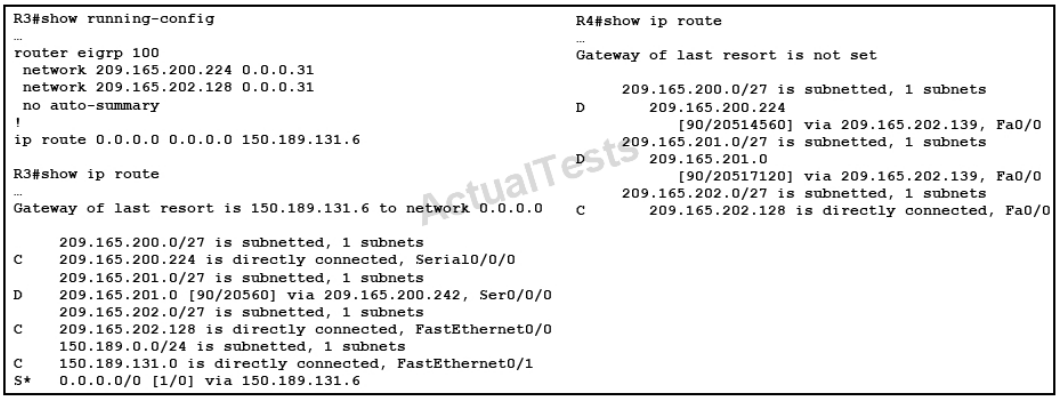
routers are on a LAN.
D. The routing metrics on R4 and R3 are different.

Answer: A
Explanation:
The passive-interface command disables the transmission and receipt of EIGRP hello packets on an interface. Unlike IGRP or RIP, EIGRP sends hello packets in order to form and sustain neighbor adjacencies. Without a neighbor adjacency, EIGRP cannot exchange routes with a neighbor.

Therefore, the passive-interface command prevents the exchange of routes on the interface. Although EIGRP does not send or receive routing updates on an interface configured with the passive-interface command, it still includes the address of the interface in routing updates sent out of other non-passive interfaces.

QUESTION NO: 427

Refer to the exhibit.



The static route to 150.189.131.6 on R3 is intended to serve as the gateway of last resort for the EIGRP network. However, while R3 installs the gateway of last resort, its EIGRP neighbor R4 does not. What might explain the problem?

A. R3 is missing the global command ip default-network 150.189.131.6.
B. Autosummary must be enabled on R3 in order for default-routes to be propagated.
C. Default-networks must be configured individually on each EIGRP router.
D. The static route on R3 must be redistributed into EIGRP.

Answer: D
Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800c2d96.shtml (see potential problems)

QUESTION NO: 428

Refer to the exhibit.

```
R4#show ip protocols
Routing Protocol is "eigrp 100"
  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=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 100
  EIGRP NSF-aware route hold timer is 240s
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    209.165.202.128/27
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: internal 90 external 170
```

```
R3#show ip protocols
Routing Protocol is "eigrp 100"
  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=255, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 100
  EIGRP NSF-aware route hold timer is 240s
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    209.165.200.224/28
    209.165.202.128/28
  Routing Information Sources:
    Gateway         Distance      Last Update
    209.165.200.242       90      00:01:26
  Distance: internal 90 external 170
```

Which statement would explain why R4 and R3 are unable to build an EIGRP adjacency?

A. The network masks on R4 (209.165.202.128/27) and R3 (209.165.202.128/28) are different.
B. The local EIGRP process on R4 and R3 are the same (but they must be unique).
C. The routing metrics on R4 and R3 are different.
D. R4 is not routing for the network 209.165.200.224/28.

Answer: C
Reference: http://packetlife.net/blog/2010/aug/9/eigrp-feasible-successor-routes/

QUESTION NO: 429

You are the network administrator of a Layer 3 switched network. Users in one VLAN are complaining that access to the server VLAN is very slow from time to time. Traffic in the local VLAN works without any issue, and users in other VLANs do not have any complaint to reach the server VLAN. What is most likely the cause of this issue?

A. routing issue
B. denial-of-service attack
C. MAC flooding
D. spanning-tree recalculation

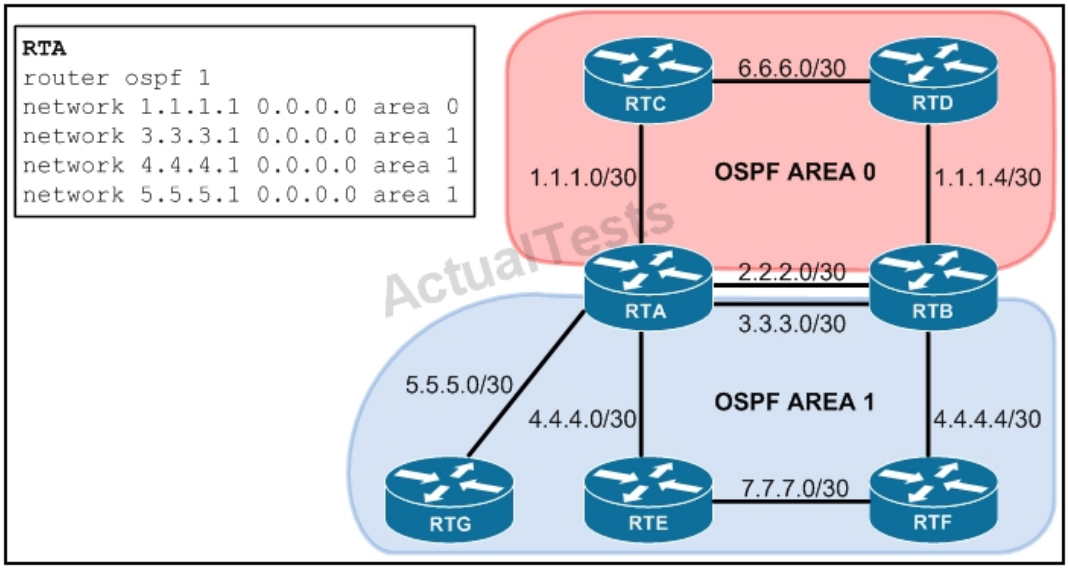
E. Layer 2 loop
F. ARP spoofing attack

Answer: F

Explanation: Reference
http://www.cisco.com/en/US/tech/tk389/tk689/technologies_tech_note09186a00801f9eb3.shtml

QUESTION NO: 430

Refer to the exhibit.



RTA and RTB are ABRs for Area 0 and Area 1.

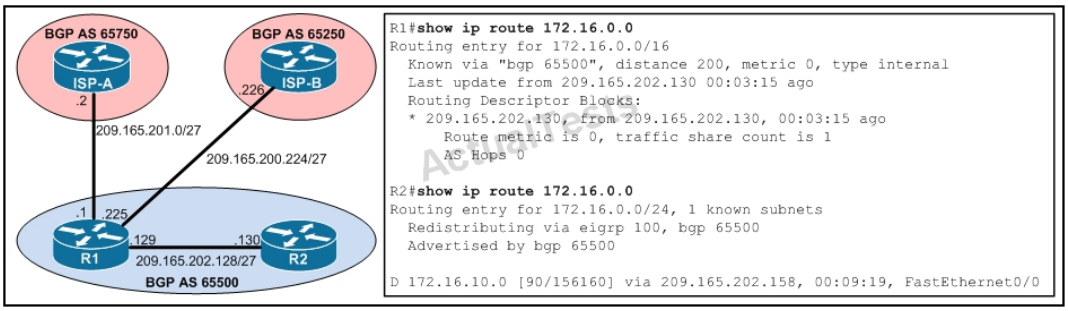
What is the issue with this configuration?

A. There is no issue because OSPF will work fine in any condition.
B. Suboptimal routing may occur because there is no Area 1 adjacency between the ABRs.
C. This is a wrong OSPF configuration because OSPF Area 0 is a must between ABRs.
D. If the link between RTC and RTD is down, an OSPF virtual link is required to connect ABR RTA and RTB.

Answer: D
Reference:
http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094e9e.shtml#t16

QUESTION NO: 431

Refer to the exhibit.



R1 is not learning about the 172.16.10.0 subnet from the BGP neighbor R2 (209.165.202.130). What should be done so that R1 will learn about this network?

A. Disable auto-summary on R2.
B. Configure an explicit network command for the 172.16.10.0 subnet on R2.
C. Subnet information cannot be passed between IBGP peers.
D. Disable auto-summary on R1.

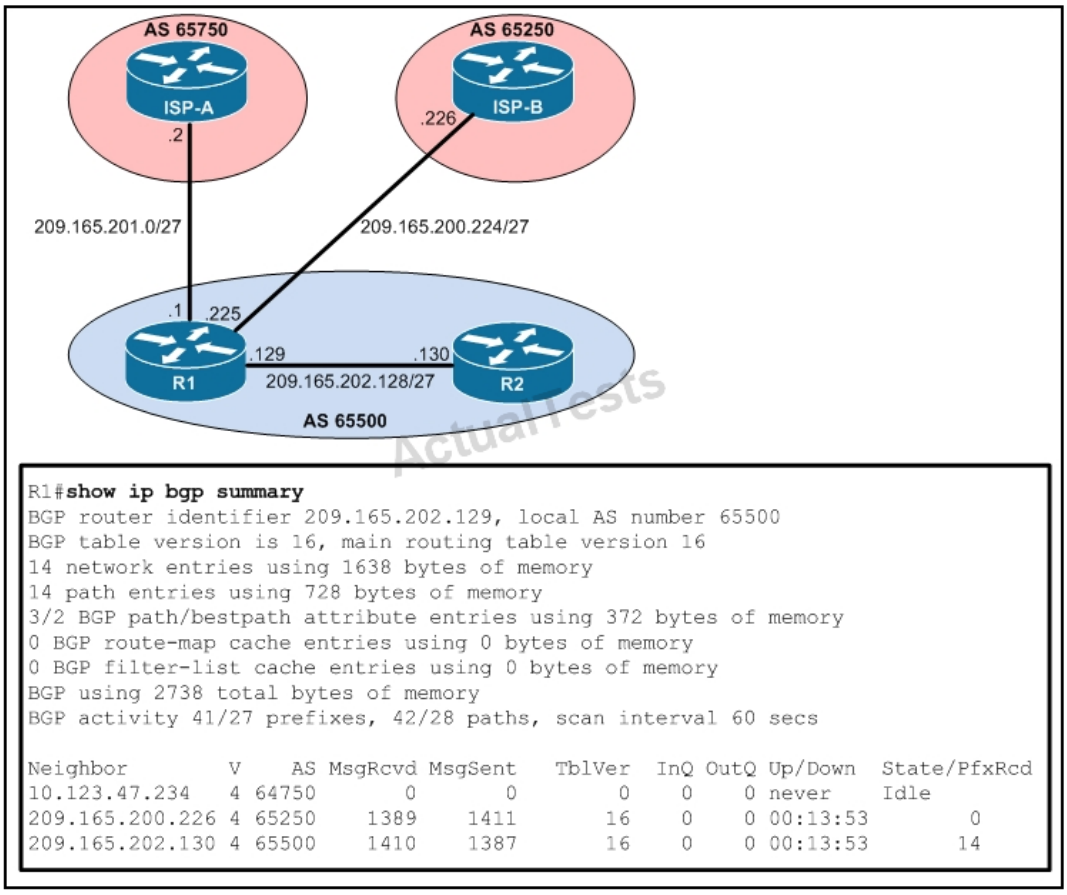
Answer: B

Explanation:

You need to configure explicit network command for 172.16.10.0 subnet on R2.

QUESTION NO: 432

Refer to the exhibit.




```
R1#show ip bgp summary
BGP router identifier 209.165.202.129, local AS number 65500
BGP table version is 16, main routing table version 16
14 network entries using 1638 bytes of memory
14 path entries using 728 bytes of memory
3/2 BGP path/bestpath attribute entries using 372 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 2738 total bytes of memory
BGP activity 41/27 prefixes, 42/28 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.123.47.234   4 64750       0       0        0    0    0 never    Idle
209.165.200.226 4 65250    1389    1411       16    0    0 00:13:53        0
209.165.202.130 4 65500    1410    1387       16    0    0 00:13:53       14
```

R1 is configured as an eBGP neighbor to ISP-A and ISP-B. While the adjacency between R1 and ISP-B is okay, R1 and ISP-A are not able to establish an adjacency.

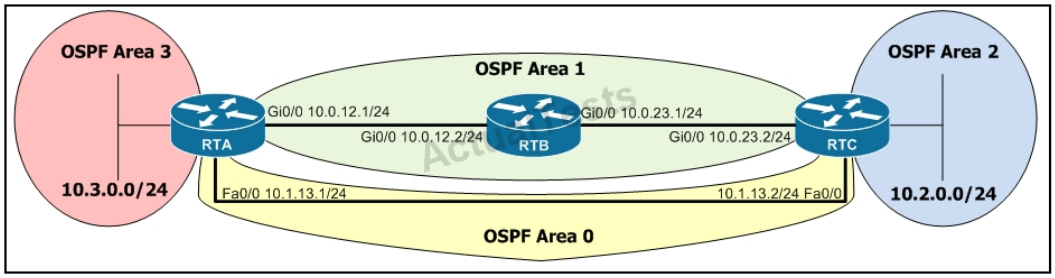
The ISP-A and ISP-B neighbor configurations on R1 are nearly identical. What is the problem between R1 and ISP-A?

A. R1 should be configured with the BGP AS 64750
B. R1 is missing the next-hop-self option under the neighbor command for ISP-A
C. R1 needs to add the ebgp-multihop option to the ISP-A neighbor command
D. R1 should redistribute static routes into the BGP process

Answer: C
Explanation:

QUESTION NO: 433

Refer to the exhibit.



You have performed multiple changes to your OSPF configuration. After these changes, you receive a lot of calls from users in OSPF area 2 complaining about application performance issues when they access servers connected to area 3. When you check the routing table of router RTC, you notice that packets to network 10.3.3.0/24 are sent out via the slower path out of the Fa0/0 interface. What can you do to solve this issue?

A. Change the OSPF configuration in OSPF area 1 to OSPF area 0
B. Change the OSPF configuration in OSPF area 0 to OSPF area 1
C. Change the OSPF configuration in OSPF area 2 to OSPF area 3
D. Change the OSPF configuration in OSPF area 3 to OSPF area 1

Answer: A

Explanation:

RTC will go through interface f0/1 to reach 10.3.0.1 as this interface is in area 0. See the below for additional information RTC

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.3.0.1, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 32/42/48 ms

RTC#traceroute 10.3.0.1

Type escape sequence to abort.

Tracing the route to 10.3.0.1

1 10.1.13.1 64 msec * 24 msec

RTC#sh ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

ia - IS-IS inter area, * - candidate default, U - per-user static route

o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C 10.1.13.0/24 is directly connected, FastEthernet0/1
O 10.0.12.0/24 [110/2] via 10.0.23.1, 00:00:35, FastEthernet0/0
O IA 10.3.0.1/32 [110/782] via 10.1.13.1, 00:01:50, FastEthernet0/1
C 10.2.0.0/24 is directly connected, Loopback0
C 10.0.23.0/24 is directly connected, FastEthernet0/0

RTC#config t
Enter configuration commands, one per line. End with CNTL/Z.
RTC(config)#router ospf 1
RTC(config-router)#network 10.0.23.0 0.0.0.255 area 0
RTC(config-router)#
*Mar 1 00:13:46.931: %OSPF-5-ADJCHG: Process 1, Nbr 10.0.23.1 on FastEthernet0/0 from FULL to DOWN,
Neighbor Down: Interface down or detached
RTC(config-router)#
*Mar 1 00:13:46.939: %OSPF-6-AREACHG: 10.0.23.0/24 changed from area 1 to area 0
RTC(config-router)#end
RTB
RTB#config t
Enter configuration commands, one per line. End with CNTL/Z.
RTB(config)#router ospf 1
RTB(config-router)#network 10.0.23.0 0.0.0.255 area 0
RTB(config-router)#
*Mar 1 00:14:38.103: %OSPF-6-AREACHG: 10.0.23.0/24 changed from area 1 to area 0
RTB(config-router)#network 10.0.12.0 0.0.0.255 area 0
RTB(config-router)#
*Mar 1 00:14:55.655: %OSPF-5-ADJCHG: Process 1, Nbr 10.3.0.1 on FastEthernet0/0 from FULL to DOWN,
Neighbor Down: Interface down or detached
RTB(config-router)#en
*Mar 1 00:14:55.663: %OSPF-6-AREACHG: 10.0.12.0/24 changed from area 1 to area 0
RTB(config-router)#end

RTA
RTA#config t
Enter configuration commands, one per line. End with CNTL/Z.
RTA(config)#
*Mar 1 00:16:02.219: %OSPF-4-ERRRCV: Received invalid packet: mismatch area ID, from backbone area must be virtual-link but not found from 10.0.12.2, FastEthernet0/0
RTA(config)#router ospf 1

*Mar 1 00:16:11.751: %OSPF-4-ERRRCV: Received invalid packet: mismatch area ID, from backbone area must be virtual-link but not found from 10.0.12.2, FastEthernet0/0

RTA(config-router)#network 10.0.12.0 0.0.0.255 area 0
RTA(config-router)#
*Mar 1 00:16:16.779: %OSPF-6-AREACHG: 10.0.12.0/24 changed from area 1 to area 0
RTA(config-router)#
*Mar 1 00:16:16.983: %OSPF-5-ADJCHG: Process 1, Nbr 10.0.23.1 on FastEthernet0/0 from LOADING to
FULL, Loading Done

RTC
RTC#clear ip ospf 1 proc
Reset OSPF process? [no]: yes
RTC#c
*Mar 1 00:18:00.267: %OSPF-5-ADJCHG: Process 1, Nbr 10.0.23.1 on FastEthernet0/0 from FULL to DOWN,
Neighbor Down: Interface down or detached
*Mar 1 00:18:00.271: %OSPF-5-ADJCHG: Process 1, Nbr 10.3.0.1 on FastEthernet0/1 from FULL to DOWN,
Neighbor Down: Interface down or detached
*Mar 1 00:18:00.491: %OSPF-5-ADJCHG: Process 1, Nbr 10.0.23.1 on FastEthernet0/0 from LOADING to

FULL, Loading Done
*Mar 1 00:18:00.495: %OSPF-5-ADJCHG: Process 1, Nbr 10.3.0.1 on FastEthernet0/1 from LOADING to
FULL, Loading Done
RTC#ping 10.3.0.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.0.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/64/76 ms
RTC#traceroute 10.3.0.1
Type escape sequence to abort.
Tracing the route to 10.3.0.1
1 10.0.23.1 32 msec 44 msec 28 msec
2 10.0.12.1 64 msec * 56 msec
RTC#show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks

C 10.1.13.0/24 is directly connected, FastEthernet0/1
O 10.0.12.0/24 [110/2] via 10.0.23.1, 00:00:43, FastEthernet0/0
O IA 10.3.0.1/32 [110/3] via 10.0.23.1, 00:00:43, FastEthernet0/0
C 10.2.0.0/24 is directly connected, Loopback0
C 10.0.23.0/24 is directly connected, FastEthernet0/0

QUESTION NO: 434

Refer to the exhibit.

```
R4#show ip ospf interface
FastEthernet0/0 is up, line protocol is up
  Internet Address 209.165.202.140/27, Area 209.165.202.128
  Process ID 100, Router ID 209.165.202.140, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 209.165.202.140, Interface address 209.165.202.140
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:01
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)

R4#show ip route
Codes: C - connected, D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2

Gateway of last resort is not set

     209.165.202.0/27 is subnetted, 1 subnets
C       209.165.202.128 is directly connected, FastEthernet0/0
```

What is preventing R4 from learning any networks via OSPF?

A. R4 does not have any OSPF neighbors.
B. There is no OSPF backup designated router on FastEthernet0/0.

C. The OSPF area number for FastEthernet0/0 does not match the interface network.
D. There are no OSPF learned routes in the network beyond R4.

Answer: A

Explanation: Explanation;
As you can see in the exhibit, no OSPF neighbors are configured so R4 will not be able to learn about networks via OSPF. There is no DR elected router so this is the only OSPF speaker on the Ethernet segment.

QUESTION NO: 435

Refer to the exhibit.

```
*Nov  1 20:37:50.943: OSPF: Rcv hello from 209.165.202.130 area 209.165.202.128 from FastEthernet0/0 209.165.202.130
*Nov  1 20:37:50.943: OSPF: Mismatched hello parameters from 209.165.202.130
*Nov  1 20:37:50.943: OSPF: Dead R 40 C 40, Hello R 10 C 10  Mask R 255.255.255.240 C 255.255.255.224
```

What problem does the debug ip ospf event output from R4 indicate?

A. a mismatched dead timer between R4 and 209.165.202.130
B. a mismatched hello timer between R4 and 209.165.202.130
C. mismatched areas between R4 and 209.165.202.130
D. mismatched masks between R4 and 209.165.202.130

Answer: D

Explanation: Explanation
The last line of the logging output shows that this is an issue with the subnet mask

QUESTION NO: 436

On a Cisco router that is in ROMMON mode, how can you set the configuration register to its default value?

