



642-891 (Composite®)

TestKing's Composite Exam

Version 22.0

Leading The Way
in IT Testing And Certification Tools

www.testking.com

Important Note Please Read Carefully

Study Tips

This product will provide you questions and answers along with detailed explanations carefully compiled and written by our experts. Try to understand the concepts behind the questions instead of cramming the questions. Go through the entire document at least twice so that you make sure that you are not missing anything.

Further Material

For this test TestKing plans to provide:

* Study Guide. Concepts and labs.

* Interactive Test Engine Examiner. Check out an Examiner Demo at <http://www.testking.com/index.cfm?pageid=724>

Latest Version

We are constantly reviewing our products. New material is added and old material is revised. Free updates are available for 90 days after the purchase. You should check your member zone at TestKing an update 3-4 days before the scheduled exam date.

Here is the procedure to get the latest version:

1. Go to www.testking.com
2. Click on Member zone/Log in
3. The latest versions of all purchased products are downloadable from here. Just click the links.

For most updates, it is enough just to print the new questions at the end of the new version, not the whole document.

Feedback

Feedback on specific questions should be send to feedback@testking.com. You should state: Exam number and version, question number, and login ID.

Our experts will answer your mail promptly.

Copyright

Each pdf file contains a unique serial number associated with your particular name and contact information for security purposes. So if we find out that a particular pdf file is being distributed by you, TestKing reserves the right to take legal action against you according to the International Copyright Laws.

Table of contents

Part 1, Building Scalable Cisco® Internetworks (BSCI®) (635 questions).....	7
Topic 1: Technology (183 questions)	7
Section 1: List the key information routers needs to route data (8 questions).....	7
Section 2: Describe classful and classless routing protocols (3 questions)	13
Section 3: Describe link-state router protocol operation.....	15
Section 4: Compare classful and classless routing protocols (3 questions).....	18
Section 5: Compare distance vector and link state routing protocols (9 questions).....	21
Section 6: Describe concepts relating to extending IP addresses and the use of VLSMs to extend IP addresses (8 questions)	28
Section 7: Describe the features and operation of EIGRP (21 questions)	34
Section 8: Describe the features and operation of single area OSPF (15 questions).....	50
Section 9: Describe the features and operation of multi-area OSPF (26 questions).....	61
Section 10: Explain basic OSI terminology and network layer protocols used in OSI (5 questions)	79
Section 11: Identify similarities and differences between Integrated IS-IS and OSPF (6 questions)	82
Section 12: List the types of IS-IS routers and their role in IS-IS area design (5 questions) ...	86
Section 13: Describe the hierarchical structure of IS-IS areas (22 questions).....	89
Section 14: Describe the concept of establishing adjacencies (4 questions)	104
Section 15: Describe the features and operation of BGP (29 questions).....	107
Section 16: Explain how BGP policy-based routing functions within an autonomous system (5 questions)	126
Section 17: Explain the use of redistribution between BGP and Interior Gateway Protocols (IGPs) (9 questions)	130
Topic 2: Implementation and Configuration (136 questions).....	138
Section 1: Given a set of network requirements, identify the steps to configure an Enhanced IGRP environment and verify proper operation (within described guidelines) of your routers (19 questions).....	138
Section 2: Given an addressing scheme and other laboratory parameters, identify the steps to configure a single-area OSPF environment and verify proper operation (within described guidelines) of your routers (8 questions)	154
Section 3: Given an addressing scheme and other laboratory parameters, identify the steps to configure a multiple area OSPF environment and verify proper operation (within described guidelines) of your routers (17 questions)	162
Section 4: Given an addressing scheme and other laboratory parameters, identify the steps to configure Cisco routers for proper Integrated IS-IS operation (10 questions)	176
Section 4: Identify the steps to select and configure the different ways to control routing update traffic (22 questions)	185
Section 5: Identify the steps to configure router redistribution in a network (22 questions)..	203
Section 6: Identify the steps to configure policy-based routing using route maps (13 questions)	221

Section 7: Given a set of network requirements, identify the steps to configure a BGP environment and verify proper operation (within described guidelines) of your routers (25 questions)	233
Section 8: Identify the steps to configure a router for Network Address Translation with overload, static translations, and route maps. (4 questions)	256
Topic 3: Design (41 questions).....	261
Section 1, Describe the three-layer hierarchical design model and explain the function of each layer: Access, Distribution and Core (4 questions)	261
Section 2: Given specific requirements, choose the correct routing protocol to meet the requirements (14 questions).....	264
Section 3: Identify the correct IP addressing scheme, including features of IPv6 (8 questions)	273
Section 4: Describe the concepts relating to route summarization and apply them to hypothetical scenarios (15 questions)	278
Topic 4: Troubleshooting (50 questions).....	287
Section 1: Identify the steps to verify OSPF operation in a single area (7 questions).....	287
Section 2: Identify verification methods which ensure proper operation of Integrated IS-IS on Cisco routers (4 questions).....	293
Section 3: Identify the steps to verify OSPF operation in multiple areas (8 questions)	296
Section 4: Describe the scalability problems associated with internal BGP (4 questions).....	303
Section 5: Interpret the output of various show and debug commands to determine the cause of route selection errors and configuration problems (20 questions).....	306
Section 6: Identify the steps to verify Enhanced IGRP operation (7 questions)	322
Topic 5: Mixed Questions (19 Questions).....	328
Topic 6: Extra practice Questions (200 questions)	341
Part 2, Building Cisco® Multilayer Switched Networks (BCMSN®) (411 questions)	439
TOPIC 1: TECHNOLOGY (205 questions).....	439
Section 1: Describe the Enterprise Composite Model used for designing networks and explain how it addresses enterprise network needs for performance, scalability and availability (1 question)	439
Section 2, Describe the physical, data-link and network layer technologies used in a switched network, and identify when to use each (10 questions)	441
Section 3: Explain the role of switches in the various modules of the Enterprise Composite Model (Campus Infrastructure, Server Farm, Enterprise Edge, Network Management) (2 questions)	448
Section 4: Explain the function of the Switching Database Manager [specifically Content Addressable Memory (CAM) and Ternary Content Addressable Memory (TCAM)] within a Catalyst switch (1 question).....	451
Section 5: Describe the features and operation of VLANs on a switched network (6 questions)	452
Section 6: Describe the features of the VLAN trunking protocols including 802.1Q, ISL (emphasis on 802.1Q) and dynamic trunking protocol (17 questions).....	456
Section 7: Describe the features and operation of 802.1Q Tunneling (802.1QinQ) within a service provider network (4 questions).....	467
Section 8: Describe the operation and purpose of managed VLAN services (3 questions) ...	471

Section 9: Describe how VTP versions 1 and 2 operate including domains, modes, advertisements, and pruning (26 questions).....	474
Section 10: Explain the operation and purpose of the Spanning-Tree Protocol (STP) on a switched network (35 questions)	489
Section 11: Identify the specific types of Cisco route switch processors, and provide implementation details (18 questions)	511
Section 12: List and describe the operation of the key components required to implement interVLAN routing (5 questions).....	523
Section 13: Explain the types of redundancy in a multilayer switched network including hardware and software redundancy (6 questions).....	526
Section 14: Explain how IP multicast operates on a multilayer switched network, including PIM, CGMP and IGMP (32 questions)	530
Section 15: Describe the quality issues with voice traffic on a switched data network, including jitter and delay (3 questions).....	550
Section 16: Describe the QoS solutions that address voice quality issues (6 questions).....	552
Section 17: Describe the features and operation of network analysis modules on Catalyst switches to improve network traffic management (8 questions)	556
Section 18: Describe Transparent LAN Services and how they are implemented in a service provider network (1 question).....	561
Topic 2: IMPLEMENTATION (135 questions).....	562
Section 1: Convert CatOS to native IOS on Catalyst switches and manage native IOS images according to best practices (2 questions)	562
Section 2: Configure access ports for static and multi-VLAN membership (20 questions)...	564
Section 3: Configure and verify 802.1Q trunks (5 questions)	578
Section 4: Configure and verify ISL trunks (6 questions)	583
Section 5: Configure VTP domains in server, client and transparent modes (23 questions)..	587
Section 6: Enable Spanning Tree on ports and VLANs (7 questions).....	603
Section 7: Configure Spanning Tree parameters including: port priority, VLAN priority, root bridge, BPDU guard, PortFast and UplinkFast (17 questions)	610
Section 8: Implement IP technology on a switched network with auxiliary VLANs (7 questions)	626
Section 9: Configure and verify router redundancy using HSRP, VRRP, GLBP, SRM, and SLB (20 questions)	631
Section 10: Configure QoS features on multilayer switched networks to provide optimal quality and bandwidth utilization for applications and data (4 questions)	646
Section 11: Configure Fast EtherChannel and Gigabit EtherChannel to increase bandwidth for interswitch connections (10 questions).....	649
Topic 3: DESIGN (28 questions).....	656
Section 1: Compare end-to-end and local VLANs, determine when to use each (5 questions)	656
Section 2: Design a VLAN configuration with VTP to work for a given specific scenario (4 questions)	659
Section 3: Plan QoS implementation within a multilayer switched network (3 questions)....	662
Section 4: Select multilayer switching architectures, given specific multilayer switching needs (15 questions).....	664

Section 5: Describe the general design models when implementing IP telephony in a switched network environment (2 questions)	673
TOPIC 4: TROUBLESHOOTING (36 questions)	675
Section 1: Verify and Troubleshoot common VLAN problems on a switched network (7 questions)	675
Section 2: Tune and troubleshoot spanning-tree protocol on a multilayer switched network to enhance network performance, prevent network loops, and minimize downtime (16 questions)	679
Section 3: Identify interVLAN routing performance and scalability issues, and propose solutions (1 question)	691
Section 4: Verify and troubleshoot interVLAN routing on a switched network (2 questions)	692
Section 5: Identify QoS implementation issues at the network access layer (3 questions)	694
Section 6: Identify QoS implementation issues at the network distribution and core layers (3 questions)	696
Topic 5, Mixed Questions (45 Questions)	698

Total number of questions: 1046

Part 1, Building Scalable Cisco® Internetworks (BSCI®) (635 questions)

Topic 1: Technology (183 questions)

Section 1: List the key information routers needs to route data (8 questions)

QUESTION NO: 1

A packet that needs to be forwarded arrives on an interface of a router. In order for a router to route data, what must that router determine? (Select the best answer)

- A. The route age of the next-hop device
- B. The subnet mask of the source network
- C. The cost metric of the path of the destination
- D. The outbound interface of the best path to the destination
- E. All of the above

Answer: D

Explanation:

In order to route the data to the correct destination the router must know the outbound interface that provides the best destination. When a receives a packet that needs to be forwarded, the router determines the destination network, looks up the best path to that destination in the routing table, and then forwards the packet out the correct interface.

Incorrect Answers:

- A. The age of the next-hop router is not consulted before forwarding information.
- B. The destination subnet mask is used by the router to determine the best path (most specific network match rule), but not the source subnet mask.
- C. The cost metric is used by routing protocols to determine the best route to a destination. This best route is then installed into the routing table. The router uses the routing table to forward packets, but does not use the metrics for that specific route before forwarding each packet.

QUESTION NO: 2

A new router is being installed into an existing network, and the routing table is being built for the first time on this network. Which of the following statements is true regarding the routing tables on a Cisco router?

- A. Entries are listed in the order of the route cost metric.
- B. Only the active link is shown for load-balanced routers.
- C. Privileged EXEC mode is required to view the routing table.
- D. The `clear ip route *` command refreshes the entire routing table.
- E. All of the above are true.

Answer: C

Explanation:

The **show ip route** command is used to display IP routing table entries. It can only be run in privileged (enable) mode.

Incorrect Answers:

- A. Entries are not listed in route cost order. Entries are ordered by destination IP address.
- B. The other links in the load balanced bundle are also shown.
- D. The `clear ip route` command is used to delete IP routing table entries. In particular the **clear ip route *** (or **clear ip route all**) command deletes IP routing table entries. This is not a refresh. Static routes will not be recreated.

QUESTION NO: 3

Which of the following phrases is the correct term for what happens to a network when a topology change causes all the routers to synchronize their routing tables?

- A. Flooding
- B. Broadcasting
- C. Convergence
- D. Summarization
- E. None of the above

Answer: C

Explanation: A topology change forces a convergence of the routers to a new routing state. Convergence is the time that it takes for all routers to agree on the network topology after a change.

Incorrect Answers:

- A. Flooding is the process of sending out routing topology information used by link state protocols.
- B. Broadcasting is used on LAN level communications. By default, routers do not forward broadcasts.
- D. Summarization is the act of taking multiple routes within the routing table, and advertising them as one less specific route.

QUESTION NO: 4

What is the purpose of configuring a router with the “IP Helper address” command?

- A. IP Helper is used to direct BOOTP clients to a BOOTP server.
- B. IP Helper is used to prevent the router from forwarding IP broadcasts.
- C. IP Helper is used to allow IPX clients to communicate with IP-based servers.
- D. IP Helper is used to accommodate compatibility routers using different IP routing protocols.

Answer: A

Explanation:

The **ip helper-address** command is used to have the Cisco IOS software forward User Datagram Protocol (UDP) broadcasts, including BOOTP, received on an interface. DHCP protocol information is carried inside of BOOTP packets. To enable BOOTP broadcast forwarding for a set of clients, configure a helper address on the router interface closest to the client. The helper address should specify the address of the DHCP server.

Note: A DHCP server can be considered to be a BOOTP server, even though a DHCP server is more advanced.

Incorrect Answers:

- B: Combined with the **ip forward-protocol** global configuration command, the **ip helper-address** command allows you to control which broadcast packets and which protocols are forwarded. However, the main purpose of the IP helper feature is not to prevent the router from forwarding IP broadcasts.
- C: IP helper does not use IPX.
- D: This is false.

QUESTION NO: 5

On router TK1 the command “ip helper address” is already configured. Which of the following commands would you use if you wanted to send SNMP broadcast packets off to a specific server?

- A. ip server udp 161
- B. ip helper-protocol 161
- C. ip forward-protocol 161
- D. ip directed-broadcast 161
- E. ip forward snmp

Answer: C

Explanation: SNMP requests are typically sent to UDP port 161.

The **ip forward-protocol** command is used to specify which protocols and ports the router forwards when forwarding broadcast packets. If an IP helper address is defined, UDP forwarding is enabled on default ports.

Note: Syntax: `ip forward-protocol {udp [port] | nd | sdns}`

Incorrect Answers:

A: The use of **ip server udp** is not correct.

B: There is no such command.

D: The **ip directed-broadcast** command is used to enable the translation of directed broadcast to physical broadcasts. It does not apply in this scenario. Furthermore, an access-list with number 161 has to be configured.

Syntax: `ip directed-broadcast [access-list-number] | [extended access-list-number]`

QUESTION NO: 6

When you execute the “ip helper-address” command on a router, which three UDP ports get enabled automatically by default? (Select three)

- A. 53 (DNS)
- B. 69 (TFTP)
- C. 515 (LPR)
- D. 161 (SNMP)
- E. 49 (TACACS)

Answer: A, B, E

Explanation:

To forward the BootP/DHCP request from the client to the DHCP server, the **ip helper-address interface** command is used. The IP helper-address can be configured to forward any UDP broadcast based on UDP port number. By default, the IP helper-address will forward the following UDP broadcasts:

- DNS (port 53), time service (port 37)
- Trivial File Transfer Protocol (TFTP) (port 69)
- Terminal Access Control Access Control System (TACACS) service (port 49)
- NetBIOS name server (port 137)
- NetBIOS datagram server (port 138)
- Boot Protocol (DHCP/BootP) client and server datagrams (ports 67 and 68)
- IEN-116 name service (port 42)

Reference: Understanding and Troubleshooting DHCP in Catalyst Switch or Enterprise Networks

<http://www.cisco.com/warp/public/473/100.html>

QUESTION NO: 7

Which administrative distance is given to EIGRP summary routes?

- A. 0
- B. 1
- C. 5
- D. 90
- E. 95
- F. 150

Answer: C

Explanation:

The following table displays the default AD for all routing protocols:

Default Administrative Distances	
Route Source	Default Distance
Connected interface	0
Static route	1
Enhanced IGRP summary route	5
External BGP	20
Internal Enhanced IGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
EGP	140
EIGRP external route	170
Internal BGP	200
Unknown	255

QUESTION NO: 8

If there is a route to the subnet 190.10.1.0/25 learned via RIP and a route to the subnet 190.10.1.0/24 learned via STATIC, which route would be preferred to reach the destination address 190.10.1.125?

Leading the way in IT testing and certification tools, www.testking.com

- A. The 190.10.1.0/25 route learned via RIP will be used because its prefix has the longest match.
- B. RIP has an administrative distance of 120, and static routes have an administrative distance of 1, so the static route would be preferred.
- C. The static route to 190.10.1.0/24 will be preferred because static routes have an administrative distance of 0 and the static route looks as though it is directly connected.
- D. A **show ip route** to the destination will show that the destination is learned from both RIP and the static route, so the traffic to 190.10.1.1.125 will be load balanced between the two paths.
- E. The information given is not sufficient to determine this.

Answer: A

Explanation:

The longest match rule always takes precedence over any other routing information. The administrative distance (AD) values are only compared for destinations with subnet masks of equal length.

Section 2: Describe classful and classless routing protocols (3 questions)**QUESTION NO: 1**

You are determining the routing protocol to use throughout your network. In doing this you compare the advantages of classless and classful protocols. Which of the following statements are true regarding classless routing protocols?

- A. A default gateway is required.
- B. Variable-length subnet masks are not supported.
- C. Routers are automatically summarized to class boundaries.
- D. Networks with different subnet masks can exist in the same address class.
- E. All of the above

Answer: D

Explanation:

Classless routing protocols understand that different routes within a major network can have different masks. They include VLSM information in the routing updates, enabling the use of different subnet mask lengths.

Incorrect Answers:

- A. There is no requirement for a default gateway.
- B. Classless routing protocols advertise the routing mask for each route. This enables VLSM (variable length subnet masks) support.
- C. With classless routing the summarization process can manually controlled and can be invoked at any point within the network. It is not limited to class boundaries.

QUESTION NO: 2

Which two of the following describe advantages of implementing a classless routing protocol, when compared to a classful routing protocol?

- A. Support for VLSM.
- B. Support for FLSM.
- C. Summarization of discontinuous subnets.
- D. Auto-summarization across network boundaries.
- E. The **ip classless** command improves convergence time.

Answer: A, C

Explanation:

- A. Classless routing protocols support VLSM, and that, in turn, leads to more efficient allocation of subnet masks to meet different host requirements on different subnetworks, resulting in better utilization of host addresses.
- C. Because subnets routes are propagated throughout the routing domain, summarization is often required to keep the routing tables at a manageable size.

Incorrect Answers:

- B, D: Fixed Length Subnet Masks (FLSM) and auto-summarized routes across network boundaries are functions of classful routing protocols, not classless.
- E. The convergence time of a network is due to numerous factors, including the timers of the routing protocol, as well as support for triggered updates. The “ip classless” command has no impact on the convergence time of any network.

Reference: Building Scalable Cisco Networks (Cisco Press) page 19-20.

QUESTION NO: 3

You wish to use a classless IP routing protocol within your network. Which of the following classless routing protocols could you use? (Select all that apply)

- A. IS-IS
- B. IGRP
- C. RIPv1
- D. OSPF
- E. EIGRP

Answer: A, D, E

Explanation:

Intermediate System to Intermediate System (IS-IS), Open Shortest Path First (OSPF) and Enhanced IGRP are all classless routing protocols.

Note: RIPv2 and BGP are also classless routing protocols.

Incorrect Answers:

- B, C: IGRP and RIPv1 are not classless. Both of these protocols do not support VLSM networks.

Section 3: Describe link-state router protocol operation (5 questions)

QUESTION NO: 1

The Testking network administrator is considering using a link state routing protocol to replace their existing distance vector protocol. Which if the following are true regarding link state routing protocols?

- A. Link-state protocols do not support summarization.
- B. Static routes must be used to accommodate link redundancy.
- C. All routers in the area know when another router joins the area.
- D. Link-state protocols utilize spanning tree to propagate routing changes.
- E. The spanning tree protocol must be enabled for link state protocols to work.

Answer: C

Explanation:

All routers within an area will have the same view of the area, meaning they will all have the same topology table. All of them will know when another router joins the area, since the routing topology table of all routers in any area must remain identical.

Incorrect Answers

- A: Link state protocols support route summarization, as well as variable length subnet masking.
- B: Link redundancy does not require the use of static routes. Redundant links can be used with link state protocols, and load balancing over redundant links is also supported.
- D, E: The spanning tree is not to propagate routing changes. Instead it used to find the shortest path to the destinations. The spanning tree protocol is used at layer 2 to determine bridging and switching loops, not routing loops.

QUESTION NO: 2

IS-IS and OSPF are two examples of link state routing protocols. Regarding the operation of link state protocols, which of the following are true?

- A. Link state protocols periodically multicast the Link State Advertisements at 90 second intervals.
- B. Link state routing protocols use Hello packets to build the link state database.
- C. Link state routing protocols use poison reverse and holddown timer to prevent routing loops.
- D. Link state routing protocols use Link State Advertisements to announce route changes.
- E. All of the above

Answer: D

Explanation:

LSA – Link State Announcement. Used by OSPF, an LSA is used to announce changes in network topology to adjacent routers

Each router periodically sends an LSA to provide information on a router's adjacencies or to inform others when a router's state changes. By comparing established adjacencies to link states, failed routers can be detected quickly and the network's topology altered appropriately. From the topological database generated from LSAs, each router calculates a shortest-path tree, with itself as root. The shortest-path tree, in turn, yields a routing table.

Incorrect Answers:

A: Although link state protocols use multicasts for routing updates, they are sent using hellos.

OSPF uses multicast IP addresses 224.0.0.5 and 224.0.0.6. LSAs are not sent every 90 seconds.

B: The Hello Protocol is responsible for establishing and maintaining neighbor relationships. It is used to build the neighbor table, not the link state database.

C: These are loop avoidance mechanisms used by distance vector routing protocols, not link state.

QUESTION NO: 3

One of the serial links in an OSPF network is experiencing problems and continuously changes from up to down. In a link-state environment, what does the router do when a route flaps?

- A. It enters the exstart state with its neighbors.
- B. It floods the area with new routing information.
- C. It generates a routing exchange using the hello protocol.
- D. It waits for the holdown timers to expire and then sends an update.

Answer: B

Explanation:

One of the drawbacks of using OSPF is that it can be prone to issues resulting in flapping routes. OSPF is a very strong protocol in terms of convergence time—each router is aware of the entire topology in the area. This results in fast convergence. However, if a link *flaps*, or changes between up and down status quickly, a flood of LSAs may be generated. This may prevent the routers in the network from converging. Administrators may use the “spf holdtime” command to force OSPF into a waiting state before computing a new route.

QUESTION NO: 4

You are an administrator of a network segment that uses RIP v2 as the routing protocol. How would the core router react if it were to detect a flapping link to a neighboring router?

- A. It recalculates the network topology.
- B. It purges that link from its routing table.
- C. It places a hold-down on the routes from that link.
- D. It sends a LSA to other router requesting an RIP update.

Answer: C

Explanation:

RIP is a distance vector routing protocol, which utilizes hold down timers and the rule of split horizons. The purpose of the hold-down state is to ensure the validity of any new routes for the same destination.

Incorrect Answers:

- A: RIP does not calculate a network topology.
- B: This would not serve any practical purpose, since the purged link would then reappear. The link would be placed in a held state, but not purged.
- D: OSPF utilizes link state advertisements (LSAs), but RIP routers do not.

QUESTION NO: 5

Which of the following routing protocols listed below use the Class D address of 224.0.0.9 to multicast its routing updates?

- A. EIGRP
- B. OSPF
- C. IGRP
- D. RIPv2

Answer: D

Explanation:

Class D addresses are multicast addresses used by many routing protocols for updating all neighbors on a link. Some of the Class D multicast addresses used by routing protocols are as follows:

- OSPF – 224.0.0.5 and 224.0.0.6
- Routing Information Protocol version 2 (RIPv2) – 224.0.0.9
- EIGRP – 224.0.0.10

Reference: Building Scalable Cisco Networks (Cisco Press) page 69

Section 4: Compare classful and classless routing protocols (3 questions)**QUESTION NO: 1**

When comparing and contrasting the differences between classless and classful routing protocols, which of the following is true?

- A. In a classful system, a router uses a bit mask to determine the network and host portions of an address and there is no class restriction.
- B. In a classless system, a router determines the class of an address and then identifies the network and host octets based on that class.
- C. In a classful system, subnet mask information is maintained and passed along with each routing update.
- D. In a classless system, when a routing update is received about a different major network as configured on the receiving interface, the default subnet mask is applied.
- E. In a classful system, when a routing update is received about the same major network as configured on the receiving interface, the router applies the subnet mask configured on the receiving interface.

Answer: E

Explanation:

Routing protocols can be divided into various categories, such as classful or classless, and distance-vector or link-state. A classful routing protocol is one that exchanges routing information based on the classful IP boundaries. A router running a classful routing protocol such as RIPv1 or IGRP that receives an update for a network to which it is not physically connected will summarize that network at the default classful boundary. If the router received a route to network 10.1.2.0, for instance, it would summarize the route to 10.0.0.0 in its routing table, because an 8-bit (255.0.0.0) subnet mask is the default classful boundary for Class A addresses. A classless routing protocol, such as RIPv2, OSPF or EIGRP, exchanges the subnet mask with the subnet information. So if the aforementioned router was running OSPF, it would receive a route to network 10.1.2.0 with a 24-bit (255.255.255.0) subnet mask included, and would not summarize the route to 10.0.0.0.

QUESTION NO: 2

Which two statements are true with regard to RIPv1 and OSPF? (Choose two)

- A. RIPv1 uses the Dijkstra algorithm while OSPF uses the Bellman-Ford algorithm for calculating best path.

- B. RIPv1 uses the Bellman-Ford algorithm, OSPF uses the Dijkstra algorithm for calculating best path.
- C. RIP forwards the entire routing table incrementally, OSPF link-state advertisements are sent out when a change occurs and every thirty minutes if no change occurs.
- D. RIPv1 maintains a 15 hop count limit while OSPF maintains a 255 hop count limit.
- E. Both RIPv1 and OSPF carry subnet mask information and therefore support VLSM.