A. set confreg 0x2102
B. confreg 0x2102
C. config-register 0x2102
D. set config-register 0x2102

Answer: B
Reference:
http://www.cisco.com/en/US/products/hw/routers/ps133/products_tech_note09186a008022493f.shtml (configreg)

QUESTION NO: 437

Refer to the exhibit.

```
Router#show int gig0/0
GigabitEthernet0/0 is up, line protocol is up
  Hardware is AmdP2, address is aabb.cc00.6400 (bia aabb.cc00.6400)
  Internet address is 10.1.1.1/30
  MTU 9000 bytes, BW 10000 Kbit/sec, DLY 1000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  ARP type: ARPA, ARP Timeout 04:00:00
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     19 packets input, 3041 bytes, 0 no buffer
     Received 9 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     39 packets output, 4653 bytes, 0 underruns
     0 output errors, 0 collisions, 1 interface resets
     0 output buffer failures, 0 output buffers swapped out

Router#show ip ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.1.2          0   EXSTART/  -     00:00:38    10.1.1.101      GigabitEthernet0/0
```

A new Ethernet link was added to an existing OSPF network. It is possible to ping the remote end of the link, but the OSPF adjacency does not form. Which two commands configured on the interface could resolve the issue? (Choose two.)

A. ip ospf mtu-ignore
B. ip ospf flood-reduction
C. ip mtu 1500
D. ip mss 1500

Answer: A,C
Explanation: Explanation
ip ospf mtu-ignore
To disable OSPF MTU mismatch detection on receiving DBD packets, use the ip ospf mtu-ignore command in interface configuration mode. To reset to default, use the no form of this command.

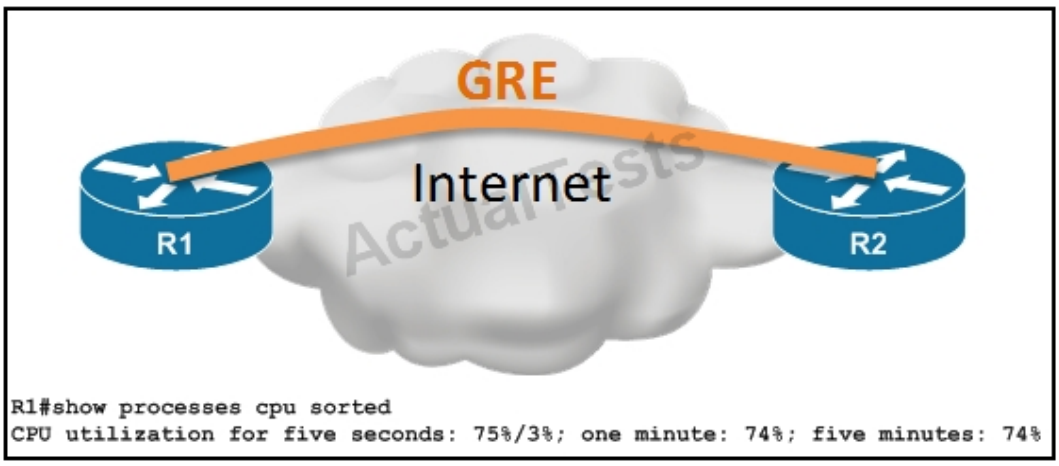
ip ospf mtu-ignore

no ip ospf mtu-ignore

Reference
http://www.cisco.com/en/US/docs/ios/12_3t/ip_route/command/reference/ip2_i1gt.html#wp1109168

QUESTION NO: 438

Refer to the exhibit.



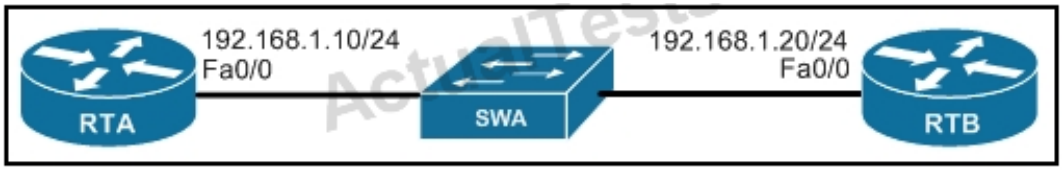
A GRE tunnel between R1 and R2 has been configured over an Internet connection. As soon as traffic is sent over the tunnel, high CPU utilization is observed on R1. Reducing the traffic shows a reduction in CPU utilization. What could be a possible cause for the high CPU utilization?

A. The BGP process is processing the complete Internet routing table.
B. An MTU issue causes GRE packets to be fragmented.
C. A GRE-based DDoS attack is overloading the router control plane.
D. All GRE packets are always process switched.

Answer: B
rence: http://blog.ine.com/2008/11/05/dealing-with-fragmented-traffic/

QUESTION NO: 439

Refer to the exhibit.



You are investigating connectivity issues on your LAN and have noticed that sometimes you cannot use Telnet from RTA to RTB. Sometimes, Telnet sessions will disconnect on RTA with the message "connection reset by peer." However, you can ping from RTA to RTB at all times. What is most likely the cause of this issue?

A. Quality of service is configured on RTA, limiting the Telnet sessions.
B. Quality of service is configured on RTB, limiting the Telnet sessions.
C. Control Plane Policing is configured on RTB, limiting the Telnet sessions.
D. On router RTB, the "login block" mode went to a "quiet period" for the IP address of RTA.
E. Another device that is connected to the switch is using the same IP address as RTB.
F. Another device that is connected to the switch is using the same MAC address as RTB.

Answer: E
Explanation:
When another device connected to the switch has the same IP address as the device from which you are pinging, the ping will go through but you will not be able to use any network services such as Telnet because the same IP address is been used by another device.

Answer: D
Reference:
https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=9&ved=0CGkQFjAI&url=http%3A%2F%2Fwww.cs.cmu.edu%2F~srini%2F15-441%2FS05%2Flectures%2F11-BGP.ppt&ei=YCbMUYKIFIfOhAeSg4Ew&usg=AFQjCNE2ioClS5UbKXhqIdnDQUt_9yNaUQ&sig2=THFuUNjA-RZEEFIsI97xHw (slide 31, 32, and 33)

QUESTION NO: 440

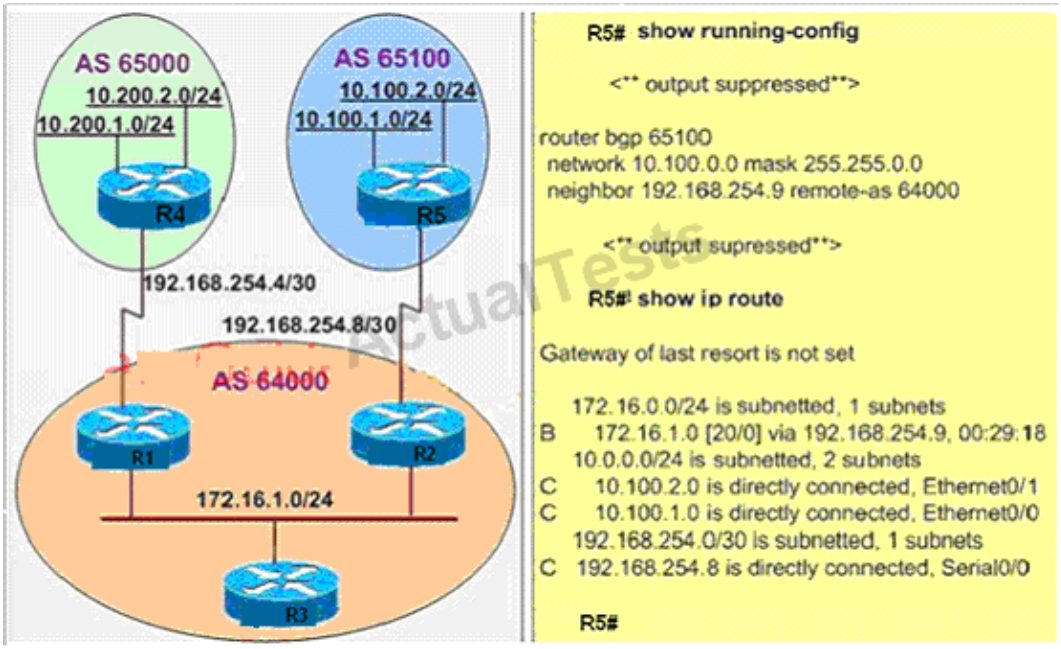
Which value should be used in a router configuration register in order to ignore the content of the NVRAM?

A. 0x2102
B. 0x2124
C. 0x2142
D. 0x2101

Answer: C
Reference:
http://books.google.com.pk/books?id=xbfe_tu9DI4C&pg=PA438&lpg=PA438&dq=Which+value+should+be+used+in+a+router+configuration+register+in+order+to+ignore+the+content+of+the+NVRAM&source=bl&ots=0KkTj9rqEk&sig=UF6KGI7ssMpv19ATDoPOYmFEh30&hl=en&sa=X&ei=WSfMUeP3BYaChQfzjIFI&redir_esc=y#v=onepage&q=Which%20value%20should%20be%20used%20in%20a%20router%20configuration%20register%20in%20order%20to%20ignore%20the%20content%20of%20the%20NVRAM&f=false (see password recovery procedure, second para)

QUESTION NO: 441

Based on the exhibit presented. R2 does not have any 10.100.x.x routes in either its routing table or its BGP table. What will you do at R5 to solve this problem?

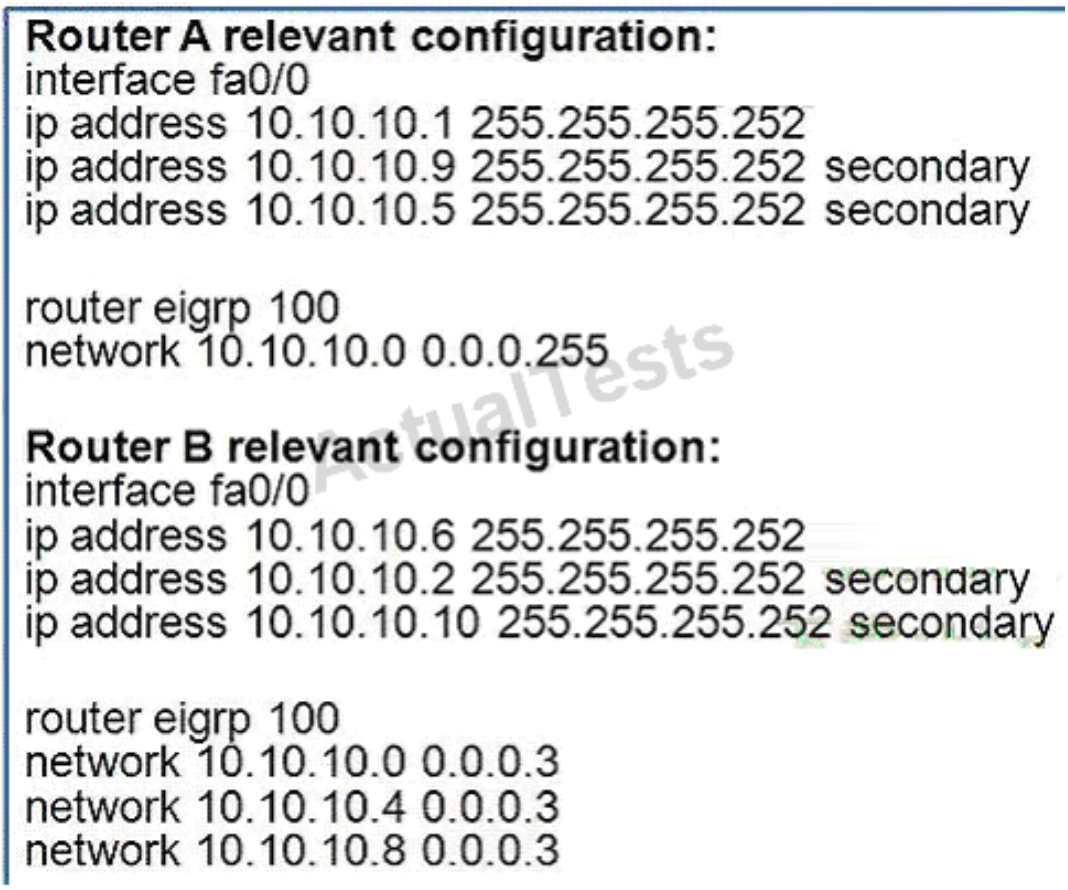


A. Disable BGP synchronization.
B. Set the BGP next-hop-self command for neighbor R2
C. Configure a static route for 10.100.0.0/16 to null0
D. Add a BGP network statement to encompass the serial link.

Answer: C
Explanation:
The null 0 interface means disregard the packet. So if I get the packet and there is a more specific match than 175.220.0.0 (which exists of course) the router will send it to the specific match otherwise it will disregard it. This is a nice way to advertise a super net.

QUESTION NO: 442

Refer to the exhibit. Routers A and B are directly connected and running EIGRP, but they are unable to form a neighbor relationship. What is the most likely cause?



A. The network statements are misconfigured.
B. The IP address statements are misconfigured
C. The autonomous system is misconfigured.
D. There is a physical issue with the cable.

Answer: B

Explanation:

To form neighbor relationship in EIGRP, these conditions must be met:

Pass the authentication process

Have the same con figured AS number

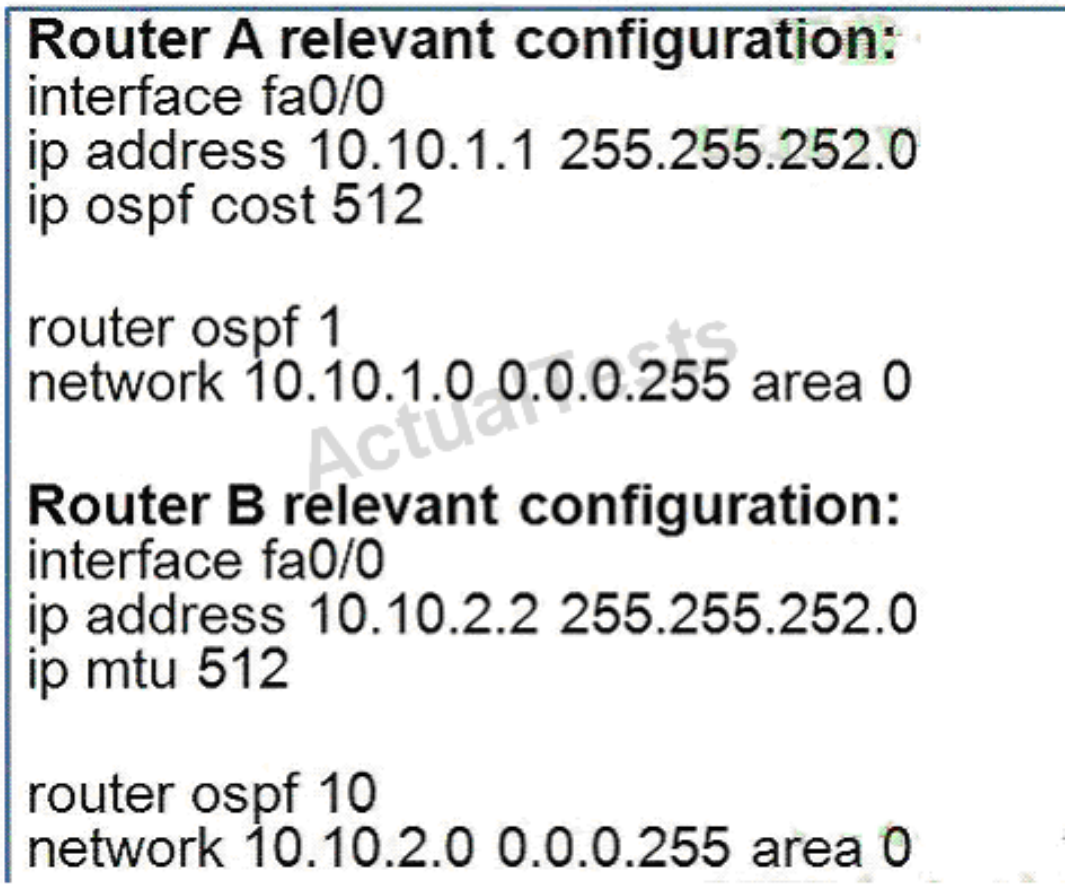
Must believe that the source IP address of a received Hello is in that router's primary connected subnet on that interface

Match K values

The third item means that the primary ip address of the neighbor must be in the same subnet with the primary ip address of the received interface. But in this case the primary ip address of router A is 10.10.10.1/30 and it is not in the same subnet with the primary ip address of router B 10.10.10.6/30 -> no EIGRP neighbor relationship is formed.

QUESTION NO: 443

Refer to the exhibit. Routers A and B are directly connected and running OSPF, but they are unable to form a neighbor relationship. What is the most likely cause?



A. The routers are not on the same network.
B. The network statements do not match.
C. The process number does not match.
D. The MTU does not match.
E. The OSPF cost does not match.
F. There is a physical issue with the cable.

Answer: D

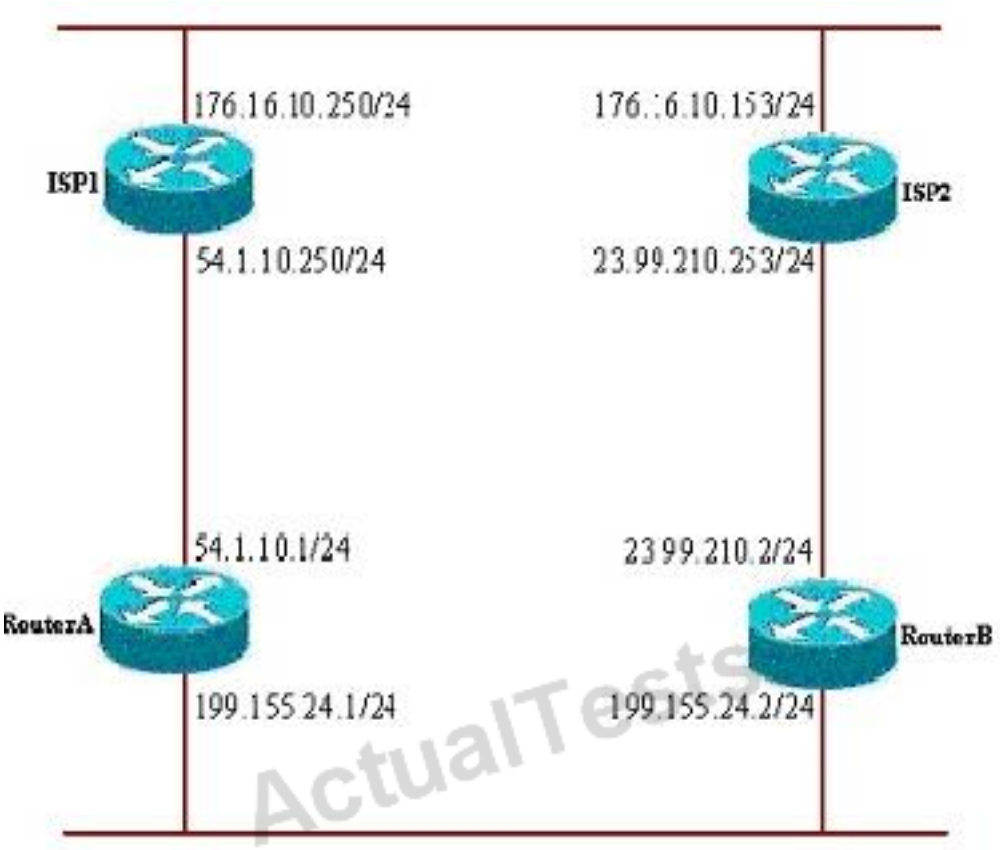
Explanation:

OSPF sends the interface MTU in a database description packet. If there is a MTU mismatch, OSPF will not form an adjacency and they are stuck in exstart/exchange state. The interface MTU option was added in RFC
2178. Previously, there was no mechanism to detect the interface MTU mismatch. This option was added in Cisco IOS Software Release 12.0.3 and later.
If the router with the higher MTU sends a packet larger that the MTU set on the neighboring router, the neighboring router ignores the packet and the neighbor state remains in exstart.
Note: By default, the MTU for Ethernet is 1500 bytes. We can check the OSPF adjacency process with the command "show ip ospf neighbor".

Reference
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0d.shtml

QUESTION NO: 444

Refer to the exhibit. Users on the 199.155.24.0 network are unable to reach the 172.16.10.0 network. What is the most likely solution?



Router A relevent configuration
router bgp 200
neighbor 54.1.10.250 remote-as 100
network 54.1.10.0 mask 255.255.255.0
network 199.155.24.0 mask 255.255.255.0

Router B relevent configuration
router bgp 200
neighbor 23.99.210.253 remote-as 300
network 23.99.210.0 mask 255.255.255.0
network 199.155.24.0 mask 255.255.255.0

Router ISP1 relevent configuration
router bgp 100
neighbor 54.1.10.1 remote-as 200
network 54.1.10.0 mask 255.255.255.0

Router ISP2 relevent configuration
router bgp 300
neighbor 23.99.210.2 remote-as 200
network 23.99.210.0 mask 255.255.255.0

A. Router ISP1 should be configured to peer with router B.
B. Router ISP2 should be configured with no synchronization.
C. Router ISP1 should be configured with no synchronization.
D. Router ISP2 should be configured with no auto-summary.
E. Router ISP1 or IPS2 should be configured with network 176.16.10.0 mask 255.255.255.0.

Answer: E

Explanation:

Neither ISP1 or ISP2 are advertising the 172.16.10.0/24 network therefore neither RouterA or RouterB are aware of how to get to these networks.

QUESTION NO: 445

Two BGP peers connected through a routed firewall are unable to establish a peering relationship. What could be the most likely cause?

A. BGP peers must be Layer 2-adjacent.
B. EBGP multihop is not configured.
C. The firewall is not configured to allow IP protocol 89.
D. The firewall is not configured to allow UDP 179.

Answer: B
Explanation:
Routed Mode Overview
In routed mode, the security appliance is considered to be a router hop in the network. It can perform NAT between connected networks, and can use OSPF or RIP (in single context mode). Routed mode supports many interfaces. Each interface is on a different subnet. You can share interfaces between contexts.
This section includes the following topics:
IP Routing Support
Network Address Translation
How Data Moves Through the Security Appliance in Routed Firewall Mode

IP Routing Support
The security appliance acts as a router between connected networks, and each interface requires an IP address on a different subnet. In single context mode, the routed firewall supports OSPF and RIP. Multiple context mode supports static routes only. We recommend using the advanced routing capabilities of the upstream and downstream routers instead of relying on the security appliance for extensive routing needs.

Passing Traffic Not Allowed in Routed Mode
In routed mode, some types of traffic cannot pass through the security appliance even if you allow it in an access list. The transparent firewall, however, can allow almost any traffic through using either an extended access list (for IP traffic) or an EtherType access list (for non-IP traffic).

Note
The transparent mode security appliance does not pass CDP packets or IPv6 packets, or any

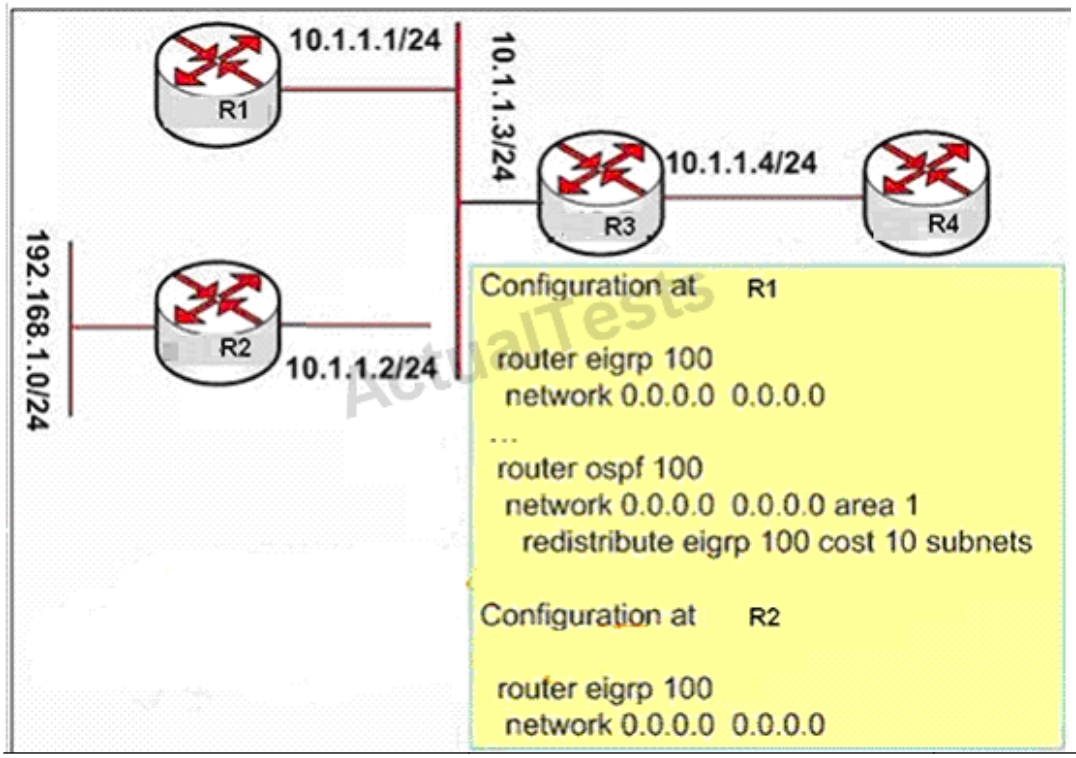
packets that do not have a valid EtherType greater than or equal to 0x600. For example, you cannot pass IS-IS packets. An exception is made for BPDUs, which are supported.
For example, you can establish routing protocol adjacencies through a transparent firewall; you can allow OSPF, RIP, EIGRP, or BGP traffic through based on an extended access list. Likewise, protocols like HSRP or VRRP can pass through the security appliance.

Non-IP traffic (for example AppleTalk, IPX, BPDUs, and MPLS) can be configured to go through using an EtherType access list.
For features that are not directly supported on the transparent firewall, you can allow traffic to pass through so that upstream and downstream routers can support the functionality. For example, by using an extended access list, you can allow DHCP traffic (instead of the unsupported DHCP relay feature) or multicast traffic such as that created by IP/TV.

Reference
http://www.cisco.com/en/US/docs/security/asa/asa72/configuration/guide/fwmode.html#wp1201691

QUESTION NO: 446

On the basis of the network provided in the exhibit, R3 and R4 are configured to run all connected links in OSPF Area 1. The network administrator is complaining that traffic destined to 192.168.1.0/24 is being routed to R2, even if R2 is not running OSPF. Which would be the cause of this problem?



A. The next hop towards 192.168.1.0/24 at R3 should be 10.1.1.2, which is R2.
B. The next hop towards 192.168.1.0/24 at R4 should be 10.1.1.1, since R1 is redistributing the route from EIGRP into OSPF. R3 is forwarding traffic incorrectly.
C. The next hop towards 192.168.1.0/24 at R4 should be 10.1.2.2, which is R3. R3 should be load-sharing between R1 and R2 for its next hop.
D. R4 does not have a route towards 192.168.1.0/24, so the network administrator is wrong in thinking any traffic is being forwarded there.

Answer: A

Explanation:

when R1 redistributes the the eigrp route into ospf , the external type 5 LSA forwarding address = 10.1.1.2, and so R3 will see it as the next hop to the network 192.168.1.0 /24.

QUESTION NO: 447

This question is about the formation of OSPF adjacency. An OSPF adjacency will not form correctly across a point-to-point link in the same area. Which would most likely cause this problem?

A. Each interface has a different OSPF cost.
B. Each interface is configured with secondary addresses as well as primary addresses.
C. Each interface has a different MTU size.
D. Each interface is configured with the ip unnumbered loopback 0 command.

Answer: C

Explanation:

Unequal MTU means stuck in EX-START

The states are Down, Attempt, Init, 2-Way, Exstart, Exchange, Loading, and Full.

Down This is the first OSPF neighbor state. It means that no information (hellos) has been received from this neighbor, but hello packets can still be sent to the neighbor in this state.

During the fully adjacent neighbor state, if a router doesn't receive hello packet from a neighbor within the RouterDeadInterval time (RouterDeadInterval = 4*HelloInterval by default) or if the manually configured neighbor is being removed from the configuration, then the neighbor state changes from Full to Down.

Attempt

This state is only valid for manually configured neighbors in an NBMA environment. In Attempt state, the router sends unicast hello packets every poll interval to the neighbor, from which hellos have not been received within the dead interval.

Init
This state specifies that the router has received a hello packet from its neighbor, but the receiving router's ID was not included in the hello packet. When a router receives a hello packet from a neighbor, it should list the sender's router ID in its hello packet as an acknowledgment that it received a valid hello packet.

2-Way
This state designates that bi-directional communication has been established between two routers. Bidirectional means that each router has seen the other's hello packet. This state is attained when the router receiving the hello packet sees its own Router ID within the received hello packet's neighbor field. At this state, a router decides whether to become adjacent with this neighbor. On broadcast media and non-broadcast multiaccess networks, a router becomes full only with the designated router (DR) and the backup designated router (BDR); it stays in the 2-way state with all other neighbors. On Point-to-point and Point-to-multipoint networks, a router becomes full with all connected routers.

At the end of this stage, the DR and BDR for broadcast and non-broadcast multiacess networks are elected.

For more information on the DR election process, refer to DR Election.
Note: Receiving a Database Descriptor (DBD) packet from a neighbor in the init state will also a cause a transition to 2-way state.

Exstart
Once the DR and BDR are elected, the actual process of exchanging link state information can start between the routers and their DR and BDR.
In this state, the routers and their DR and BDR establish a master-slave relationship and choose the initial sequence number for adjacency formation. The router with the higher router ID becomes the master and starts the exchange, and as such, is the only router that can increment the sequence number. Note that one would logically conclude that the DR/BDR with the highest router ID will become the master during this process of master-slave relation. Remember that the DR/BDR election might be purely by virtue of a higher priority configured on the router instead of highest router ID. Thus, it is possible that a DR plays the role of slave. And also note that master/slave election is on a per-neighbor basis.

Exchange
In the exchange state, OSPF routers exchange database descriptor (DBD) packets. Database descriptors contain link-state advertisement (LSA) headers only and describe the contents of the entire link-state database. Each DBD packet has a sequence number which can be incremented only by master which is explicitly acknowledged by slave. Routers also send link-state request packets and link-state update packets (which contain the entire LSA) in this state. The contents of

the DBD received are compared to the information contained in the routers link-state database to check if new or more current link-state information is available with the neighbor.

Loading
In this state, the actual exchange of link state information occurs. Based on the information provided by the DBDs, routers send link-state request packets. The neighbor then provides the requested link-state information in link-state update packets. During the adjacency, if a router receives an outdated or missing LSA, it requests that LSA by sending a link-state request packet. All link-state update packets are acknowledged.

Full
In this state, routers are fully adjacent with each other. All the router and network LSAs are exchanged and the routers' databases are fully synchronized.
Full is the normal state for an OSPF router. If a router is stuck in another state, it's an indication that there are problems in forming adjacencies. The only exception to this is the 2-way state, which is normal in a broadcast network. Routers achieve the full state with their DR and BDR only. Neighbors always see each other as 2-way.

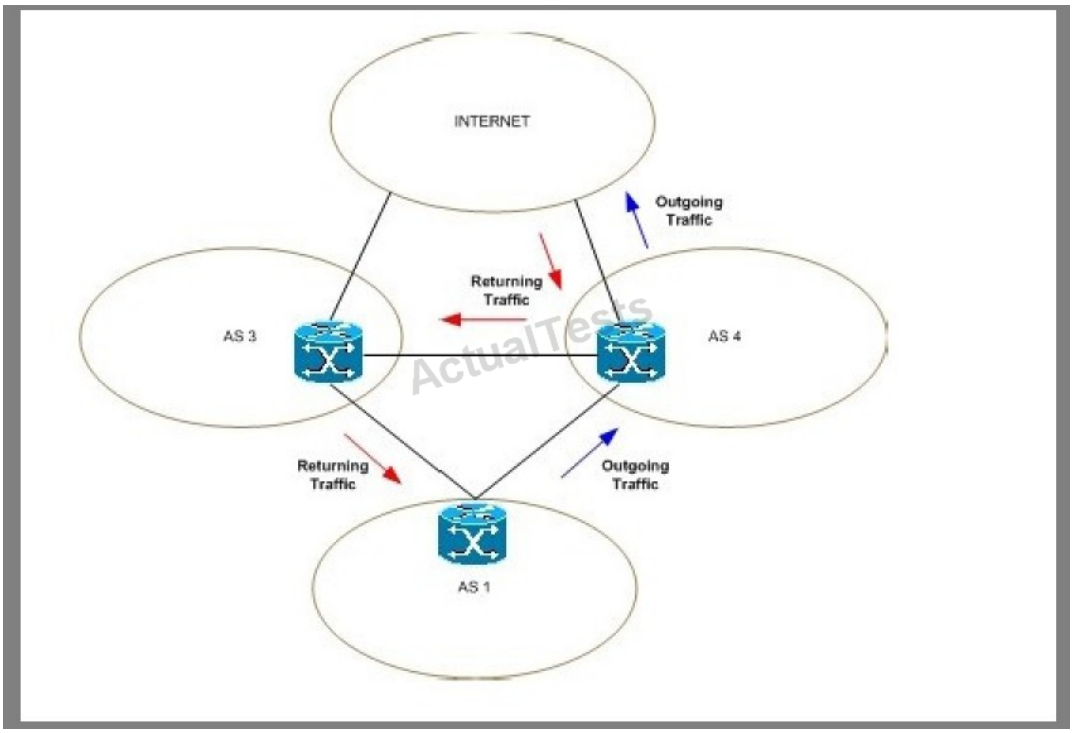
Troubleshooting OSPF Neighbor Relationships
OSPF Neighbor List is Empty
OSPF not enabled properly on appropriate interfaces
Layer 1 or 2 not functional
Passive interface configured
Access list(s) blocking Hello packets in multiple directions
Error in IP address or subnet mask configuration
Hello or dead interval mismatch
Authentication configuration error
Area ID mismatch
Stub flag mismatch
OSPF adjacency exists with secondary IP addressing or asynchronous interface
Incorrect configuration type for nonbroadcast multiaccess (NBMA) environment
OSPF Neighbor Stuck in Attempt
Misconfigured neighbor statement
Unicast nonfunctional in NBMA environment
OSPF Neighbor Stuck in init
Access list or Layer 2 problem blocking Hellos in one direction
Multicast nonfunctional on one side
Authentication configured on only one side
Broadcast keyword missing from the map command
OSPF Neighbor Stuck in Two-Way
Priority 0 configured on all routers

OSPF Neighbor Stuck in Exstart/Exchange
Mismatch interface maximum transmission unit (MTU)
Duplicate router IDs on routers
Broken unicast connectivity
Network type of point-to-point between Primary Rate Interface (PRI) and Basic Rate Interface (BRI)/dialer
OSPF Neighbor Stuck in Loading
Mismatched MTU
Corrupted link-state request packet

Reference
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0e.shtml
Cisco General Networking Theory Quick Reference Sheets

QUESTION NO: 448

Refer to the exhibit. BGP-4 routing to the Internet, in normal behavior, may create asymmetrical routing for different prefixes. The BGP routing table indicates that traffic should follow the paths indicated in the exhibit, but packets are not going further than the border router in AS 4. What could be the cause of this problem?



A. TCP Intercept is configured in AS 4.

B. Unicast Reverse Path Forwarding is configured in loose mode in this router.
C. Packets may be leaving AS 1 without the BGP routing flag set to 1.
D. Unicast Reverse Path Forwarding is configured in strict mode in this router.
E. There is a missing Unicast Reverse Path Forwarding configuration.

Answer: D
Explanation: When administrators use Unicast RPF in strict mode, the packet must be received on the interface that the router would use to forward the return packet. Unicast RPF configured in strict mode may drop legitimate traffic that is received on an interface that was not the router's choice for sending return traffic.

Dropping this legitimate traffic could occur when asymmetric routing paths are present in the network.

Reference
http://www.cisco.com/web/about/security/intelligence/unicast-rpf.html

QUESTION NO: 449

You replaced your Layer 3 switch, which is the default gateway of the end users. Many users cannot access anything now, including email, Internet, and other applications, although other users do not have any issues. All of the applications are hosted in an outsourced data center. In order to fix the problem, which one of these actions should you take?

A. Clear the MAC address table in the switch.
B. Clear the ARP cache in the switch.
C. Clear the ARP cache in the end devices.
D. Clear the ARP cache in the application servers.

Answer: C
Explanation:
Each workstation has its own arp cache. To delete the arp cache on a windows desktop do the following:

1. Open the "Command Prompt," the Windows application that enables running Windows commands and software applications, by clicking the Windows "Start" button, clicking "Programs," clicking "Accessories" and then clicking "Command Prompt."
2. Type "netsh interface ip delete arpcache" in the Command Prompt to clear your ARP cache.
3. Type "arp -a" in the Command Prompt to verify that the ARP cache was cleared. The output of this command should be "No ARP Entries Found."

Reference
http://www.ehow.com/how_5097488_clear-arp-cache.html#ixzz1lTxv4dRg

QUESTION NO: 450

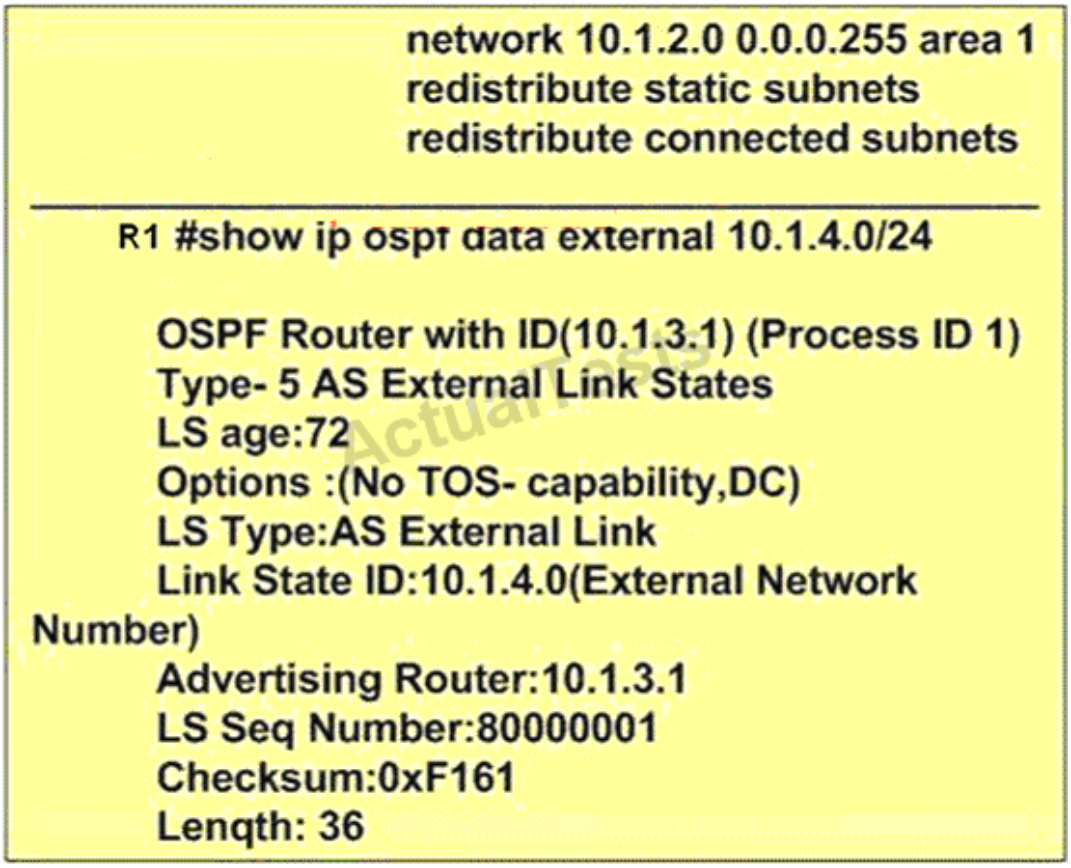
Half of your network uses RIPv2 and the other half runs OSPF. The networks do not communicate with each other. Which two of these factors describe the impact of activating EIGRP over each separate part? (Choose two.)

A. EIGRP will not be accepted when configured on the actual RIPv2 routers.
B. OSPF will no longer be used in the routing table, because you only have EIGRP internal routes running.
C. OSPF will no longer be used in the routing table, because you only have EIGRP external routes running.
D. RIPv2 will populate its RIP database but not its routing table, because you only have EIGRP external routes running.
E. RIPv2 will populate its RIP database but not its routing table, because you only have EIGRP internal routes running.
F. OSPF database will have RIPv2 routes.

Answer: B,E
Reference:
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800943c5.shtml

QUESTION NO: 451

Look at the following exhibit carefully; there is no route to 10.1.4.0/24 in the local routing table. According to the output of R1 in the exhibit, can you tell me why 10.1.4.0/24 is not in R1's routing table?



A. The forwarding address, 10.1.3.2, is also redistributed into OSPF, and an OSPF external route cannot use another OSPF external as its next hop
B. R3 is not redistributing 10.1.4.0/24 properly.
C. R2 is not properly configured as an Area Border Router
D. Area 1 is a stub area, and external routes cannot be originated in a stub area.

Answer: A

Explanation:

Network Missing from the Routing Table
R2515 has a RIP (R) derived route for network 200.1.1.0/24. R2515 is the ASBR and redistributes the RIP protocol into OSPF. R2504 learns about network 200.1.1.0/24 from R2515 and installs it in its routing table as an OSPF external type 2 (E2) route. The problem is that R2507 does not have network 200.1.1.0/24 in its routing table.