Answer: B, C

Explanation:

RIP is a distance vector routing protocol, which uses the Bellman-Ford algorithm for calculating the best path to a destination. OSPF is a link state routing protocol, which means that it uses the Dijkstra algorithm to determine the best path. RIP also periodically advertises the entire routing protocol to all neighbors, while OSPF uses LSA information to flood the area with routing information, but only after a topology change has occurred. To maintain stability, OSPF also floods out the entire OSPF table every thirty minutes if no change has occurred.

Incorrect Answers:


A: The reverse is true.

D: RIP maintains a hop limit of 15 (16 means it is unreachable) but OSPF does not have any hop limitations.

E: RIPv1 does not carry the subnet mask information in the routing updates, and so it does not support VLSM. RIPv2 does, however.

QUESTION NO: 3

The TestKing IGRP network is displayed in the diagram below:



TestKing1 and TestKing2 are running IGRP. TestKing1 receives an update on its Fa0/0 interface from TestKing2 about the 199.10.50.64/26 network. Which mask will be applied to the update?

- A. 255.255.255.0
- B. 255.255.255.192
- C. 255.255.0.0
- D. 255.255.255.255

Answer: A

Explanation:

The ability to specify a different subnet mask for the same network number on different subnets is called Variable-Length Subnet Mask (VLSM). RIPv1 and IGRP are classful protocols and are incapable of carrying subnet mask information in their updates. Before RIPv1 or IGRP sends out an update, it performs a check against the subnet mask of the network that is about to be advertised and, in case of VLSM, the subnet gets dropped.

In this example, first router TestKing2 checks to see whether 199.10.50.64/26 is part of the same major net as 10.10.0.0/16, which is the network assigned to the interface that will be sourcing the update. It is not, and so TestKing2 will summarize at the network boundary, which is 199.10.50.0/24 since this is a class C network.

Section 5: Compare distance vector and link state routing protocols (9 questions)**QUESTION NO: 1**

You're a systems administrator who's just chosen OSPF over RIP version 1. Your junior administrator Britney is confused and asks you why you didn't choose RIP. What would you tell her? (Choose all that apply)

- A. OSPF maintains smaller routing tables than RIP
- B. OSPF cost metric is based on number of hops.
- C. OSPF only sends routing updates only when necessary.
- D. OSPF supports VLSM and allows more efficient use of IP addresses.

Answer: C, D

Explanation:

C: RIP use periodic broadcast of the entire routing table, while OSPF use event-triggered announcements. OSPF uses

D: RIP Version 1 does not support VLSM, while OSPF does.

Incorrect Answers:

- A: The size of the routing table depends on the number of routes. Since RIP version 1 automatically summarizes at the network boundary, the routing tables of RIP networks are generally smaller than OSPF networks.
- B: RIP uses hop as cost metric. OSPF uses a metric based on the bandwidth of the links to the destination.

QUESTION NO: 2

OSPF is the popular choice as the standard open source routing protocol. Why is it more popular than both version 1 and version 2 of RIP?

- A. OSPF uses greater CPU overhead.
- B. OSPF has greater router memory requirements.
- C. OSPF allows for greater scalability.
- D. OSPF uses a simpler distance vector algorithm.
- E. OSPF allows for a simpler router configuration
- F. OSPF uses a simpler route selection process.

Answer: C

Explanation:

OSPF allows for greater flexibility than both RIP versions 1 and 2. OSPF allows for the use of areas, with each area connecting into the backbone area. This provides for a tiered, hierarchical design. With both RIP versions, the entire network is seen as one flat network. In addition, there is no limitation on the number of hops that a route can take with OSPF. With both RIP versions, there is a 16 hop limit. Lastly, as advantage that OSPF has over RIP version 1 is that it supports VLSM information while RIPv1 does not. However, this advantage does not apply to RIP version 2 since it also supports VLSM.

Incorrect Answers:

- A, B: Although true, these are considered disadvantages to using OSPF.
- D. OSPF is a link state routing protocol, while RIP is considered to be distance vector.
- E. The OSPF router configuration is more complex to implement than a RIP router configuration.
- F. OSPF uses the cost metrics of each link in determining routes, which is relatively simple. However, RIP uses a simple hop count as the metric, which is even simpler.

Reference: Building Scalable Cisco Networks (Cisco Press) pages 99-100.

QUESTION NO: 3

The Testking network is expanding, and there is a growing concern as to the total number of router hops that are used. Which two of the routing protocols below use hop count to calculate their distance metric? (Select all that apply)

- A. IGRP
- B. OSPF
- C. EIGRP
- D. RIPv2
- E. RIPv1
- F. Static routes

Answer: D; E

Explanation:

Both RIP version 1 and version 2 use the hop count to measure the distance to a destination.

Incorrect Answers

- A, C: IGRP and EIGRP use a composite metric that is calculated by factoring weighted mathematical values for internetwork delay, bandwidth, reliability, and load. By default, only bandwidth and delay are used in the calculation.
- B: OSPF uses cost as the metric, which is found by taking the value 100,000,000 and dividing it by the bandwidth of each link.

QUESTION NO: 4

Leading the way in IT testing and certification tools, www.testking.com

A network designer is considering the best routing protocol to use in a new network. What are the advantages of using OSPF over RIPv1 in a large network? (Select two)

- A. OSPF has a faster convergence time.
- B. OSPF requires less router memory.
- C. OSPF manages fewer internal tables.
- D. OSPF consumes less bandwidth through the use of incremental updates.

Answer: A, D

Explanation:

A: OSPF has a faster convergence time than RIP. OSPF supports triggered updates while RIP does not.

D: OSPF use incremental updates, while RIP broadcasts the entire routing table to all neighbors periodically.

Incorrect Answers

B: OSPF is more processor intensive and requires more memory. With faster processors this is not a big issue, but it can still be considered a disadvantage of OSPF.

C: OSPF use a more complex algorithm. These algorithms require more internal tables, such as routing, topology, and neighbor tables.

QUESTION NO: 5

The Bellman-Ford algorithm is used by the grand majority of distance vector routing protocols in calculating routes. One sophisticated routing protocol doesn't use this algorithm; instead it uses the DUAL (diffusing update algorithm). Which of the following routing protocols is it?

- A. IGRP
- B. OSPF
- C. EIGRP
- D. RIP v.2
- E. RIP v.1

Answer: C

Explanation:

DUAL is used by EIGRP. The Diffusing Update Algorithm (DUAL) is the algorithm used to obtain loop-freedom at every instant throughout a route computation. This allows all routers involved in a topology change to synchronize at the same time. Routers that are not affected by topology changes are not involved in the recomputation. The DUAL finite state machine embodies the decision process for all route computations. It tracks all routes advertised by all neighbors. The distance information, known as a metric, is used by DUAL to select efficient loop free paths. DUAL selects routes to be inserted into a routing table based on feasible successors. A successor is a neighboring router used for packet forwarding that has a least cost path to a

destination that is guaranteed not to be part of a routing loop. When there are no feasible successors but there are neighbors advertising the destination, a recomputation must occur.

QUESTION NO: 6

Regarding distance vector and link state routing protocols, match each characteristic on the left to the appropriate routing protocol type on the right:

- Gains knowledge of the network based only on information learned from directly connected neighbors.
- Based on the Dijkstra algorithm.
- Routing updates occur only when a change occurs.
- Proprietary protocol type offered by Cisco.
- Requires more memory and processing power.
- Based on the Bellman-Ford algorithm.
- Sends complete routing tables at regular intervals

Distance Vector

Place here

Place here

Place here

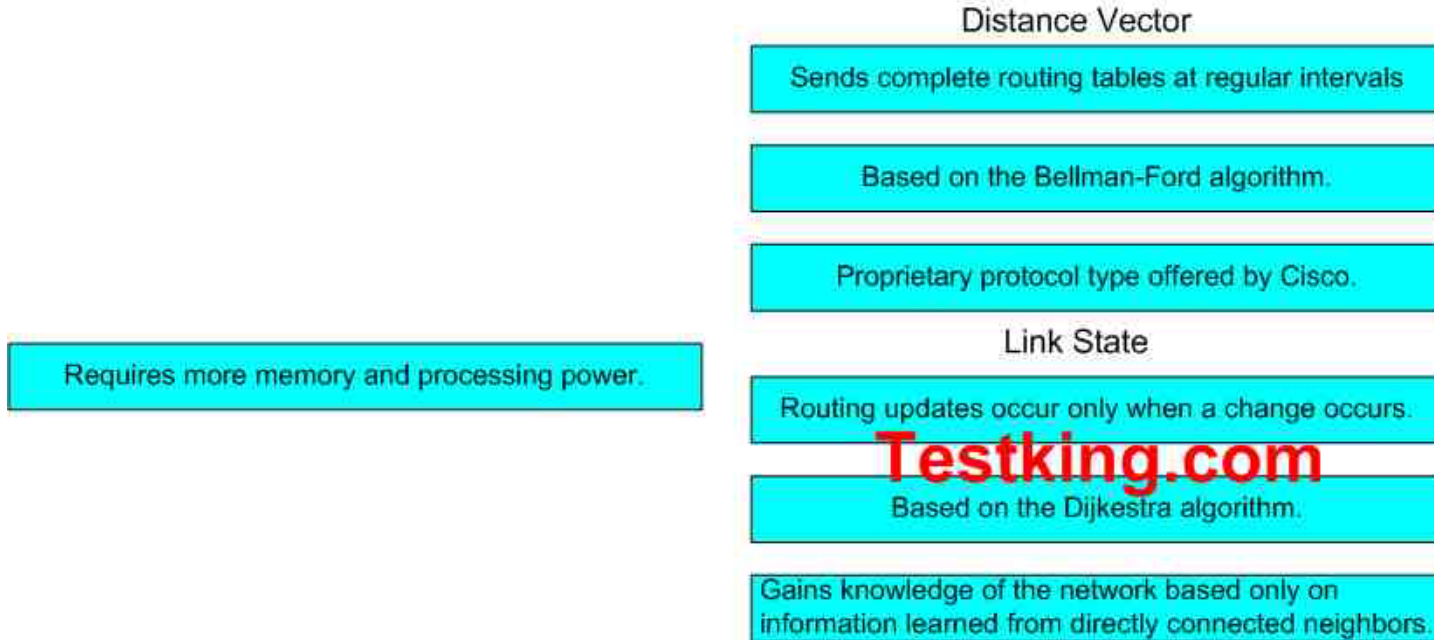
Link State

Place here

Place here

Place here

Answer:

**QUESTION NO: 7**

With regards to RIP version 1 networks, which of the following two statements are true? (Choose two)

- A. RIP v.1 route updates have a subnet mask field.
- B. RIP v.1 networks are known as classful networks.
- C. RIP v.1 networks are known as classless networks.
- D. RIP v.1 route updates do not have a subnet mask field.

Answer: B, D

Explanation:

There are two versions of RIP, namely RIPv1 and RIPv2. RIPv1 is only capable of classful routing. The routing updates do not carry subnet information, which means that a network's size is determined solely by the network class of its IP address, and there is no way to split a network into smaller subnets, each routed along a different path.

RIP cannot handle classless routing. RIP v1 summarizes all routes it knows on classful network boundaries, so it is impossible to subnet a network properly via VLSM if you are running RIP v1, which is a Classful routing protocol.

Incorrect Answers:

A, C: RIP version 2 is a classless routing protocol that supports VLSM through the use of including subnet mask information within the routing updates. However, RIP version 1 does not.

QUESTION NO: 8

In an effort to increase network security, RIP routing authentication is being configured on a network. Which of the following types of authentication does RIPv2 support?

- A. MD5 only
- B. PAP and CHAP
- C. clear text and PAP
- D. clear text and MD5
- E. MD5 and CHAP

Answer: D

Explanation:

Authentication

Authentication can prevent communication with any RIP routers that are not intended to be part of the network, such as UNIX stations running **routed**. Only RIP updates with the authentication password are accepted. RFC 1723 defines simple plain-text authentication for RIPv2.

MD5 Authentication

In addition to **plain-text passwords**, the Cisco implementation provides the ability to use **Message Digest 5 (MD5)** authentication, which is defined in RFC 1321. Its algorithm takes as input a message of arbitrary length and produces as output a 128-bit fingerprint or message digest of the input, making it much more secure than plain-text passwords.

Incorrect Answers:

A: RIP version 2 supports plain text password authentication in addition to MD5.

B, C, E: PAP and CHAP are authentication methods used by the Point to Point Protocol (PPP), not RIP.

QUESTION NO: 9

OSPF and IS-IS are examples of link state protocols. Which two features are characteristics of link-state routing protocols? (Choose two)

- A. Link-state protocols are based on the Dijkstra algorithm.
- B. Any incremental updates contain the full routing table.
- C. Link-state protocols flood the area with LSAs when a network change occurs.
- D. Routing updates are broadcast to neighbors.
- E. Link-state protocols use split-horizon updates to avoid routing loops.

Answer: A, C

Explanation:

Link-state routing protocols are based on the Dijkstra algorithm, meaning that routers exchange topology information with their nearest neighbors. The topology information is flooded throughout the AS, so that every router within the AS has a complete picture of the topology of

the AS. This picture is then used to calculate end-to-end paths through the AS, normally using a variant of the Dijkstra algorithm. Therefore, in a link-state routing protocol, the next hop address to which data is forwarded is determined by choosing the best end-to-end path to the eventual destination.

Incorrect Answers:

B, D: Incremental updates and neighbor broadcasts are not utilized in link state routing protocols. Updates are triggered by a change in the network topology.

E: Split horizons are used by distance vector routing protocols, not link state protocols.

Section 6: Describe concepts relating to extending IP addresses and the use of VLSMs to extend IP addresses (8 questions)

QUESTION NO: 1

When an administrator uses VLSM to subnet a network address, which of the following rules must be adhered to? (Select three)

- A. A subnet is made up of all aggregated routes.
- B. A subnet can address hosts, or be further subnetted.
- C. A routing protocol must carry the subnet mask in updates.
- D. Summarized networks must have the same high-order bits.

Answer: B, C, D

Explanation:

All these statements are true regarding VLSM subnetting.

B: A subnet can address hosts, or be further subnetted.

C: The subnet mask must be included in routing updates, since the subnet mask information needs to be distributed throughout the entire network.

D: Summarized routes must have the same high order bits. All subnet masks start with a series of contiguous 1's, followed by 0's.

Incorrect Answers:

A: Route aggregation is not subnetting. Supernetting is the term being described in this choice.

QUESTION NO: 2

The Testking network is migrating to a classless routing protocol to support VLSM.

Which two of the statements below correctly describe features of Variable Length Subnet Masking (VLSM)? (Select two)

- A. It supports IP version 4 and 6.
- B. It accommodates overlapping address ranges.
- C. It allows for better route summarization within routing tables.
- D. It allows subnetworks to be further divided into smaller subnets.

Answer: C, D

Explanation:

C: VLSM allows multiple levels of subnetted IP addresses within a single network, and thus allow for better route summarization within routing tables.

D: VLSMs provide the ability to include more than one subnet mask within a network, and the ability to subnet an already subnetted network address.

Incorrect Answers:

- A: VLSM was constructed for IP version 4.
 B: Overlapping address ranges should not be used with VLSM.

QUESTION NO: 3

The CTO of TestKing INC. has called a meeting with all the system administrators to update them on the latest expansion plan of establishing multiple remote offices, and connecting those remote offices to the central office by a WAN. Management is insisting that an addressing scheme using VLSM be used, and they need six useable host addresses at each remote location. Assuming that TestKing is using a class C IP range, which variable length subnet mask should be used to support the 6 hosts at each location?

- A. /24
- B. /28
- C. /29
- D. /30
- E. /31
- F. /32

Answer: C**Explanation:**

To provide for six usable host addresses three host bits must be used. This would give exactly six host address ($2^3-2=6$). This would leave 29 ($32-3$) bits for the network mask. For any subnet, we must remember to subtract 2 from the usable host portion; 1 for the network IP address and one for the broadcast IP address.

Incorrect Answers:

- A: This would result in 254 hosts and only 1 network.
 B: This would result in 14 hosts, with 14 networks.
 D: This results in only 2 host IP addresses. This subnet mask is often used on point to point WAN links.
 E, F: This would not allow for 6 hosts at each network.

QUESTION NO: 4

You're boss has assigned you the difficult task of subnetting a Class B network, so you choose 172.29.100.0/28. Which of the following answer choices reflect a valid VLSM subnet that can be obtained from your Class B network?

- A. 172.29.100.8
- B. 172.29.100.32
- C. 172.29.100.60

- D. 172.29.100.104
- E. 172.29.100.126

Answer: B**Explanation:**

We have the network 172.29.100.0/28 which is a subnetted Class B network. We want to identify a valid VLSM subnet mask of this network, which subnets this network further.

We recalculate the subnetted network in binary to better study it. We must make sure that only the host bits, and not the network bits, are used in the fourth octet, so when the number is shown in binary the last 4 numbers should be 0.

First octet	Second octet	Third Octet	Fourth Octet	Comments
172	29	100	0	
10101100	00011101	01100100	00000000	This is the network address
172	29	100	8	
10101100	00011101	01100100	00000100	This is a host on the 172.29.100.0/28 subnet
172	29	100	32	
10101100	00011101	01100100	00100000	This is a valid subnet address.
172	29	100	60	
10101100	00011101	01100100	00111100	This is a host on the 172.29.100.48/28 subnet.
172	29	100	104	
10101100	00011101	01100100	01101000	This is a host on the 172.29.100.96/28 subnet.

With a /28 network mask, the network boundaries will always be a multiple of 16 (172.29.100.0, 172.29.100.16, 172.29.100.32, 172.29.100.48, etc).

QUESTION NO: 5

You are using the address space of 172.28.100.0/24. You want to divide use this address space for your WAN links by dividing the address space using a subnet mask of 172.28.100.0/30

How many subnets will this provide for WAN links?

- A. 6
- B. 14
- C. 30

- D. 62
- E. 126
- F. 254

Answer: D

Explanation:

When sub-dividing a /24 network with a subnet mask of /30 (255.255.255.252) we have 62 usable networks, with 2 usable hosts on each subnet.

First octet	Second octet	Third Octet	Fourth Octet	Comments
172	28	100	0	
10101100	00011110	01100100	00000000	This is the entire address space
172	28	100	0	
10101100	00011110	01100100	00000000	This is the 30-bit VLSM mask

We can use 6 bits in the fourth octet for subnets. This gives us $2^6 - 2 = 62$ subnets

QUESTION NO: 6

TestKing is using the private IP address range of 192.168.100.0/24 for use on all of the WAN links. If this network is divided using a /30 mask, how many total WAN subnets can be supported?

- A. 14
- B. 30
- C. 62
- D. 126
- E. 254

Answer: C

Explanation:

When sub-dividing a class C network with a subnet mask of /30 (255.255.255.252) we have 62 usable networks, with 2 usable hosts on each subnet.

First octet	Second octet	Third Octet	Fourth Octet	Comments
192	168	100	0	
10000000	10101000	01100100	00000000	This is the entire address space
192	168	100	0	

11000000 10101000 01100100 **00000000** This is the 30-bit VLSM mask

We can use 6 bits in the fourth octet for subnets. This gives us $2^6 - 2 = 62$ subnets

Reference: Building Scalable Cisco Networks (Cisco Press) page 73

QUESTION NO: 7

You've secured the subnet and mask of 186.37.100.0/26 to use for some of your companies remote offices. Which VLSM mask below could provide you with six useable host addresses in each branch subnet?

- A. /24
- B. /28
- C. /29
- D. /30
- E. /31
- F. None of the above

Answer: C

Explanation:

3 bits (32-29) for the host would provide for $2^3 - 2 = 6$ hosts as is required. In this case, using the /29 mask on the existing /26 network would create 6 different networks with 6 hosts per network.

QUESTION NO: 8

Routers TK1 and TK2 are connected via a point to point serial line as shown in the diagram below:



Which command will configure a static route from the corporate network to 172.27.6.0?

- A. TK1(config)# ip route 172.27.6.0 255.255.255.0 172.17.8.2
- B. TK2(config)#ip route 172.27.6.0 255.255.0.0 172.17.8.2

- C. TK2(config)#ip route 172.27.6.0 255.255.255.0 172.17.8.2
- D. TK1(config)#ip route 172.27.6.0 255.255.0.0 172.17.8.1
- E. TK2(config)#ip route 172.27.6.0 255.255.255.0 172.17.8.1

Answer: E

Explanation:

Only answer choice uses the correct syntax on the correct router. The static route must be placed on router TK2, and the value for the next hop should be the serial IP address of the connected router, which in this case is 172.17.8.1 on router TK1.

Section 7: Describe the features and operation of EIGRP (21 questions)

QUESTION NO: 1

The TestKing network is using EIGRP as the network routing protocol. Which of the following statements correctly describe features and characteristics of routing using EIGRP? (Select three)

- A. It sends periodic updates every 60 seconds.
- B. EIGRP uses DUAL to achieve rapid convergence.
- C. Adjacencies exist between master routers (MRs) in each domain.
- D. It uses multicast to discover other EIGRP routers on an internetwork.
- E. EIGRP provides support for multiple network layer protocols: IPX, AppleTalk, and IP.

Answer: B, D, E

Explanation:

B: Enhanced IGRP uses the Diffusing Update based algorithm (DUAL).

D: EIGRP use hello multicast packets for neighbor discovery/recovery.

E: EIGRP supports IPX, Appletalk, and IP. Separate EIGRP instances can be created for each routed protocol.

Incorrect Answers:

A: EIGRP use triggered updates, not periodic updates.

C: Each EIGRP router maintains a neighbor table that lists adjacent routers. However, there is no concept of master routers in EIGRP.

Reference: "Introduction to Enhanced IGRP (EIGRP)"

<http://www.cisco.com/warp/public/459/7.html>

QUESTION NO: 2

While comparing the tables associated with EIGRP and OSPF, you notice some similarities. Which EIGRP table is similar to OSPF's Link State Adjacency table?

- A. Neighbor table
- B. Routing table
- C. Topology table
- D. Successor table
- E. None of the above

Answer: A

Explanation:

Each EIGRP router maintains a neighbor table that lists adjacent routers. This table is comparable to the neighbors (adjacency) database used by OSPF.

Incorrect Answers:

B: The routing table of EIGRP would compare to the routing table of OSPF

C: Both EIGRP and OSPF contain topology tables, which would compare to each other.

D: EIGRP does not have a successor table. Successors are maintained in the routing and topology tables of EIGRP.

Reference: Building Scalable Cisco Networks (Cisco Press) page 250

QUESTION NO: 3

Match the correct EIGRP term on the left to the slot in the middle that is adjacent to the correct description on the right.

EIGRP term, use these		Description
Successor	place here	lists adjacent routers
Routing table	place here	route entries for all destinations
Neighbor table	place here	primary route to destination
Feasible successor	place here	best routes to destinations
Topology table	place here	backup route to destination

Answer:

Neighbor table	- lists adjacent routers
Topology Table	- route entries for all destinations.
Routing table	- best routes to a destination
Successor	- primary route used to reach a destination
Feasible successor	- backup route to the destination.

Explanation:

- Neighbor Table – Each EIGRP router maintains a neighbor table that lists adjacent routers. This table is comparable to the neighbors (adjacency) database used by OSPF.
- Topology Table – An EIGRP router maintains a topology table for each network protocol configured: IP, IPX, and AppleTalk. All learned routes to a destination are maintained in the topology table.

- Routing Table – EIGRP choose the best routes to a destination from the topology table and places these routes in the routing table. The router maintains one routing table for each network protocol.
- Successor – This is the primary route used to reach a destination. Successors are kept in the routing table.
- Feasible Successor – This is a neighbor that is downstream with respect to the destination, but it is not the least-cost path and thus is not used for forwarding data. In other words, this is a backup route to the destination. These routes are selected at the same time as successors, but are kept in the topology table.

Reference: Building Scalable Cisco Networks (Cisco Press) page 250

QUESTION NO: 4

Different types of EIGRP transmissions are sent between neighboring routers. Which three of the following terms are known as 'reliable packets' in EIGRP? (Select three)

- A. hello
- B. ACK
- C. reply
- D. query
- E. update

Answer: C, D, E

Explanation:

Updates are used to convey the 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.

EIGRP reliable packets are: Update, Query and Reply.

EIGRP unreliable packets are: Hello and Ack.

Incorrect Answers:

A, B: 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 non-zero acknowledgment number.

Reference: Cisco BSCN version 1.0 study guide, pages 6-18.

QUESTION NO: 5

Which one of the following statements best describes the way EIGRP advertises subnet mask information to its destination networks?

- A. EIGRP advertises a prefix length for each destination network.
- B. EIGRP advertises a fixed length subnet mask for each destination network.
- C. EIGRP advertises only a classful subnet mask for each destination network.
- D. EIGRP, like IGRP and RIP, does not advertise a subnet mask for each destination network.

Answer: A

Explanation:

EIGRP advertises a prefix length to each destination network. It supports VLSM which allows subnet masks to be automatically summarized at the network boundary, and EIGRP can also be configured to summarize on any bit boundary at any interface.

Incorrect Answers:

- B: This is incorrect because EIGRP supports Variable Length subnet masks.
- C: This is incorrect because EIGRP is considered to be a classless protocol, not classful.
- D: EIGRP is more sophisticated than IGRP and RIP and actually does advertise a subnet mask to each destination network.

QUESTION NO: 6

EIGRP was designed specifically not to use excessive bandwidth for routing updates and functions across WAN links. If EIGRP is configured on an interface, what is the maximum percent of its bandwidth EIGRP would use, by default?

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

Answer: C

Explanation:

By default, EIGRP will limit itself to using no more than 50% of the available bandwidth. The maximum bandwidth that can be used is a configurable parameter.

Reference: Cisco, Configuration Notes for the Enhanced Implementation of EIGRP.
<http://www.cisco.com/warp/public/103/12.html>

QUESTION NO: 7

EIGRP supports Variable Length Subnet Masking (VLSM). Which two features of EIGRP are true regarding EIGRP support for VLSM? (Select two)

- A. It advertises a routing mask
- B. It is a classful routing protocol
- C. It is a classless routing protocol
- D. It does not advertise a routing mask.

Answer: A, C

Explanation:

EIGRP includes the routing mask when it advertises routes. Subnet mask information must be included in all routing updates for classless protocols. EIGRP is a classless protocol, meaning that it does not solely rely on the network class in order to find the subnet mask. These two features enable EIGRP to support VLSM.

Incorrect Answers

- B: Classful routing protocols do not support VLSM.
- D: The subnet mask must be included in the route to support VLSM.

QUESTION NO: 8

In an effort to decrease the size of the routing tables in the TestKing network, summarization is being configured on all routers. Where in an EIGRP network is it most appropriate to implement route summarization?

- A. At area border routers
- B. At autonomous system boundary routers
- C. It is done automatically by the DR router in each area
- D. Manually at any interface of any EIGRP router within the network
- E. In backbone area routers.

Answer: B

Explanation:

Enhanced IGRP performs route summarization at classful network boundaries by default. Automatic route summarization occurs at major network boundaries. ASBRs (autonomous system boundary routers) are used at the major network boundaries.

Incorrect Answers

- A, C, E: Area Border Router and DR are used in OSPF, but not in EIGRP. Backbone area routers are also concepts used in OSPF, but not in EIGRP.

D: Manual summarization can be done in any interface at any router within network, but this is not the preferred solution. Summarizing manually at each router can result in black hole routing. However, done properly this is an acceptable practice. If the question calls for more than one answer then the best choice would be B and D.

QUESTION NO: 9

The TestKing network is using route summarization to decrease the size of the routing tables. In an EIGRP network, which of the following describes the best method for implementing summarization?

- A. At WAN interfaces.
- B. Manually at major network boundaries.
- C. Dynamically at discontinuous interfaces.
- D. Dynamically at major network boundaries.

Answer: D

Explanation:

EIGRP supports arbitrary route summarization. Route summarization takes place automatically at major network boundaries (where network are supernetted).

Note: Route summarization is the consolidation of advertised addresses.

Incorrect Answers

- A: Manually configuring summarization at WAN boundaries is not necessarily recommended in all situations. If the routers belong to the same major network boundary then summarization should not be configured.
- B: Route summarization is automatic in EIGRP, and it is recommended to keep the defaults.
- C: Route summarization should only be applied at contiguous interfaces. Choosing this option would most likely lead to black hole routing, making many networks unreachable.

QUESTION NO: 10

The TestKing network is running EIGRP on the T1 links within their NBMA frame relay network. By default what would the hold time be for hello packets across these WAN links?

- A. 30 seconds
- B. 60 seconds
- C. 90 seconds
- D. 180 seconds
- E. 5 Seconds

Answer: D

Explanation:

EIGRP sends hello packets every 5 seconds on high bandwidth links and every 60 seconds on low bandwidth multipoint links. The hold time is typically three times the hello interval. In this scenario, on slow NBMA media, hold time will be 180 seconds.

Incorrect Answers:

A, B, C: Although the hello timers can be manually configured for these values, they are not the default time.

E: This is the default for high bandwidth links, such as Ethernet networks.

Reference: Enhanced Interior Gateway Routing Protocol

<http://www.cisco.com/warp/public/103/eigrp1.html>

QUESTION NO: 11

EIGRP is being used as the routing protocol within the TestKing network. Which two of the following features allow EIGRP to support classless routing? (Select two)

- A. Dijkstra's algorithm
- B. discontinuous subnets
- C. variable length subnet masks
- D. periodic update announcements
- E. unequal path-cost load balancing

Answer: B, C

Explanation:

B: Discontinuous subnets are supported by EIGRP. This is a classless routing protocol feature.

C: EIGRP support variable length subnet masks (VLSM). This is a classless routing protocol feature.

Incorrect Answers

A: EIGRP use the DUAL algorithm, not the Dijkstra's algorithm. The Dijkstra algorithm is used by distance vector protocols. EIGRP is considered to be a hybrid routing protocol.

D: EIGRP use periodic hello messages, not update announcements. EIGRP use update packets to convey reachability of destinations. When a new neighbor is discovered, Update packets are sent so the neighbor can build up its. These are not periodic, however.

E: IGRP and EIGRP support unequal cost path load balancing, which is known as variance. However, this is not a classless feature of EIGRP.

Note: Classless routing protocols include the routing mask with the route advertisement. This enables discontinuous subnets and variable length subnet masks.

Reference: Introduction to Enhanced IGRP (EIGRP)

<http://www.cisco.com/warp/public/103/1.html>

How Does Unequal Cost Path Load Balancing (Variance) Work in IGRP and EIGRP?

<http://www.cisco.com/warp/public/103/19.html>

QUESTION NO: 12

EIGRP is being used as the routing protocol on the TestKing network. While troubleshooting some network connectivity issues, you notice a large number of EIGRP SIA (Stuck In Active) messages. What causes these SIA routes? (Select two)

- A. The neighboring router stops receiving ACK packets from this router.
- B. The neighboring router starts receiving route updates from this router.
- C. The neighboring router is too busy to answer the query (generally caused by high CPU utilization).
- D. The neighboring router is having memory problems and cannot allocate the memory to process the query or build the reply packet.

Answer: C, D

Explanation:

SIA routes are due to the fact that reply packets are not received. This could be caused by a router which is unable to send reply packets. The router could have reached the limit of its capacity, or it could be malfunctioning.

Incorrect Answers

A: Missing replies, not missing ACKs, cause SIA.

B: Routes updates do not cause SIA.

Notes: If a router does not receive a reply to all outstanding queries within 3 minutes, the route goes to the stuck in active (SIA) state. The router then resets the neighbors that fail to reply by going active on all routes known through that neighbor, and it re-advertises all routes to that neighbor.

Reference: Enhanced Interior Gateway Routing Protocol

<http://www.cisco.com/warp/public/103/eigrp3.html>

QUESTION NO: 13

EIGRP is being configured on the TestKing network. After the initial configuration, the EIGRP routers form relationships with its peer routers. Which three of the following statements correctly describe these peer relationships? (Select three)

- A. EIGRP will form neighbors if the routers are not adjacent
- B. EIGRP will not form neighbors if the metric K-values do not match.
- C. EIGRP will not form neighbors if the router AS numbers do not match.
- D. EIGRP will form neighbors over primary and secondary interface addressing.
- E. EIGRP will form neighbors even though hello and hold timers do not match on the peering interfaces.

Answer: B, C, E

Explanation:

B: To become neighbors the routers metric must be comparable.

C: EIGRP neighbors must have the same AS number.

E: It is possible for two routers to become EIGRP neighbors even though the hello and hold timers don't match.

Incorrect Answers:

A: Neighbor discovery/recovery is the process that routers use to dynamically learn of other routers on their directly attached networks.

D: EIGRP doesn't build peer relationships over secondary addresses.

Reference: Cisco White Paper, Enhanced Interior Gateway Routing Protocol
<http://www.cisco.com/warp/public/103/eigrp1.html>

QUESTION NO: 14

The TestKing network is an EIGRP network consisting of four links (Link A, Link B, Link C, & Link D). The costs associated with each link is A=1, B=3, C=2, D=4. According to the variables below, which link is going to be the feasible successor?

- A. Link A
- B. Link B
- C. Link C
- D. Link D

Answer: C

Explanation:

Link A is the best link and it therefore the successor. Link C is the next best link, and it is therefore the feasible successor.

Incorrect Answers

A: Link A is the successor, not the feasible successor.

B: Both Link A and Link C are better than Link B.

D: The feasible successor must have Advertised Destination that is less than the Feasible Distance.

Note: Feasible distance is the best metric along a path to a destination network, including the metric to the neighbor advertising that path. Reported distance, or advertised distance, is the total metric along a path to a destination network as advertised by an upstream neighbor. A feasible successor is a path whose reported distance is less than the feasible distance.

Reference: Enhanced Interior Gateway Routing Protocol
<http://www.cisco.com/warp/public/103/eigrp1.html>

QUESTION NO: 15

Part of the routing table of router TK1 is displayed below:

S	62.99.153.0/24 [1/0] via 209.177.64.130 172.209.12.0/32 is subnetted, 1 subnets
D EX	172.209.1 [170/2590720] via 209.179.2.114, 06:47:28, Serial0/0/0.1239 62.113.17.0/24 is variably subnetted, 2 subnets, 2 masks
D EX	99.3.215.0/24 [170/27316] via 209.180.96.45, 09:52:10, FastEthernet11/0/0 [170/27316] via 209.180.96.44, 09:52:10, FastEthernet11/0/0
	25.248.17.0/24 [90/1512111] via 209.179.66.25, 10:33:13, Serial0/0/0.1400001 [90/1512111] via 209.179.66.41, 10:33:13, Serial0/0/0.1402001 62.113.1.0/24 is variably subnetted, 12 subnets, 2 masks
D	62.113.1.227/32 [90/2611727] via 209.180.96.45, 10:33:13, FastEthernet1/0/0 [90/2611727] via 209.180.96.44, 10:33:13, FastEthernet1/0/0
S*	0.0.0.0/0 [1/0] via 209.180.96.14

From analyzing the above command output, what is the administrative distance of the external EIGRP routes?

- A. 24
- B. 32
- C. 90
- D. 170
- E. 27316
- F. None of the above

Answer: D

Explanation:

By default an external EIGRP route has a value of 170. By examining the exhibit we see that this default value of the external EIGRP routes (see D-EX in exhibit) indeed is set to 170. The first value within the brackets display the AD, so with a value of [170/27316] the AD is 170 and the metric of the route is 27316.

Incorrect Answers:

A, B: This is the subnet mask used for some of the routes in the table.

C: This is the AD of the internal EIGRP routes, which is the default


E: This is the EIGRP metric of the external EIGRP routes.

Reference: What Is Administrative Distance?

http://www.cisco.com/warp/public/105/admin_distance.html

QUESTION NO: 16

The TestKing network is shown below, along with the relevant router configurations:

**TK1# show run**

```
interface Loopback0
  ip address 10.10.10.1 255.255.255.0
!
interface Ethernet0
  ip address 172.29.1.1 255.255.255.0
  media-type 10BaseT
!
!
router eigrp 100
  redistribute connected
  network 172.29.0.0
  auto-summary
  no eigrp log-neighbor-changes
!
ip classless
no ip http server
```

TK2# show run

```
interface Ethernet0
  ip address 172.29.1.2 255.255.255.0
  media-type 10BaseT
!
interface Ethernet1
  ip address 172.19.2.2 255.255.255.0
  media-type 10BaseT
!
router eigrp 100
  network 172.19.0.0
  network 172.29.0.0
!
ip classless
no ip http server
```

TK3# show run

Leading the way in IT testing and certification tools, www.testking.com

```

interface Ethernet1/0
 ip address 172.19.2.3 255.255.255.0
 !
router eigrp 100
 network 172.19.0.0
 auto-summary
 no eigrp log-neighbor-changes
 !
ip classless
ip http server

```

With the topology found in the graphic, what will the TK1 loopback 0 be in the TK3 routing table?

- A. It will show up in the routing table as D 10.0.0/8.
- B. It will show up in the routing table as D EX 10.0.0.0/8.
- C. It will show up in the routing table as D 10.0.0./24.
- D. It will not show up in TK3 routing table because there is no network command on TK1.

Answer: B

Explanation:

Because router TK1 is configured with route redistribution, it will redistribute the connected loopback network into EIGRP. Because redistributed routes will show up as external EIGRP routes in the routing table, choice B is correct. Although the loopback interface is using a /24 subnet mask, EIGRP summarizes at network boundaries by default so the network will appear as the class A network of 10.0.0.0/8 in the routing table of the other routers.

Incorrect Answers:

- A: The route will be external, since it was redistributed into EIGRP.
- C: It will be external because of redistribution, and it will also be summarized since that is the default behavior of EIGRP.
- D: Although it was not configured under the EIGRP network command, it would be redistributed because it is a connected route.

QUESTION NO: 17

EIGRP has been configured as the routing protocol on the TestKing network. Which statements are true regarding EIGRP? (Choose three)

- A. By default, EIGRP performs auto-summarization across classful network boundaries.
- B. EIGRP uses an area hierarchy to increase network scalability
- C. To speed convergence, EIGRP attempts to maintain a successor and feasible successor path for each destination.
- D. EIGRP uses hellos to establish neighbor relationships.

- E. By default, EIGRP uses the Dijkstra algorithm to determine the best path to destination network based on bandwidth and delay.

Answer: A, C, D

Explanation:

The default behavior of EIGRP routers is to automatically summarize routes at the network boundary.

EIGRP routers maintain information regarding the successors and the feasible successors to each network destination. This information is useful in the convergence time whenever any of the links fail.

EIGRP neighbors periodically use hellos to establish the relationship. Should any of the neighbors fail, triggered updates are sent to update the network regarding the topology change.

Incorrect Answers:

B: EIGRP networks are inherently flat, as each router is perceived as a peer to every other EIGRP router. IS-IS and OSPF utilize hierarchical network topologies.

E: The Dijkstra algorithm is used by link state routing protocols. EIGRP is considered to be a hybrid routing protocol.

QUESTION NO: 18

Router Testking1 is the headquarters router in a hub and spoke topology supporting 24 remote offices. Point-to-multipoint Frame Relay EIGRP network is deployed between the headquarters and the remote offices. There is no bandwidth command configured under either the major serial interface or the subinterface on router Testking1.

What is the bandwidth of each Frame Relay connection perceived by the EIGRP process?

- A. 64 kbps
- B. 128 kbps
- C. 1.544 Mbps
- D. 1.536 Mbps

Answer: A

Explanation:

By default in point-to-multipoint Frame Relay EIGRP network all subinterfaces share the bandwidth and each one receives $(\text{throughput of the channel})/(\text{number of subinterfaces})$ throughput. $1544\text{kbps}/24 = 64\text{kbps}$. 1544 – bandwidth of the serial interface.

QUESTION NO: 19

What administrative distance is given to EIGRP summary routes?

- A. 0

- B. 1
- C. 5
- D. 90
- E. 95
- F. 170

Answer: C

Explanation:

Default administrative Distance

Connected interface or static route that identifies the outgoing interface rather than the next hop	0
Static route	1
EIGRP summary route	5
External BGP	20
EIGRP	90
IGRP	100
OSPF	110
RIP	120
External EIGRP	170
Internal BGP	200
An unknown network	255 or infinity

QUESTION NO: 20

**From the following choices, identify the characteristics of EIGRP feasible successors.
(Select three)**

- A. A feasible is selected by comparing the advertised distance of a non-successor router to the feasible distance of the best route.
- B. If the advertised distance of the non-successor route is less than the feasible of best route, then that route is identified as a feasible successor.
- C. If the successor becomes unavailable, the feasible successor can be used immediately without recalculating for a lost route.
- D. The feasible successor can be found in the routing table.
- E. Traffic will be load balanced between feasible successors with the same advertised distance.

Answer: A, B, C

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

Once a neighbor relationship has been formed, called an **Adjacency**, the routers exchange routing update information and each router builds its own topology table. The Updates contain all the routes known by the sender. For each route, the receiving router calculates a distance for that route based on the distance that is conveyed and the cost to that neighbor that advertised the particular route. If the receiving router sees several routes to a particular network with different metrics, then the route with the lowest metric becomes the **Feasible Distance (FD)** to that network. The Feasible Distance is the metric of a network advertised by the connected neighbor plus the cost of reaching that neighbor. This path with the best metric is entered into the routing table because this is the quickest way to get to that network.

With the other possible routes to a particular network with larger metrics, the receiving router also receives the **Reported Distance (RD)** to this network via other routers. The Reported Distance being the total metric along a path to a destination network as advertised by an upstream neighbor. The Reported Distance for a particular route is compared with the Feasible Distance that it already has for that route. If the Reported Distance is larger than the Feasible Distance then this route is not entered into the Topology Table as a Feasible Successor. This prevents loops from occurring. If the Reported Distance is smaller than the Feasible Distance, then this path is considered to be a Feasible Successor and is entered into the Topology table. The **Successor** for a particular route is the neighbor/peer with the lowest metric/distance to that network.

If the receiving router has a Feasible Distance to a particular network and it receives an update from a neighbor with a lower advertised distance (Reported Distance) to that network, then there is a **Feasibility Condition**. In this instance, the neighbor becomes a **Feasible Successor** for that route because it is one hop closer to the destination network. There may be a number of Feasible Successors in a meshed network environment, up to 6 of them are entered into the Topology table thereby giving a number of next hop choices for the local router should the neighbor with the lowest metric fail. What you should note here, is that the metric for a neighbor to reach a particular network (i.e. the Reported Distance) must always be less than the metric (Feasible Distance) for the local router to reach that same network. This way routing loops are avoided. This is why routes that have Reported Distances larger than the Feasible Distance are not entered into the Topology table, so that they can never be considered as successors, since the route is likely to loop back through that local router.

Incorrect Answers:

D: The feasible successors are found in the topology table, but not the active routing table.

E: With EIGRP, traffic is load balanced across equal cost links in the routing table, but not between feasible successors.

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

QUESTION NO: 21

Which of the following statements are true about EIGRP operation? (Choose three)

- A. When summarization is configured, the router will add a route to null 0
- B. The summary route remains in the route table, even if there are no more specific routes to the network.

- C. Summarization is configured on a per-interface level.
- D. The maximum metric for the specific routes is used as the metric for the summary route.
- E. Automatic summarization across major network boundaries is enabled by default.

Answer: A, C, E

Explanation:

Using the **ip summary-address eigrp** summarization command will cause the creation of an EIGRP summary default route to the null 0 interface with an administrative distance of 5. Caution should be taken when using this as a default route. 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.

By default, Cisco routers will automatically summarize EIGRP routes across major network boundaries. With EIGRP, 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.

Incorrect Answers:

- B: When all of the specific routes used within a summary route are deleted from the routing table, the summary route will also be deleted.
- D: The metric used in a summary route is the best metric from among the summarized routes, not the maximum (worst).

**Section 8: Describe the features and operation of single area OSPF
(15 questions)**

QUESTION NO: 1

OSPF has recently been configured on router TK1. A router running OSPF can be in one of many different states. Which state must router TK1 be in with a neighbor before it will begin to route traffic?

- A. full state
- B. active state
- C. two-way state
- D. forwarding state

Answer: A

Explanation:

Full is the normal state for an OSPF router. In this state, routers are fully adjacent with each other and the router is routing traffic. The OSPF neighbor states are described below:

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. Bi-directional 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.

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.

Incorrect Answers

B: An IGRP or EIGRP router can be in an active state, but an OSPF cannot be in an active state.

C: At the two-way state, a router decides whether to become adjacent with this neighbor.

However, routes are not exchanged between the routers.

D: An interface (link) can be in a forwarding state, but not a router.

Note: An OSPF router can be in one of the following states: down, attempt, init, 2-way, exstart, exchange, loading, or full.

Reference: OSPF Neighbor States

<http://www.cisco.com/warp/public/104/13.html>

RFC 2328, OSPF Version 2

QUESTION NO: 2

One of the WAN links on an OSPF has just gone down, causing a network topology change. What does an OSPF network do to maintain route consistency when it notices a change in the network?

- A. It enters the exstart state with its neighbors.
- B. It floods the area with new routing information.
- C. It generates a routing exchange using the hello protocol.
- D. It waits for the holddown timers to expire, then sends the update.

Answer: B

Explanation:

Link state routing protocols generate routing updates only when there is a change in the topology. When a link changes state, a link-state advertisement (LSA) concerning that link (route) is created by the device that detected the change and propagated to all neighboring devices using a special multicast address. This process is called flooding.

Incorrect Answers:

A: The exstart state is not entered when a change in the topology occurs.

Note: Once the DR and BDR have been elected, the routers are considered to be in the exstart state and are ready to discover the link-state information about the internetwork and create their link-state databases.

C: The hello protocol is not used to exchange routing information. The purpose of hellos is for identifying neighbors and verifying periodically that the neighbors exist.

D: OSPF does not use holddown timers like RIP. Topology changes are flooded immediately.

QUESTION NO: 3

Certain types of OSPF networks require the election of a Designated Router (DR) as well as a Backup Designated Router (BDR). In which of the following network types does this election take place? (Select two)

- A. Point-to-point
- B. Nonbroadcast multi-access
- C. Point-to-multipoint
- D. Broadcast multi-access

Answer: B, D

Explanation:

Mode

NBMA


Adjacency

Manual Configuration DR/BDR elected

Broadcast Automatic DR/BDR elected
Reference: Building Scalable Cisco Networks (Cisco Press) page 124

QUESTION NO: 4

Three TestKing routers and connected via a frame relay network as shown below:



Examine the s0 interface on router TestKing1. What is the OSPF network type by default?

- A. Nonbroadcast
- B. Broadcast
- C. Point-to-multipoint
- D. Point-to-point
- E. Point-to-multipoint nonbroadcast

Answer: A

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

Based on the information above, the physical serial 0 interface is used, with no subinterfaces. When OSPF is configured on a physical interface (for example, interface S0) of a nonbroadcast multiaccess technology such as Frame Relay, the default network type of NON_BROADCAST is assigned. When OSPF is configured on point-to-point subinterfaces, the default interface type of POINT_TO_POINT is assigned. When OSPF is configured on multipoint subinterfaces, the default interface type of NON_BROADCAST is assigned.

Reference: http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080094054.shtml

QUESTION NO: 5

Router TK1 has been configured for OSPF on all of the interfaces. A logical loopback interface has not been created in TK1. By default, what will determine the OSPF router ID of router TK1?

- A. It defaults to 255.
- B. The name set by the hostname command.
- C. The highest IP address configured in the router.
- D. The priority number of the router set by the priority command.
- E. The lowest IP address configured in the router.

Answer: C

Explanation:

The **show ip ospf interface** command verifies that interfaces have been configured in the intended areas. If no loopback address is specified, the interface with the highest address is the taken router ID. It also gives the timer intervals, including the hello interval, and shows the neighbor adjacencies.

Note: the router ID used is not pre-emptive, meaning that if a loopback address is configured or another interface is configured with a higher IP address, the router ID will not change until it is restarted.

Reference: Building Scalable Cisco Networks (Cisco Press) page 134

QUESTION NO: 6

While troubleshooting some connectivity issues, you issue the “show ip ospf database” in order to examine the link state database. Which three of the statements below are true regarding the OSPF link state database? (Select three)

- A. Each router has an identical link state database.
- B. External routes are imported into a separate link state database.
- C. Synchronization of link state databases is maintained via flooding of LSAs.
- D. Information in the link state database is used to build a routing table by calculating a shortest-path tree.

- E. By default, link state databases are refreshed every 10 minutes in the absence of topology changes.

Answer: A, C, D

Explanation:

The Link state database is a collection of link state advertisement for all routers and networks. Each router in the OSPF network maintains an identical database. LSA flooding occurs whenever there is a change in the OSPF topology, ensuring that the databases are synchronized. OSPF also uses the SPF algorithm to build the database tables.

Incorrect Answers:

- B. Only one link state database is maintained, and it is used for all OSPF routes.
E. The default refresh time is 30 minutes.

Reference: Building Scalable Cisco Networks (Cisco Press) page 178.

QUESTION NO: 7

When designing OSPF networks, it is not recommended to have an excessive number of routers in any single area. Which two of the following statements describe the consequences of having too many routers in an area? (Select two)

- A. Excessive LSA traffic.
B. Frequent routing table recalculation.
C. Frequent adjacencies table recalculation.
D. Unacceptable number of reachability errors.

Answer: A, B

Explanation:

In OSPF, if too many routers are in the same area they spend too much time communicating with each other, because anytime a topology change in any single route occurs, LSAs are flooded to all other routers within the area. In OSPF the amount of overhead used increases exponentially with the number of routers in an area. Too many routers can cause an overwhelming amount of LSA and OSPF database calculations.

Incorrect Answers:

- C: The number of routers in an area has little or no impact with the number of neighbor adjacencies that must remain established for any given router.
D. An excessive number of routers in any single area will not alone increase the number of errors within the network.

QUESTION NO: 8

Router TK1 has OSPF configured on its fast Ethernet interface. What is the OSPF cost associated with this 100Mbps Ethernet connection?

- A. 1
- B. 6
- C. 10
- D. 100

Answer: A

Explanation:

The OSPF cost of an interface is inversely proportional to the bandwidth of that interface. A higher bandwidth indicates a lower cost. There is a higher cost and time delay involved in crossing a 56k serial line than crossing a 10M ethernet line. The formula used to calculate the cost is of any single OSPF network is:

Cost= 100,000,000/bandwith in bps.

Since a 100M Ethernet link is 100,000,000 bps we have:

$10000\ 0000 / 10000\ 0000 = 1$

Reference:

<http://www.cisco.com/warp/public/104/2.html>

QUESTION NO: 9

OSPF is a routing protocol that supports variably subnetted IP networks. By what means is it able to support VLSM?

- A. By using route summarization.
- B. By maintaining a topological database.
- C. By carrying subnet mask information in the route updates.
- D. By allocating addresses in groups to support multiple areas.

Answer: C

Explanation:

Each route update includes subnet mask information, which is passed on to the other OSPF routers within the network.

Incorrect Answers

- A: VLSM allows route summarization, but VLSM does not use route summarization.
- B: Although a topology table is maintained, a topological database does not contain VLSM information.
- D: This does not describe the way OSPF supports VLSM..

QUESTION NO: 10

Which of the following statements regarding OSPF and VLSM support are true? (Select two)

- A. A multi-area OSPF design dictates that VLSM be used throughout the AS.
- B. Summarization can be performed with different prefix lengths throughout the network.
- C. The use of VLSM enables a truly hierarchical addressing scheme to be developed.
- D. OSPF's support for VLSM does not compensate for its inability to handle discontinuous subnets.

Answer: B, C

Explanation:

B: Summarization can be performed with different prefix lengths throughout the network, as long as the ranges don't overlap.

C: VLSM enables a true hierarchical addressing scheme.

Incorrect Answers

A: VLSM can be used locally only. VLSM is not required to be used throughout multiple areas.

D: In order to take advantage of summarization, network numbers in areas should be assigned in a contiguous way to be able to lump these addresses into one range.