R2507 has external routes for networks 3.3.4.0/24, 3.22.88.0/24 and 3.44.66.0/24, even though all of these networks should be included in the summary of 3.0.0.0/8.
The reason these external routes show up is that the ASBR, which redistributes RIP into OSPF, has RIP running on these three subnets. It therefore redistributes the subnets as external routes into OSPF. Since these subnets are external routes, they are not summarized by the ABR (R2504). External OSPF routes can only be summarized by the ASBR. In this case, R2515. The ABR summarizes only internal OSPF routes from area 1 into area 0.
Note: With the fix of Cisco bug ID CSCdp72526 (registered customers only) , OSPF does not

generate a type-5 link-state advertisement (LSA) of an overlapped external network. R2507 only has a summary inter-area route of 3.0.0.0/8. Then, R2507 installs 200.1.1.0/24 as the forwarding address and it is reachable via interarea route 3.0.0.0/8. This is in compliance with RFC 2328.
This output shows the external LSA for network 200.1.1.0/24 in the OSPF database of R2507:
R2507# show ip ospf data external 200.1.1.0

OSPF Router with ID (7.7.7.1) (Process ID 1)
Type- 5 AS External Link States
LS age: 72
Options: (No TOS- capability, DC)
LS Type: AS External Link
Link State ID: 200.1.1.0 (External Network Number)
Advertising Router: 3.44.66.3
LS Seq Number: 80000001
Checksum: 0xF161
Length: 36
Network Mask: /24
Metric Type: 2 (Larger than any link state path)
TOS: 0
Metric: 20
Forward Address: 3.3.4.4
External Route Tag: 0
OSPF allows the ASBR to specify another router as the forwarding address to external routes. In this case, the
ASBR (R2515) has specified 3.3.4.4 as the forwarding address for the external network 200.1.1.0.
RFC 2328 , section 16.4 (Calculating AS external routes), states:
"If the forwarding address is non-zero, look up the forwarding address in the routing table. The matching routing table entry must specify an intra-area or inter-area path; if no such path exists, do nothing with the LSA and consider the next in the list."
In this example, the route to the forwarding address 3.3.4.4 is shown here:

R2507#
show ip route 3.3.4.4
Routing entry for 3.3.4.0/ 24
Known via "ospf 1", distance 110, metric 20,type extern 2, forward
metric 70
Redistributing via ospf 1
Last update from 1.1.1.2 on Serial0, 00: 00: 40 ago
Routing Descriptor Blocks:
* 1.1.1.2, from 3.44.66.3, 00: 00: 40 ago, via Serial0
Route metric is 20, traffic share count is 1

The forwarding address of 3.3.4.4 is matched by the external route 3.3.4.0/24 instead of the inter-area summary route 3.0.0.0/8 due to the longest match rule. Because the router does not have an internal OSPF route to the forwarding address, it does not install the external route 200.1.1.0/24 in the routing table. The use of an external route to reach another external route may lead to loops. Therefore OSPF does not permit it.

Reference
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009405a.shtml

QUESTION NO: 452

While troubleshooting a network, you need to verify the liveness of hosts in the subnet 192.168.1.64/26. All of the hosts are able to reply to ping requests. How would you confirm the existing nodes using one single command?

A. ping 192.168.1.255
B. ping with sweep option
C. ping 192.168.1.127
D. ping 192.168.1.64
E. P>ng with broadcast option

Answer: C
Explanation:
The 192.168.1.27 is the broadcast address of the 192.168.1.64/26 sub-network so by sending a ping request to this address all the hosts in this subnet will reply (to the broadcast address). But it is not quite right nowadays as all the Casco's routers which have IOS version 12.0 or above will simply drop these pings. If you wish to test this function then you have to turn on the ip directed-broadcast function (which is disabled by default from version 12.0).

The purpose of the ip directed-broadcast command is to enable forwarding of directed broadcasts. When this is turned on for an interface, the interface will respond to broadcast messages that are sent to its subnet. Cisco introduced this command in IOS version 10 (and it is enabled by default) but they soon realized this command was being exploited in denial of service attacks and disabled it from version 12.0. As you can guess, a ping to the broadcast address requires all hosts in that subnet to reply and it consumes much traffic if many are sent. A type of this attack is smurf attack, in which the attacker tries to borrow the victims IP address as the source address and sends ICMP packets to the broadcast address of the network. When all the hosts in that subnet hear the ICMP request, they will reply to the computer which the attacker borrowed the IP address from. You can

try this function by enabling ip directed-broadcast command in interface mode. Then from the directly connected router issue the ping to the broadcast address of that subnet (or ping 255.255.255.255).

QUESTION NO: 453

Two routers are connected by a serial link, and are configured to run EIGRP on all interfaces. You examine the EIGRP neighbor table on both routers (using the show ip eigrp neighbor command) and see that the router connected over the serial link is listed as a neighbor for a certain amount of time, but is periodically removed from the neighbor table.

None of the routes from the neighbor ever seem to be learned, and the neighbor transmission statistics (SRTT, RTO, and Q Count) seem to indicate that no packets are being transmitted between the neighbors. Which would most likely cause this problem?

A. While multicast packets are being successfully sent over the link, unicast packets are not
B. There is a bug in the EIGRP code that needs to be fixed.
C. This is correct behavior for the first few minutes of EIGRP neighbor formation. After four or five cycles, it should straighten itself out and the neighbor
D. The hello or hold intervals are set differently on the two routers.

Answer: A
Explanation:
EIGRP uses five packet types:
•Hello/Acks
•Updates
•Queries
•Replies
•Requests

Hellos are multicast for neighbor discovery/recovery. They do not require acknowledgment. A hello with no data is also used as an acknowledgment (ack). Acks are always sent using a unicast address and contain a nonzero acknowledgment number.

Updates are used to convey reachability of destinations. When a new neighbor is discovered, update packets are sent so the neighbor can build up its topology table. In this case, update packets are unicast. In other cases, such as a link cost change, updates are multicast. Updates are always transmitted reliably. Queries and replies are sent when destinations go into Active state. Queries are always multicast unless they are sent in response to a received query. In this case, it is unicast back to the successor that originated the query. Replies are always sent in

response to queries to indicate to the originator that it does not need to go into Active state because it has feasible successors. Replies are unicast to the originator of the query. Both queries and replies are transmitted reliably.

Request packets are used to get specific information from one or more neighbors. Request packets are used in route server applications. They can be multicast or unicast. Requests are transmitted unreliably.

QUESTION NO: 454

Study the exhibit carefully. Two directly connected routers are configured with OSPF. The output presented in the exhibit can be seen on the console of one router. What most likely cause this problem?



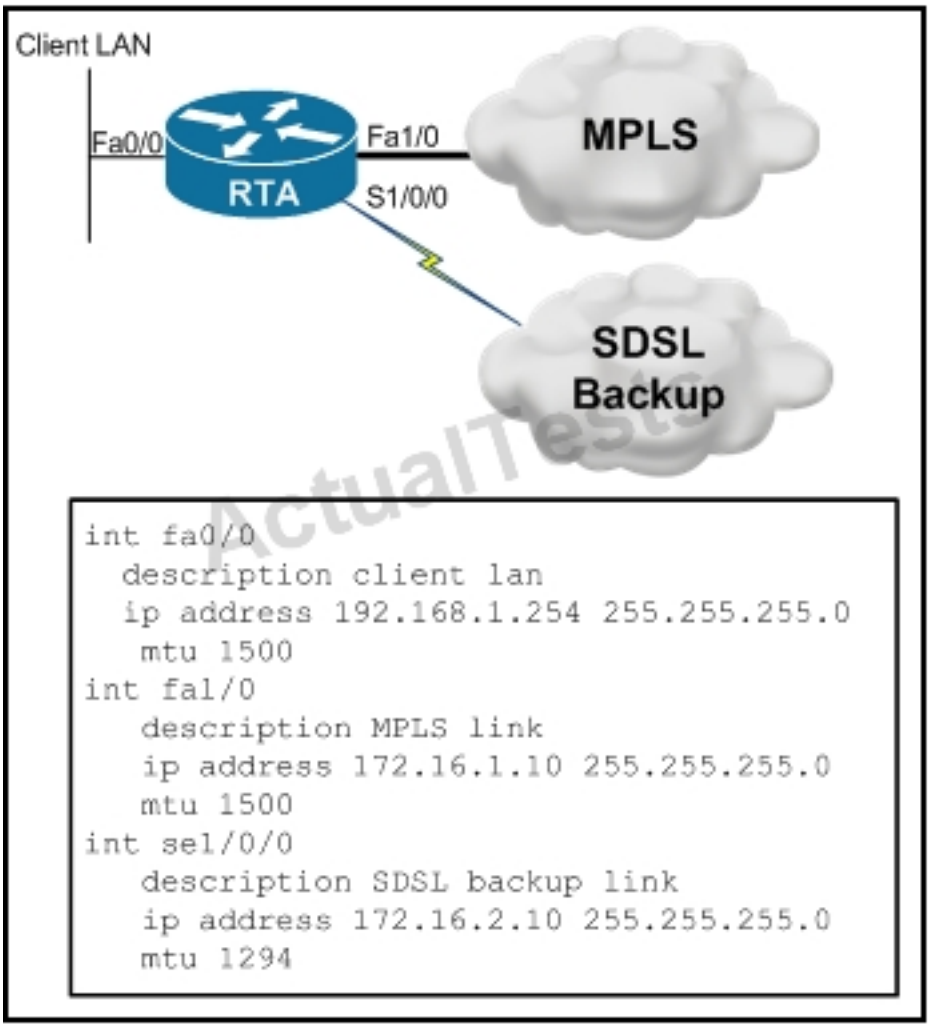
A. The maximum transmission unit on either side of the link is not the same
B. This debug is wrong, OSPF does not exchange DBD packets
C. This is normal for OSPF running over an FDDI ring
D. OSPF has received a packet that will not fit in its local buffer, so the packet has been discarded.

Answer: A

Explanation: Explanation
The debug output shows that the interface on the other side has a large MTU than what is currently being used

QUESTION NO: 455

Refer to the exhibit.



You get reports that Windows users cannot log in anymore when the main MPLS link is down. The message that users receive is, "There are currently no login servers available to service the login request." However, if you modify the MTU of the client machine to 1294, the user can log in. What would be a reliable fix for this issue?

A. add "ip tcp mss-adjust 1254" to se1/0/0
B. add "ip tcp mss-adjust 1254" to se1/0/0 and fa1/0
C. force Windows to use TCP for Kerberos
D. add "ip icmp error-send" to se1/0/0

Answer: C

Explanation:

You can force windows to use TCP for Kerberos so that even if you modify MTU of the client machine to 1294, use cannot log in.

QUESTION NO: 456

Refer to the exhibit. You would like to guarantee 7 Mb/s for FTP traffic in your LAN, as it seems that peer-to-peer traffic is taking up a large amount of bandwidth. When testing the configuration, you notice that FTP traffic doesnl reach 7 Mb/ s. What is the problem?

```
ip cef

class-map match-all ftp
 match protocol ftp
class-map match-all voice
 match precedence 5

policy-map mark
 class voice
  priority 200
 class ftp
  set dscp af32
  bandwidth 7000
  random-detect dscp-based
 class class-default

interface Tunnel1
 ip address 20.1.1.1 255.255.255.252
 load-interval 30
 qos pre-classify
 tunnel source 4.4.4.1
 tunnel destination 4.4.4.2
!
interface Tunnel2
 ip address 20.2.2.1 255.255.255.252
 load-interval 30
 qos pre-classify
 tunnel source 4.4.4.1
 tunnel destination 4.4.4.2

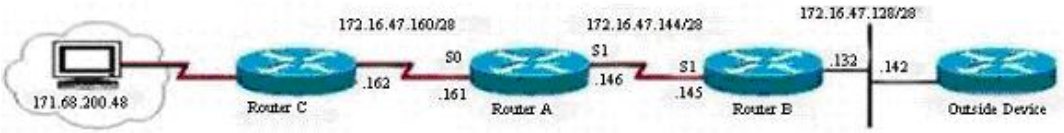
interface Ethernet0/1
 ip address 4.4.4.1 255.255.255.0
 no ip proxy-arp
 load-interval 30
 half-duplex
 no keepalive
 no cdp enable
 service-policy output mark
```

A. The Ethernet interface should have keepalives enabled.
B. The duplex settings are wrong on the Ethernet interface.
C. The qos pre-classify command should be removed from the tunnel interfaces.
D. the priority queue for the voice class is probably taking all the bandwidth
E. there are probably not enough interface buffers; they should be tuned.

Answer: B
Reference: http://en.wikipedia.org/wiki/Duplex_mismatch

QUESTION NO: 457

Refer to the exhibit. This exhibit shows the NAT configuration for Router A and the output for a ping issued from device 171.68.200.48 and destined to 172.16.47.142. Based on this information, what change must be made to Router A in order for the ping to work?




```
!
ip nat inside source static 171.68.200.48 172.16.47.150
interface Serial0
 ip address 172.16.47.161 255.255.255.240
 ip nat inside
!
interface Serial1
 ip address 172.16.47.146 255.255.255.240
 ip nat outside
!
no ip classless
ip route 0.0.0.0 0.0.0.0 172.16.47.145
ip route 171.68.200.0 255.255.255.0 172.16.47.162

Router-A#show ip nat translation
Pro Inside global      Inside local       Outside local      Outside global
--- 172.16.47.150      171.68.200.48      ---                ---

Router-A#debug ip packet detail
Router-A#debug ip nat
Router-A#IP: s=171.68.200.48 (Serial0), d=172.16.47.142, len 100, unroutable
    ICMP type=8, code=0
Router-A#IP: s=172.16.47.161 (local), d=171.68.200.48 (Serial0), len 56, sending
    ICMP type=3, code=1
Router-A#IP: s=171.68.200.48 (Serial0), d=172.16.47.142, len 100, unroutable
    ICMP type=8, code=0
Router-A#IP: s=171.68.200.48 (Serial0), d=172.16.47.142, len 100, unroutable
    ICMP type=8, code=0
Router-A#IP: s=172.16.47.161 (local), d=171.68.200.48 (Serial0), len 56, sending
    ICMP type=3, code=1
Router-A#IP: s=171.68.200.48 (Serial0), d=172.16.47.142, len 100, unroutable
    ICMP type=8, code=0
Router-A#IP: s=171.68.200.48 (Serial0), d=172.16.47.142, len 100, unroutable
    ICMP type=8, code=0
Router-A#IP: s=172.16.47.161 (local), d=171.68.200.48 (Serial0), len 56, sending
    ICMP type=3, code=1
```

A. reload the router
B. clear the route cache
C. add a static route
D. configure IP as classless
E. load a newer IOS image

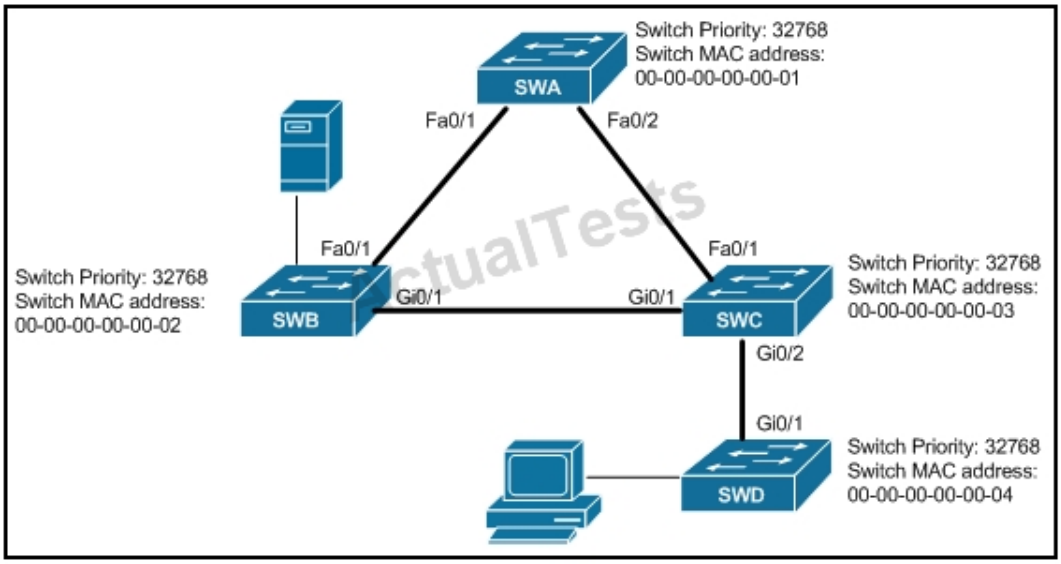
Answer: D

Explanation:

The assumption with classful routing is that all of a classful address belongs to one organization. The idea of a component subnet would be any subnet of a classful network. When you look at the routing table it lists the subnets based on the classful networks like this: Network 192.168.1.0/27 is subnetted 2 subnets 192.168.1.32 …. 192.168.1.64 …. The 192.168.1.32 and 192.168.1.64 subnets would be the "component subnets" of the 192.168.1.0 network. So the point is that if you tried to route a packet to 192.168.1.151 it would be dropped even if there was a default route in the routing table. The reason fo this is that the router knows of subnets that belong to the major network (192.168.1.0 in this case), but does not have an entry for the specific subnet.

QUESTION NO: 458

Refer to the exhibit.



Users that are connected to switch SWD are complaining about slow performance when they are doing large file transfers from a server connected to switch SWB. All switches are running PVST+. Which option will improve the performance of the file transfers?

A. Reconnect the clients from switch SWD to switch SWA.
B. Reconnect the clients from switch SWD to switch SWC.
C. Change PVST+ to RSTP.
D. Change the STP root switch from switch SWA to switch SWB.
E. Configure an EtherChannel between switch SWB and switch SWC.

Answer: D

Explanation:

When SWA is the root bridge the traffic path will be SWD <-> SWC <-> SWA <-> SWB

When SWB is the root bridge the traffic path will be SWD <-> SWC <-> SWB

QUESTION NO: 459

Refer to the exhibit.

```
Switch#show interface | include Vlan 10 is
 interface vlan 10 is down, line protocol is down
```

Clients in VLAN 10 complain that they cannot access network resources and the Internet. When you try to ping the default gateway from one of the affected clients, you get ping timeouts. What is the most likely cause of this issue?

A. VLAN 10 is only enabled on trunk interfaces.
B. VLAN 10 is not created in the switch database.
C. STP is not running on the switch.
D. IP routing is disabled on the switch.
E. The switch CAM table is corrupted.

Answer: B

Explanation: Explanation

Troubleshooting the Autostate Feature on IOS Based Switches

Perform these troubleshooting steps if the VLAN interface is down.

This is the symptom of a VLAN interface being in up/down status.

Corgon-6000#sh int vlan 151

Vlan151 is up, line protocol is down

!--- Line protocol on interface VLAN 151 is down.
!--- You need to investigate why this line protocol is not up
!--- (at least one L2 port exists, and there should be a
!--- link up on this VLAN).

Check to make sure that VLAN 151 exists in the VLAN database and is active. The command below shows that the VLAN exists and is active on the switch.

Corgon-6000#sh vlan 151 | i 151
151 VLAN151 active Gi4/10
151 enet 100151 1500 - - - - - 0 0
Corgon-6000#
!--- VLAN 151 exists in VLAN database and is active.
!--- L2 port Gig4/10 is assigned to VLAN 151.

Check the status of interface gig 4/10 assigned to VLAN 151.
Corgon-6000#sh int gig 4/10
GigabitEthernet4/10 is up, line protocol is down (notconnect)
Corgon-6000#sh run int gig 4/10

Building configuration...
Current configuration : 182 bytes
!
interface GigabitEthernet4/10
no ip address
logging event link-status
logging event bundle-status
switchport
switchport access vlan 151
switchport mode access
end

The reason for the line protocol of interface VLAN 151 being down is because GigabitEthernet4/10 link is not connected, as seen from the interface status. It is possible that no device is connected to the interface or that the link has cabling or auto-negotiation issues preventing the link from being up.
Connect the device to GigabitEthernet4/10 to bring the interface link up.
Mar 11 12:10:52.340: %LINK-3-UPDOWN: Interface GigabitEthernet4/10, changed state to up
Mar 11 12:10:53.156: %LINEPROTO-5-UPDOWN: Line protocol on Interface

GigabitEthernet4/10, changed state to up
Corgon-6000#
Corgon-6000#
Corgon-6000#sh int vlan 151
Vlan151 is up, line protocol is down

Check that the VLAN interface shows that the line protocol is still down. You need to investigate why this line protocol is not up. Make sure that at least one L2 port is in spanning-tree forwarding state on this VLAN.

```
Corgon-6000#sh spanning-tree vlan 151
VLAN0151
  Spanning tree enabled protocol rstp
  Root ID     Priority    32768
              Address     00d0.003f.8897
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
  Bridge ID   Priority    32768
              Address     00d0.003f.8897
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Gi4/10           Desg LRN 4         128.202  P2p
Corgon-6000#
```

Actual Tests

The Spanning-tree port status is LRN, which means learning state. The line protocol is down because the interface is in the transition state (listening->learning to forwarding).
Corgon-6000# Mar 11 12:11:23.406: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan151, changed state to up
Note: Time stamp difference between logs when the line protocol on GigabitEthernet4/10 went up, and Interface Vlan151 is around 30 seconds, which represents 2xforwarding delay in STP (listening-> learning->forwarding)
Corgon-6000#sh int vlan 151
Vlan151 is up, line protocol is up
The line protocol is up. You need to verify spanning-tree port status on the L2 port (should be forwarding).
Corgon-6000#sh spanning-tree vlan 151

```
VLAN0151
  Spanning tree enabled protocol rstp
  Root ID     Priority    32768
              Address     00d0.003f.8897
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
  Bridge ID   Priority    32768
              Address     00d0.003f.8897
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Gi4/10           Desg FWD 4         128.202  P2p
```

Actualtests

!--- Verified spanning-tree port status on L2 port
!--- is FWN = forwarding.