Reference: OSPF Design Guide

<http://www.cisco.com/warp/public/104/3.html>

QUESTION NO: 11

In order to reduce the overall size of the routing table, summarization is being configured throughout the TestKing network. What is true about route summarization in OSPF? (Select the best answer)

- A. Type 3 and type 4 LSAs carry external summarized routes.
- B. Summarization prevents type 1 link LSAs from being propagated into the backbone area0.
- C. Route summarization can be performed at any point in the network where enough contiguous addresses are present.
- D. Route summarization reduces the amount of bandwidth, CPU, and memory resources consumed by the OSPF process.

Answer: D

Explanation:

Summarization reduces the overall number of routes within the network, which will increase performance by reducing the workload on the router.

QUESTION NO: 12

OSPF is considered to be an extremely scalable routing protocol. Which of the following would be reasons for this? (Select two)

- A. Neighbor adjacencies control distribution of routing protocol updates.
- B. Routing table information does not flood the network until holddown timers have expired.
- C. The hello protocol is a more efficient means of sending routing updates than table exchange used in RIPv1.
- D. Topological database is maintained with incremental updates, with full exchange occurring only every 30 minutes.

Answer: C, D

Explanation:

- C: The Hello Protocol used by OSPF to establish and maintain neighbor relationship. These periodic hellos are much more scalable than the entire routing updates used by RIP.
- D: Incremental updates are used to maintain the OSPF topology. By default, the entire SPF algorithm is only run every 30 minutes.

Incorrect Answers:

- A: LSA flooding occurs whenever a topology change occurs in the OSPF network. The DR controls this distribution, not the neighbor relationship.
- B: LSA flooding occurs immediately, after each topology change within an area.

QUESTION NO: 13

Router TK1 is configured for OSPF on all of its interfaces. One of the links on this router goes down, causing a link state change. Assuming that router TK1 lies within an Ethernet segment, how will it respond to this link state change?

- A. It will broadcast the LSA on each of its interfaces.
- B. It multicasts the link-state changes to the DR and BDR.
- C. It will immediately flood the change to all outer routers on the same segment.
- D. It will update its routing table, and then floods the change to all other routers on the same segment.

Answer: B

Explanation:

The steps for OSPF convergence are as follows:

1. When a router detects a link failure, the router sends an LSA to its neighbors. If the router is on a multi-access link, it sends the update to the designated router (DR) and the backup designated router (BDR), not to all neighbors.
 2. The path is removed from the originating router's tables.
 3. On receipt of the LSA, all routers update the topology table and flood the LSA out its interfaces.
 4. The routing protocol runs the Dijkstra algorithm to rebuild the routing table.
- For OSPF, convergence is detection time, plus LSA flooding, plus 5 seconds before computing the topology table. This amounts to a few seconds.

QUESTION NO: 14

With which other routers on an Ethernet LAN should the DROTHER routers form OSPF full adjacency?

- A. the DR and BDR only
- B. all OSPF routers on the Ethernet LAN
- C. all OSPF routers in the area
- D. all other DROTHER routers on the Ethernet LAN
- E. the DR only

Answer: A

Explanation:

DROTHER – The router was not chosen as the DR or the BDR. If the priority on the interface had been set to zero, the state would always be DROTHER because the router could not be elected as a DR or a BDR. Each router should form the full adjacency only with DR and BDR to receive the OSPF updates from them.

QUESTION NO: 15

What does an OSPF router do if it receives an LSA with a lower sequence number than what is already in its link-state database?

- A. Updates its database with the LSA information.
- B. Ignores the LSA
- C. Forwards the LSA to the DR
- D. Ignores the LSU and sends the source router the more current LSA info in its own database.

Answer: D

Explanation:

If the LSA entry was in the database, but the LSA that has just been received has an older sequence number, the router asks whether the information in the database is the same.

If the information is the same and the new LSA has an older sequence number, the process discards the packet. It might be old news, but there is no inconsistency in the database.

If the information is different and the newly received LSA has an older sequence number, however, the receiving router discards the LSA update. It issues a copy of the LSA it has in its database, sending it out of the receiving interface to the source address of the out-of-date LSA. The logic is that the sending router has bad or old information and must be updated because its topological database is obviously not synchronized with the rest of the area.

Section 9: Describe the features and operation of multi-area OSPF (26 questions)**QUESTION NO: 1**

Which of the three areas below are valid OSPF area types? (Select three)

- A. Stub
- B. Active
- C. Remote
- D. Backbone
- E. Ordinary or standard

Answer: A, D, E

Explanation:

There is no such thing as an active area or a remote area in OSPF. The correct OSPF area types are:

- ordinary area (standard area)
- stub area
- totally stubby area
- not so stubby area (NSSA)
- backbone area (area 0)

QUESTION NO: 2

OSPF networks create Link State Advertisements (LSAs). What kind of LSA is conceived by an Area Border Router (ABR) containing route summarization of area subnets?

- A. Router link, type 1.
- B. Router link, type 8.
- C. Network summary link, type 3.
- D. AS external summary link, type 4.

Answer: C

Explanation:

The network summary link entry is a LSA type 3 that is originated by ABRs. It describes the links between the ABR and the internal routers of a local area. These entries are flooded throughout the backbone area to the other ABRs.

Incorrect Answers

- A: Router link, LSA type 1 is generated by each router for each area it belongs to. It describes the states of the router's link to the area.
- B: Router LSA is type 1.
- D. External link LSA is type 5.

QUESTION NO: 3

Router TK1 has an interface in one OSPF area and another interface in a different OSPF area. What kind of router is TK1?

- A. ABR
- B. ASBR
- C. internal router
- D. backbone router

Answer: A

Explanation:

An ABR (Area Border Router) shares an interface with at least one other OSPF area.

Incorrect Answers

- B: An ASBR (Autonomous System Border Router) have at least one interface in a non-OSPF network.
- C: An internal router is only connected to routers in the internal area.
- D: A backbone router has at least one interface in the backbone, also known as Area 0.

QUESTION NO: 4

When an OSPF ASBR announces external networks, what kind of LSA type does it use?

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

Answer: E

Explanation:

In OSPF, Type 5 LSA is used to announce external networks.

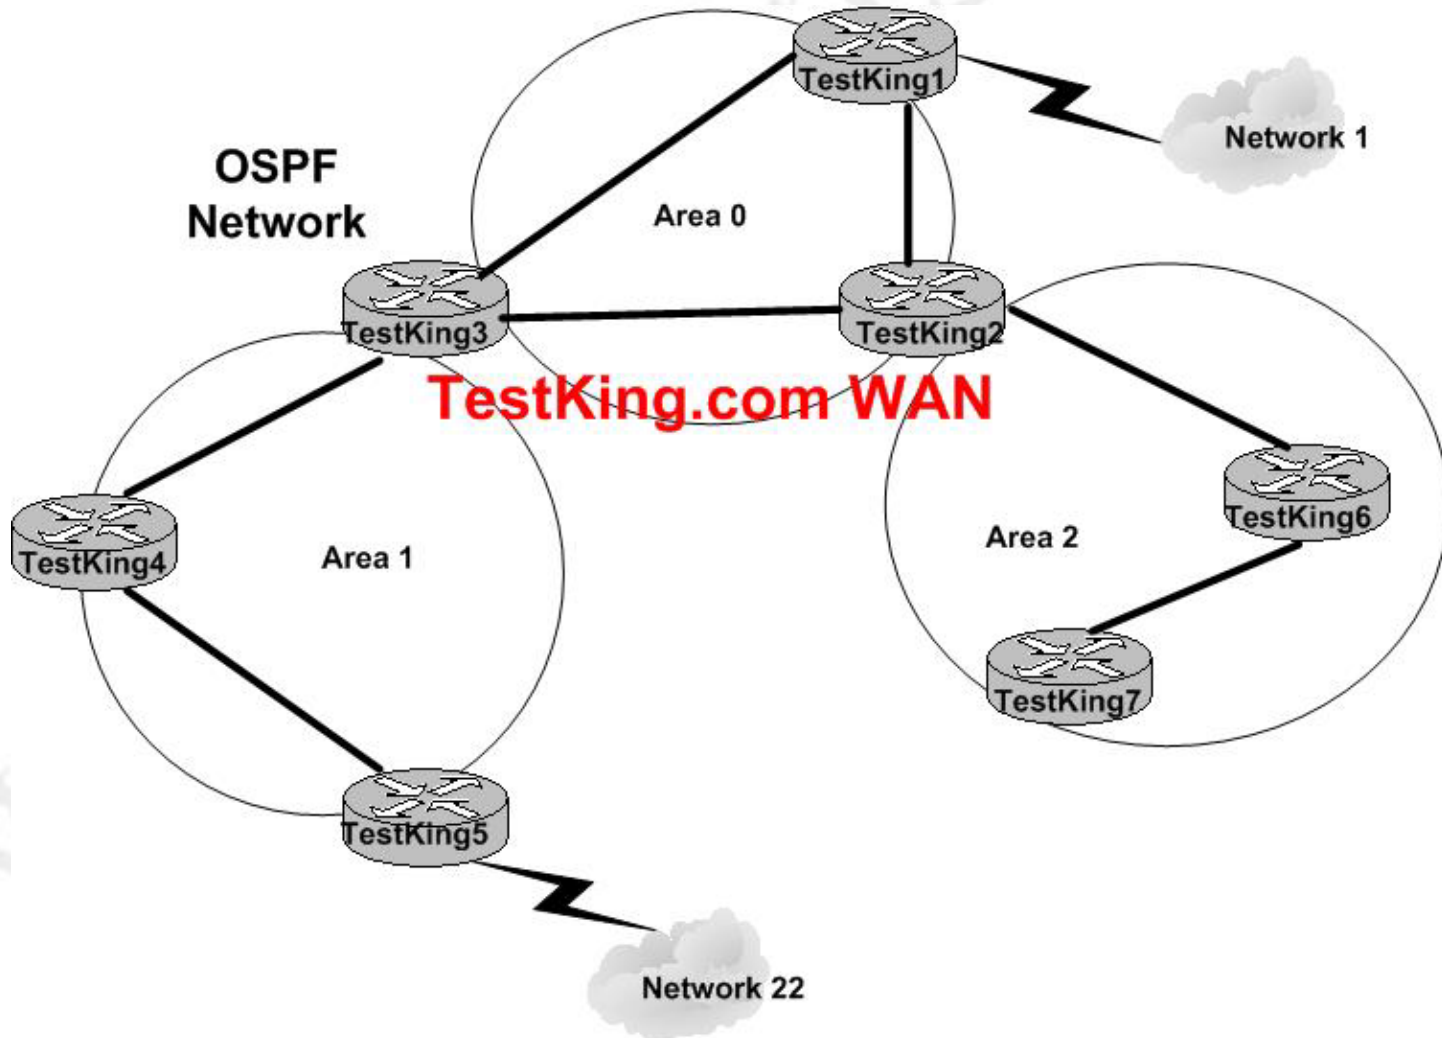
The OSPF Router LSA types are described here:

- *Router LSA* (type 1) - information about the router and its directly connected links; flooded only w/in the area

- *Network LSA* (type 2) - information about a LAN and the routers connected to it; originated by the DR; flooded only w/in the area
- *Summary LSA* (type 3) - describe networks reachable outside the area; originated by the area border routers (ABR)
- *ASBR Summary LSA* (type 4) – reachability information to the Autonomous System Border Routers outside the area; originated by the ABRs
- *External LSA* (type 5) - information about destinations outside the OSPF domain; originated by the ASBRs; flooded throughout the whole network
- Other LSA types exist for multicast information (type 6) and other extensions, including NSSA area (type 7).

QUESTION NO: 5

The TestKing OSPF network is displayed in the diagram below:



The status of external Network 22 has just changed. The router TestKing5 uses a type 7 LSA to update area 1 of this change. In OSPF terminology, what kind of area is area 1 classified as?

- A. Backbone area
- B. Transit area
- C. Stubby area
- D. Totally stubby area
- E. Not-so-stubby area

Answer: E

Explanation:

NSSA use type 7 LSAs. The different OSPF area types are described below.

- **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).
- **Normal Areas:** These areas can either be standard areas or transit (backbone) areas. Standard areas are defined as areas that can accept intra-area, inter-area and external routes.
- **Backbone area** is the central area to which all other areas in OSPF connect.
- **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.
- **NSSA:** This type of area allows the flexibility of importing a few external routes into the area while still trying to retain the stub characteristic. Assume that one of the routers in the stub area is connected to an external AS running a different routing protocol, it now becomes the ASBR, and hence the area can no more be called a stub area. However, if the area is configured as a NSSA, then the ASBR generates a NSSA external link-state advertisement (LSA) (Type-7) which can be flooded throughout the NSSA area. These Type-7 LSAs are converted into Type-5 LSAs at the NSSA ABR and flooded throughout the OSPF domain

Reference:

http://www.cisco.com/en/US/tech/tk365/tk480/technologies_tech_note09186a0080094a74.shtml

QUESTION NO: 6

What kind of information can you find in a type 3 OSPF network summary link LSA?

- A. Summary of routes in the AS.
- B. Summary of link states in an OSPF area.
- C. Summary of IP subnets in an OSPF area.

- D. Summary of metric coast from ABR to ASBR.
- E. None of the above.

Answer: C

Explanation:

A type 3 area LSA is Inter-area prefix LSA for ABR. It advertises internal networks to routers in other areas (interarea routes). Type 3 LSAs may represent a single network or a set of networks summarized into one advertisement. Only ABRs generate summary LSAs. In OSPF for IPv6, addresses for these LSAs are expressed as *prefix, prefix length* instead of *address, mask*. The default route is expressed as a prefix with length 0.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide_chapter09186a0080145c56.html

QUESTION NO: 7

Drag the OSPF area type icons on the left, and drop them into the centre target area that adjoins the corresponding description on the right.

OSPF area type, use these		Descriptions
Stub Area	place here	interconnects areas
Transit Area	place here	uses type 7 LSAs
Totally Stubby area	place here	does not accept external LSAs
Backbone Area	place here	interconnects discontinuous backbones
Not-so-stubby Area	place here	does not accept external or summary LSAs

Answer:

OSPF area type, use these

	Descriptions
Backbone Area	interconnects areas
Not-so-stubby Area	uses type 7 LSAs
Totally Stubby area	does not accept external or summary LSAs
Transit Area	interconnects discontinuous backbones
Stub Area	does not accept external LSAs

Explanation:

- 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).
- Normal Areas: These areas can either be standard areas or transit (backbone) areas. Standard areas are defined as areas that can accept intra-area, inter-area and external routes.
- The Backbone area is the central area to which all other areas in OSPF connect.
- 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.
- NSSA: This type of area allows the flexibility of importing a few external routes into the area while still trying to retain the stub characteristic. Assume that one of the routers in the stub area is connected to an external AS running a different routing protocol, it now becomes the ASBR, and hence the area can no more be called a stub area. However, if the area is configured as a NSSA, then the ASBR generates a NSSA external link-state advertisement (LSA) (Type-7) which can be flooded throughout the NSSA area. These Type-7 LSAs are converted into Type-5 LSAs at the NSSA ABR and flooded throughout the OSPF domain

Reference:

http://www.cisco.com/en/US/tech/tk365/tk480/technologies_tech_note09186a0080094a74.shtml

QUESTION NO: 8

In a multi-area OSPF network inter-area summarization produces type's 3 & 4 Link State Advertisements (LSAs). At which router are these LSA types produced?

- A. At the ABR.
- B. At the ASBR.

- C. At the backbone DR.
- D. Each internal router calculates and sends network updates.
- E. None of the above.

Answer: A

Explanation: Network summary link entries, LSA type 3 or LSA type 4 are originated by ABRs. They describe the links between the ABR and the internal routers of a local area. These entries are flooded throughout the backbone area to the other ABRs.

Note: Type-3 describes routes to networks within the local area and are sent to the backbone area. Type-4 describes reachability to ASBRs. These link entries are not flooded through totally stubby areas.

Incorrect Answers:

B: Autonomous system external link entries, LSA type 5, are originated by ASBRs. They describe routes to destinations external to the autonomous system.

C: Backbone DRs does not generate these LSAs.

D: Internal routers produce router link entries, LSA type 1. They describe the states of the router's link to the area.

QUESTION NO: 9

What does Cisco recommend for those who set up networks with multiple OSPF areas? (Select two)

- A. There should not be more than three areas per route.
- B. Area 0 must be larger than any subsequent OSPF area.
- C. A router can be a DR or BDR for more than one LAN.
- D. You should not run more than one instance of the OSPF process on an ABR.

Answer: C, D

Explanation:

C: A router can be a DR or BDR for more than one LAN. Since the loopback address is normally used as the router ID, this ID is used for all LAN segments.

D: Running multiple OSPF processes on the same router is not recommended because it creates multiple database instances that add extra overhead.

Incorrect Answers:

A: There is no such limit. The hierarchical structure could be larger.

B: There is no requirement that Area 0 must be the largest OSPF area.

QUESTION NO: 10

When designing and implementing OSPF in multiple areas; what should be taken in consideration? (Select two)

- A. Each area must connect to area 0.
- B. Each area must have a unique AS number.
- C. Remote areas must be configured as stub or NSSA areas.
- D. Traffic between two areas must travel across the backbone area.

Answer: A, D

Explanation:

In OSPF, all areas must be connected to a backbone area, which is area 0. It is also required that traffic from one area to another non-backbone area must traverse area 0, since all areas must connect to area 0.

Incorrect Answers

B: OSPF routers within the same Autonomous System should be configured with the same AS number. The AS number is also called the OSPF process ID.

C: Remote areas are not required to be either stub or not so stubby areas.

QUESTION NO: 11

In OSPF, a Designated Router (DR) floods LSAs to all the routers on the segment that it has a neighbor relationships with. Over what kind of link does this happen and what kind of LSA type is sent?

- A. Router link, type 1.
- B. Network link, type 2.
- C. External link, type 5.
- D. Network summary link, type 3.
- E. None of the above

Answer: B

Explanation:

A network Link LSA type 2 is generated by DRs in multi-access networks. They describe the set of routers attached to a particular network. Flooding occurs within the area that contains the network only.

Incorrect Answers:

A: Router link, LSA type 1 is generated by each router for each area it belongs to. It describes the states of the router's link to the area.

C: External link, LSA type 5 is originated by the ASBR. It describes routes to destinations external to the autonomous system. It is flooded throughout an OSPF autonomous system except for stub and totally stubby areas.

D: Network summary link entry, LSA type 3 is originated by ABRs. It describes the links between the ABR and the internal routers of a local area. These entries are flooded throughout the backbone area to the other ABRs.

QUESTION NO: 12

In an OSPF network, neighbor relationships allow the topology of the network to scale. What are two reasons as to why this can help a network topology scale? (Select two)

- A. Neighbor adjacencies control distribution of routing protocol updates.
- B. Routing table information does not flood the network until holddown timers have expired.
- C. The hello protocol is a more efficient means of sending routing updates than table exchanges used in RIPv1.
- D. Topological database is maintained with incremental updates, with full exchange occurring only every 30 minutes.

Answer: A, C

Explanation:

When a network topology changes and becomes more sophisticated, fewer routers are burdened with the responsibility of having to relearn the entire network. The hello protocol is also efficient because rather than focusing on knowing what all the routers know, the routers only focus on who is still around, therefore available for routing.

QUESTION NO: 13

In an OSPF network; which IP address multicasts all OSPF Designated Routers and Backup Designated Routers?

- A. 224.0.0.5
- B. 224.0.0.6
- C. 224.0.0.9
- D. 224.0.0.11
- E. 224.0.0.2

Answer: B

Explanation:

224.0.0.6 is the IP address used by all OSPF DRs and BDRs.

Incorrect Answers:

- A: This is the multicast address used by all OSPF routers.
- C: This address is used by RIP version 2 routers.

D: This multicast address is used by mobile agents.

E: This multicast IP address is used by all routers on the subnet.

Reference: Building Scalable Cisco Networks (Cisco Press) page 114

QUESTION NO: 14

What kind of OSPF Link State Advertisement (LSA) is found in a type 3 network summary link?

- A. Summary of routes in the AS.
- B. Summary of link states in an OSPF area.
- C. Summary of IP subnets in an OSPF area.
- D. Summary of metric costs from ABR to ASBR.

Answer: C

Explanation:

Type 3 LSAs are INTER-AREA advertisements and they are not used within the local area.

The various LSA types used in OSPF are described below:

LS Type	Link State ID
1	The originating router's Router ID.
2	The IP interface address of the network's Designated Router.
3	The destination network's IP address.
4	The Router ID of the described AS boundary router.
5	The destination network's IP address.

QUESTION NO: 15

You are the administrator of the large TestKing OSPF and are considering dividing it to multiple areas. What are two advantages of configuring multiple OSPF areas? (Select two)

- A. It eliminates the need for a DR or BDR in the OSPF network.
- B. It eliminates security concerns by segregating portions of the network.
- C. Type 1 and 2 LSAs are confined to a single area, reducing routing overhead.
- D. Area members have smaller topological databases than if the network was one large area, requiring less CPU to derive routes.

Answer: C, D

Explanation:

C: Link State update (LSU) traffic is reduced. Rather than send an LSU about each network within an area, you can advertise a single or fewer summarized routes between areas to reduce the overhead associated with linkstate updates when they are crossing areas.

Note: All LSA types, except the AS-external-LSAs (LS type = 5), are flooded throughout a single area only.

D: Topological database size would decrease and fewer CPU cycles would be required to calculate routes.

Incorrect Answers:

A: A DR and BDR for each area would still be required.

B: Security concerns would not be eliminated. There are no security advantages in using more than one area in an OSPF network.

QUESTION NO: 16

OSPF has been configured on router TK1, which lies on a multi point network. On this multi-access network, what type of router floods type 2 Link State Advertisements (LSAs) into an OSPF area?

- A. DR
- B. ABR
- C. BDR
- D. ASBR

Answer: A

Explanation: Type 2 LSAs are generated by Designated Routers (DRs) in multi-access networks. They describe the set of routers attached to a particular network and are flooded within the area that contains the network only.

Note: Instead of routers exchanging link-state information with every other router on the segment, each router sends the link-state information to the DR and Backup Designated Router (BDR). The DR sends each router's link-state information to all other routers in the network. This flooding process significantly reduces the router-related traffic on a segment.

Incorrect Answers

B: Area Border Routers generate Type 3 and Type 4 LSAs.

C: BDRs are just used for backup, should the DR fail.

D: Autonomous System Boundary Routers (ASBRs) generate Type 5 LSAs.

QUESTION NO: 17

You want router one of the OSPF areas on your network to be able to accept type 7 external routers, but not type 5 external routes. What kind of area should be configured to accomplish this?

- A. A not-so stubby area
- B. A stubby area
- C. A totally stubby area
- D. A backbone area
- E. An on-demand area

Answer: A

Explanation:

The NSSA external LSA —Identified as Type 7, these LSAs are created by the ASBR residing in a not so stubby area (NSSA). This LSA is similar to an autonomous system external LSA, except that this LSA is contained within the NSSA area and is not propagated into other areas, but it is converted into a Type 5 LSA by the ABR.

An NSSA is an area that is seen as a stub area but can receive external routes, which it will not propagate into the backbone area and thus the rest of the OSPF domain. Another LSA, Type 7, is created specifically for the NSSA. This LSA can be originated and communicated throughout the area, but it will not be propagated into other areas, including Area 0. If the information is to be propagated throughout the AS, it is translated into an LSA Type 5 at the NSSA ABR.

QUESTION NO: 18

Match the OSPF area type on the left to the type of OSPF routes (as would be seen in the routing table) it will support on the right hand side:

O, O* IA	Stub	Place here
O, O IA, O N2, O* N2	Totally Stubby	Place here
O, O IA, O*IA	NSSA	Place here

Answer:

Stub	O, O IA, O*IA
Totally Stubby	O, O* IA
NSSA	O, O IA, O N2, O* N2

QUESTION NO: 19

Which of the following OSPF routes are supported by the Totally Stubby type? (Select two)

- A. O
- B. O IA
- C. O N2
- D. O* N2
- E. O* IA

Answer: A, E

Explanation:

The only routes that are seen in a totally stubby area are intra-area routes (O) and the default Intra area route (O*IA).

Example routing table of a totally stubby router:

TK1#**show ip route**

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate
```

default

Gateway of last resort is not set

```
      203.250.15.0 255.255.255.252 is subnetted, 1 subnets
C       203.250.15.0 is directly connected, Serial0
      131.108.0.0 255.255.255.240 is subnetted, 1 subnets
O       131.108.79.208 [110/74] via 203.250.15.1, 00:31:27, Serial0
O*IA 0.0.0.0 0.0.0.0 [110/74] via 203.250.15.1, 00:00:00, Serial0
```

Note that the only routes that show up are the intra-area routes (O) and the default-route 0.0.0.0. The external and inter-area routes have been blocked. The cost of the default route is now 74 (64 for a T1 line + 10 advertised by RTC). No configuration is needed on RTE in this case. The area is already stub, and the **no-summary** command does not affect the Hello packet at all as the stub command does.

QUESTION NO: 20

Which type of OSPF network requires manual configuration of OSPF neighbors?

- A. broadcast multi-access
- B. nonbroadcast multi-access
- C. point-to-point

D. point-to-multipoint

Answer: B

Explanation:

Non-broadcast multi-access networks, as found in frame-relay and ATM, require the neighbors to be explicitly specified. This is the case where the PVCs are configured on the physical interface using frame map commands. When sub-interfaces are used, the neighbor commands do not need to be configured.

Note: On the physical interface of an NBMA network, the command “ip ospf network-type point-to-multipoint” command can be used instead of specifying the neighbors individually.

QUESTION NO: 21

Which of the following OSPF routes are supported by the Stub Area type? (Select three)

- A. O
- B. O IA
- C. O N2
- D. O* N2
- E. O* IA

Answer: A, B, E

Explanation:

When an area is configured as a stub area, it will learn about all routes, except for the external routes.

Example routing table of a stub router:

TK2#**show ip route**

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate
default
```

```
Gateway of last resort is 203.250.15.1 to network 0.0.0.0
```

```
203.250.15.0 255.255.255.252 is subnetted, 1 subnets
C      203.250.15.0 is directly connected, Serial0
O IA   203.250.14.0 [110/74] via 203.250.15.1, 00:26:58, Serial0
       128.213.0.0 255.255.255.252 is subnetted, 1 subnets
O IA   128.213.63.0 [110/84] via 203.250.15.1, 00:26:59, Serial0
       131.108.0.0 255.255.255.240 is subnetted, 1 subnets
O      131.108.79.208 [110/74] via 203.250.15.1, 00:26:59, Serial0
O*IA   0.0.0.0 0.0.0.0 [110/65] via 203.250.15.1, 00:26:59, Serial0
```

Leading the way in IT testing and certification tools, www.testking.com

Note that all the routes show up except the external routes which were replaced by a default route of 0.0.0.0.

QUESTION NO: 22

Which of the following OSPF routes are supported by the Not-So-Stubby-Area (NSSA) type? (Select three)

- A. O N1
- B. O IA
- C. O N2
- D. O* N2
- E. O* IA

Answer: B, C, D

Explanation:

The various route types used by OSPF are as follows:

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

An OSPF NSSA will receive inter-area, external type 2, and default external type 2 routes.

QUESTION NO: 23

Routers TK1, TK2, and TK3 are part of OSPF area 100 as shown below:



In the TestKing network shown above, router TK1 was elected the DR and TK2 was elected the BDR. Assume that TK1 fails, and that TK2 takes the place of the DR while TK3 becomes the new BDR. What will happen when TK1 comes back online?

- A. TK1 will take the place of DR immediately upon establishing its adjacencies.
- B. TK1 will take the place of DR only if TK2 fails.
- C. TK1 will take the place of DR only if both TK2 and TK3 fail.
- D. A new election will take place establishing an all new DR and BDR based on configured priority levels and MAC addresses.

Answer: C

Explanation:

If a router with a higher priority value gets added to the network, it does not preempt the DR and BDR. The only time a DR and BDR changes is if one of them is out of service. If the DR is out of service, the BDR becomes the DR, and a new BDR is selected. If the BDR is out of service, a new BDR is elected. In a multi-access network, the router that is powered on first will generally become the DR, since the DR/BDR process is not pre-emptive.

Reference:: CCNP Self-Study Second Edition P.243

QUESTION NO: 24

Which of the following are features are characteristics of an OSPF ABR? (Choose three)

- A. maintains one link-state database for all areas
- B. routes traffic to and from other areas
- C. maintains a separate link-state database for each area
- D. originates type 3 link-state advertisements which can be summarized

Leading the way in IT testing and certification tools, www.testking.com

- E. has at least one physical interface in area 0.
- F. has at least one interface in an external internetwork.

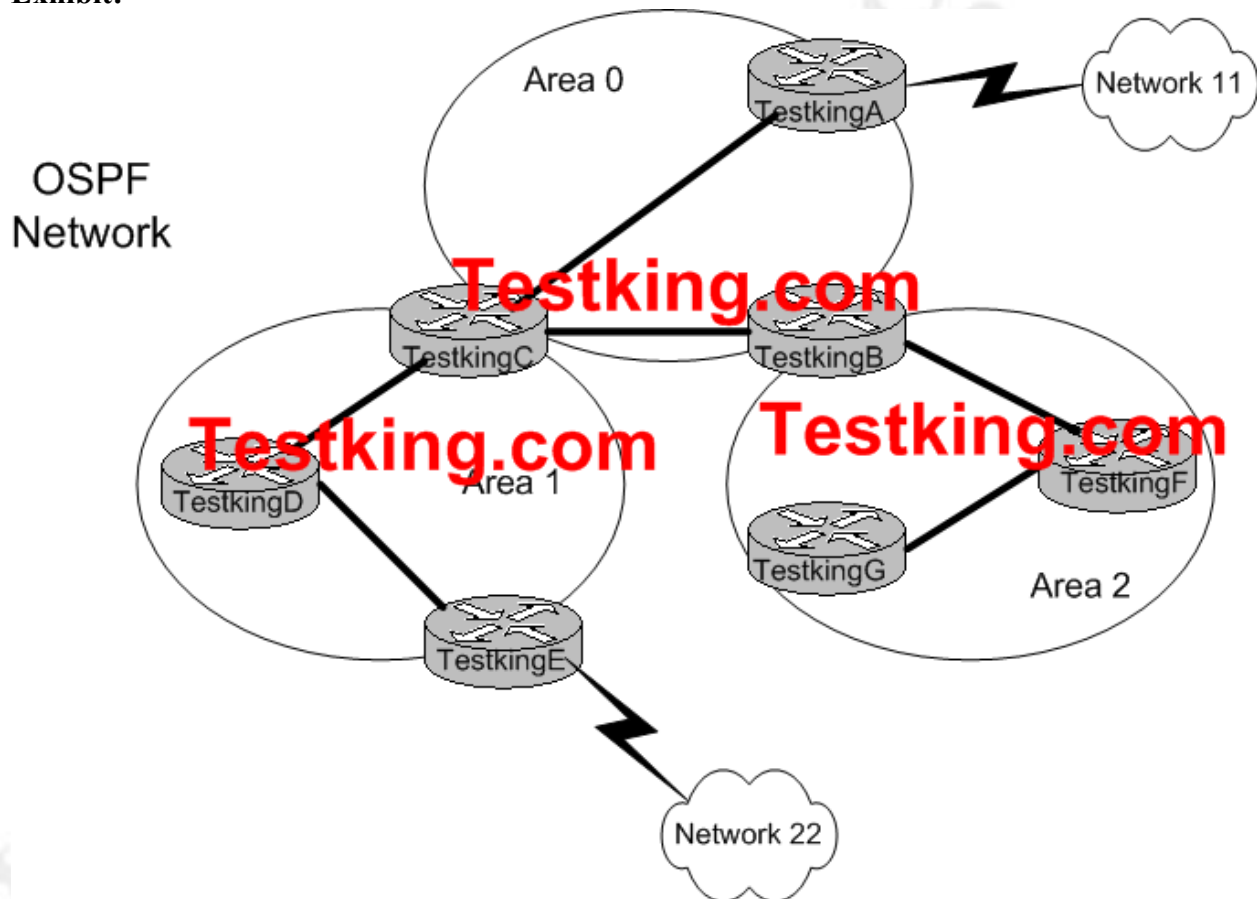
Answer: B, C, D

Explanation:

ABR forwards summary LSAs. It forwards both Type 3 LSAs and Type 4 LSAs. Type 3 LSAs are forwarded to the other ABRs, and Type 4 LSAs are forwarded to the ASBRs. ABR also forwards Type 3 LSAs from other areas into its own area. If the ABR has multiple links in the same area, it also forwards Type 1 and Type 2 LSAs in its capacity as an internal router.

QUESTION NO: 25

Exhibit:



Which type of Link State Announcement will TestkingE use to announce changes in network 22 to Area 1?

- A. Type 1 LSA
- B. Type 2 LSA

- C. Type 3 LSA
- D. Type 4 LSA
- E. Type 5 LSA

Answer: E

Explanation:

The external link LSA is originated by AS boundary routers and is flooded throughout the AS. Each external advertisement describes a router to a destination in another autonomous system. Default routes for the AS can also be described by AS external advertisements. This is identified as a Type 5 LSA.

QUESTION NO: 26

What is the OSPF cost for the 100 Mbps Ethernet media type?

- A. 1
- B. 2
- C. 6
- D. 10
- E. 100

Answer: A

Section 10: Explain basic OSI terminology and network layer protocols used in OSI (5 questions)

QUESTION NO: 1

Using OSI terms, what type of system is a router?

- A. end system
- B. routing system
- C. inter-domain system
- D. intermediate system
- E. intra-domain system

Answer: D

Explanation:

Data communication sessions generally occur between end stations, or end hosts. Routers are considered to be intermediate systems, since they are used as a means for transporting data between end stations. An easy way to remember this is the OSI routing protocol IS-IS, which stands for "Intermediate System to Intermediate System."

QUESTION NO: 2

The network layer of the OSI protocol suite defines two of the routing protocols below. Which two are they? (Select two)

- A. End System-to-End System
- B. Routing Information Protocol
- C. Interior Gateway Routing Protocol
- D. Enhanced Interior Gateway Routing Protocol
- E. Intermediate System-to-Intermediate System
- F. End System-to-Intermediate System

Answer: E, F

Explanation:

Both are defined by ISO in the OSI protocol suite.

Intermediate System-to-Intermediate System (IS-IS) Protocol is an intradomain Open System Interconnection (OSI) dynamic routing protocol specified in International Organization for Standardization (ISO) 10589. The protocol is designed to operate in OSI Connectionless Network Service (CLNS). Data is carried using the protocol specified in ISO 8473.

OSI CLNS is a network layer service similar to bare IP service. A CLNS entity communicates over Connectionless Network Protocol (CLNP) with its peer CLNS entity.

Incorrect Answers:

A: This is not a routing protocol.

B: RIP is a standards based routing protocol, but it is not defined by the OSI or ISO organizations.

C, D: These answer choices refer to the Cisco proprietary routing protocols, IGRP and EIGRP.

Reference:

http://www.cisco.com/en/US/tech/tk365/tk381/technologies_white_paper09186a00800a3e6f.shtml

QUESTION NO: 3

Which of the following routing protocols belongs in the network layer, as defined by the OSI protocol suite?

- A. End System-to End System
- B. Routing Information Protocol
- C. Interior Gateway Routing Protocol
- D. Enhanced Interior Gateway Routing Protocol
- E. Intermediate System-to-Intermediate System
- F. None of the above

Answer: E**Explanation:**

Intermediate System-to-Intermediate System (IS-IS) Protocol is an intradomain Open System Interconnection (OSI) dynamic routing protocol specified in International Organization for Standardization (ISO) 10589. The protocol is designed to operate in OSI Connectionless Network Service (CLNS). Data is carried using the protocol specified in ISO 8473.

OSI CLNS is a network layer service similar to bare IP service. A CLNS entity communicates over Connectionless Network Protocol (CLNP) with its peer CLNS entity.

Incorrect Answers:

A: This is not a routing protocol.

B: RIP is a standards based routing protocol, but it is not defined by the OSI or ISO organizations.

C, D: These are the Cisco proprietary routing protocols, IGRP and EIGRP.

Reference:

http://www.cisco.com/en/US/tech/tk365/tk381/technologies_white_paper09186a00800a3e6f.shtml

QUESTION NO: 4

In network engineering and computer science the term domain can be interpreted in multiple contexts. How is a domain defined in OSI terminology?

- A. A set of non-routing network nodes.
- B. A contiguously connected area that can reach all other areas.
- C. All devices configured for OSI protocols within an internetwork.
- D. A contiguous set of routers and hosts and the data links that connect them.
- E. Any portion of an OSI network that is under a common administrative authority.
- F. None of the above

Answer: E

Explanation: A domain is any portion of an OSI network that is under a common administrative authority. A domain is a collection of End Systems, Intermediate Systems, and subnetworks operated by a single organization or administrative authority. The components which make up the domain are assumed to interoperate with a significant degree of mutual trust among them, but interoperate with other Administrative Domains in a mutually suspicious manner

Reference: ISO TR 9575

QUESTION NO: 5

How is network layer addressing accomplished in the OSI protocol suite?

- A. Internet Protocol address
- B. Media Access Control address
- C. Packet Layer Protocol address
- D. Network Service Access Point address
- E. Authority and Format Identifier address
- F. None of the above

Answer: D

OSI network layer addressing is implemented by using two types of hierarchical addresses: network service access point addresses and network entity titles.

A *network service access point (NSAP)* is a conceptual point on the boundary between the network and the transport layers. The NSAP is the location at which OSI network services are provided to the transport layer. Each transport layer entity is assigned a single NSAP, which is individually addressed in an OSI internetwork using NSAP addresses.

A network entity title (NET) is used to identify the network layer of a system without associating that system with a specific transport layer entity (as an NSAP address does). NETs are useful for addressing intermediate systems (ISs), such as routers, that do not interface with the transport layer. An IS can have a single NET or multiple NETs, if the IS participates in multiple areas or domains.

Reference: http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/osi_prot.htm#1022230

Section 11: Identify similarities and differences between Integrated IS-IS and OSPF (6 questions)

QUESTION NO: 1

Drag the routing protocol characteristics on the left to the corresponding OSPF box in the middle or the corresponding IS-IS box on the right:

IS-IS term	Drop OSPF Characteristics here	Drop IS-IS Characteristics here
1000 routers per area	place here	place here
50 routers per area	place here	place here
routers belong to one level 2 area	place here	place here
flexible extension of backbone		
backup designated router	place here	place here
boundaries lie inside routers	place here	place here
boundaries lie on links		
forms adjancies with all neighbors		

Answer:

IS-IS term	Drop OSPF Characteristics here	Drop IS-IS Characteristics here
	50 routers per area	1000 routers per area
	backup designated router	routers belong to one level 2 area
	boundaries lie inside routers	boundaries lie on links
	place here	forms adjancies with all neighbors
	place here	flexible extension of backbone

QUESTION NO: 2

The similar routing protocols OSPF and IS-IS share many characteristics. Which four of the characteristics below do they share? (Select four)

- A. link-state database

- B. backbone design using area 0
- C. Shortest Path First (SPF) algorithm
- D. Update, Decision, and Flooding Process
- E. Hello protocol to establish and maintain adjacencies

Answer: A, C, D, E

Explanation:

- A: IS-IS uses a link-state database, similar to OSPF. Both OSPF and IS-IS are considered to be link state routing protocols.
- C: The IS-IS decision process runs shortest-path-first (SPF) algorithm on the link-state database, and creates the forwarding database.
- D: The Update, Decision, and Flooding Process of IS-IS and OSPF are similar.
- E: Routers running IS-IS will send hello packets out all IS-IS-enabled interfaces to discover neighbors and establish adjacencies. This is similar to OSPF.

Incorrect Answers:

- B: IS-IS does not have a backbone area like the OSPF area 0. The IS-IS backbone is a contiguous collection of Level 2-capable routers, each of which can be in a different area.

Reference: Introduction to Intermediate System-to-Intermediate System Protocol

http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm

QUESTION NO: 3

OSPF and IS-IS are similar in many ways. In OSPF, the backbone area is area 0. How is an IS-IS backbone area numbered?

- A. Area 0
- B. Area 1
- C. Area 51
- D. Any legal area number.
- E. There is no backbone area number.

Answer: E

Explanation:

IS-IS does not have a backbone area like the OSPF area 0. The IS-IS backbone is a contiguous collection of Level 2-capable routers, each of which can be in a different area.

Reference: Cisco, Introduction to Intermediate, System-to-Intermediate System Protocol

http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm

QUESTION NO: 4

OSPF are similar in many ways, but there are some important differences between them. Which three characteristics apply to IS-IS but not to OSPF? (Choose three)

- A. Encapsulates PDUs directly into a data-link frame.
- B. Uses a DIS and a backup DIS to present the pseudo-node on the LAN.
- C. Uses stubby areas to improve network scalability.
- D. Uses a default IOS metric of 10 on each interface.
- E. Runs PRC (Partial Calculations) to calculate IP reachability information.
- F. Uses an on-demand circuit to reduce the hello and LSA flooding across switched WAN links, such as ISDN.

Answer: A, D, E

Explanation:

IS-IS encapsulates protocol data units (PDUs) directly into a data link frame, as outlined by the OSI standard. With IS-IS each link has a metric of 10, regardless of the actual bandwidth of the link. Finally, the PRC is the software's process of calculating routes without performing an SPF calculation.

Incorrect Answers:

B: On broadcast multi-access networks, a single router is elected as the DIS. There is no backup DIS elected. The DIS is the router that creates the pseudonode and acts on behalf of the pseudonode.

C: Stubby areas are only used by OSPF.

F: On demand circuits are only used by OSPF, to suppress the LSA hello packets from triggering DDR calls.

QUESTION NO: 5

IS-IS is often considered an alternative to OSPF in the IP world. Which two statements identify similarities between IS-IS and OSPF? Select two.

- A. support for designated intermediate systems and backup designated intermediate systems
- B. support for multiple areas per router
- C. support for classless routing
- D. support for address summarization between area
- E. support for both DIS adjacencies and neighbor adjacencies

Answer: C, D

Explanation:

Similarities Between Integrated IS-IS and OSPF

They are both link-state protocols and are based on the Dijkstra algorithm of Shortest Path First (SPF). In addition, they both have a two-level hierarchy. OSPF tends to be deployed mostly as an enterprise solution, whereas Integrated IS-IS is used for IP routing in some ISP networks. Both support address summarization between areas.

QUESTION NO: 6

**Which of the following statements are correct about the differences in IS-IS and OSPF?
(Choose three)**

- A. IS-IS LSP contains TLV fields and OSPF LSU contains the LSAs.
- B. New additions to the protocol are easily implemented in OSPF but not with IS-IS
- C. For greater fine tuning there are more options available in IS-IS.
- D. OSPF has more area types than does IS-IS.
- E. IS-IS is more CPU-intensive than is OSPF.

Answer: A, C, D

Explanation:

TLV's are a main strength of IS-IS. TLV's provide flexibility and extend the functionality of the protocol. In IS-IS, the TLV fields are variable in length and appended to the various packet formats. This means that IS-IS has advantages over OSPF and the protocol can adapt to the changing needs and advances to technology by simply defining a new TLV.

Section 12: List the types of IS-IS routers and their role in IS-IS area design (5 questions)

QUESTION NO: 1

There are several varieties of IS-IS routers. Two types of IS-IS provide intra-area routing services, meaning that they are used to route traffic within the same area.. Which of the following two match this description? (Select two)

- A. L1 IS
- B. L1 ES
- C. L2 IS
- D. L2 ES
- E. L1/L2 IS

Answer: A, E

Explanation:

L1 IS and L1/L2 IS routers provide intra-area routing services. A Level 1 router knows the topology only of its own area and has Level 1 or Level1/Level2 neighbors in this area. It has a Level 1 link-state database with all the information for intra-area routing. It uses the closest Level 2-capable router in its own area to send packets out of the area, a scenario that may result in suboptimal routing.

Incorrect Answers:

B, D: End systems are OSI terms used for end hosts, such as PCs or servers. They do not provide for any type of routing functionality.

C: L2 routers make up the backbone of the IS-IS network, and they are used to provide for inter-area routing.

QUESTION NO: 2

In OSPF, the backbone area is area 0. How could you describe the backbone area of an IS-IS domain?

- A. A contiguous collection of L1 routers.
- B. A contiguous collection of L1 and L1/L2 routers.
- C. A contiguous collection of L1 and L2 routers.
- D. A contiguous collection of L2 and L1/L2 routers.
- E. None of the above.

Answer: D

Explanation:

IS-IS does not have a backbone area like the OSPF area 0. The IS-IS backbone is a contiguous collection of Level 2-capable routers, each of which can be in a different area. L2 routers are used for inter-area routing, and a L1 router can only route traffic to other areas via an L2 or and L1/L2 IS-IS router.

QUESTION NO: 3

Which one of the following statements correctly describes a characteristic of an IS-IS backbone?

- A. A router in all areas.
- B. All routers in one area.
- C. A chain of L2 and L1/L2 IS-IS routers.
- D. An unbroken string of L1 and LS IS-IS routers.
- E. A central area to which all other areas are attached.
- F. None of the above

Answer: C

Explanation:

IS-IS does not have a backbone area like the OSPF area 0. The IS-IS backbone is a contiguous collection of Level 2-capable routers, each of which can be in a different area. L1 routers are used for internal routing within the area and so only L2 and L1/L2 routers form the backbone in IS-IS.

QUESTION NO: 4

The routing protocol IS-IS uses areas. What must an administrator connect each IS-IS area to?

- A. Area 0
- B. Area 1
- C. Level-1 backbone
- D. Level-2 backbone
- E. External IS-IS areas

Answer: D

Explanation:

Small IS-IS networks are built as a single area that includes all the routers in the network. As the network grows larger, it is usually reorganized into a backbone area made up of the connected set of all Level 2 routers from all areas, which are in turn connected to local areas. Routers that are solely used as L1 routers must route all traffic that is destined outside of the area to a L1/L2

or a L2 router. The contiguous set of L2 routers make up the OSPF backbone, similar to area 0 in OSPF networks.

Incorrect Answers:

- A, B: Area 0 or Area 1 has no special significance in IS-IS. In IS-IS, areas are broken down into levels and are not numbered like OSPF areas.
- C: Level 2, not Level 1.
- E: This is not a requirement.

QUESTION NO: 5

Which packet type is used to acknowledge LSPs on point-to-point links and to request missing pieces of information in the IS-IS link-state database?

- A. complete SNP (CSNP)
- B. partial SNP (PSNP)
- C. hello
- D. database query
- E. database description packet (DDP)

Answer: B

Explanation:

Partial SNP (PSNP) – Includes a subset of LSPs, used to request individual LSPs and to acknowledge receipt of these LSPs: Level1,Level2.

Reference: CCNP Self-Study CCNP BSCI Exam Certification Guide p.361

Section 13: Describe the hierarchical structure of IS-IS areas (22 questions)**QUESTION NO: 1**

In your network you have two Level-1/ Level-2 IS-IS routers named TK1 and TK2. The two routers are connected together via a WAN link, and they both belong to the same area. What kind of an adjacency can you expect them to establish?

- A. Level-1 only
- B. Level-2 only
- C. Both Level-1 and Level-2
- D. None unless statically configured
- E. WAN links do not support IS-IS adjacencies

Answer: C

Explanation:

IS-IS routers that are configured as L1/L2, which is the default, will run two separate SPF algorithms, one for each level. Neighbors on point-to-point networks always become adjacent unless they do not see themselves in their neighbors' hello PDU and match on certain parameters. On broadcast networks and nonbroadcast multiaccess (NBMA) networks, the DIS (Designated Intermediate System) will become adjacent with its neighbors.

Two routers will become neighbors if the following parameters are agreed upon:

- Level 1—The two routers sharing a common network segment must have their interfaces configured to be in the same area if they are to have a Level 1 adjacency.
- Level 2—The two routers sharing a common network segment must be configured as Level 2 if they are in different areas and want to become neighbors.

Within an Area Level 1 routing is used. Routing between areas is referred to as Level 2 routing. A Level 1/Level 2 router may have neighbors in any area. It has two link-state databases: a Level 1 link-state database for intra-area routing and a Level 2 link-state database for inter-area routing. A Level 1/Level 2 router runs two SPFs and may require more memory and processing as a result.

QUESTION NO: 2

You have a network segment with an IS-IS level 1-2 router, and you're thinking of upgrading your network. What is the maximum number of Level-2 routing processes that this router can belong to?

- A. 1
- B. 2
- C. 30
- D. 40
- E. limited only by the router's resources
- F. Depends on the total number of IDB blocks.

Answer: A

Explanation:

The Cisco IOS software can handle simultaneous operation of up to 30 dynamic IP routing processes. The combination of routing processes on a router or access server consists of the following protocols (with the limits noted):

- Up to 30 IGRP routing processes
- Up to 30 OSPF routing processes
- One RIP routing process
- **One IS-IS process**
- One BGP routing process

Up to 30 EGP routing processes

QUESTION NO: 3

Which IOS command would you enter to if you wanted to see the IS-IS level-2 routing database table?

- A. `show isis route`
- B. `show clns route`
- C. `show isis database`
- D. `show clns neighbors`

Answer: C

Explanation:

The **show isis database (detail)** command displays the contents of the IS-IS database. By default, this command will display all of the L1 and L2 database information.

QUESTION NO: 4

Which of the following types of routers can a Level-2 IS establish adjacencies with? (Select three)

- A. Any Level-1 IS in any area
- B. Any Level-2 IS in any area

- C. Any Level-1 IS in the same area
- D. Any Level-1/Level-2 IS in any area
- E. Any Level-1/Level-2 IS in the same area

Answer: B, D, E

Explanation:

The various adjacency types that an IS-IS router can be are shown below:

level-1 A Level 1 adjacency may be established if there is at least one area address in common between this system and its neighbors. Level 2 adjacencies will never be established over this interface.

level-1-2 A Level 1 and Level 2 adjacency is established if the neighbor is also configured as **level-1-2** and there is at least one area in common. If there is no area in common, a Level 2 adjacency is established. This is the default.

level-2-only Level 2 adjacencies are established if the other routers are L2 or L1L2 routers and their interfaces are configured for L1L2 or L2. Level 1 adjacencies will never be established over this interface.

Reference:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/cs/csprtn1/csisis.htm>

QUESTION NO: 5

Which IOS command displays the Level-1 routing table in Integrated IS-IS?

- A. show isis route
- B. show clns route
- C. show isis database
- D. show clns neighbors

Answer: A

Explanation:

The “show isis route” command displays the IS-IS Level 1 forwarding table for IS-IS-learned routes.

QUESTION NO: 6

One of the characteristics of the routing protocol IS-IS is its ability to use level-1 routing. Where in an IS-IS network does this happen?

- A. Between domains.

- B. Between areas.
- C. Between intermediate systems in the same area.
- D. Between end systems and intermediate systems in the same area.

Answer: C

Explanation:

A two-level hierarchy is used to support large routing domains. A large domain may be administratively divided into areas. Each system resides in exactly one area.¹ **Routing within an area is referred to as Level 1 routing.** Routing between areas is referred to as Level 2 routing. A Level 2 Intermediate System (IS) keeps track of the paths to destination areas. A Level 1 IS keeps track of the routing within its own area. For a packet destined for another area, a Level 1 IS sends the packet to the nearest Level 2 IS in its own area, regardless of what the destination area is. Then the packet travels via Level 2 routing to the destination area, where it may travel via Level 1 routing to the destination. It should be noted that selecting an exit from an area based on Level 1 routing to the closest Level 2 IS might result in suboptimal routing.²

Reference:

http://www.cisco.com/en/US/tech/tk365/tk381/technologies_white_paper09186a00800a3e6f.shtml

QUESTION NO: 7

Your newly appointed TestKing trainee wants to know why Level-3 area routing is not a supported feature of integrated IS-IS on Cisco routers.

What will your reply be?

- A. The system ID on a Cisco router is limited to 6 bytes.
- B. The NET on a Cisco router is restricted to a maximum of 8 bytes.
- C. The lack of a Domain portion of the NSAP only accommodates for 2 levels of routing hierarchy.
- D. Cisco routers cannot route CLNS data that use the ISO/IEC 10589 standard of NSAP addressing.
- E. Since the NSAP service identifier (N-SEL) must always be set to 00, no other service types are available.