Reference
http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a0080160b14.shtml

QUESTION NO: 460

Refer to the exhibit.

```
R4#show ip ospf virtual-links
Virtual Link OSPF_VL1 to router 209.165.202.130 is down
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 0, Cost of using 65535
  Transmit Delay is 1 sec, State DOWN,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

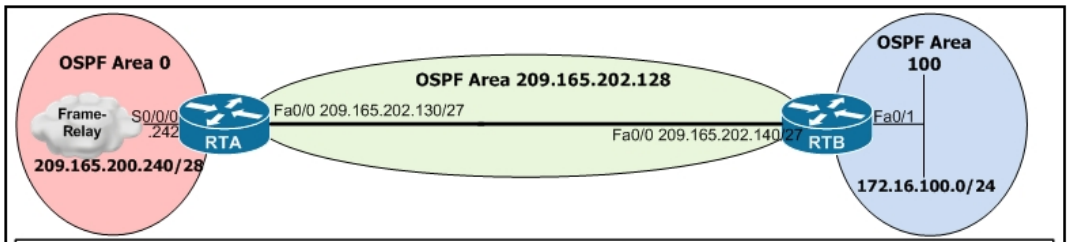
An OSPF virtual link is configured between R4 and R3. Based upon the show command output, why is the virtual-link down?

A. Virtual links cannot transit area 0.
B. The cost of the virtual link needs to be configured as 1.
C. The timer intervals for virtual links need to be aggressive (2, 8, 8, and 1).
D. The virtual interface VL1 is shut down.

Answer: A
Reference:
http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a00801ec9ee.shtml

QUESTION NO: 461

Refer to the exhibit.




```
RTA#show ip ospf virtual-links
Virtual Link OSPF_VL1 to router 209.165.202.140 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 209.165.202.128, via interface FastEthernet0/0, Cost of using 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01

RTB#show ip ospf virtual-links
Virtual Link OSPF_VL4 to router 209.165.200.242 is down
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 209.165.202.128, Cost of using 65535
  Transmit Delay is 1 sec, State DOWN,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

RTB# *Nov  4 19:11:05.058: %OSPF-4-ERRRCV: Received invalid packet: mismatch area ID, from backbone area
must be virtual-link but not found from 209.165.202.130, FastEthernet0/0
```

An OSPF virtual link is configured between RTB and RTA. Based upon the exhibit, why is the

virtual link on RTB down?

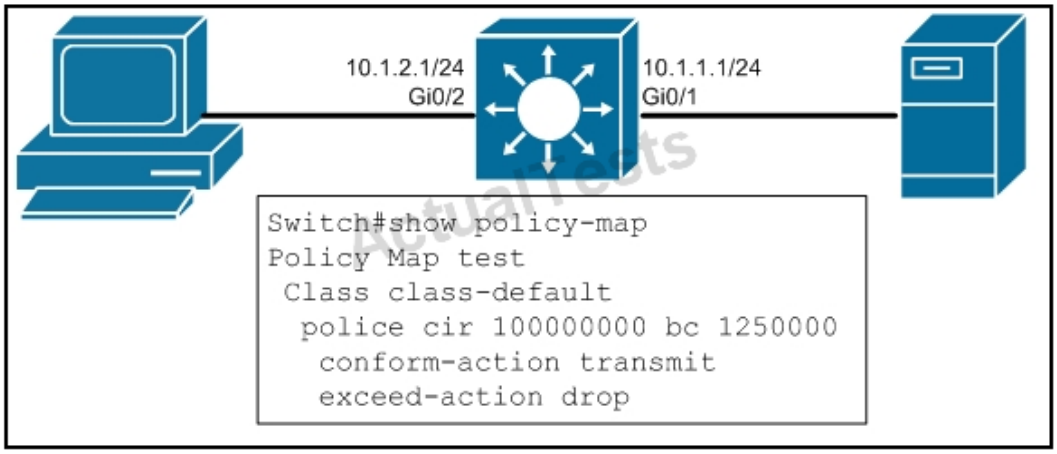
A. The cost on each end of the OSPF virtual link must be identical.
B. There is a unidirectional physical layer issue from RTB to RTA.
C. The OSPF virtual link neighbor IP address on RTB is incorrect.
D. The virtual link state on FastEthernet0/0 of RTA must be point-to-multipoint.

Answer: C
Explanation:

Refer to the diagram above and see that the OSPF virtual link neighbor IP address on RTB is incorrect in the IOS screen.

QUESTION NO: 462

Refer to the exhibit.



You are trying to police down to 100 Mb/s. While testing, you notice that you rarely exceed 70-80 Mb/s.

What do you need to change in your MQC configuration to allow for 100 Mb/s speeds?

A. Change the CIR value from 100 Mb/s to 200 Mb/s.
B. Change the Bc value to allow for a large enough burst.
C. Change the QoS queue from default to priority.
D. Change the exceed-action to transmit.

Answer: B

Explanation:
Burst size—Also called the Committed Burst (Bc) size, it specifies in bits (or bytes) per burst how much traffic can be sent within a given unit of time to not create scheduling concerns. (For a shaper, such as GTS, it specifies bits per burst; for a policer, such as CAR, it specifies bytes per burst.)

QUESTION NO: 463

Refer to the exhibit.

```
Switch#show interface | include Vlan 10 is
 interface vlan 10 is up, line protocol is down
```

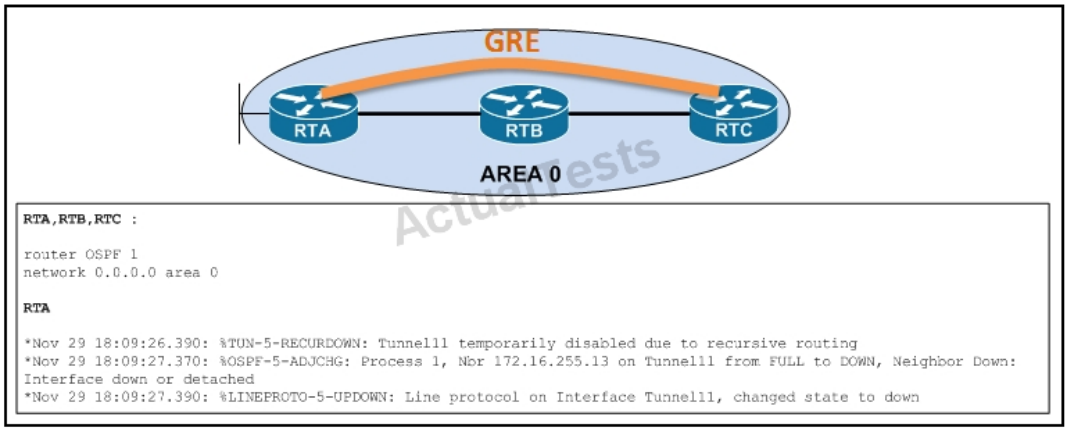
Clients in VLAN 10 complain that they cannot access network resources and the Internet. When you try to ping the default gateway from one of the affected clients, you get ping timeouts. What is most likely the cause of this issue?

A. VLAN 10 is only enabled on trunk interfaces.
B. VLAN 10 is not created in the switch database.
C. STP is not running on the switch.
D. IP routing is disabled on the switch.
E. The switch CAM table is corrupted.
F. The no shutdown command is issued under the VLAN 10 interface configuration.

Answer: D
Explanation:
When the line protocol is down, it means IP routing is disabled on the switch.

QUESTION NO: 464

Refer to the exhibit.



As soon as the OSPF neighbors are established across the GRE tunnel between RTA and RTC, the GRE tunnel immediately goes down/down. What could be a possible cause?

A. The route to the tunnel destination is preferred via the tunnel.
B. Autosummary is not configured.
C. GRE tunnels cannot be in area 0.
D. A misconfigured access list is on the router C tunnel interface.
E. A misconfigured access list is on the router A tunnel interface.

Answer: A
Explanation:
The error message, %TUN-5-RECURDOWN: Tunnel0 temporarily disabled due to recursive routing, means that the generic routing encapsulation (GRE) tunnel router has discovered a recursive routing problem. This condition is usually due to one of the following 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.
Tunnel interface status depends on the IP reachability to the tunnel destination. When the router detects a recursive routing failure for the tunnel destination, it shuts the tunnel interface down for a few minutes so that the situation causing the problem can resolve itself as routing protocols converge. If the problem is caused by misconfiguration, the link may oscillate indefinitely.

Another symptom of this problem is continuously flapping Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), or Border Gateway Protocol (BGP) neighbors, when the neighbors are over a GRE tunnel.
This document shows an example of troubleshooting an oscillating tunnel interface that is running EIGRP.

Reference
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094690.shtml

QUESTION NO: 465

When you are troubleshooting duplex mismatches, which two errors are seen on the full-duplex end? (Choose two.)

A. runts
B. FCS errors
C. interface resets
D. late collisions

Answer: A,B

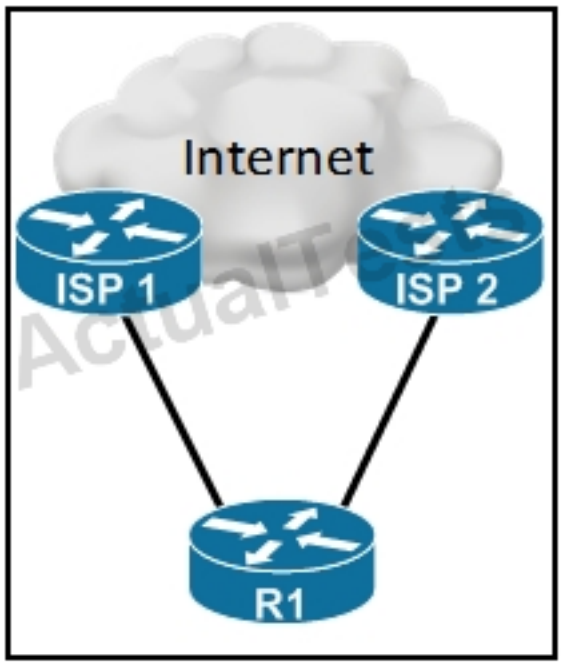
Explanation:

n half-duplex environments, it is possible for both the switch and the connected device to sense the wire and transmit at exactly the same time and result in a collision. Collisions can cause runts, FCS, and alignment errors, caused when the frame is not completely copied to the wire, which results in fragmented frames. In full duplex, collisions are not possible.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a00800a7af0.shtml

Topic 10, Optimize the Network

QUESTION NO: 466

Refer to the exhibit.



A small enterprise connects its office to two ISPs, using separate T1 links. A static route is used for the default route, pointing to both interfaces with a different administrative distance, so that one of the default routes is preferred.

Recently the primary link has been upgraded to a new 10 Mb/s Ethernet link.

After a few weeks, they experienced a failure. The link did not pass traffic, but the primary static route remained active. They lost their Internet connectivity, even though the backup link was operating.

Which two possible solutions can be implemented to avoid this situation in the future? (Choose two.)

A. Implement HSRP link tracking on the branch router R1.
B. Use a track object with an IP SLA probe for the static route on R1.
C. Track the link state of the Ethernet link using a track object on R1.
D. Use a routing protocol between R1 and the upstream ISP.

Answer: B,D

Explanation:

Interface Tracking

Interface tracking allows you to specify another interface on the router for the HSRP process to monitor in order to alter the HSRP priority for a given group.

If the specified interface's line protocol goes down, the HSRP priority of this router is reduced, allowing another HSRP router with higher priority can become active (if it has preemption enabled).

To configure HSRP interface tracking, use the standby [group] track interface [priority] command.

When multiple tracked interfaces are down, the priority is reduced by a cumulative amount. If you explicitly set the decrement value, then the value is decreased by that amount if that interface is down, and decrements are cumulative. If you do not set an explicit decrement value, then the value is decreased by 10 for each interface that goes down, and decrements are cumulative.
The following example uses the following configuration, with the default decrement value of 10.

Note: When an HSRP group number is not specified, the default group number is group 0.

interface ethernet0
ip address 10.1.1.1 255.255.255.0
standby ip 10.1.1.3
standby priority 110
standby track serial0
standby track serial1

The HSRP behavior with this configuration is:
0 interfaces down = no decrease (priority is 110)
1 interface down = decrease by 10 (priority becomes100)
2 interfaces down = decrease by 10 (priority becomes 90)

Reference
http://www.cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a0080094a91.shtml#intracking

QUESTION NO: 467

To troubleshoot network issues more accurately, milliseconds should be included in the syslog of the router. Which command will achieve this?

A. service timestamps log datetimec msec
B. logging timestamps msec
C. syslog timestamps hour minute second miliseconds
D. service logging timestamp msec
E. logging service timestamp msec

Answer: A
Explanation: Explanation

Logging is a quite essential part of a secure network configuration. Logging not only helps the Network Administrators to identify the issue while troubleshooting, also enables them to react to intrusion attempts or Denial-of –Service attacks.

By default on Cisco IOS , no timestamp information is included; however, you can enable timestamps and also modify the format of the timestamp attached to SYSLOG message by using the service timestamps log global configuration commands as follows:

ITKE(Config)# service timestamps log {uptime |datetime [msec |localtime |show-timezone]}
http://itknowledgeexchange.techtarget.com/network-technologies/what-is-service-timestamps-logging-and-howit-can-be-configured-cisco-switch-or-a-router/

QUESTION NO: 468

Refer to the exhibit.

```
process cpu threshold type total rising 80 interval 5 falling 50 interval 5
event manager applet CCIE
event syslog pattern "SYS-1-CPUFALLINGTHRESHOLD"
action 1.1 cli command "enable"
action 1.2 cli command "show log | a flash:LOG"
```

Based on the above commands, when will the output of the show log command be saved?

A. Each time the total CPU utilization goes below 50 percent
B. Each time the total CPU utilization goes above 80 percent
C. Every 5 minutes while the total CPU utilization is above 80 percent
D. Every 5 seconds while the total CPU utilization is above 80 percent
E. Every 5 minutes while the total CPU utilization is below 50 percent
F. Every 5 seconds while the total CPU utilization is below 50 percent

Answer: A
Explanation: The cpu threshold generates syslog messages when it goes above 80 % and when it comes back down below 50% after being above 80%.
It checks cpu utilization every 5 seconds.

When the cpu has been above 80%, and has come back below 50%, the syslog message SYS-1-CPUFALLINGTHRESHOLD is generated
-thats when the "show log" command is triggered

The closest answer is "Each time the total CPU goes below 50 percent"

QUESTION NO: 469

Which configuration would make an IP SLA probe use a precedence value of 5?

A.
ip sla 1
icmp-echo 1.1.1.1
tos 160
B.
ip sla 1
icmp-echo 1.1.1.1
tos 20
C.
ip sla 1
icmp-echo 1.1.1.1
precedence 5
D.
ip sla 1
icmp-echo 1.1.1.1
dscp 20

Answer: A
Explanation: Explanation

SUMMARY STEPS
1. enable
2. configure terminal
3. ip sla monitor operation-number
4. type echo protocol ipIcmpEcho {destination-ip-address | destination-hostname} [source-ipaddr {ip-address |
hostname} | source-interface interface-name]
5. buckets-of-history-kept size
6. distributions-of-statistics-kept size
7. enhanced-history [interval seconds] [buckets number-of-buckets]
8. filter-for-history {none | all | overThreshold | failures}
9. frequency seconds
10. hours-of-statistics-kept hours
11. lives-of-history-kept lives

12. owner owner-id
13. request-data-size bytes
14. statistics-distribution-interval milliseconds
15. tag text
16. threshold milliseconds
17. timeout milliseconds
18. tos number
19. verify-data
20. vrf vrf-name
21. exit
22. ip sla monitor schedule operation-number [life {forever | seconds}] [start-time {hh:mm[:ss] [month day | day
month] | pending | now | after hh:mm:ss] [ageout seconds] [recurring]
23. exit
24. show ip sla monitor configuration [operation-number] tos number

Example:
Router(config-sla-monitor-echo)# tos 160
(Optional) Defines a type of service (ToS) byte in the IP header of an IP SLAs operation.

Reference
http://www.cisco.com/en/US/docs/ios/12_4/ip_sla/configuration/guide/hsicmp.html

QUESTION NO: 470

Refer to the exhibit.

```
event manager applet MONITOR
 event interface name FastEthernet0/0 parameter receive_rate_pps entry-op ge entry-val
 10000 entry-type value exit-op le exit-val 1000 exit-type value poll-interval 5
```

When would the EEM applet be triggered?

A. every time that the input errors counter is higher than 100
B. every time that the input errors counter is slower than 10 errors per 10 seconds
C. every time that the input errors counter is lower than 100
D. every time that the input errors counter is faster than 100 errors per 10 seconds

Answer: A
Explanation:

Input errors includes runts, giants, no buffer, cyclic redundancy checksum (CRC), frame, overrun, and ignored counts. Other input-related errors can also cause the input errors count to be increased. Some datagrams may have more than one error.

QUESTION NO: 471

Refer to the exhibit.

```
event manager applet TIMER
 event timer cron cron-entry "0 * * * *"
```

When would the EEM applet be triggered?

A. once a month
B. once a day
C. once an hour
D. once a minute

Answer: D
Explanation:
Interface Counter Event Detector-Rate Based Trigger
The interface counter Event Detector (ED) adds the ability for an interface event to be triggered based on a rate of change over a period of time. A rate can be specified both for the entry value and the exit value.

The "event interface" applet CLI command has been modified to accept three new keywords:

[entry-type {value | increment | rate}]
[exit-type {value | increment | rate}]
[average-factor <average-factor-value>]
Similar commands are added for Tcl scripts:
[entry_type {value | increment | rate}]
[exit_type {value | increment | rate}]
[average_factor <average-factor-value>]
Applet syntax :
[no] event [<ev-label>] interface name <interface-name> parameter <counter-name> entry-val <entry-val> entry-op {gt|ge|eq|ne|lt|le} [entry-type {value | increment | rate}]

[exit-comb {or | and}]
[exit-val <exit-val> exit-op {gt|ge|eq|ne|lt|le} exit-type {value | increment | rate}]
[exit-time <exit-time-val>]
poll-interval <poll-int-val>

The following is an example of the rate based trigger in action. This applet monitors for errors on an interface. If the rate of change averages to two or more over three 60 second polling cycles, then the interface is reset by doing a shut/no shut. The policy will re-arm after the rate has dropped below 1.

event manager applet int-rate-test
event interface name FastEthernet0/24 parameter input_errors entry-op ge entry-val 2 entry-type rate exit-op lt
exit-val 1 exit-type rate average-factor 3 poll-interval 60
action 1.0 syslog msg "Interface input error rate for $_interface_name is $_interface_value; resetting..."
action 2.0 cli command "enable"
action 3.0 cli command "interface $_interface_name"
action 4.0 cli command "shut"
action 5.0 cli command "no shut"
action 6.0 cli command "end"

QUESTION NO: 472

Refer to the exhibit.

```
event manager applet CCIE
event timer cron name CCIE cron-entry  */5 * * * *
action 1 cli command "en"
action 2 cli command "show log"
```

ActualTests

Which output will the EEM applet in the exhibit produce?

A. The output of show version will be executed every 5 hours.
B. The output of show log will be executed every 5 hours.
C. The output of show log will be executed every Friday.
D. The output of show log will be executed every 5 minutes.

Answer: B

Explanation:

http://www.cisco.com/en/US/docs/ios/netmgmt/command/reference/nm_06.html#wp1157622

event timer

cron-entry

Specifies the first five fields of a UNIX crontab entry as used with the UNIX CRON daemon.

cron-entry

Text string that consists of five fields separated by spaces. The fields represent the times and dates when CRON timer events will be triggered. Fields and corresponding values are as follows:

–minute—A number in the range from 0 to 59 that specifies when a CRON timer event is triggered.

–hour—A number in the range from 0 to 23 that specifies when a CRON timer event is triggered.

–day-of-month—A number in the range from 1 to 31 that specifies the day of the month when a CRON timer event is triggered.

–month—A number in the range from 1 to 12 or the first three letters (not case-sensitive) of the name of the month in which a CRON timer event is triggered.

–day-of-week—A number in the range from 0 to 6 (Sunday is 0) or the first three letters (not case-sensitive) of the name of the day when a CRON timer event is triggered.

Instead of the first five fields, special strings can be entered. See the "Usage Guidelines" section for details.

==========

Usage Guidelines

For the cron-entry argument, the following special strings also are allowed in syntax:

•Range of numbers—The specified range is inclusive, and a hyphen separates the numbers. For example, 8-11 after the hour field specifies execution of a CRON timer event at hours 8, 9, 10, and 11.

•Asterisk (*)—Indicates that a field is not specified and can be any value.

•List—A list is a set of numbers or ranges separated by a comma but no space. For example, 1,2,5,9 or 0-4,8-12.

•Step value in conjunction with a range—Following a range with /number specifies skips of the number value through the range. For example, 0-23/2 in the hour field specifies that an event is triggered every second hour. Steps are permitted after an asterisk, for example */2 means every two hours.

QUESTION NO: 473

Refer to the exhibit.

```
7609-1#show int g5/2
GigabitEthernet5/2 is up, line protocol is down (monitoring)
  Hardware is C7600 1Gb 802.3, address is c471.fe02.9100 (bia c471.fe02.9100)
  Internet address is 172.16.1.33/24
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 100Mb/s
  Media-type configured as RJ45 connector
  input flow-control is off, output flow-control is off
  Clock mode is auto
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:01:00, output 00:01:00, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/28236/26546 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 88000 bits/sec, 25 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  L2 Switched: ucast: 2340906 pkt, 194394394 bytes - mcast: 89580706 pkt, 38164979127 bytes
  L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes mcast
  L3 out Switched: ucast: 0 pkt, 0 bytes mcast: 0 pkt, 0 bytes
```

Why is the interface in the up/down state?

A. There is no physical layer connection.
B. A span session has been configured with g5/2 as the destination.
C. There were too many input drops on the interface.
D. A span session has been configured with g5/2 as the source

Answer: B
Explanation:
SPAN Destination Port Up/Down
When ports are spanned for monitoring, the port state shows as UP/DOWN.
When you configure a SPAN session to monitor the port, the destination interface shows the state down (monitoring), by design. The interface shows the port in this state in order to make it evident

that the port is currently not usable as a production port. The port as up/down monitoring is normal.

Reference
http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a008015c612.shtml#topic8-8

QUESTION NO: 474

Refer to the exhibit.

```
router#sh ip cache flow
IP packet size distribution (331 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .050 .001 .002 .001 .003 .001 .006 .003 .001 .003 .001 .000 .000 .000

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .000 .000 .900 .000 .000 .000 .000 .000 .000 .000 .000
```

ActualTests

You are investigating a performance problem between two hosts. You have enabled NetFlow.

What is most likely the cause of this issue?

A. A firewall is stripping the TCP MSS option.
B. A firewall is stripping the IP MSS option.
C. An IPS is stripping the TCP MSS option.
D. There is a VPN link causing low MTU.
E. You must configure the MTU on the links on the router.

Answer: A
Explanation:

QUESTION NO: 475

Which of the following is the encryption algorithm used for priv option when using SNMPv3?

A. HMAC-SHA
B. HMAC-MD5
C. CBC-DES
D. AES

E. 3DES

Answer: C

Explanation:

SNMPv3

Feature Summary

Simple Network Management Protocol Version 3 (SNMPv3) is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting packets over the network. The security features provided in SNMPv3 are:

Message integrity—Ensuring that a packet has not been tampered with in-transit.

Authentication—Determining the message is from a valid source.

Encryption—Scrambling the contents of a packet prevent it from being seen by an unauthorized source.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the group in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level will determine which security mechanism is employed when handling an SNMP packet. Three security models are available:

SNMPv1, SNMPv2c, and SNMPv3. Table 1 identifies what the combinations of security models and levels mean:

Table 1 SNMP Security Models and Levels

Model	Level	Authentication	Encryption	What Happens
v1	noAuthNoPriv	Community String	No	Uses a community string match for authentication.
v2c	noAuthNoPriv	Community String	No	Uses a community string match for authentication.
v3	noAuthNoPriv	Username	No	Uses a username match for authentication.
v3	authNoPriv	MD5 or SHA	No	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms.
v3	authPriv	MD5 or SHA	DES	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Provides DES 56-bit encryption in addition to authentication based on the CBC-DES (DES-56) standard.

ActualTests

Reference
http://www.cisco.com/en/US/docs/ios/12_0t/12_0t3/feature/guide/Snmp3.html#wp4363
http://www.cisco.com/en/US/docs/ios/12_0t/12_0t3/feature/guide/Snmp3.html
http://www.cisco.com/en/US/docs/ios/12_4t/12_4t2/snmpv3ae.html

QUESTION NO: 476

Refer to the exhibit.

```
event manager applet CCIE
event timer cron name CCIE cron-entry * * */5 * *
action 1 cli command "en"
action 2 cli command "show log"
```

ActualTests

Which output will the EEM applet in the exhibit produce?

A. The output of show version will be executed every 5 hours.
B. The output of show log will be executed every 5 hours.
C. The output of show log will be executed every 5 days.
D. The output of show log will be executed every 5 minutes.

Answer: C
Explanation:
event_register_timer
Create a timer and register for a timer event as both a publisher and a subscriber. Use this keyword when there is a need to trigger a policy that is time specific or timer based. This event timer is both an event publisher and a subscriber. The publisher part indicates the conditions under which the named timer is to go off. The subscriber part identifies the name of the timer to which it is subscribing.
Note Both the CRON and absolute time specifications work on local time.
Syntax

event_register_timer watchdog|countdown|absolute|cron
[name ?] [cron_entry ?] (for cron timer)
[time ?] (for other types of timer)
[priority low|normal|high] [maxrun ?]
[nice 0|1]

Timer event types (one must be specified):	
watchdog	Mandatory. Watchdog timer.
countdown	Mandatory. Countdown timer.
absolute	Mandatory. Absolute timer.
cron	Mandatory. Cron timer.

Other arguments:	
name	Mandatory. Name of the timer.
cron_entry	Mandatory. Must be specified if cron timer type is specified. Must not be specified if any other timer type is specified. A cron_entry is a partial Unix crontab entry (the first five fields) as used with the Unix cron daemon. A cron_entry specification consists of a text string with five fields. The fields are separated by spaces. The fields represent the time and date when cron timer events will be triggered. The fields are described in Table 4: Ranges of numbers are allowed. Ranges are two numbers separated with a hyphen. The specified range is inclusive. For example, 8-11 for an "hour" entry specifies execution at hours 8, 9, 10 and 11. A field may be an asterisk (*), which always stands for "first-last". Lists are allowed. A list is a set of numbers (or ranges) separated by commas. Examples: "1,2,5,9", "0-4,8-12". Step values can be used in conjunction with ranges. Following a range with "/<number>" specifies skips of the number's value through the range. For example, "0-23/2" can be used in the "hour" field to specify an event that is triggered every other hour. Steps are also permitted after an asterisk, so if you want to say "every two hours", use "*/2". Names can also be used for the "month" and "day of week" fields. Use the first three letters of the particular day or month (case doesn't matter). Ranges or lists of names are not allowed. The day that a timer event is triggered can be specified by two fields -- day of month, and day of week. If both fields are restricted (ie, are not *), an event will be triggered when either field matches the current time. For example, "30 4 1,15 * 5" would cause an event to be triggered at 4:30 am on the 1st and 15th of each month, plus every Friday. Instead of the first five fields, one of seven special strings may appear. These seven special strings are described in Table 5.

Actual Tests

	Example 1: "0 0 1,15 * 1" would trigger an event at midnight on the first and fifteenth of each month, as well as on every Monday. To specify days by only one field, the other field should be set to *; "0 0 * * 1" would trigger an event at midnight only on Mondays. Example 2: "15 16 1 * *" would trigger an event at 4:15PM on the first day of each month. Example 3: "0 12 * * 1-5" would trigger an event at noon on Monday through Friday of each week. Example 4: "@weekly" would trigger an event at midnight once a week on Sunday.
time	Mandatory. Must be specified if timer type other than cron is specified. Must not be specified if cron timer type is specified. For watchdog and countdown timers, the number of seconds and milliseconds until the timer expires; for absolute timer, the calendar time of the expiration time. Specified in SSSSSSSSSS[.MMM] format where SSSSSSSSSS must be an integer number representing seconds between 0 and 4294967295 inclusive. MMM represents milliseconds, and must be an integer number between 0 and 999. An absolute expiration date is the number of seconds and milliseconds since January 1, 1970. If the date specified has already passed, the timer expires immediately.
priority	Optional. Priority level the script will be queued. Priority not specified means normal priority, the default.
maxrun	Optional. Maximum runtime of the script (specified in SSSSSSSSSS[.MMM] format where SSSSSSSSSS must be an integer number representing seconds between 0 and 4294967295 inclusive. MMM represents nanoseconds, and must be an integer number between 0 and 999) If not specified, use the default 20 seconds runtime limit.
nice	Optional. When nice is set to 1, the policy is run at a runtime priority that is less than the default priority. The default value is 0.

Table 4 Time and Date When Cron Events Will Be Triggered

Field	Allowed Values
minute	0-59
hour	0-23
day of month	1-31
month	1-12 (or names, see below)
day of week	0-7 (0 or 7 is Sun, or names, see below)

Table 5 Special Strings for cron_entry

String	Meaning
@yearly	Trigger once a year, "0 0 1 1 *".
@annually	(same as @yearly)
@monthly	Trigger once a month, "0 0 1 * *".
@weekly	Trigger once a week, "0 0 * * 0".
@daily	Trigger once a day, "0 0 * * *".
@midnight	(same as @daily)
@hourly	Trigger once an hour, "0 * * * *".

QUESTION NO: 477

On a router that is configured with multiple IP SLA probes, which command can be used to manage the CPU load that was created by the periodic probing?

A. ip sla monitor low-memory
B. ip sla group schedule
C. ip sla reaction-trigger
D. ip sla enable timestamp

Answer: B
Reference: http://www.cisco.com/en/US/docs/ios/ipsla/command/reference/sla_02.html (see usage guidelines)

QUESTION NO: 478

Which configuration would make an IP SLA probe use a precedence value of 3?

A. ip sla 1
icmp-echo 1.1.1.1
tos 12
B. ip sla 1
icmp-echo 1.1.1.1
tos 96
C. ip sla 1
icmp-echo 1.1.1.1
precedence 3
D. ip sla 1
icmp-echo 1.1.1.1
dscp 12

Answer: B
Explanation:
SUMMARY STEPS

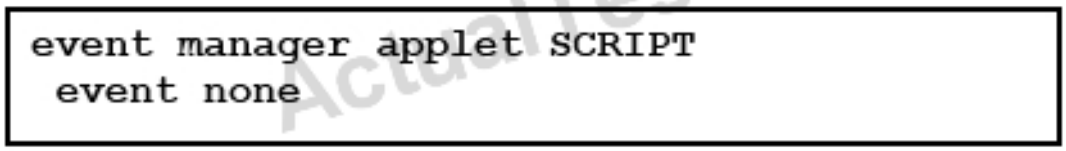
1. enable
2. configure terminal
3. ip sla monitor operation-number
4. type echo protocol ipIcmpEcho {destination-ip-address | destination-hostname} [source-ipaddr {ip-address |
hostname} | source-interface interface-name]

5. frequency seconds
6. exit
7. ip sla monitor schedule operation-number [life {forever | seconds}] [start-time {hh:mm[:ss] [month day | day
month] | pending | now | after hh:mm:ss] [ageout seconds] [recurring]
8. exit

Reference
http://www.cisco.com/en/US/docs/ios/12_4/ip_sla/configuration/guide/hsicmp.html

QUESTION NO: 479

Refer to the exhibit.



Which of the following options will trigger the applet?

A. an external Cisco IOS event
B. a manually run policy event
C. a preconfigured timer
D. an automated RPC call

Answer: B
Explanation:
There are two ways to manually run an EEM policy. EEM usually schedules and runs policies on the basis of an event specification that is contained within the policy itself. The event none command allows EEM to identify an EEM policy that can either be run manually or be run when an EEM applet is triggered. To run the policy, use either the action policy command in applet configuration mode or the event manager run command in global configuration mode.

QUESTION NO: 480

Refer to the exhibit.

```
event manager applet CCIE
event timer cron name CCIE cron-entry  * */5 * * *
action 1 cli command "en"
action 2 cli command "show log"
```

Which result will the EEM applet in the exhibit produce?

A. The output of show version will be executed every 5 hours.
B. The output of show log will be executed every 5 hours.
C. The output of show log will be executed every Friday.
D. The output of show log will be executed every 5 minutes.

Answer: B

Explanation:

The cron entry indicates 5 hours. So the output of show log will be executed every 5 hours.

QUESTION NO: 481

Refer to the exhibit.

```
process cpu threshold type total rising 80 interval 5 falling 50 interval 5
event manager applet CCIE
event syslog pattern "SYS-1-CPURISINGTHRESHOLD"
action 1.1 cli command "enable"
action 1.2 cli command "show log | a flash:LOG"
```

Based on the above commands, when will the output of the show log command be saved?

A. Each time the total CPU utilization goes below 50 percent
B. Each time the total CPU utilization goes above 80 percent
C. Every 5 minutes while the total CPU utilization is above 80 percent
D. Every 5 seconds while the total CPU utilization is above 80 percent
E. Every 5 minutes while the total CPU utilization is below 50 percent
F. Every 5 seconds while the total CPU utilization is below 50 percent

Answer: B

Explanation: Here is a reference:

http://www.cisco.com/en/US/docs/ios/netmgmt/configuration/guide/nm_cpu_thresh_notif.pdf

QUESTION NO: 482

Refer to the exhibit.