Answer: C

Explanation:

Integrated IS-IS is a version of the OSI IS-IS routing protocol that uses a single routing algorithm to support more network layer protocols than just CLNP. Integrated IS-IS sometimes is called Dual IS-IS, named after a version designed for IP and CLNP networks. Only one IS-IS process is allowed whether you run it in integrated mode, ISO CLNS only or IP only.

QUESTION NO: 8

**Which of the following three statements correctly describe the characteristics of IS-IS?
(Select three)**

- A. L1 routers have no knowledge about routes outside their area.
- B. L1/L2 routers maintain a separate Level 1 link-state database and a Level 2 link-state database; they do not advertise L2 routes to L1 routers.
- C. To route packets to another area, L1 routers must forward the packets to the L2 router of the destination area.
- D. To route packets to another area, L1 routers must forward the packets to an L1/L2 router within their area.
- E. L2 routers form adjacencies with L1 and L1/L2 neighbors.

Answer: A, B, D

Explanation:

Intermediate System-to-Intermediate System (IS-IS) is an OSI link-state hierarchical routing protocol that floods the network with link-state information to build a complete, consistent picture of network topology. To simplify router design and operation, IS-IS distinguishes between Level 1 and Level 2 ISs. Level 1 ISs communicate with other Level 1 ISs in the same area. Level 2 ISs route between Level 1 areas and form an intradomain routing backbone. Hierarchical routing simplifies backbone design because Level 1 ISs need to know only how to get to the nearest Level 2 IS. The backbone routing protocol also can change without impacting the intra-area routing protocol.

Reference: http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/osi_rout.htm

QUESTION NO: 9

On router TK1, IS-IS is configured on all Ethernet interfaces. In IS-IS, what network type identifies a LAN interface?

- A. broadcast
- B. point-to-point
- C. pseudo-node
- D. non-broadcast
- E. point-to-multipoint

Answer: A

Explanation:

The types of networks that IS-IS defines include Point-to-point networks and Broadcast networks. Broadcast networks are used on all LAN interfaces.

Reference: http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm

Incorrect Answers:

- B: Point-to-point networks, such as serial lines, connect a single pair of routers.
- C: A Designated Intermediate System (DIS) creates a pseudonode (a virtual node), and all the routers on a LAN, including the DIS, form an adjacency with the pseudonode instead of forming $n*(n-1)$ order adjacencies with each other in a full mesh. DISs are not used by default however.
- D: Non-broadcast is not used by IS-IS.
- E: Point-to-multipoint is not used by IS-IS.

QUESTION NO: 10

The TestKing network is utilizing IS-IS for the routing protocol. By default, what is the metric used by the TestKing routers on IS-IS interfaces?

- A. Delay
- B. Error
- C. Default
- D. Expense
- E. Bandwidth
- F. Hops

Answer: C

Explanation: The default metric used on Cisco IOS for IS-IS is cost. For each IS-IS interface, the range is from 0 to 63 with the default value being 10.

QUESTION NO: 11

In the IS-IS routing protocol, what metric is supported exclusively by the Cisco IOS?

- A. Cost
- B. Delay
- C. Expense
- D. Error
- E. None of the above


Answer: A**Explanation:**

The original IS-IS specification defines four different types of metrics. Cost, being the default metric, is supported by all routers. Delay, expense, and error are optional metrics. The delay metric measures transit delay, the expense metric measures the monetary cost of link utilization, and the error metric measures the residual error probability associated with a link.

The Cisco implementation uses cost only. If the optional metrics were implemented, there would be a link-state database for each metric and SPF would be run for each link-state database. IS-IS uses a single required default metric with a maximum path value of 1024. The metric is arbitrary and typically is assigned by a network administrator. Any single link can have a maximum value of 64, and path links are calculated by summing link values. Maximum metric values were set at these levels to provide the granularity to support various link types while at the same time ensuring that the shortest-path algorithm used for route computation will be reasonably efficient. IS-IS also defines three optional metrics (costs): delay, expense, and error.
Reference:http://www.cisco.com/en/US/tech/tk365/tk381/technologies_white_paper09186a00800a3e6f.shtml

QUESTION NO: 12

The TestKing network is displayed in the diagram below:



In the network segment of the above exhibit, which router is going to become the level-2 DIS (Designated Intermediate System)?

- A. TestKing1
- B. TestKing2
- C. TestKing3
- D. TestKing4
- E. None

Answer: B

Explanation:

Since the Level 2 priority of TestKing2 is higher, it will become the DIS. The DIS election process is described below:

On a LAN, one of the routers elects itself the DIS, based on interface priority (the default is 64). If all interface priorities are the same, the router with the highest subnetwork point of attachment (SNPA) is selected. The SNPA is the MAC address on a LAN, and the local data link connection identifier (DLCI) on a Frame Relay network. If the SNPA is a DLCI and is the same at both sides of a link, the router with the higher system ID becomes the DIS. Every IS-IS router interface is assigned both a L1 priority and a L2 priority in the range of 0 to 127.

The DIS election is preemptive (unlike OSPF). If a new router boots on the LAN with a higher interface priority, the new router becomes the DIS. It purges the old pseudonode LSP and floods a new set of LSPs.

Reference: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094b42.shtml#subfirstthree

QUESTION NO: 13

On a point-to-point IS-IS network connection, the PSNP (Partial Sequence Number Packet) is responsible for which of the following?

- A. Acknowledging LSPs.
- B. Replacing IIH packets.
- C. Establishing adjacencies.
- D. Sending link-state changes.

Answer: A

Explanation:

Partial sequence number PDUs (PSNPs) are used to request an LSP (or LSPs) and acknowledge receipt of an LSP (or LSPs).

Reference: Cisco, Introduction to Intermediate, System-to-Intermediate System Protocol http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm

QUESTION NO: 14

Which of the following metrics can be used by IS-IS when determining the best path to a destination? (Choose four)

- A. default
- B. load
- C. delay
- D. expense
- E. reliability
- F. error
- G. uptime

Answer: A, C, D, F

Explanation:**IS-IS Metrics**

IS-IS uses a single required default metric with a maximum path value of 1024. The metric is arbitrary and typically is assigned by a network administrator. Any single link can have a maximum value of 64, and path links are calculated by summing link values. Maximum metric values were set at these levels to provide the granularity to support various link types while at the same time ensuring that the shortest-path algorithm used for route computation will be reasonably efficient. IS-IS also defines three optional metrics (costs): delay, expense, and error. The delay cost metric reflects the amount of delay on the link. The expense cost metric reflects the communications cost associated with using the link.

The error cost metric reflects the error rate of the link. IS-IS maintains a mapping of these four metrics to the quality of service (QoS) option in the CLNP packet header. IS-IS uses these mappings to compute routes through the network.

QUESTION NO: 15

Which packet type is used to acknowledge LSPs on point-to-point links and to request missing pieces of information in the IS-IS link state database?

- A. complete SNP (CSNP)
- B. partial SNP (PSNP)
- C. hello
- D. database query
- E. database description packet (DDP)

Answer: B

Explanation:

The different IS-IS packet types are described below:


CSNP—Complete Sequence Number PDUs are used to tell other routers about all the LSPs in one router database; similar to an OSPF database descriptor packet

PSNP—Partial Sequence Number PDUs are used to request an LSP and acknowledge receipt of an LSP

PDUs (PSNP's) every 3 seconds and complete sequence number PDUs (CSNPs) every 10 seconds that describe the full database of the DIS. Another router can ask the DIS for a missing LSP using a partial sequence number packet (PSNP), or in turn give the DIS a new LSP.

QUESTION NO: 16

The TestKing network is shown below, along with the relevant information on router TK1:



```
TK1#show clns interface ethernet 0
Ethernet0 is up, line protocol is up
Checksums enabled, MTU 1497, Encapsulation
SAP
Routing Protocol: ISIS
Circuit Type: level-1-2
Interface number 0x0, local circuit ID 0x1
Level-1 Metric: 10, Priority: 64, Circuit ID: TK2.02
Number of active level-1 adjacencies: 1
Level-2 Metric: 10, Priority: 64, Circuit ID: TK2.01
Number of active level-2 adjacencies: 1
Next ISIS LAN Level-1 Hello in 5 seconds
Next ISIS LAN Level-2 Hello in 2 seconds
```

Based on the information above, which statement is correct about IS-IS?

- A. Router TK1 is the DIS for both L1 and L2
- B. Router TK1 is the backup DIS for both L1 and L2.
- C. Router TK2 is the DIS for both L1 and L2.
- D. Router TK1 is the DIS for L1 and a different router (TK2) is the DIS for L2.
- E. Router TK1 is the DIS for L2 and a different router (TK2) is the DIS for L1.

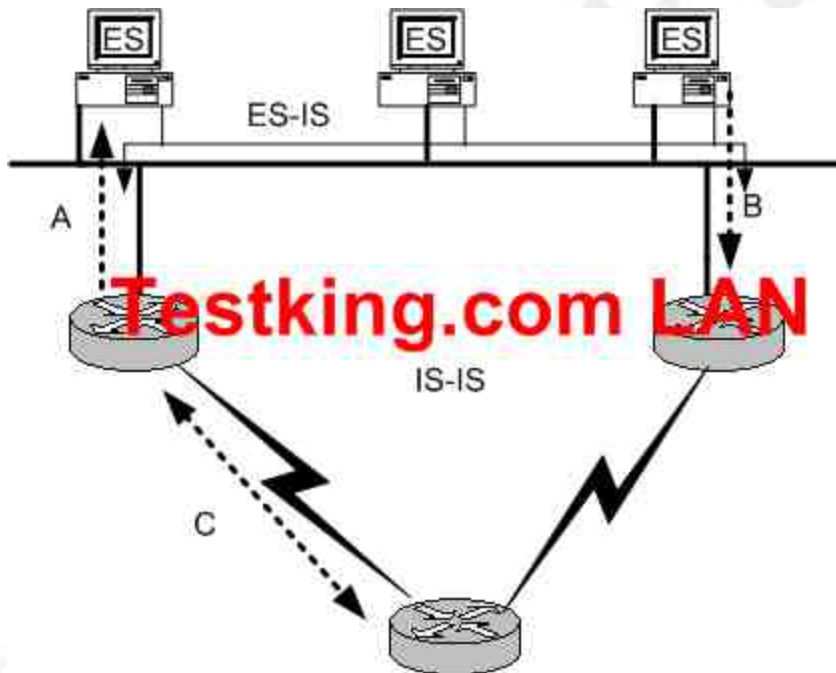
Answer: C

Explanation:

It is the DIS that generates the pseudonode Link State Packet (LSP) and is denoted with a non-zero LSP-ID - R2.01

QUESTION NO: 17

The TestKing IS-IS network is displayed below:



Given the above diagram, which statement correctly matches the letter to the type of IS-IS hello PDU?

- A. A=ISH; B=IIH; C:ESH
- B. A=IIH; B=ESH; C=ISH
- C. A=ISH; B=ESH; C:IIH

- D. A=ESH; B=ISH; C=IIH
E. A=IIH; B=ISH; C=ESH

Answer: C

Explanation:

OSI CLNS is a network layer service similar to bare IP service. A CLNS entity communicates over Connectionless Network Protocol (CLNP) with its peer CLNS entity.

In the OSI architecture there are "systems": Routers are ISs, and hosts are End Systems (ESs). ESs themselves have no routing information; they discover ISs (routers) by listening to Intermediate System Hellos (ISHs) and sending traffic to any random router. ESs send End System Hellos (ESHs); they do not choose a designated router to handle all traffic, and optimal routing is accomplished via redirects.

ISs discover ESs by listening to ESHs, and ISs send ISHs to ESs.

Intermediate System-to-Intermediate System Hello (IIH)—Used by routers to detect neighbors and form adjacencies. In addition to the IIH, which is an IS-IS protocol data unit (PDU), there is an ISH and an ESH, which are End System-to-Intermediate System (ES-IS) PDUs.

QUESTION NO: 18

Drag and Drop

As a TestKing.com network technician you are required to drag the appropriate IS-IS routing them to its correct definition.

Between ES's and IS's on the same subnet. OSI routing begins at this level.

Routing between separate domains. Similar to BGP.

Between IS's on the same area. Also called area routing.

Called inter-area routing.

Drop answer here

Drop answer here

Drop answer here

Drop answer here

Select from these

Level 1/2 routing

Level 2 routing

Level 3 routing

Level 0 routing

Level 1 routing

Answer:

As a TestKing.com network technician you are required to drag the appropriate IS-IS routing them to its correct definition.

Between ES's and IS's on the same subnet. OSI routing begins at this level.	Level 0 routing
Routing between separate domains. Similar to BGP.	Level 3 routing
Between IS's on the same area. Also called area routing.	Level 1 routing
Called inter-area routing.	Level 2 routing

Select from these

Level 1/2 routing

QUESTION NO: 19

What does IS-IS use to establish and maintain neighbor relationships between ISs?

- A. IHH
- B. LSP
- C. CLNS
- D. CLNP
- E. ISH

Answer: A

Explanation:

IHH (IS-IS Hello) – These packets create and maintain neighbor relationships and adjacencies. There are three types of Integrated IS-IS Hello packet. The type of packet is defined in the fixed header under the Type field and allows the packet to be handed off to the appropriate process.

Reference: CCNP BSCI Exam Certification Guide p.370

QUESTION NO: 20

In an IS-IS environment, what happens when the designated IS router crashes?

- A. The elected backup designated router takes the place of the DIS indefinitely without the necessity of a new election.

- B. The elected backup designated router takes the place of the DIS indefinitely without the necessity of a new election, until the original DIS comes back online.
- C. A new election process occurs immediately, establishing a new DIS until a router with a higher priority or MAC address establishes an adjacency.

Answer: C

Explanation:

There is no backup designated router in IS-IS. Therefore, if the DIS meets an untimely death, a new DIS would be elected, based on priority or highest MAC address. If another router comes online with a higher priority, it will dislodge the existing DIS and rule in its place. This behavior is different from that of OSPF. Once a new DIS is elected, the link-state databases are purged and new LSPs are flooded.

Reference: CCNP Self-Study CCNP BSCI Exam certification guide p.806

QUESTION NO: 21

IS-IS Level 1 routers use LSPs for which purpose?

- A. to build topology database for the local area only
- B. to build topology database for the local and remote areas
- C. to build topology database for remote areas
- D. to build topology database for areas outside the AS
- E. to establish with L1 IS's
- F. to establish adjacency with ES's

Answer: A

IS-IS Level 1 routers shares LSAs with L1 routers in the same area.

QUESTION NO: 22

In IS-IS routing, when would the use of Level 1 IS routing be required?

- A. When routing between end systems and intermediate systems.
- B. When routing between Level 1 areas in different domains.
- C. When routing between Level 1 areas in the same domain.
- D. When routing between intermediate systems in the same area.
- E. When routing between intermediate systems in different areas.

Answer: D

Explanation:

Layer 1 links connect the routers in the same area, so layer 1 routing takes place between intermediate systems in the same area. Layer 2 routing takes place between intermediate systems in different areas.

www.testking.com

Section 14: Describe the concept of establishing adjacencies (4 questions)

QUESTION NO: 1

Router TK1 is configured as a Level-1 IS-IS router. Which of the following routers can router TK1 establish an adjacency with? (Select two)

- A. Any Level-1 IS in any area.
- B. Any Level-2 IS in any area.
- C. Any Level-1 IS in the same area.
- D. Any Level-2 IS in the same area.
- E. Any Level-1/Level-2 IS in the same area.

Answer: C, E

Explanation:

A Level-1 IS router can establish adjacencies with other routers Level-1 and Level-1/Level-2 IS routers within the same area.

Incorrect Answers:

- A: Level-1 adjacencies can only be established within the same area.
B, D: Level-2 adjacencies require Level-2 IS routers.

QUESTION NO: 2

Router TK1 is an IS-IS Level-1/Level2 router. Which of the following type of router will TK1 NOT be able to establish an adjacency with?

- A. Any Level-1 IS in any area.
- B. Any Level-2 IS in any area.
- C. Any Level-1 IS in the same area.
- D. Any Level-1/Level-2 IS in any area.
- E. Any Level-1/Level-2 IS in the same area

Answer: A

Explanation:

Level-1 routers can only establish adjacencies with other Level-1 routers within the same area. This is because L1 routers in different areas must connect through a L2 capable router, so two L1 routers must be in the same area in order to become adjacent.

Incorrect Answers:

B, C, D, E: IS-IS has a two-level hierarchy. Contiguous Level 2-capable routers form the backbone. Both Level 2 and Level 1 routers live in areas. Routers can be Level 1 (L1), Level 2 (L2), or both (L1/L2). Within Cisco IOS[®] Software, the default configuration is both Level 1 and Level 2 at the same time which allows an IS-IS network to run with minimal configuration in a plug-and-play fashion. L1/L2 routers act as 2 separate routers in terms of adjacency building, so it will become adjacent with other L2 routers, as well as other L1 routers that are in the same area.

QUESTION NO: 3

What is periodically sent by a DIS on a LAN to ensure that all adjacent neighbors' IS-IS link-state databases are synchronized?

- A. complete SNP (CSNP)
- B. partial SNP (PSNP)
- C. database query
- D. database description packet (DDP)
- E. link-state summary
- F. hello

Answer: A

Explanation:

All ISs maintain adjacencies with all other ISs on a broadcast network. DIS sends CSNP to all ISs. Periodic CSNPs ensure the databases are synchronized.

Reference: CCNP Self-Study CCNP BSCI Exam Certification Guide p.350

QUESTION NO: 4

Which of the following items must match in an OSPF hello packet in order for two routers to become OSPF neighbors? (Choose four)

- A. Neighbors
- B. Hello/dead intervals
- C. Area ID
- D. Address of DR/BDR
- E. Stub area flag
- F. Authentication password

Answer: B, C, E, F

Explanation:

The following hello parameters must match in order for two OSPF routers to become adjacent neighbors:

1. OSPF area number

2. OSPF area type, such as stub or NSSA flag
3. Subnet and subnet mask
4. OSPF HELLO and Dead timer values
5. Authentication values/passwords

www.testking.com

Section 15: Describe the features and operation of BGP (29 questions)

QUESTION NO: 1

Which three of the statements below correctly describe the characteristics of Autonomous Systems in routed networks? (Select three)

- A. Within an AS, all routers must run either BGP or IBGP.
- B. An AS uses exterior gateway protocols (EGPs) to exchange information with other autonomous systems.
- C. An AS is a group of routers under the same technical administration.
- D. Within an AS, routes learned through BGP can be redistributed using interior gateway protocols.
- E. Within an AS, routes learned through an interior protocol cannot be redistributed using BGP to other autonomous systems.

Answer: B, C, D

Explanation:

- **Exterior Gateway Protocol (EGP)** routing protocol used to connect between autonomous systems.
- The use of the term autonomous system in connection with BGP stresses the fact that the administration of an autonomous system appears to other autonomous systems to have a single coherent interior routing plan, and presents a consistent picture of those networks that are reachable through it.
- BGP is used between autonomous systems

Incorrect Answers:


A: Within an AS, all routers need not run either interior BGP or exterior BGP. While all IBGP peers need to be fully meshed, all routers within the network do not need to run BGP.

E: One of the major reasons for running BGP is so that the interior public IP networks within an AS can be distributed to other systems in the Internet via BGP.

Reference: Building Scalable Cisco Networks (Cisco Press) page 313

QUESTION NO: 2

The TestKing network is displayed in the following diagram:



On the assumption that there is no IGP running in AS 64700 and synchronization is OFF, which router in AS 64700 is going to advertise the route to 183.32.0.0?

- A. TK2 only
- B. TK5 only
- C. TK2 and TK5 only
- D. TK2, TK3, and TK4 only
- E. TK2, TK3, TK4, and TK5

Answer: A

Explanation: We need to run IGP on AS 64700 to control the exit points from the AS. When router TK2 receives the EBGP route, it will advertise it to all IBGP routers. However, IBGP routers do not advertise the IBGP learned route unless it is also in their IGP table, or synchronization is turned off.

Reference: Using the Border Gateway Protocol for Interdomain Routing
<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/icsbgp4.htm>

QUESTION NO: 3

Which of the following statements about IBGP routers are true? (Select one.)

- A. They must be fully meshed.
- B. They can be in a different AS.
- C. They must be directly connected.
- D. They do not need to be directly connected.

Answer: D

Explanation:

Since BGP runs over TCP, the IBGP speakers need to only be able to reach the other IBGP speakers via a TCP connection across the network and do not need to be directly connected. Even if two IBGP neighbors are separated by other routers, they will still become neighbors as long as they are reachable via TCP.

Incorrect Answers:

A: Within an AS, all IBGP speakers must be configured as a full mesh. The only exceptions to this rule occur when router reflection or confederations are used.

B: IBGP routers need to be configured with the same AS number. The only exception to this rule is when BGP confederations are in use.

C: The routers only need to be reachable via a TCP connection.

QUESTION NO: 4

Given the following choices, what kind of BGP router advertises routes to other IBGP neighbors?

- A. Client
- B. EBGP peer
- C. Route reflector
- D. cluster of clients
- E. None of the above

Answer: C

Explanation:

Ordinarily, with no route reflector, a full mesh of IBGP peers is required. Route reflectors modify the BGP split horizon rule by allowing the router configured as the route reflector to propagate routes learned by IBGP to other IBGP peers. Route reflectors reduce the number of BGP neighbor relationships in an AS.

QUESTION NO: 5

BGP has been configured on some of the TestKing routers. The BGP routing process relies on two different types of tables. What are they? (Select two)

- A. An IP routing table.
- B. A BGP topology table.
- C. A BGP attribute table.
- D. A table that contains BGP information received from and sent to other routers.
- E. A combined table that contains both IP routes and BGP information received from and sent to other routers.

Answer: B, C

Explanation:

BGP uses one table for topology and another for attribute. The attribute table assigns values to various attributes (weight, local preference, multi-exit discriminator, origin, AS path, next hop, & community) and the topology table matches the values of these attributes to the various routes it can select.

QUESTION NO: 6

Attributes are a means of sending additional route information over BGP. Which of the following statements are true regarding BGP attributes? (Select three)

- A. MED is an optional attribute.
- B. Origin is an optional attribute.
- C. Next-hop is an optional attribute.
- D. Local Preference is a discretionary attribute.
- E. AS-Path is a well-known mandatory attribute.
- F. Community is a well-known mandatory attribute.

Answer: A, D, E

Explanation:

A: The Multiple Exit Discriminator (MED) attribute is optional.

D: The local preference attribute is used to prefer an exit point from the local autonomous system (AS). It is a well-known discretionary attribute.

E: 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. The AS-Path attribute is mandatory.

Note: BGP metrics are called path attributes. Optional attributes are recognized by some implementations, but are expected not to be recognized by everyone. A well-known mandatory attributes must be present in all update messages.

The attributes defined by BGP include:

Well-known mandatory attributes:

- AS-path
- Next-hop
- Origin

Well-known discretionary attributes:

- Local preference
- Atomic aggregate

Optional transitive attributes:

- Aggregator
- Communities

Optional non-transitive attribute:

- Multi-Exit-Discriminator (MED)

Incorrect Answers:

- B: The origin attribute indicates how BGP learned about a particular route. It can have three values: IGP, EGP or incomplete. It is a required attribute.
- C: The EBGP next-hop attribute is the IP address that is used to reach the advertising router. It is a required attribute.
- F: The community attribute provides a way of grouping destinations, called communities, to which routing decisions can be applied. It is not mandatory.

Reference: Border Gateway Protocol

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

QUESTION NO: 7

The BGP community attribute tags route for the sake of consistency in route filtering. Which of the answer choices below correctly describe the community attribute?

- A. Optional and transitive.
- B. Optional and non-transitive.
- C. Well-known and mandatory.
- D. Well-known and discretionary.

Answer: A

Explanation:

The community attribute is an optional transitive attribute that can be in the range 0 to 4,294,967,200. Each network can be a member of more than one community.

The attributes defined by BGP include:

Well-known mandatory attributes:

- AS-path
- Next-hop
- Origin

Well-known discretionary attributes:

- Local preference
- Atomic aggregate

Optional transitive attributes:

- Aggregator
- Communities

Optional non-transitive attribute:

- Multi-Exit-Discriminator (MED)

QUESTION NO: 8

Leading the way in IT testing and certification tools, www.testking.com

You are a guest lecturer at the TestKing Academy teaching some CCNP hopefuls about BGP. One of the students, Bob asks you what a BGP community is. How would you respond?

- A. Communities are tagged by default in outgoing updates.
- B. Communities can only be used within one autonomous system.
- C. Communities are a means of tagging routes to ensure consistent filtering.
- D. Communities perform summarization of blocks of contiguous network prefixes.

Answer: C

Explanation:

A community is a group of destinations which share some common property. No tag is used by default. Communities are a means of tagging routes to ensure consistent filtering or route-selection policy.

Incorrect Answers

- A: By default, all destinations belong to the general Internet community and are not tagged.
- B: Each autonomous system administrator may define which communities a destination belongs to. Community information is passed on between different autonomous systems.
- D: No summarization is performed by communities. Communications can be aggregated, however.

Reference: RFC 1997, BGP Communities Attribute

QUESTION NO: 9

The TestKing IBGP routers are configured in a fully meshed fashion. Which of the following are reasons why IBGP routers should peer with each other within a single AS? (Select two)

- A. IBGP routes are not propagated to other EBGP peers.
- B. IBGP routes that a router originates are propagated to other IBGP peers.
- C. IBGP routes are propagated to other IBGP speakers in the AS that are not peers.
- D. IBGP routes that are learned from an IBGP neighbor are propagated to only EBGP peers.

Answer: A, B

Explanation:

IBGP routes are propagated to all IBGP peers and only the IBGP peers. They are not exchanged with EBGP peers, which is why all IBGP routers must be configured in a full mesh.

Note: You can configure Border Gateway Protocol (BGP) either within an autonomous system or between different autonomous systems. When run within an autonomous system, it's called internal BGP (IBGP). When run between different autonomous systems, it's called external BGP (EBGP).

QUESTION NO: 10

Routers TK1 and TK2 belong to the same BGP peer group. Which two of the statements below are true regarding this BGP peer group? (Select two)

- A. The peer group name is passed to other routers in the peer group.
- B. A peer group is a group of BGP neighbors with different update policies.
- C. The peer group name is only local to the router on which it is configured.
- D. A peer group allows options that affect outbound updates to be overridden.
- E. A peer group is a more efficient way to update BGP than configuration individual neighbors.