```
event manager applet MONITOR
 event interface name FastEthernet0/0 parameter receive_rate_pps entry-op ge entry-val
 10000 entry-type value exit-op le exit-val 1000 exit-type value poll-interval 5
```

When would the EEM applet be triggered?

A. every time that the input packet per second counter is below 10, 000
B. every time that the input packet per second counter has increased by 1, 000
C. every time that the input packet per second counter is above 10, 000
D. every time that the input packet per second counter has decreased by 1, 000

Answer: C
Reference: http://www.cisco.com/en/US/docs/ios-xml/ios/eem/command/eem-cr-e1.html

QUESTION NO: 483

Which two statements are correct about Nonstop Forwarding? (Choose two.)

A. It allows the standby RP to take control of the device after a hardware or software fault on the active RP.
B. It is a Layer 3 function that works with SSO to minimize the amount of time a network is unavailable to users following a switchover.
C. It is supported by the implementation of EIGRP, OSPF, RIPv2, and BGP protocols.
D. It synchronizes startup configuration, startup variables, and running configuration.
E. The main objective of NSF is to continue forwarding IP packets following a switchover.
F. Layer 2 802.1w or 802.1s must be used, as 802.1d cannot process the Layer 2 changes.
G. Routing protocol tuning parameters must be the same as the NSF parameters, or failover will be inconsistent.

Answer: B,E
Explanation: Explanation
Cisco Nonstop Forwarding (NSF) works with the Stateful Switchover (SSO) feature in Cisco IOS

software. NSF works with SSO to minimize the amount of time a network is unavailable to its users following a switchover. The main objective of Cisco NSF is to continue forwarding IP packets following a Route Processor (RP) switchover.

Reference
http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fsnsf20s.html

QUESTION NO: 484

On a router that is configured with multiple IP SLA probes, which command can be used to manage the CPU load that was created by the periodic probing?

A. ip sla monitor low-memory
B. ip sla group schedule
C. ip sla reaction-trigger
D. ip sla enable timestamp

Answer: B
Reference: http://www.cisco.com/en/US/docs/ios/ipsla/command/reference/sla_02.html (see usage guidelines)

QUESTION NO: 485

Which configuration would make an IP SLA probe use a precedence value of 3?

A. ip sla 1
icmp-echo 1.1.1.1
tos 12
B. ip sla 1
icmp-echo 1.1.1.1
tos 96
C. ip sla 1
icmp-echo 1.1.1.1
precedence 3
D. ip sla 1
icmp-echo 1.1.1.1
dscp 12

Answer: B
Explanation:
SUMMARY STEPS

1. enable
2. configure terminal
3. ip sla monitor operation-number
4. type echo protocol ipIcmpEcho {destination-ip-address | destination-hostname} [source-ipaddr {ip-address |
hostname} | source-interface interface-name]
5. frequency seconds
6. exit
7. ip sla monitor schedule operation-number [life {forever | seconds}] [start-time {hh:mm[:ss] [month day | day
month] | pending | now | after hh:mm:ss] [ageout seconds] [recurring]
8. exit

Reference
http://www.cisco.com/en/US/docs/ios/12_4/ip_sla/configuration/guide/hsicmp.html

QUESTION NO: 486

Which NetFlow version should be used to collect accounting data for IPv6 traffic?

A. version 1
B. version 5
C. version 7
D. version 8
E. version 9

Answer: E
Reference: http://www.cisco.com/en/US/docs/ios-xml/ios/netflow/configuration/15-s/ip6-nflw-v9.html

QUESTION NO: 487

You are responsible for network monitoring and need to monitor traffic over a routed network from

a remote source to an IDS or IPS located in the headquarters site. What would you use in order to accomplish this?

A. VACLs and VSPAN
B. RSPAN
C. ERSPAN
D. NetFlow

Answer: C
Explanation:
ERSPAN Overview
ERSPAN supports source ports, source VLANs, and destinations on different switches, which provides remote monitoring of multiple switches across your network (see Figure 68-3). ERSPAN uses a GRE tunnel to carry traffic between switches.

ERSPAN consists of an ERSPAN source session, routable ERSPAN GRE-encapsulated traffic, and an ERSPAN destination session. You separately configure ERSPAN source sessions and destination sessions on different switches.

To configure an ERSPAN source session on one switch, you associate a set of source ports or VLANs with a destination IP address, ERSPAN ID number, and optionally with a VRF name. To configure an ERSPAN destination session on another switch, you associate the destinations with the source IP address, ERSPAN ID number, and optionally with a VRF name.

ERSPAN source sessions do not copy locally sourced RSPAN VLAN traffic from source trunk ports that carry RSPAN VLANs. ERSPAN source sessions do not copy locally sourced ERSPAN GRE-encapsulated traffic from source ports.

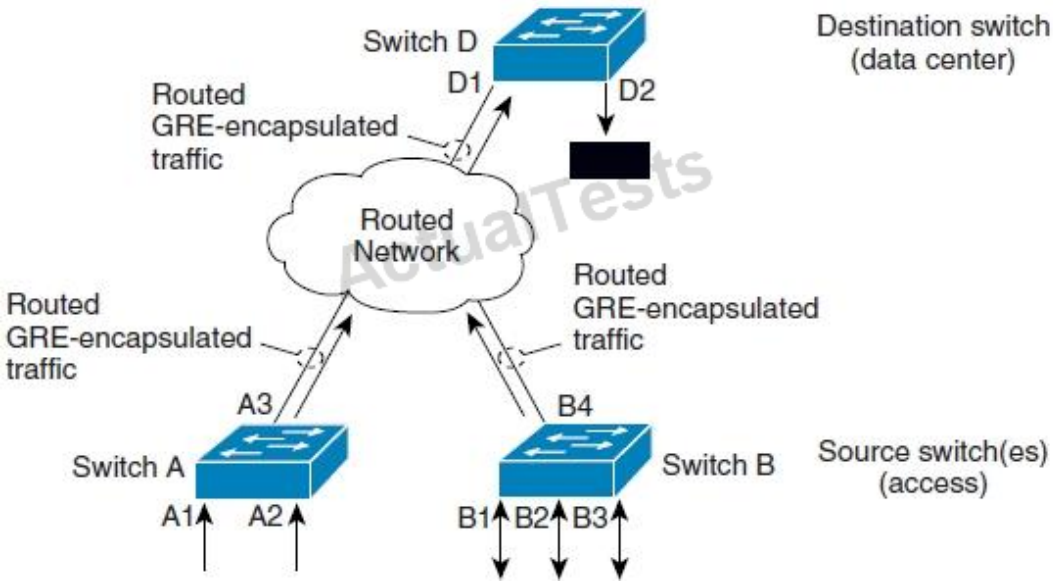
Each ERSPAN source session can have either ports or VLANs as sources, but not both.
The ERSPAN source session copies traffic from the source ports or source VLANs and forwards the traffic using routable GRE-encapsulated packets to the ERSPAN destination session. The ERSPAN destination session switches the traffic to the destinations. supports source ports, source VLANs, and destinations on different switches, which provides remote monitoring of multiple switches across your network. ERSPAN uses a GRE tunnel to carry traffic between switches.

ERSPAN consists of an ERSPAN source session, routable ERSPAN GRE-encapsulated traffic, and an ERSPAN destination session. You separately configure ERSPAN source sessions and destination sessions on different switches. To configure an ERSPAN source session on one switch, you associate a set of source ports or VLANs with a destination IP address, ERSPAN ID number, and optionally with a VRF name. To configure an ERSPAN destination session on another switch, you associate the destinations with the source IP address, ERSPAN ID number, and optionally with a VRF name. ERSPAN source sessions do not copy locally sourced RSPAN

VLAN traffic from source trunk ports that carry RSPAN VLANs. ERSPAN source sessions do not copy locally sourced ERSPAN GRE-encapsulated traffic from source ports.

Each ERSPAN source session can have either ports or VLANs as sources, but not both. The ERSPAN source session copies traffic from the source ports or source VLANs and forwards the traffic using routable GREencapsulated packets to the ERSPAN destination session. The ERSPAN destination session switches the traffic to the destinations.

Figure 68-3 ERSPAN Configuration



Reference
http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/span.pdf

QUESTION NO: 488

What keywords do you need to the access-list to provide to the logging message like source address and source mac address?

A. Log
B. Log-input
C. Log-output
D. Logging

Answer: B

Explanation: The log-input keyword exists in Cisco IOS Software Release 11.2 and later, and in certain Cisco IOS Software

Release 11.1 based software created specifically for the service provider market. Older software does not support this keyword. Use of this keyword includes the input interface and source MAC address where applicable.

Reference
http://www.cisco.com/web/about/security/intelligence/acl-logging.html

QUESTION NO: 489

Which option is true when calculating round-trip delay in IP SLA operations?

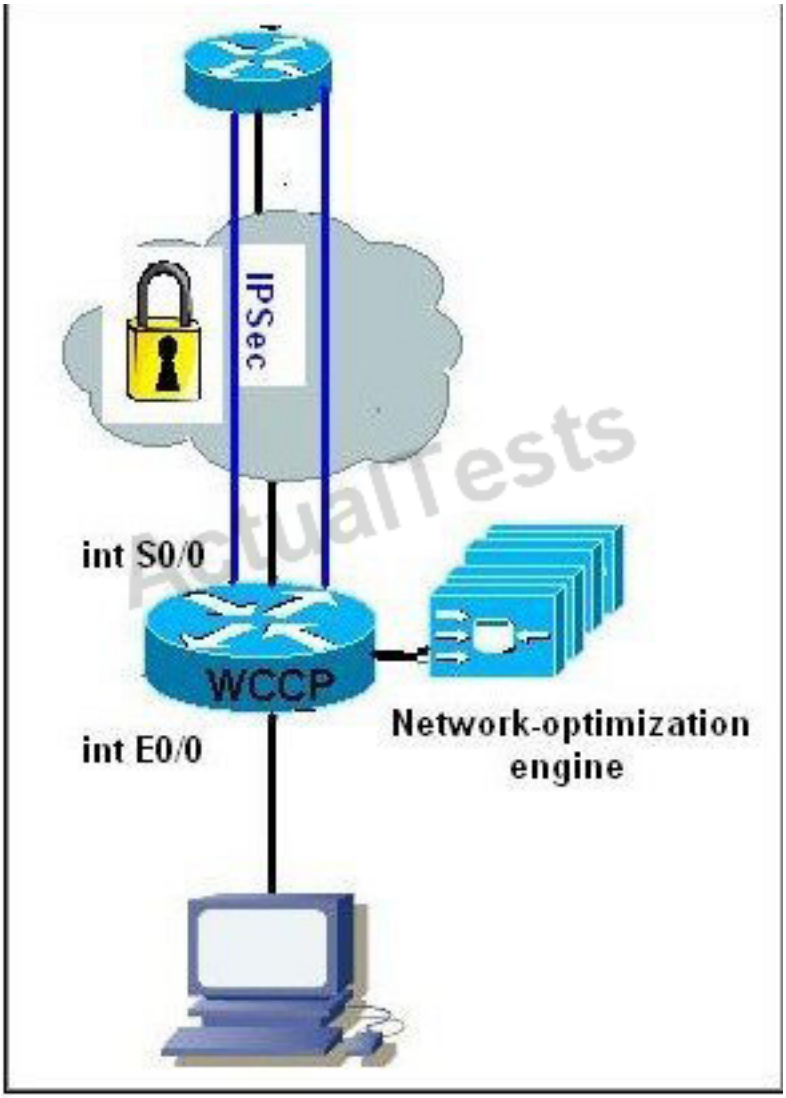
A. The processing time on the end routers is only assessed for operations that involve the responder.
B. The processing time on the end routers is only assessed for operations that involve the transmitter.
C. The processing time on the end routers is only assessed for operations that involve both the respond.
D. The processing time on the end routers is not assessed for neither the responder nor the transmitter.

Answer: A
Explanation:
The Cisco IOS IP SLAs Responder is a component embedded in the destination Cisco routing device that allows the system to anticipate and respond to Cisco IOS IP SLAs request packets. The Cisco IOS IP SLAs Responder provides an enormous advantage with accurate measurements without the need for dedicated probes and additional statistics not available via standard ICMP-based measurements. The patented Cisco IOS IP SLAs Control Protocol is used by the Cisco IOS IP SLAs Responder providing a mechanism through which the responder can be notified on which port it should listen and respond. Only a Cisco IOS device can be a source for a destination IP SLAs Responder.
Figure 2 shows where the Cisco IOS IP SLAs Responder fits in relation to the IP network. The Cisco IOS IP SLAs Responder listens on a specific port for control protocol messages sent by a Cisco IOS IP SLAs operation. Upon receipt of the control message, the responder will enable the specified UDP or TCP port for the specified duration. During this time, the responder accepts the requests and responds to them. The responder disables the port after it responds to the Cisco IOS IP SLAs packet, or when the specified time expires. For added security, MD5 authentication for control messages is available.

Enabling the Cisco IOS IP SLAs Responder on the destination device is not required for all Cisco IOS IP SLAs operations. For example, if services that are already provided by the destination router (such as Telnet or HTTP) are chosen, the Cisco IOS IP SLAs Responder need not be enabled. For non-Cisco devices, the Cisco IOS IP SLAs Responder cannot be configured and Cisco IOS IP SLAs can send operational packets only to services native to those devices.

QUESTION NO: 490

Refer to the exhibit. You are asked to enable redirection for a network optimization engine that will be connected directly to your company CPE. What is the correct configuration to enable redirection for traffic optimization?



A. (config)#interface s0/0
(config-if)#ip wccp 61 out

(config)#interface e0/0
(config-if)#ip wccp 62 out
B. (config)#intetface s0/0
(config-if)#ip wccp 62 in
(config)#interface e0/0
(config-if)#ip wccp 61 in
C. (config)#interface s0/0
(config-if)#ip wccp 61 in
(config-if)#ip wccp 62 out
D. (config)#interface e0/0
(config-if)#ip wccp 61 in
(config-if)#ip wccp 62 out

Answer: D

Explanation:

Configuring WCCP for redirection for inbound traffic on interfaces allows you to avoid the overhead associated with CEF forwarding for outbound traffic. Setting an output feature on any interface results in the slower switching path of the feature being taken by all packets arriving at all interfaces. Setting an input feature on an interface results in only those packets arriving at that interface taking the configured feature path; packets arriving at other interfaces will use the faster default path. Configuring WCCP for inbound traffic also allows packets to be classified before the routing table lookup, which translates into faster redirection of packets.

QUESTION NO: 491

When using IP SLA FTP operation, which two FTP modes are supported? (Choose two.)

A. Only the FTP PUT operation type is supported.
B. Active mode is supported.
C. Passive FTP transfer modes are supported.
D. FTP URL specified for the FTP GET operation is not supported.

Answer: B,C

Explanation:

Both active and passive FTP transfer modes are supported. The passive mode is enabled by default. Only the FTP GET (download) operation type is supported. The URL specified for the FTP GET operation must be in one of the following formats:
If the username and password are not specified, the defaults are anonymous and test, respectively.
FTP carries a significant amount of data traffic and can affect the performance of your network. The results of an IP SLAs FTP operation to retrieve a large file can be used to determine the

capacity of the network but retrieve large files with caution because the FTP operation will consume more bandwidth. The FTP operation also measures your FTP server performance levels by determining the RTT taken to retrieve a file.

QUESTION NO: 492

When running IP SLA, which application type should be used if you want to know round-trip delay, jitter, and packet loss for the full path?

A. ICMP path echo
B. UDP echo
C. ICMP path jitter
D. Application Performance Monitor
E. TCP connect

Answer: C
Explanation:
Before configuring any IP SLAs application, you can use the show ip sla application command to verify that the operation type is supported on your software image. In contrast with other IP SLAs operations, the IP SLAs Responder does not have to be enabled on either the target device or intermediate devices for Path Jitter operations. However, the operational efficiency may improve if you enable the IP SLAs Responder. The IP SLAs ICMP Path Jitter operation is ICMP-based. ICMP-based operations can compensate for source processing delay but cannot compensate for target processing delay. For more robust monitoring and verifying, use of the IP SLAs UDP Jitter operation is recommended. The jitter values obtained using the ICMP Path Jitter operation are approximates because ICMP does not provide the capability to embed processing times on routers in the packet. If the target router does not place ICMP packets as the highest priority, then the router will not respond properly. ICMP performance also can be
affected by the configuration of priority queueing on the router and by ping response.
The path jitter operation does not support hourly statistics and hop information.
Unlike other IP SLAs operations, the ICMP Path Jitter operation is not supported in the RTTMON MIB. Path Jitter operations can only be configured using Cisco IOS commands and statistics can only be returned using the show ip sla commands. The IP SLAs Path Jitter operation does not support the IP SLAs History feature (statistics history buckets) because of the large data volume involved with Jitter operations.

Topic 11, Evaluate proposed changes to a Network

QUESTION NO: 493

You are about to migrate a customer network to use a VSS. Which of these statements is true about a VSS?

A. The VSS switch must be the root bridge for all VLANs and is automatically designated.
B. The VSS switch is defined in RFC 4318 as a managed object.
C. The PAgP+ or LACP protocols are used to maintain the operational state of the VSS devices.
D. A VSS interoperates with a virtual port channel.
E. The 802.1 Q or ISL protocols are used to maintain the operational state of the VSS devices.
F. A VSS increases the size of the spanning-tree domain.

Answer: C
Explanation:
Root Switch and Root Guard Protection
The root of the STP should always be the VSS. Use a statically-defined, hard-coded value for the spanning tree root so that no other switches in the network can claim the root for a given spanning tree domain. Use either Root Guard on a link of VSS-facing access-layer switch or enable it at access-layer switch user port (although the later does not prevent someone from replacing access-layer switch with another switch that can take over as root). The root change might not affect forwarding in non-looped designs (root selection matter only when alternate path (loop) is presented to STP); however, the loss of BPDU or inconstancies generated by a noncompliant switch becoming root could lead to instability in the network.
By default, the active switch's base MAC address is used as the root address of the VSS. This root address does on change during SSO switchover so that an access-layer switch does see the root change.

VSL EtherChannel
Since VSL EtherChannel uses LMP per member link, the link-aggregation protocols, such as PAgP and LACP, are not required; each member link must be configured in unconditional EtherChannel mode using the channelgroup group-number mode on command. Once the VSL configuration is completed, using the switch convert mode virtual CLI command at the enable prompt will start the conversion process. The conversion process includes changing the interface naming convention from slot/interface to switch_number/slot/interface, saving the configuration, and rebooting. During switch rebooting, the systems recognize the VSL configuration and proceeds with their respective VSL ports initialization processes Trunking Configuration Best Practices In a traditional multilayer design featuring standalone switches, when Dynamic Trunking

Protocol (DTP) and 802.1Q or Inter-Switch Link (ISL) negotiation are enabled, considerable time can be spent negotiating trunk settings when a node or interface is restored. During negotiation, traffic is dropped because the link is operational from a Layer-2 perspective. Up to two seconds can be lost depending on where the trunk interface is being brought up.

However, in this configuration, DTP is not actively monitoring the state of the trunk and a misconfigured trunk is not easily identified. There is a balance between fast convergence and your ability to manage your configuration and change control.
In VSS, trunk mode of a port-channel interface being either desirable or undesirable does not exhibit the behavior of standalone node. In VSS, each access-layer is connected via port-channel (MEC), where a link member when brought on line is not a separate negotiation; rather it is an addition to EtherChannel group. The node-related restoration losses are also not an issue when compared to a standalone dual- node design in which each node has a separate control plane that negotiates a separate trunking event. As with VSS, when the node is restored, the link-up event is an additional member link of the MEC and not a trunk interface VSS Virtual Switching System (VSS) is a network virtualization technology that allows two physical Cisco Catalyst 6500 series switches to act as a single logical virtual switch. The VSS increases operational efficiencies and scales bandwidth up to 1.4 Tb/s. This technology is very similar to StackWise technology used with the Cisco Catalyst 3750 series product line, which enables switches stacked together to operate as one and use a single command-line interface (CLI) for management. However, VSS is limited to two physical chassis connected together.

vPC
Virtual Port Channel (vPC) technology works by combining two Cisco Nexus 7000 series switches or two Cisco Nexus 5000 series switches with 10GE links, which are then represented to other switches as a single logical switch for port channeling purposes. With vPC, the spanning-tree topology appears loop- free, although multiple redundant paths are present in the physical topology.

RFC 4318
This memo defines an SMIv2 MIB module for managing the Rapid Spanning Tree capability defined by the IEEE P802.1t and P802.1w amendments to IEEE Standard 802.1D-1998 for bridging between Local Area Network (LAN) segments. The objects in this MIB are defined to apply both to transparent bridging and to bridges connected by subnetworks other than LAN segments.

References
Cisco CCDA Official Certification Guide Fourth Edition
VSS Enabled Campus Design
http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/VSS-dg_ch3.html#wpxref89818

Virtual Switching System (VSS) Q&A
http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps9336/prod_qas0900aecd806ed74b.html
Cisco Catalyst 6500 Virtual Switching System Deployment Best Practices
http://www.cisco.com/en/US/products/ps9336/products_tech_note09186a0080a7c837.shtml
http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/VSS-dg_ch3.html
http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/VSS-dg_ch3.html#wp1079784

QUESTION NO: 494

You have done a partial migration from 802.1D STP to 802.1w STP. Which of the following is true?

A. 802.1 D and 802.1w intemperate only when the 802.1 D STP domain supports rapid convergence.
B. Ports leading to 802.1D devices will run in compatibility mode, while the rest of the ports will run in 802.1w mode.
C. This is an invalid configuration and a partial migration cannot be done.
D. The bridge timers will be set to match the 802.1 D devices.
E. A secondary root bridge will always be populated within the 802.1 D domain.
F. If the root bridge is selected within the 802.1 D domain, the whole STP domain will run in 802.1 D compatibility mode.
G. In partially migrated 802.1w networks, it is recommended to keep the STP diameter below 4.

Answer: B
Explanation:
IEEE 802.1w RSTP is designed to be compatible with IEEE 802.1d STP. Even if all the other devices in your network are using STP, you can enable RSTP on your switch, and even using the default configuration values, your switch will interoperate effectively with the STP devices. If any of the switch ports are connected to switches or bridges on your network that do not support RSTP, RSTP can still be used on this switch. RSTP automatically detects when the switch ports are connected to non-RSTP devices in the spanning tree and communicates with those devices using 802.1d STP BPDU packets.

QUESTION NO: 495

Which two are effects of connecting a network segment that is running 802.1D to a network segment that is running 802.1w? (Choose two.)

A. The entire network switches to 802.1D and generates BPDUs to determine root bridge status.
B. A migration delay of three seconds occurs when the port that is connected to the 802.1D bridge comes up.
C. The entire network reconverges and a unique root bridge for the 802.1D segment, and a root bridge for the 802.1w segment, is chosen.
D. The first hop 802.1w switch that is connected to the 802.1D runs entirely in 802.1D compatibility mode and converts the BPDUs to either 802.1D or 802.1w BPDUs to the 802.1D or 802.1w segments of the network.
E. Classic 802.1D timers, such as forward delay and max-age, will only be used as a backup, and will not be necessary if point-to-point links and edge ports are properly identified and set by the administrator.

Answer: B,E

Explanation:

Each port maintains a variable that defines the protocol to run on the corresponding segment. A migration delay timer of three seconds also starts when the port comes up. When this timer runs, the current STP or RSTP mode associated to the port is locked. As soon as the migration delay expires, the port adapts to the mode that corresponds to the next BPDU it receives. If the port changes its mode of operation as a result of a BPDU received, the migration delay restarts.
802.1D works by the concept that the protocol had to wait for the network to converge before it transitioned a port into the forwarding state. With Rapid Spanning Tree it does not have to rely on any timers, the only variables that that it relies on is edge ports and link types.

Any uplink port that has an alternate port to the root can be directly placed into the forwarding state (This is the Rapid convergence that you speak of "restored quickly when RSTP is already in use?"). This is what happened when you disconnected the primary look; the port that was ALT, moved to FWD immediately, but the switch also still needs to create a BDU with the TC bit set to notify the rest of the network that a topology has occurred and all non-edge designated ports will transition to BLK, LRN, and then FWD to ensure there are no loops in the rest of the network. This is why if you have a host on a switchport, and you know for a fact that it is only one host, enable portfast to configure the port as an edgeport so that it does not have to transition to all the STP states.

Reference
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_white_paper09186a0080094cfa.shtml

QUESTION NO: 496

You add the following commands into a routed topology: router eigrp 1 variance 3 traffic-share min across-interfaces. Users now complain about voice quality in your VoIP system. What should be done?

A. Add the command: router eigrp 1 traffic-share voice interface fast 0/0.
B. Reconfigure EIGRP to recognize voice packets.
C. Remove the variance from the configuration.
D. Reconfigure the VoIP system to use RTP sequence number headers.
E. Use an H.323 gatekeeper for your VoIP system to negotiate an H.245 uneven packet buffer.
F. Reconfigure EIGRP to version 2.

Answer: C

Explanation:

Traffic-share min command causes EIGRP to divide traffic only among the routes with the best metric. When the traffic-share min command is used with the across-interfaces keyword, an attempt is made to use as many different interfaces as possible to forward traffic to the same destination.

Therefore with the configuration above, EIGRP will only use equal-cost load-balancing feature even when the variance command is used. However, if you use both the traffic-share min command and variance command, even though traffic is sent over the minimum-cost path only, all feasible routes get installed into the routing table, which decreases the convergence times.

QUESTION NO: 497

You are the network administrator of a medium-sized company, and users are complaining that they cannot send emails to some organizations. During your troubleshooting, you notice that your DNS MX record is blacklisted by several public blacklist filters. After clearing these listings for your IP address, and assuming that your email server has the right virus protection in place, what are two possible solutions to prevent this from happening in the future? (Choose two.)

A. Change your Internet provider.
B. Change your public IP address.
C. Allow the email server to send traffic only to TCP port 25.
D. Put your email server in a DMZ.
E. Use a separate public IP address for your email server only.

Answer: C,E

Reference: http://www.parkansky.com/tutorials/dmz.htm

QUESTION NO: 498

A router that acts as an Internet border gateway has multiple upstream connections that are used in a load-sharing setup. The NOC has identified a DDoS attack from a specific source entering its network via interface GigabitEthernet0/1.

The NOC wants to block this suspicious traffic on the border router in a scalable way and without major changes to the different interface configurations.

Which configuration would block the DDoS attack from the known source (194.90.1.5)?

A.
interface GigabitEthernet0/1
ip address 192.168.1.1 255.255.255.252
ip verify unicast source reachable-via any
!
ip route 194.90.1.5 255.255.255.255 Null0
B.
interface GigabitEthernet0/1
ip address 192.168.1.1 255.255.255.252
ip verify unicast source reachable-via any
!
ip route 194.90.1.5 255.255.255.255 192.168.1.2
C.
interface GigabitEthernet0/1
ip address 192.168.1.1 255.255.255.252
ip verify unicast source reachable-via rx
!
ip route 194.90.1.5 255.255.255.255 Null0
D.
interface GigabitEthernet0/1
ip address 192.168.1.1 255.255.255.252
ip verify unicast source reachable-via rx
!
ip route 194.90.1.5 255.255.255.255 192.168.1.2

Answer: A

Explanation:

QUESTION NO: 499

Refer to exhibits 1 and 2. A company uses a Metro Ethernet (Gigabit Ethernet) dedicated circuit to communicate between users (subnet B) and servers (subnet A) as shown in Exhibit 1. Both routers use OSPF to advertise the subnets. During a weekly management meeting, they realize

that the WAN link is oversize. They have been using only 2 Mb/s in the worst-case scenario. So they propose a new, cheaper WAN connection using a 2-Mb/s Frame-Relay point-to-point link to interconnect both sites (Exhibit 2). The Frame Relay service provider informs them that multicast traffic is not allowed to run over the service provider network. Which one of these options is best to enable the company to establish the OSPF neighbor adjacency?

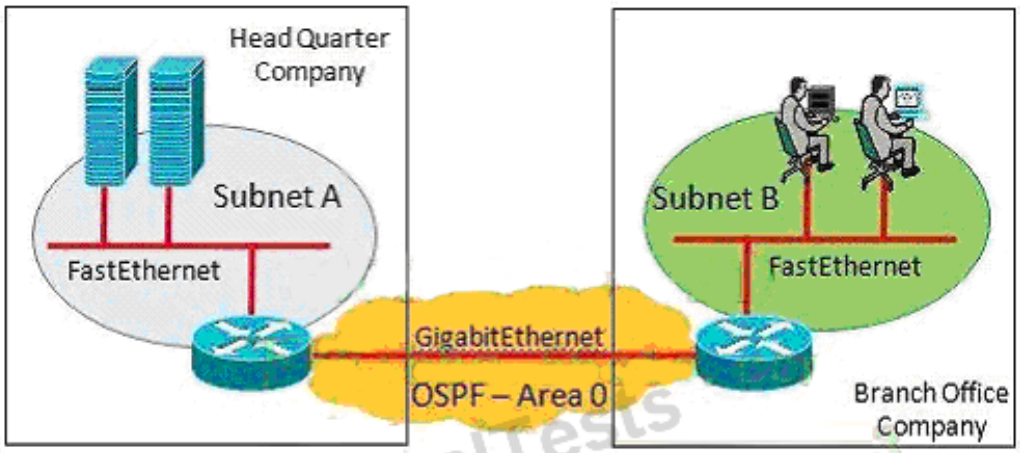
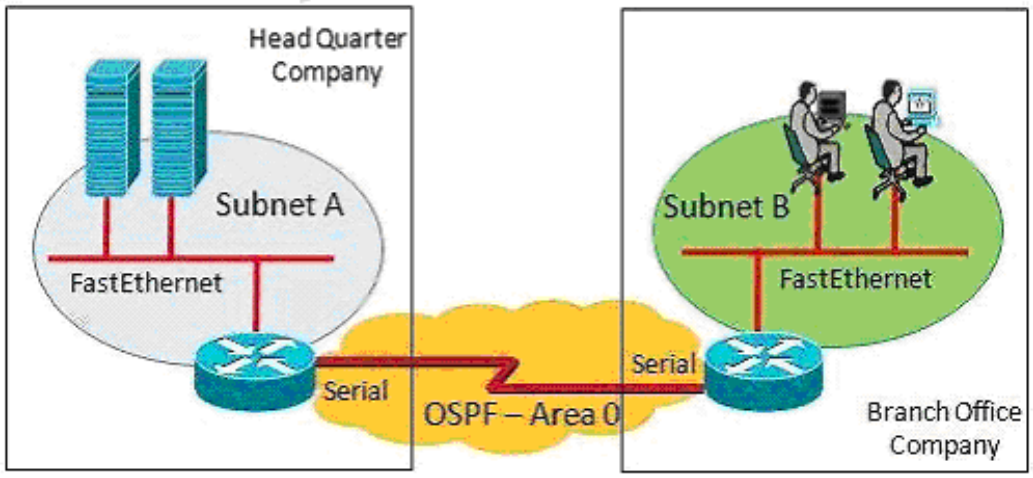


Exhibit 1



A. Use OSPF network broadcast, because it uses unicast to establish a neighbor relationship.
B. Use OSPF network point-to-multipoint, because it uses unicast to establish a neighbor relationship.
C. Use OSPF network point-to-point, because it uses unicast to establish a neighbor relationship.
D. Use OSPF network point-to-multipoint nonbroadcast, because it establishes a neighbor relationship using unicast packets.
E. Use OSPF network nonbroadcast, because it establishes a neighbor relationship using multicast.

Answer: D

Explanation:

OSPF Point-to-Multipoint Network with Separate Costs per Neighbor

Description

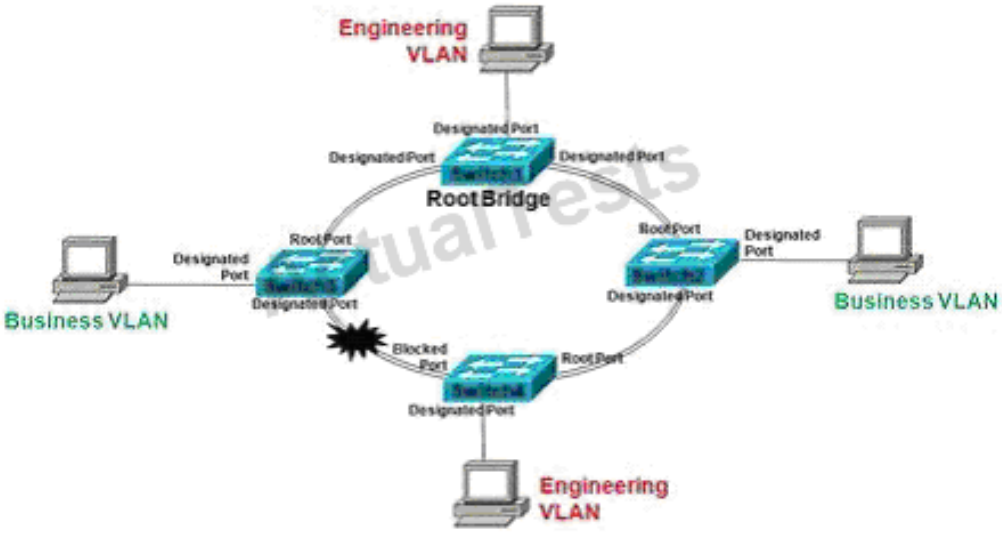
OSPF has two new features related to point-to-multipoint networks. One feature applies to broadcast networks; the other feature applies to nonbroadcast networks.
On point-to-multipoint, broadcast networks, there is no need to specify neighbors. However, you can specify neighbors with the neighbor command, in which case you should specify a cost to that neighbor.

On point to multipoint, nonbroadcast networks, you now use the neighbor command to identify neighbors.
Assigning a cost to a neighbor is optional.
Before this feature, some OSPF point-to-multipoint protocol traffic was treated as multicast traffic. Therefore, the neighbor command was not needed for point-to-multipoint interfaces because multicast took care of the traffic. Hellos, updates and acknowledgments were sent using multicast. In particular, multicast hellos discovered all neighbors dynamically.
However, some customers were using point-to-multipoint on nonbroadcast media (such as classic IP over ATM), so their routers could not dynamically discover their neighbors. This feature allows the neighbor command to be used on point-to-multipoint interfaces.
On any point-to-multipoint interface (broadcast or not), the Cisco IOS software assumed the cost to each neighbor was equal. The cost was configured with the ip ospf cost command. In reality, the bandwidth to each neighbor is different, so the cost should be different. With this feature, you can configure a separate cost to each neighbor. This feature applies to point-to-multipoint interfaces only.

Benefits
You can now configure neighbors on point-to-multipoint interfaces and assign a cost to each neighbor. These capabilities allow the router to dynamically discover neighbors over nonbroadcast media and to prefer some routes over others by assigning different costs to neighbors.
Reference
http://www.cisco.com/en/US/docs/ios/11_3/feature/guide/ospfpmp.html

QUESTION NO: 500

Refer to the exhibit. Users from the Engineering VLAN complain that every time Business VLAN users have a network connectivity issue, the Engineering VLAN users usually have problems experiencing slow response or network connectivity problems. After troubleshooting, an unauthorized switch 2 was found. This unauthorized switch has been a regular problem, assuming the root bridge function under the spanning-tree domain and causing the Engineering VLAN to be unstable. Which three of these actions could be suggested to fix the problem?



A. Upgrade Spanning Tree Protocol to Rapid Spanning Tree Protocol.
B. Change Business VLAN PCs to switch 1 and switch 4.
C. Force the root bridge to be switch 2, instead.
D. Adjust spanning-tree timers (max-age and forward-delay).
E. Shut down all unused ports.
F. Use MSTP to separate the Engineering VLAN from the Business VLAN to optimize spanning-tree convergence time within each VLAN

Answer: A,E,F

Explanation:

The best action is to shut down unused ports and use MSTP to separate Engineering VLAN from the business VLAN so that spanning-tree convergence time can be optimized within each VLAN. Also upgrade spanning tree protocol to Rapid spanning tree protocol (RSTP).

QUESTION NO: 501

You are deploying two core switches, one in each building, 50 km away from each other. The cross-connection between them will be a Layer 2 2-gigabit EtherChannel with an 802.1Q trunk. You configured it correctly but the link does not come up. The port is in the "admin up" state, and the line protocol is in the "down" state. The fiber link is OK. What would be the most likely reason for the link not to come up?

A. The switches are not the same model.
B. You are not using the correct SFP.
C. You are not using correct optical media converters.
D. Configuration should be modified, because the distance is longer.

Answer: B

Explanation:

Verifying the Line Protocol Is Up

In the output from the show interfaces fastethernet, show interfaces gigabitethernet or show interfaces

tengigabitethernet command, verify that the line protocol is up. If the line protocol is down, the line protocol software processes have determined that the line is unusable.

Perform the following corrective actions:

Replace the cable.

Check the local and remote interface for misconfiguration.

Verify that a hardware failure has not occurred. Observe the LEDs to confirm the failure. See the other troubleshooting sections of this chapter, and refer to the Cisco 7600 Series Router SIP, SSC, and SPA Hardware Installation Guide. If the hardware has failed, replace the SPA as necessary.

Reference

http://www.cisco.com/en/US/products/hw/routers/ps368/module_installation_and_configuration_guides_chapter09186a0080523f70.html#wp1038764

QUESTION NO: 502

Your company is researching a new application that runs over IPv6, but part of it must still have IPv4 support. Your company uses a traditional IPv4 network. Your plan is not to run IPv6 over the whole network, but to segment parts of the network or even to operate simultaneously with IPv6 and IPv4. You must make a brief presentation about IPv6 technology to the board of technical directors. Which three of these items could be part of your presentation? (Choose three.)

A. Tunnel IPv6 over IPv4 to connect far-end IPv6 networks.
B. Explain why configuring IPv4 and IPv6 at the same time over the same LAN interface is not possible.
C. Explain why configuring IPv4 and IPv6 at the same time over the same LAN interface is possible.
D. What is the meaning of EUI-64 and how does it work?
E. Tunnel IPv4 over IPv6 to connect far-end IPv4 networks.

Answer: A,C,D

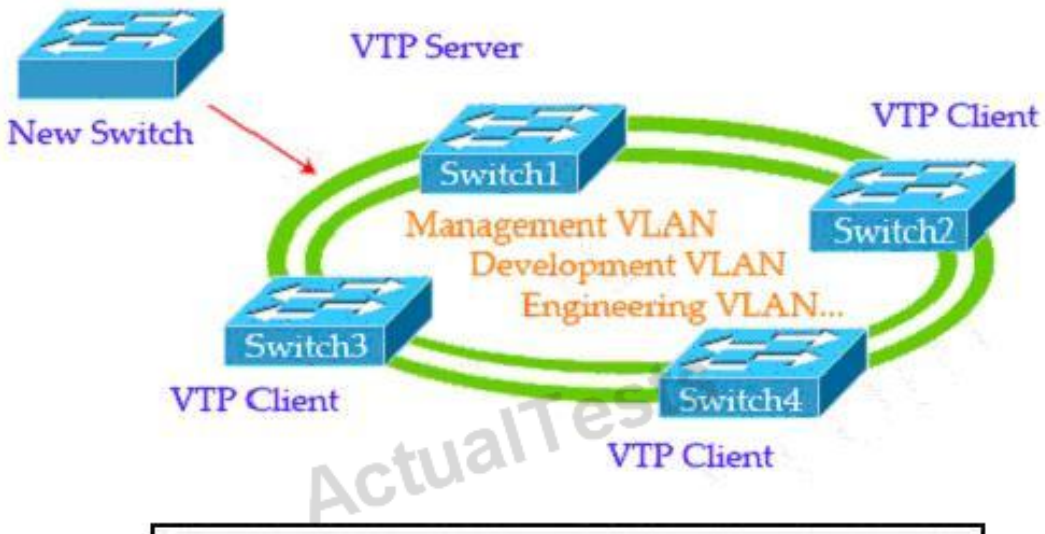
Explanation:

An interface ID is used to identify interfaces on a link. The interface ID must be unique to the link. It may also be unique over a broader scope. In many cases, an interface ID will be the same as or

based on the link-layer address of an interface. Interface IDs used in aggregatable global unicast and other IPv6 address types must be 64 bits long and constructed in the modified EUI-64 format. Interface IDs are constructed in the modified EUI-64 format in one of the following ways:
For all IEEE 802 interface types (for example, Ethernet, and FDDI interfaces), the first three octets (24 bits) are taken from the Organizationally Unique Identifier (OUI) of the 48-bit link-layer address (the Media Access Control [MAC] address) of the interface, the fourth and fifth octets (16 bits) are a fixed hexadecimal value of FFFE, and the last three octets (24 bits) are taken from the last three octets of the MAC address. The construction of the interface ID is completed by setting the Universal/Local (U/L) bit the seventh bit of the first octet a value of 0 or 1. A value of 0 indicates a locally administered identifier; a value of 1 indicates a globally
unique IPv6 interface identifier.
For other interface types (for example, serial, loopback, ATM, Frame Relay, and tunnel interface types except tunnel interfaces used with IPv6 overlay tunnels), the interface ID is constructed in the same way as the interface ID for IEEE 802 interface types; however, the first MAC address from the pool of MAC addresses in the router is used to construct the identifier (because the interface does not have a MAC address).
For tunnel interface types that are used with IPv6 overlay tunnels, the interface ID is the IPv4 address assigned to the tunnel interface with all zeros in the high-order 32 bits of the identifier.
An IPv4-compatible IPv6 address is an IPv6 unicast address that has zeros in the high-order 96 bits of the address and an IPv4 address in the low-order 32 bits of the address. The format of an IPv4- compatible IPv6 address is 0:0:0:0:0:0:A.B.C.D or ::A.B.C.D. The entire 128-bit IPv4-compatible IPv6 address is used as the IPv6 address of a node and the IPv4 address embedded in the low-order 32 bits is used as the IPv4 address of the node. IPv4-compatible IPv6 addresses are assigned to nodes that support both the IPv4 and IPv6 protocol stacks and are used in automatic tunnels.

QUESTION NO: 503

Refer to the exhibit.




```
Switch# show vtp status

VTP Version                     : 7
Configuration Revision          : 15
Maximum VLANs supported local   : 1015
Number of existing VLANs        : 10
VTP Operating Mode              : Server
VTP Domain Name                 : corp
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x22 0xF3 0x1A
Configuration last modified by 172.18.22.15 at 5-28-03 11:53:20
```

The Layer 2 network uses VTP to manage its VLAN database. A network designer created all VLANs on the VTP server (switch 1) and it has been advertised through VTP to all other VTP clients (switches 2 through 4). Due to network growth, a network operator decided to add a new switch between switch 1 and switch 3. The network operator has been instructed to use a refurbished switch and use a VTP client. Which three of these has been instructed to use a refurbished switch and use a VTP client. Which three of these factors should the network operator consider to minimize the impact of adding a new switch? (Choose three.)

A. Pay special attention to the VTP revision number, because the higher value takes the priority.
B. Configure all VLANs manually on the new switch in order to avoid connectivity issues.
C. A trunk should be established between the new switch and switches 1 and 3 as VTP only runs over trunk links.
D. Set at least the VTP domain name and password to get the new switch synchronized.
E. An ISL trunk should be established between the new switch and switches 1 and 3, because VTP only runs over ISL.
F. Pay special attention to the VTP revision number, because the lower value takes the priority.

Answer: A,C,D

Explanation:

VTP should be used whenever we have more than 1 switch with multiple VLANs. It helps us save much time so configuring all VLANs manually is just a waste of time -> B is not correct.

VLAN Trunking Protocol (VTP) can operate over 802.1q or ISL on FastEthernet link.+ On ISL: **Switch(config-if)#**switchport trunk encapsulation isl + On 802.1q: **Switch(config-if)#**switchport trunk encapsulation dot1q

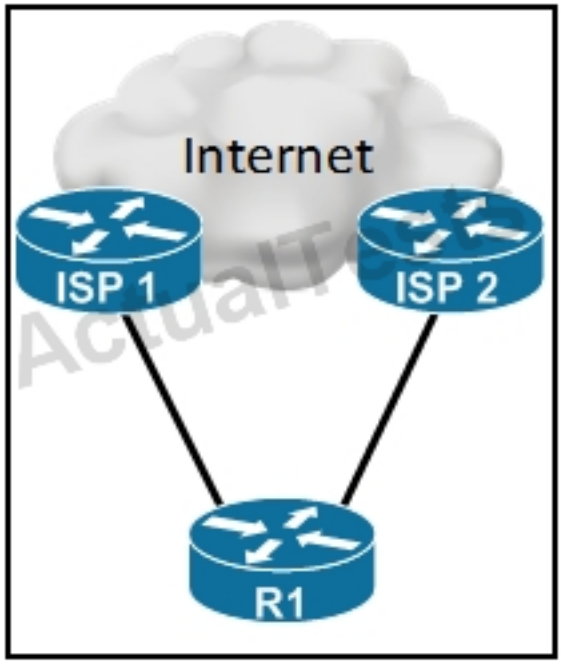
-> E is not correct

Note: The 2940/2950 switches only support 802.1q encapsulation with the **switchport mode trunk** command. The switch will automatically use 802.1q encapsulation.

Each time a VTP updates are sent out, the revision number is increased by 1. Any time a switch sees a higher revision number, it knows the information that it's receiving is more current, and it will overwrite the current database with that new information.

QUESTION NO: 504

Refer to the exhibit.



R1 has two eBGP sessions to ISP1 and ISP2 (one to each ISP router), and R1 receives the same prefixes through both links.

Which configuration should be applied for the link between R1 and ISP2 to be preferred for incoming traffic (ISP2 to R1)?

A. increase local preference on R1 for routes advertised to ISP2
B. decrease local preference on R1 for routes advertised to ISP2
C. increase MED on R1 for routes advertised to ISP2
D. decrease MED on R1 for routes advertised to ISP2

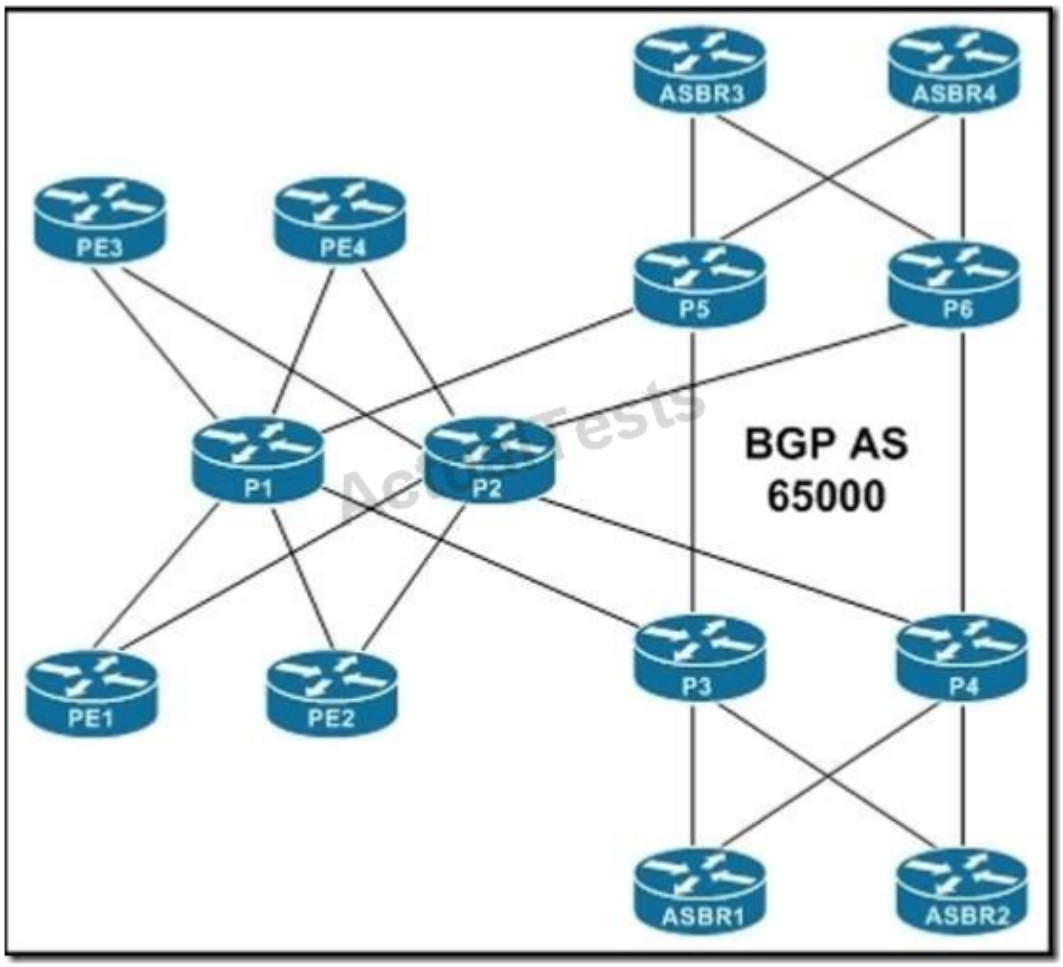
Answer: D

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.
Reference:
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094934.shtml

QUESTION NO: 505

Refer to the exhibit.



AS65000 has core network P1 - P6. The eBGP peers to another AS are through ASBR1-2 and ASBR3-4. All business and residential customer POPs are connected to the P1 - P2 core.

AS65000 has decided to enable MPLS LDP on all BGP routers, and disable BGP on the core network (P1 - P6). What is the drawback of this BGP core-free network?

A. MPLS LDP is slower compared to BGP.
B. All IP addresses of nodes and links within AS65000 are exposed to any external AS.
C. The BGP configuration must be full-mesh between the ASBR and PE routers.
D. The core router will not be involved in the BGP path selection process.

Answer: D

Explanation:

Although removing the core routers so that they are not burdened with running BGP is a goal of many service providers, one drawback is that they are then no longer participating in the BGP path selection process.