Answer: C, E

Explanation:

C: The peer group name is only local to the router it is configured on, it is not passed to any other router.

E: Neighbors with the same update policies can be grouped into peer groups to simplify configuration and make update calculation more efficient.

Incorrect Answers:

A: A BGP peer group does not necessarily have a name, it might have a number.

B: BGP peers use the same update policy.

D: A peer group does not have this option.

Reference: RFC 1771, A Border Gateway Protocol 4 (BGP-4)

QUESTION NO: 11

Routers TK1 and TK2 are configured as BGP peers. Which of the following is true regarding this peering relationship? (Select two)

- A. Periodic keepalives are used to verify connectivity.
- B. Incremental keepalives are used to verify connectivity.
- C. It provides a reliable connection between two BGP routers.
- D. It provides a “best effort” connection between two BGP routers.

Answer: A, C

Explanation:

A: The default keepalive frequency is 60 seconds.

C: BGP peering provides a reliable connection between BGP routers. BGP utilizes the inherent reliability of TCP, since all BGP sessions are maintained over TCP port 179.

Incorrect Answers:

B: Keepalives messages are not incremental.

D: BGP peers form a reliable connection.

Reference: BGP Peer Groups

<http://www.cisco.com/warp/public/459/29.html>

QUESTION NO: 12

Router TK1 is configured as a route reflector. What would happen if router TK1 would receive an update from a peer in a different autonomous system?

- A. It discards the update.
- B. It sends the update to all IBGP peers.
- C. It sends the update only to non clients.
- D. It sends the update only to route reflector clients.
- E. It sends the update to all routers in the autonomous system.

Answer: B

Explanation:

When any BGP router receives an update from an EBGP peer, it will forward this routing information to all IBGP peers. This is true for all EBGP routers, including route reflectors.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_configuration_guide_chapter09186a00800ca571.html#5155

QUESTION NO: 13

Which of the following statements is true about BGP peer communications?

- A. Communication between BGP peers runs over RIP.
- B. Communication between BGP peers runs over TCP.
- C. Communication between BGP peers runs over UDP.
- D. Communication between BGP peers runs over ICMP.

Answer: B

Explanation:

BGP communicate through a TCP connection. TCP port 179 is reserved for BGP sessions.

Reference: Border Gateway Protocol

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

QUESTION NO: 14

The TestKing network is considered to be a single autonomous system (AS). Which three of the following statements are characteristics of an autonomous system? (Select three)

- A. It uses only Interior Gateway Protocols (IGPs).
- B. EGPs are used to connect different autonomous systems.
- C. It is a set of routers under a single technical administration.
- D. It uses EGPs to route packets to other autonomous systems and IGPs to route packets within the autonomous system.
- E. It uses IGPs to route packets to other autonomous systems and EGPs to route packets within the autonomous system.

Answer: B, C, D

Explanation:

A BGP autonomous system is a group of routers all administered as a single team working on mutual packet switching goals. EGP's (Exterior Gateway Protocol's) handle routing to outside autonomous systems which are on the exterior of that particular autonomous system, and IGP's (Interior Gateway Protocol's) route within, on the interior of autonomous systems.

Incorrect Answers:

- A: This is incorrect because an autonomous system can use an exterior gateway protocol.
- E: This is incorrect because the roles of IGP's and EGP's are reversed in this answer choice.

QUESTION NO: 15

Which of the following statements is true about the BGP synchronization command? (Select two)

- A. Synchronization must be enabled when implementing a multi-homed BGP connection to multiple ISPs.
- B. If it is turned ON, a prefix learned from IBGP neighbor is valid only if a non-BGP (IGP) route exists for that prefix.
- C. Synchronization is necessary when peering with an EBGP neighbor.
- D. Synchronization improves BGP routing convergence.
- E. Synchronization can be turned off if all the transit routers in an Autonomous system are running a fully meshed IBGP.

Answer: B, E

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. Only if all routers in the transit path in the AS are running BGP it is safe to turn synchronization off.

Reference: Building Scalable Cisco Networks (Cisco Press) page 33

http://www.cisco.com/en/US/products/sw/iosswrel/ps1826/products_configuration_guide_chapter09186a00800877b5.html

QUESTION NO: 16

The TestKing BGP network is shown in the diagram below:



Router TK1 receives BGP routing advertisements from router TK2 about the network 165.50.0.0. By default, what is the value of the next hop attribute?

- A. 40.1.1.1
- B. 40.1.1.2
- C. 40.1.1.3
- D. Router TK1 does not accept the advertisement from Router TK2 because Router TK1 is not peering with Router TK3 via BGP.
- E. Router TK2 does not advertise network 165.50.0.0 to Router TK1 because the network is not directly connected to Router TK2.

Answer: A

Explanation:

The BGP next-hop attribute is a well-known mandatory attribute that indicates the next hop IP address that is to be used to reach a destination. For EBGP, the next hop is the IP address of the neighbor specified who sent the update, Router TK2 in this scenario.

Incorrect Answers:

B: Router TK2 learned this route through IBGP with the next-hop of 40.1.1.2. This value will not be used instead of 40.1.1.1 because the next hop information is not preserved across AS 64300 when being sent to a different AS.

C: This is the local interface on Router TK1 itself.

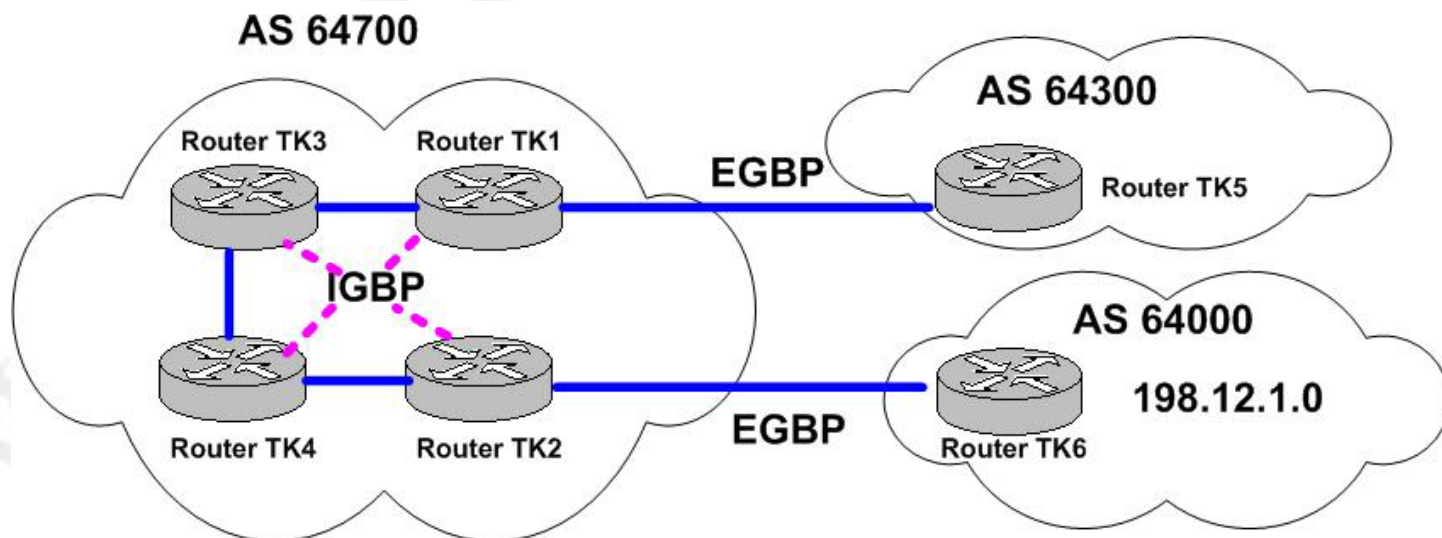
D: The advertisement is accepted.

E: The route is advertised.

Reference: Reference: Cisco Press's BSCI 2nd Edition pg 529

QUESTION NO: 17

The TestKing network is shown in the following exhibit:



- all routers are using BGP

- synchronization is OFF in AS 647000

Regarding the route to 198.12.1.0; which routers will this route propagate to?

- Router TK6
- Routers TK2 and TK6
- Routers TK2, TK4, and TK6
- Routers TK2, TK3, TK4, and TK6
- Routers TK1, TK2, TK3, TK4, and TK6
- Routers TK1, TK2, TK3, TK4, TK5, and TK6
- It will not be propagated to any other router


Answer: F

Explanation:

Router TK 6 will advertise the 192.12.1.0 route to Router TK2 using EBGP. This route would propagate to Router TK1, TK3, and TK4 through the fully meshed IGBP sessions. Since synchronization is turned off Router TK1 will advertise the internally learned route to Router TK5 using EBGP.

QUESTION NO: 18

The TestKing BGP network is shown below:



- synchronization is OFF
- NO IGP is running in AS 65500

Which router in AS 65500 is going to end up advertising a route to network 172.16.0.0?

Leading the way in IT testing and certification tools, www.testking.com

- A. B only
- B. E only
- C. B and E only
- D. B, C and D only
- E. B, C, D, and E

Answer: C

Explanation:

Router B learns about the 172.16.0.0 network route via the EBGP session from router A. With synchronization turned off, it will pass this route to all of its fully meshed IBGP peer routers. Although all routers will receive the route, only routers B and E will advertise this route to other routers. Router B will advertise the route to all of the fully meshed IBGP peers, and router E will advertise it to AS 64500.

Reference: Using the Border Gateway Protocol for Interdomain Routing
<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/icsbgp4.htm>

QUESTION NO: 19

The TestKing network administrator is planning to use a private AS number. Which of the following are choices for a private AS number?

- A. 10080
- B. 48512
- C. 64128
- D. 64524

Answer: D

Explanation:

This autonomous system designator is a 16-bit number, with a range of 1 to 65535. RFC 1930 provides guidelines for the use of AS numbers. A range of AS number, 64512 through 65535, is reserved for private use, much like the private Internet Protocol (IP) addresses defined in RFC 1918.

Reference: Building Scalable Cisco Networks (Cisco Press) page 312

QUESTION NO: 20

The router that is being used for the Internet is low on memory. Because multiple ISP's are being used, BGP is needed for Internet routing. Which of the following types of routes should be received from the Internet providers if the goal is to consume as little memory as possible?

- A. Only default routes.
- B. Only partial BGP routes.
- C. Only internal BGP routes.
- D. Only redistributed routes.

Answer: A

Explanation:

Today, the Internet consists of over 100,000 routes, so accepting the full routing table from an ISP can consume a great deal of router memory. By only accepting default routes from the ISPs ensures a minimal impact on the system. There is no risk of having your autonomous system (AS) becoming a transit AS. By receiving only default routes, the entire Internet will be reachable while only using a small amount of memory for the routing table.

Note: One recommendation of Cisco, not listed here, is to use AS_PATH filters for accepting only routes originated from an ISP and its directly connected Autonomous Systems (AS), instead of receiving the full BGP routing table from an ISP.

Reference: Achieve Optimal Routing and Reduce BGP Memory Consumption
<http://www.cisco.com/warp/public/459/41.shtml>

QUESTION NO: 21

Routers TK1, TK2, and TK3 are all IBGP peers within the TestKing network. Why should these peer sessions be fully meshed within the AS?

- A. Because BGP requires redundant TCP sessions between iBGP peers.
- B. Because a full mesh allows for optimal routing within the Transit AS.
- C. Because routes learned via iBGP are never propagated to other eBGP peers.
- D. Because routes learned via iBGP are never propagated to other iBGP peers.
- E. Because routes learned via eBGP are never propagated to other iBGP peers.

Answer: D

Since BGP does not relay routing traffic from one Interior BGP session to another (only from an Exterior BGP session to an IBGP session or another EBGP session), BGP speakers must be fully meshed. When a BGP routing update is received from a neighboring AS, it must be relayed directly to all other BGP speakers in the AS. Do not expect to relay BGP paths from one router, through another, to a third, all within the same AS.

Incorrect Answers:

A: BGP only requires a single TCP session to each peer.

B: The full mesh is needed so that all routers receive the routing information, but it does not provide for any level of route optimization.

C, E: Routes from EBGP peers are propagated to all IBGP peers, and routes learned via IBGP are sent to EBGP peers by default.

QUESTION NO: 22

Which BGP attribute is used by BGP to prevent routing loops?

- A. AS-path
- B. Next-hop
- C. MED
- D. Weight
- E. Local preference
- F. Origin

Answer: A

Explanation:

BGP uses path information to determine routing loops throughout the network. A simplified view of AS_PATH is that it is the list of Autonomous Systems that a route goes through to reach its destination. Loops are detected and avoided by checking for your own AS number in AS_PATH's received from neighboring Autonomous Systems. If your AS number appears in the AS path the second time, then a loop has occurred.

QUESTION NO: 23

When the BGP path selection process is being performed on a Cisco router, which BGP attribute is used first when determining the best path?

- A. Local preference
- B. MED
- C. Weight
- D. Origin
- E. Next-hop
- F. AS-path

Answer: C

Explanation:

How the Best Path Algorithm Works:

BGP assigns the first valid path as the current best path. It then compares the best path with the next path in list, until it reaches the end of the list of valid paths. The following is a list of rules used to determine the best path.

1. Prefer the path with the highest WEIGHT.

Note: WEIGHT is a Cisco-specific parameter, local to the router on which it's configured.

2. Prefer the path with the highest LOCAL_PREF. Note the following:

- Path without LOCAL_PREF is considered as having the value set with the bgp default local-preference command, or 100 by default.
- 3. Prefer the path that was locally originated via a network or aggregate BGP subcommand, or through redistribution from an IGP. Local paths sourced by network or redistribute commands are preferred over local aggregates sourced by the aggregate-address command.
- 4. Prefer the path with the shortest AS_PATH. Note the following:
 - The AS_CONFED_SEQUENCE and AS_CONFED_SET are not included in the AS_PATH length.
- 5. Prefer the path with the lowest origin type: IGP is lower than EGP, and EGP is lower than INCOMPLETE.
- 6. Prefer the path with the lowest multi-exit discriminator (MED).
- 7. Prefer external (eBGP) over internal (iBGP) paths. If bestpath is selected, go to Step 9 (multipath).
- 8. Prefer the path with the lowest IGP metric to the BGP next hop. Continue, even if bestpath is already selected.
- 9. Check if multiple paths need to be installed in the routing table for BGP Multipath. Continue, if bestpath is not selected yet.
- 10. When both paths are external, prefer the path that was received first (the oldest one). This step minimizes route-flap, since a newer path will not displace an older one, even if it would be the preferred route based on the next decision criteria (Steps 11, 12, and 13).
- 11. Prefer the route coming from the BGP router with the lowest router ID. The router ID is the highest IP address on the router, with preference given to loopback addresses. It can also be set manually using the bgp router-id command.
- 12. If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length. This will only be present in BGP route-reflector environments. It allows clients to peer with RRs or clients in other clusters. In this scenario, the client must be aware of the RR-specific BGP attribute.
- 13. Prefer the path coming from the lowest neighbor address. This is the IP address used in the BGP neighbor configuration, and corresponds to the remote peer used in the TCP connection with the local router.

QUESTION NO: 24

Which one of the following statements about BGP is FALSE?

- A. BGP uses TCP port 179.
- B. BGP ensures reliability of updates by using the reliable transport services of TCP.
- C. The network command with the mask option never installs a prefix into the BGP table unless there is a matching prefix exists in the IP route table.
- D. A TCP connection is required before exchanging updates.
- E. BGP uses notification and the update messages to establish and maintain the BGP neighbor relationship.

Answer: E

Explanation:

BGP relies on TCP to provide for a reliable connection between peer routers. The statement in E is false, as the notification message is used to terminate the BGP session, not to maintain or establish the peering session. BGP systems send notification messages when an error condition is detected. After the message is sent, the BGP session and the TCP connection between the BGP systems are closed. Notification messages consist of the BGP header plus the error code and subcode, and data that describes the error.

Incorrect Answers:

A: This statement is true.

B: BGP uses TCP port 179. TCP is a reliable, connection oriented protocol.

C: This statement is true. The route must actually exist in the routing table before it will be advertised using the network BGP configuration command.

D: These are both true statements regarding BGP. BGP relies on TCP to maintain sessions.

QUESTION NO: 25

Which BGP state is a router in when it has found the IP address in the neighbor statement and has created and sent out a BGP packet but the router has not received a response back?

- A. Update
- B. Connect
- C. Active
- D. Established

Answer: B

Explanation:

The following table describes the various states that a BGP router can be in:

BGP State Machine

BGP goes through six states to establish an adjacency.

Idle - incoming connections are refused, and the system gets ready to start speaking BGP. After this is done (by way of a Start event), move to Connect.

Connect - a connection is made to the peer. Send a BGP OPEN message, and go to OpenSent.

Active - a connection comes in from a peer. Send a BGP OPEN message, and go to OpenSent.

OpenSent - Wait for an OPEN message from the peer. When received, send a KEEPALIVE and go to OpenConfirm.

OpenConfirm - Wait for the KEEPALIVE from the peer, then move to Established.

Established - Bidirectional communication is established. Start sending UPDATE and KEEPALIVE messages as required

QUESTION NO: 26

Which BGP attribute will not be advertised in routing updates to its neighboring routers?

- A. weight
- B. local preference
- C. origin
- D. AS_path
- E. next hop

Answer: A

Explanation:

Weight is proprietary to Cisco and is used in route selection. It is local to the router, and because it is not propagated to other routers, there is no problem with compatibility. When there are multiple paths, it selects a path to a destination with different next hops to the same destination. Note that the weight attribute has no code. Because it is a local attribute and is not propagated to other routers, no code is needed.

Reference: Cisco Press 642-801 p.527

QUESTION NO: 27

Place the BGP attributes in the correct order used for determining a route.

originate route	1st
AS_Path	2nd
weight	3rd
local preference	4th
MED	5th

Answer:

**QUESTION NO: 28**

Which attribute must exist in the BGP update packet?

- A. LOCAL_PREF
- B. AGGREGATOR
- C. AS_Path
- D. Weight

Answer: C

Explanation:

C) AS PATH is a well known mandatory attribute. Mandatory means that it is required by all routers. These attributes are required and are therefore recognized by all BGP implementations.

QUESTION NO: 29

Which three conditions can cause BGP neighbor establishment to fail? Select three.

- A. There is an access list blocking all TCP traffic between the two BGP neighbors.
- B. The EBGP neighbor is not directly connected, and the **ebgp-multihop** option is set to the default value.
- C. The IBGP neighbor is not directly connected.
- D. BGP synchronization is enabled in a Transit AS with fully-meshed IBGP neighbors.
- E. The BGP update interval is different between the two BGP neighbors.
- F. The BGP neighbor is referencing an incorrect AS number in its **neighbor** statement.

Answer: A, B, F

Section 16: Explain how BGP policy-based routing functions within an autonomous system (5 questions)

QUESTION NO: 1

The TestKing network administrator is considering the use of policy based routing. What would happen if BGP policy-based routing was implemented and the next hop router goes down, leaving no alternative path?

- A. If the next-hop router goes down and no alternative path is in place, policy routing will route to null 0.
- B. If the next-hop router goes down and no alternative path is in place, policy routing will default to another BGP path.
- C. If the next-hop router goes down and no alternative path is in place, policy routing will deny all traffic to that destination.
- D. If the next-hop router goes down and no alternative path is in place, policy routing will default to dynamic routing decisions.
- E. None of the above.

Answer: D

Explanation:

A backup path should be in place in case the defined next-hop router goes down. If there is no alternative defined, policy routing will default to dynamic routing decisions.

With any policy based routing implementation, when a packet is unable to match any of the conditions or if the next hop is not unreachable, the default routing behavior will occur.

Reference: CCNP Self-Study CCNP BSCI Exam Certification Guide by Clare Gough, page 499.

QUESTION NO: 2

BGP Policy Based routing has been configured on router TK1. Which of the following statements are true regarding policy-based routing on this router? (Select two)

- A. Policy routing can be configured on the inbound interface.
- B. Policy routing can be used to alter the final destination of the packet.
- C. Policy routing can be used to alter the next hop in the path to the destination.
- D. Policy routing does not allow traffic to be directed based on the source address.

Answer: A, C

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

Rules of Policy-Based Routing:

BGP can implement any of the following rules associated with the hop-by-hop paradigm. This paradigm is the capability to influence which router will be the next-hop router, potentially dictating it at every router and thus influencing the entire path of the traffic, hop by hop. The following rules associated with policy-based routing seem repetitive, but, in fact, each point raises a subtly different nuance:

- Traffic can be directed on either the source address or both the source and destination addresses, as well as inbound or outbound traffic on an interface. These are only some of the criteria that can be used.
 - Policy-based routing affects only the next hop in the path to the destination.
 - Policy-based routing does not affect the destination of the packet. It affects the path used to get to the destination.
 - Policy-based routing does not allow traffic sent into another autonomous system to take a different path from the one that would have been chosen by that autonomous system.
 - It is possible to influence only how traffic will get to a neighboring autonomous system, not how it will be routed within that autonomous system.
 - As policy-based routing examines the source address, it is configured on the inbound interface.
- Reference: Self-Study CCNP BSCI Exam Certification Guide Third Edition P.523

QUESTION NO: 3

Which one of the following BGP prefix lists would you use if you wanted to deny the default route 0.0.0.0?

- A. `ip prefix-list abc deny 0.0.0.0/0`
- B. `ip prefix-list abc permit 0.0.0.0/32`
- C. `ip prefix-list abc deny 255.255.255.255/0`
- D. `ip prefix-list abc permit 255.255.255.255/32`

Answer: A

Explanation:

To deny the default route 0.0.0.0/0 use the following command:

```
ip prefix-list abc deny 0.0.0.0/0
```

Incorrect Answers:

B, D: We are required to deny it, not permit it.

C: This will effectively deny all routes, not just the default route.

QUESTION NO: 4

Which of the following prefix-list statements would you enter if you wanted to permit all prefixes that fall between /8 and /24 on the 192.0.0.0 network?

- A. ip prefix-list abc permit 192.0.0.0/8 ge 8 le 24
- B. ip prefix-list abc permit 192.0.0.0/8 ge 24 le 8
- C. ip prefix-list abc permit 192.0.0.0/24 ge 24 le 8
- D. ip prefix-list abc permit 192.0.0.0/24 ge 8 le 24

Answer: A

Explanation:

When used with prefix lists, the “ge” means “greater than or equal to” and the term “le” means “less than or equal to.” Greater than (ge) 8 and less than (le) 24.

An 8 bit subnet mask should also be used, since we are looking to match all prefixes beginning on the entire 192.X.X.X range.

QUESTION NO: 5

The TestKing IBGP network is displayed below:



AS 65100

It is desired to configure IBGP in the best, most stable manner. Which of the following options show the best configuration choice for router TestKing1?

- A. router bgp 65100
neighbor 2.2.2.2 remote-as 65100
neighbor 2.2.2.2 update-source loopback0
- B. router bgp 65100
neighbor 10.1.1.2 remote-as 65100
neighbor 10.2.2.2 remote-as 65100
neighbor 10.1.1.2 update-source loopback0
neighbor 10.2.2.2 update-source loopback0
- C. router bgp 65100
neighbor 10.1.1.2 remote-as 65100
neighbor 10.2.2.2 remote-as 65100

```
neighbor 2.2.2.2 remote-as 65100
```

- D. router bgp 65100
neighbor 10.1.1.2 remote-as 65100
neighbor 10.2.2.2 remote-as 65100
neighbor 10.1.1.2 update-source s0
neighbor 10.2.2.2 update-source s1

Answer: A

Explanation:

Option A displays the best choice, and shows the correct syntax for the use of sourcing BGP updates from the loopback interface. Using the loopback interface provides the best option, since this interface will remain up and reachable as long as one of the serial interfaces remains active.

Note: Configuring the routers in this manner can be used for both IBGP and EBGP, and can also be used for load balancing traffic across the two serial links

Section 17: Explain the use of redistribution between BGP and Interior Gateway Protocols (IGPs) (9 questions)**QUESTION NO: 1**

When you're running an Interior Gateway Protocol (like OSPF or EIGRP); why would you want to redistribute these routes into BGP?

- A. So BGP can propagate this information to other IGP neighbors.
- B. So BGP can propagate this information to other IBGP neighbors.
- C. So BGP can propagate this information to other EBGP neighbors.
- D. So BGP can propagate this information to other OSPF neighbors.

Answer: C

Explanation:

When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). In order to advertise the IP subnets used within your network to other autonomous systems, these routes must be injected into BGP.

Note: Customer networks usually employ an Interior Gateway Protocol (IGP) such as RIP or OSPF for the exchange of routing information within their networks

Reference: Border Gateway Protocol

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

QUESTION NO: 2

You are the administrator of an internal network configured for OSPF. Two gateway routers are learning external routes from BGP and inserting those learned routes into your network via redistribution. One of the gateway routers is called TestKing2 and it is inserting the subnets 128.213.64.0 through 128.213.95.0. Your task is to configure the gateway router Testking1 to summarize the subnets into one range before injecting them into OSPF. The first two lines have already been configured. Select the boxes below and place the three correct boxes in the right order to complete the configuration.

TestKing1(config)#router ospf 100

TestKing1(config-router)#redistribute bgp 50 metric 1000 subnets

TestKing1(config-router)#

Drop keyword here

Drop keyword here

Drop keyword here

Select from these

128.213.64.0

summary-
address

128.213.0.0

255.255.224.0

network

auto-summary

0.0.31.255

Answer: Summary-address 128.213.64.0 255.255.224.0**Explanation:**

External route summarization, external routes that are injected into OSPF via redistribution, is done via the following **router ospf** subcommand:

summary-address *ip-address mask*

This command is effective only on ASBRs doing redistribution into OSPF as in this scenario.

QUESTION NO: 3

Your network is running BGP as well as EIGRP. You are considering redistributing your BGP routes into your EIGRP. What factors must you consider before redistributing BGP routes to IGP?

- A. IGP's are limited to 250 routes
- B. A full BGP routing table may contain 100,000+ routes
- C. Because of possible routing loops, Cisco router configuration does not allow BGP routes to be redistributed into an IGP.
- D. Because BGP routes are not advertised unless they are known by the IGP, Cisco automatically redistributes routes into BGP.

Answer: B**Explanation:**

The size of the BGP routing table must be taken into consideration. It is generally not recommended to redistribute BGP routes into your IGP due to the large number of routes that can be obtained from BGP.

Incorrect Answers:

- A: There is no such limit.

- C: BGP routes can very well be distributed into an IGP.
- D: BGP routers are not redistributed automatically.

QUESTION NO: 4

Why is it sometimes necessary to redistribute IGP protocols like OSPF into BGP?

- A. So that BGP can propagate this information to other IGP neighbors.
- B. So that BGP can propagate this information to other IBGP neighbors.
- C. So that BGP can propagate this information to other EBGP neighbors.
- D. So that BGP can propagate this information to other OSPF neighbors.

Answer: C

Explanation:

Customer networks usually employ an Interior Gateway Protocol (IGP) such as RIP or OSPF for the exchange of routing information within their networks. In order to advertise these networks to other network across the Internet, they must be redistributed into BGP.

Reference: Border Gateway Protocol

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

QUESTION NO: 5

According to Cisco's official recommendations, what is the preferred method for populating the BGP table on your router with IGP routes?

- A. Use the `network` command.
- B. Redistribute EBGP routes into BGP.
- C. Redistribute dynamic routes into BGP.
- D. Redistribute static routes into the IGP.

Answer: A

Explanation:

For BGP the **network** command allows BGP to advertise a network that is already in the IP table.

Note: An alternate solution, not listed here, is to redistribute the IGP routes into BGP.

Incorrect Answers:

- B: We want to distribute internal routes into BGP, not external routes.
- C: Redistribution from the IGP is NOT advised because there is a great reliance on the IGP table. It could cause instability.
- D: We want to populate the BGP table, not the IGP table.

QUESTION NO: 6

You are the midst of configuring a gateway route named RouterTestK. Your task is to configure the router to redistribute the various subnets it learns via BGP AS75 into OSPF, all while using the default metric. Select three of the correct command phrases below and place them in the correct order in the command sequence.

RouterTestK(config-router)#

Select from these

bgp 75	ospf	subnets	200
default-metric	metric	area	redistribute

Answer:

RouterTestK(config-router)#

Select from these

	ospf	200
default-metric	metric	area

Explanation: We want to redistribute BGP AS 75 into OSPF.
Syntax of the redistribution command:

```
router(config-router)#redistribute protocol [process-id] [metric [metric-type type-value]
[route-map map-tag] [subnets] [tag tagvalue]
```

We specify the BGP protocol and the Autonomous System number 75. We also use the **subnets** keyword, which is an optional OSPF parameter that specifies that subnetted routes should also be redistributed. Only routes that are not subnetted are redistributed if the **subnets** keyword is not specified.

Incorrect Answers:

Leading the way in IT testing and certification tools, www.testking.com

ospf: We want to redistribute the routes learned from BGP AS 75 into OSPF so we should specify the BGP protocol.

area: Is not used in the redistribute command.

default-metric: Default-metric is a separate command.

metric, 200: We don't have to specify a metric since we want to use a default metric.

QUESTION NO: 7

You are the administrator of the TestKing network which is configured for OSPF. Two gateway routers are learning external routes from BGP and inserting those learned routes into your network via redistribution. One of the gateway routers is called TestKing1 and it is inserting the subnets 142.14.64.0 through 142.14.95.0. Your task is to configure router Testking1 to summarize the subnets into one range before injecting them into OSPF. The first two lines have already been configured. Select the boxes below and place the three correct boxes in the right order to complete the configuration.

TestKing1(config)#router ospf 200

TestKing1(config-router)#redistribute bgp 80 metric 1100 subnets

RouterTestKing1(config-router)#

Select from these

auto-summary	summary-address	network
142.14.0.0	255.255.224.0	0.0.31.255
	142.14.64.0	

Answer:

RouterTestKing1(config-router)#

Select from these

auto-summary	network
142.14.0.0	0.0.31.255

Explanation:

The **summary-address** router configuration command is used to create aggregate addresses for OSPF.

Simplified syntax: **summary-address** *address mask*

We study the subnets that the TestKing1 router is injecting in binary to decide where to summarize:

Decimal	1 st octet	2 nd octet	3 rd octet	4 th octet
142.14.64.0	10001110	00001110	01000000	00000000
142.14.95.0	10001110	00001110	01011111	00000000

255.255.224.0	11111111	11111111	11100000	00000000
---------------	----------	----------	----------	----------

The subnet we can summarize on is marked with red. This is the 142.14.64.0/19 network. It has a 255.255.224.0 subnet mask.

QUESTION NO: 8

You are the administrator of an internal network configured for OSPF. Two gateway routers (TestKing1 & TestKing2) are redistributing external BGP routes. TestKing1 is redistributing subnets ranging from 163.150.64.0 – 163.160.95.0 and TestKing2 is doing the same with 163.150.96.0 – 163.150.127.0. Your task is to configure the gateway router Testking1 to summarize these subnets into one range before redistributing them into OSPF. The first two lines have already been configured. Select the boxes below and place the three correct boxes in the right order to complete the configuration.

Exhibit:

```
TestKing1(config)#router ospf 100
```

```
TestKing1(config-router)#redistribute bgp 70 metric 500 subnets
```

TestKing1(config-router)#

Drop keyword here

Drop keyword here

Drop keyword here

Select from these

163.150.0.0

255.255.224.0

0.0.31.256

163.150.64.0

network


summary-
address

auto-summary

Answer:

TestKing1(config-router)#	summary-address	163.150.64.0	255.255.224.0
---------------------------	-----------------	--------------	---------------

Select from these



Explanation:

External route summarization, external routes that are injected into OSPF via redistribution, is done via the following router ospf subcommand:

summary-address *ip-address mask*

This command is effective only on ASBRs doing redistribution into OSPF as in this scenario.. We should summarize most matching leftmost bits as indicated below.

Decimal	1 st Octet	2 nd Octet	3 rd Octet	4 th Octet	
163.50.64.0	10100011	00110010	01000000	00000000	
163.50.95.0	10100011	00110010	01011111	00000000	
163.50.64.0	10000000	00110010	01000000	00000000	IP address
255.255.224.0	11111111	11111111	11000000	00000000	Subnet mask

Incorrect Answers

Network:

The **network** command is not used to summarize routes.

auto-summary:

The auto-summary command is used to turn auto summarization on.

128.213.0.0:

We should use a more specific summarization.

0.0.31.256:

For summarization we should specify a network mask, not a wildcard.

QUESTION NO: 9

It isn't always in your best interests to redistribute dynamically learned routes from an interior gateway protocol to a border gateway protocol. Which of the following are potential issues that can occur when doing so? (Select two)

- A. Routing loops can occur.
- B. The routes are automatically summarized.
- C. External IGP learned routes might not necessarily have originated in this AS.
- D. The BGP process will ignore the external IGP learned routes.

Answer: A, C

Explanation:

If redistribution is used, care must be taken that only local routes are redistributed. For example, routes learned from other autonomous systems (that were learned by redistributing BGP into the IGP) must not be sent out again from the IGP, or routing loops could result.

Incorrect Answers:

B: With redistribution into exterior protocols such as BGP, no routes are automatically summarized.

D: All routes will be included in the redistribution, including externally learned routes.

Reference: Building Scalable Cisco Networks (Cisco Press) page 408.

Topic 2: Implementation and Configuration (136 questions)

Section 1: Given a set of network requirements, identify the steps to configure an Enhanced IGRP environment and verify proper operation (within described guidelines) of your routers (19 questions)

QUESTION NO: 1

You are a network administrator on a Cisco router that's running EIGRP. You want to aggregate some of the routes. Which of the following commands below would you enter to configure the summary route you need?

- A. `ip auto-summary as-number address mask`
- B. `ip summary-address as-number address mask`
- C. `ip auto-summary eigrp as-number address mask`
- D. `ip summary-route eigrp as-number address mask`
- E. `ip summary-address eigrp as-number address mask`

Answer: E

Explanation:

The **ip summary-address eigrp** command is used to configure a summary aggregate address for a specified interface.

Syntax:

ip summary-address eigrp *autonomous-system-number address mask*

Incorrect Answers:

A, C, D: No such command exists using this syntax.

B: We must specify the appropriate protocol with the **eigrp** keyword.

QUESTION NO: 2

Router TK1 has a 256kbps serial interface link to another TestKing location. The configuration file of this interface is shown below:

```
interface serial 0/0
bandwidth 56
```

Leading the way in IT testing and certification tools, www.testking.com

```
ip bandwidth-percent eigrp 1 200
```

From the command output above, how many kbps of bandwidth is allocated for EIGRP traffic?

- A. 56 kbps
- B. 112 kbps
- C. 128 kbps
- D. 256 kbps

Answer: B

Explanation:

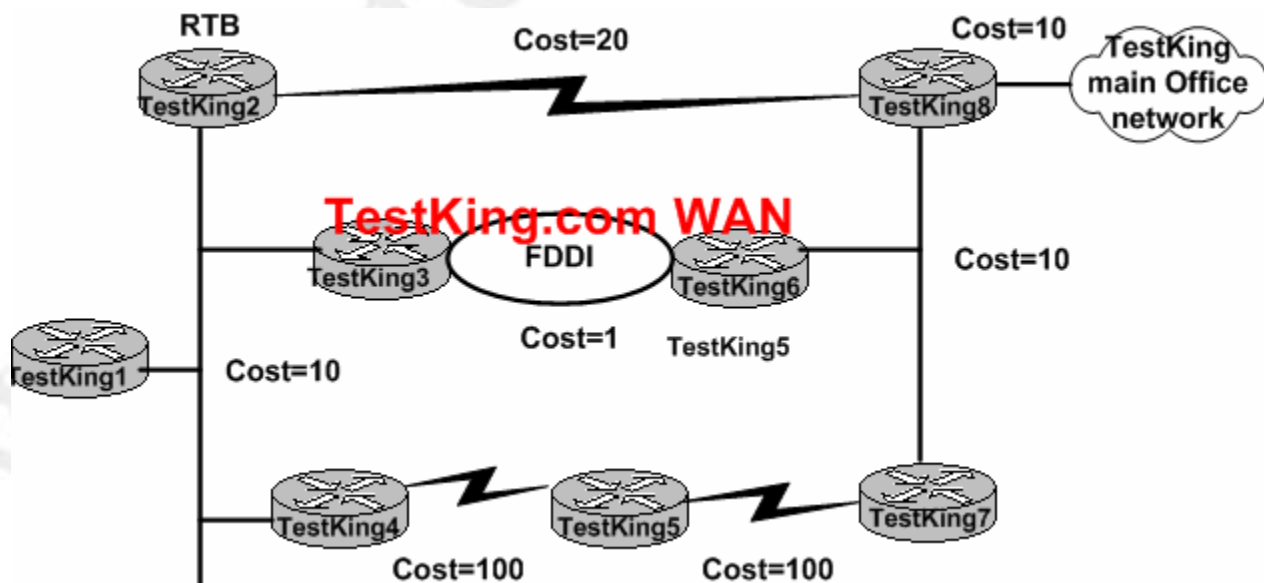
The **bandwidth-percent** command tells EIGRP what percentage of the configured bandwidth it may use. The default is 50 percent. Since the **bandwidth** command is also used to set the routing protocol metric, it may be set to a particular value for policy reasons. In this case, it is set to 56 kbps even though the actual link is 256 kbps. The **bandwidth-percent** command can have values greater than 100 if the bandwidth is configured artificially low due to such policy reasons. In this specific case, the bandwidth configured on the interface is 56kbps, so 200 percent of this value is 112 kbps.

Reference:

http://www.cisco.com/en/US/tech/tk365/tk207/technologies_tech_note09186a0080094063.shtml

QUESTION NO: 3

The TestKing EIGRP network is displayed below:



Leading the way in IT testing and certification tools, www.testking.com

All of the routers in the above network segment are configured with EIGRP. If both routers TestKing2 and Testking3 were to fail, how would TestKing1 react?

- A. TestKing1 will automatically route packets via TestKing4 to the TestKing Main Office Network.
- B. TestKing1 will place the route via TestKing4 into the hold down state.
- C. TestKing1 will go into the active state for all routers.
- D. TestKing1 will go into the active state for the route to TestKing Main Office Network.

Answer: D

Explanation:

First, TestKing1 it checks the topology table for an alternate route then it sends a query to neighboring routers for other routes to the destination. In this case, since both the optimal route and the feasible successor failed, TestKing1 must go into the active state first before routing around the failure via TestKing4. The re-route will happen only after TestKing1 transitions into the active state and learns about the route through TestKing4.

QUESTION NO: 4

You are the network engineer at TestKing. The TestKing network has five paths from Router TK1 to a given destination. The local EIGRP metric on Router TK1 for these paths are as follows:

```

TK1
-----
Path 1: 1500
Path 2: 1500
Path 3: 2000
Path 4: 4000
Path 5: 4000

```

Variance 3 is configured on Router TK1. Which paths would be included in Router TK1's routing table?

- A. path 1, 2, and 4
- B. path 1, 2, and 3
- C. path 1, 2, and 5
- D. path 1, 2, 4 and 5
- E. path 1, 2, 3, 4, and 5
- F. None of the above

Answer: E

Explanation:

If the variance number is higher than the default 1, the EIGRP process multiplies the best (lowest) cost or metric value for a path by the number stated as the variance multiplier. All paths to the same destination that have metrics within this new range are now included in load balancing. Per the question this means: $1500 * 3 = 4500$. All available options are within 1500-4500.

Reference: Self-Study CCNP BSCI Exam Certification Guide Third Edition P.485, Topic: Load Balancing in EIGRP.

QUESTION NO: 5

Which one of the following commands would configure EIGRP for IP on your Cisco router?

- A. `ip eigrp routing`
- B. `router eigrp process-id`
- C. `ip eigrp autonomous-system-number`
- D. `router eigrp autonomous-system-number`

Answer: D

Explanation:

Perform the following steps to configure EIGRP for IP:

Step 1 Enable EIGRP and define the autonomous system.

```
routerTK(config)#router eigrp autonomous-system-number
```

Step 2 Indicate which networks are part of the EIGRP autonomous system.

```
routerTK(config-router)#network network-number
```

Step 3 Define bandwidth of a link for the purposes of sending routing update traffic on the link.

```
routerTK(config-if)#bandwidth kilobits. This step is optional.
```

QUESTION NO: 6

The syntax for configuring variance on an EIGRP router is:

```
variance number
```

In the EIGRP command sequence above, what is the function of the *number* variable?

- A. It acts as a multiplier.
- B. It defines the limit for how far the metrics can be separated.
- C. It indicates how many paths can be used for load balancing.

D. It indicates how many paths can be used for unequal load balancing.

Answer: A

Explanation:

Every routing protocol supports equal cost path load balancing. IGRP and EIGRP also support unequal cost path load balancing, which is known as variance. The **variance number** command instructs the router to include routes with a metric less than or equal to **number** times the minimum metric route for that destination. Thus **number** is used as multiplier, even though it defines a limit for the metrics. This limit is used for unequal load balancing.

Incorrect Answers:

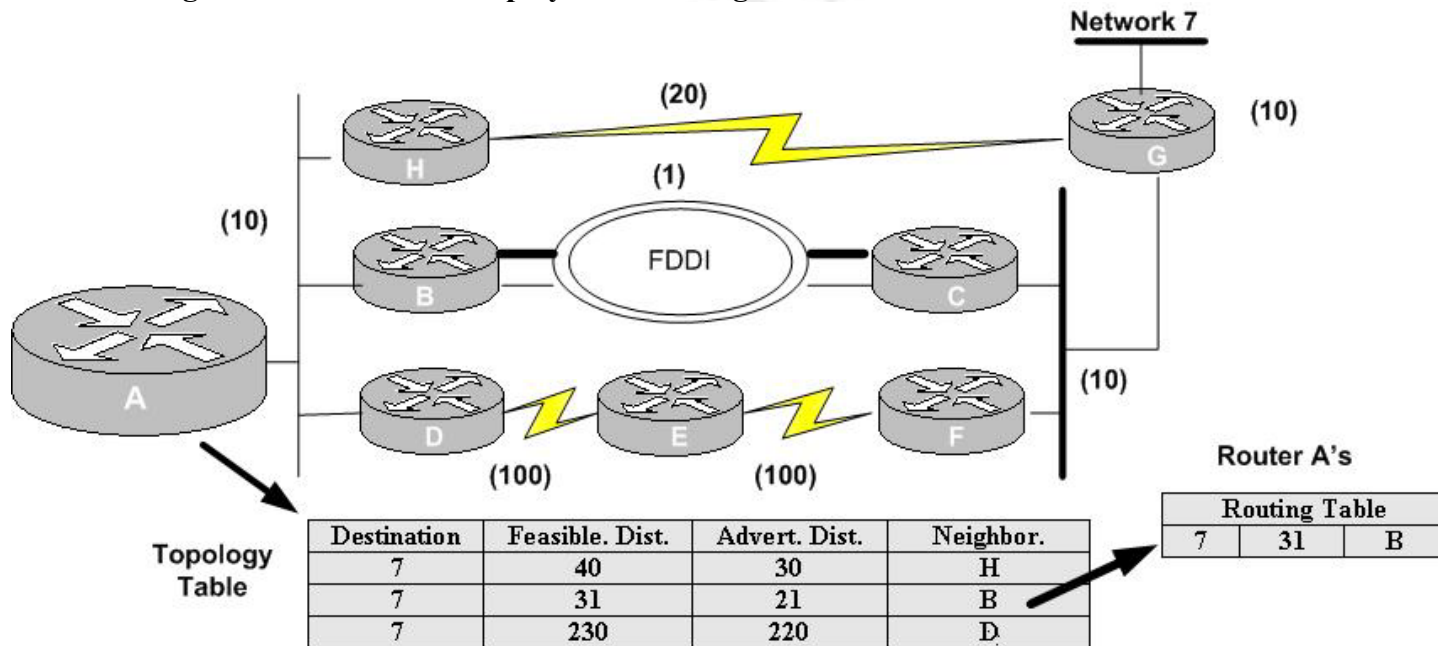
B: *Number* does not directly define a limit.

C, D: *Number* does not directly indicate the number of paths that can be used.

Reference: <http://www.cisco.com/warp/public/103/19.html>

QUESTION NO: 7

The TestKing EIGRP network is displayed in the diagram below:



You are the systems administrator of the above EIGRP network and you've elected to shut down router B's FDDI interface. After you execute the shutdown, which router will become the feasible successor to network 7 for router A?

- A. B
- B. D
- C. H

D. There is no feasible successor

Answer: D

Explanation:

Router H will be the successor, and that route will be placed in the Routing table.

Router A detects the link failure between Router B and network 7. It checks the topology table for a successor. It finds that H is the successor since the advertised distance for H (30) is less than the feasible distance for B (31).

However, there is no next best route – no feasible successor. The candidate route through D has an advertised distance (220) that is higher than the feasible distance of the successor route (40).

Note: Successor is a route selected as the primary route to use to reach a destination. Successors are the entries kept in the routing table.

A feasible successor is a backup route. These routes are selected at the same time the successors are identified, but they are kept in a topology table,

Incorrect Answers:

A: The FDDI interface of B is down.

B: The candidate route through D has an advertised distance (220) that is higher than the feasible distance of the successor route (40). It cannot be used as a feasible successor.

C: Router H is the successor, not the feasible successor.

QUESTION NO: 8

You are the network engineer at TestKing. You want to configure EIGRP to run across a 56 Kbps serial PPP link on the TestKing network. You also want to ensure the proper convergence of EIGRP routes.

What command should you issue on the serial interface?

- A. bandwidth 56
- B. bandwidth 56000
- C. ip bandwidth-percent eigrp 1 56
- D. ip bandwidth-percent eigrp 1 56000

Answer: A

Explanation:

The bandwidth is a logical construct whose value can have wide-reaching implications on the function of your network. It does not affect the actual speed of the link. In fact, it is practical to configure the **bandwidth** command only on serial lines, where the speed of the link will vary considerably. This command is useful for routing protocols that use bandwidth as part of the metric. For serial interfaces, this value defaults to a full T1, so it may appear to be a better link than what is actually being used.

Router (config) # interface S0

Router (config-if) # bandwidth speed-of-line

The bandwidth command is always specified in kbps.

EIGRP limits itself to 50 percent of the value specified in the **bandwidth** command, or if the bandwidth command is not set, the interface defaults. If you need to limit this percentage further, the upper limit that EIGRP uses can be stated as a percentage of the **bandwidth** command.

The **ip bandwidth-percent-eigrp** command interacts with the **bandwidth** command on the interface. You would use this command primarily because in your network, the **bandwidth** command does not reflect the true speed of the link. The **bandwidth** command might have been altered to manipulate the routing metric and path selection of a routing protocol, such as IGRP or OSPF. It might be better to use other methods of controlling the routing metric and return the bandwidth to a true value. Otherwise, the **ip bandwidth-percent eigrp** command is available. It is possible to set a bandwidth percent that is larger than the stated bandwidth. This is with the understanding that, although the bandwidth might be stated to be 56 kbps, the link is in fact 256 kbps. The following shows the structure of the **ip bandwidth-percent eigrp** command:

Router (config) #interface S0

Router (config-if) #ip bandwidth-percent eigrp autonomous-system-number percent

QUESTION NO: 9

The TestKing Network topology is displayed in the diagram below:



Based on the information above, what optional EIGRP configuration will be required in order to achieve full network connectivity?

- A. Use the EIGRP **no auto-summary** command on TK1 and TK2.
- B. Use the EIGRP **no auto-summary** command on TK3 and TK4.

- C. Use the **passive interface** on the TK1 and TK2 interface that connects to the 10.1.1.0/24 and 10.1.2.0/24 subnet respectively.
- D. Use the **passive interface** command between the TK3 and T1 connection and between the TK3 and TK2 connection.
- E. Use the **variance** command on TK3.

Answer: A

Explanation:


A useful feature of EIGRP is automatic route summarization; this summarizes subnets to the classful network boundary. This is enabled by default, you can turn this off per AS by using the following command in router configuration mode:

```
Router(config-router)#no auto-summary
```

In the example above, this is needed, otherwise both TK1 and TK2 will advertise the summarized network 10.0.0.0/8 to router TK3, making their networks unreachable.

QUESTION NO: 10

You are the network engineer at TestKing. The TestKing network is shown in the following graphic:



Router RTX is configured with the following commands:

```
Interface serial 0  
Encapsulation frame relay
```

You want to complete an EIGRP hybrid multipoint configuration on RTX S0 by configuring it with the correct bandwidth parameter.

Which command should you issue?

- A. bandwidth 56
- B. bandwidth 224
- C. bandwidth 256
- D. bandwidth 768
- E. Non of the above

Answer: B

Explanation:

If the multipoint network has different speeds allocated to the VCs, take the lowest CIR and simply multiply it by the number of circuits. This is applied to the physical interface. In the case of the question $4 \times 56 = 224$ – bandwidth 224

Reference: Self-Study CCNP BSCI Exam Certification Guide Third Edition P. 490.

QUESTION NO: 11

The TestKing multipoint frame relay network is using EIGRP for the routing protocol. In this network, how is the bandwidth determined for each multipoint neighbor?

- A. Bandwidth command per neighbor.
- B. The configured CIR per subinterface.
- C. The configured CIR divided by the number of neighbors on that interface.
- D. Bandwidth of the main interface divided by the number of neighbors on that interface.

Answer: D

Explanation:

When configuring multipoint interfaces, especially for Frame Relay, remember that all neighbors share the bandwidth equally, regardless of the actual CIR of each individual PVC. For multipoint interfaces, the bandwidth used by the IOS is taken from the configured bandwidth of the main interface, divided by the total number or neighbors on that interface.

QUESTION NO: 12

Your network consists of a router named R1 that's configured in a hub and spoke topology, supporting 24 remote office via a point-to-multipoint Frame Relay EIGRP network deployment. The bandwidth command was not manually entered on the frame relay main interface or the sub-interfaces. What is the perceived bandwidth (as seen by the EIGRP process) of each Frame Relay connection?

- A. 64 kbps
- B. 128 kbps
- C. 1.544 Mbps
- D. 1.536 Mbps
- E. None of the above

Answer: A

Explanation:

The default bandwidth for all serial WAN interfaces is 1.544 Mbps. EIGRP by default assumes the bandwidth of a serial interface is a full T1 (1.544 Mbps) if not specified. This includes sub-interfaces. For multipoint frame relay interfaces, the perceived bandwidth used by the IOS is

taken from the configured bandwidth, divided by the number of configured neighbors. Since there are 24 remote offices, there are 24 configured frame relay neighbors, so the bandwidth used by the IOS is the full T1 divided by 24, which is 64kbps.

QUESTION NO: 13

EIGRP has been configured on the WAN links of router TK1. On one of the serial T1 interfaces, the bandwidth statement was not specified. By default, how will the EIGRP process perceive this T1 link?

- A. 256 Kbps
- B. 1.544 Mbps
- C. It depends as it is set by the PVC
- D. It is set by the DLCI

Answer: B

Explanation:

The default EIGRP link speed is 1.544 Mbps for non-high speed WAN links

Note: The enhanced code uses the "bandwidth" subcommand on interfaces and sub-interfaces in order to determine the rate at which to generate EIGRP packets. This parameter is automatically set on fixed-bandwidth interfaces (such as LANs), but defaults to T1 (1544 Kbps) for all serial media.

Incorrect Answers:

- A: The default speed is 1.544 Mbps, not 256 Kbps.
C, D: The default speed is not set by the PVC or the DLCI.

QUESTION NO: 14

You are in the midst of configuring an NBMA network with EIGRP as the routing protocol. Which of the options below would you use to configure the interface bandwidth for a point-to-point interface?

- A. The DLCI assigns the bandwidth for the interface.
- B. The sliding window size determines interface bandwidth.
- C. You should use the default bandwidth assigned to the interface.
- D. You should manually configure bandwidth as the CIR of the PVC.

Answer: D

Explanation:

The bandwidth can be configured separately on each subinterface. Since this is NBMA we can assume that Frame Relay is used. For Frame Relay on point-to-point the bandwidth should be set it to the CIR of the PVC.

Note 1: **NBMA** (Non-broadcast Multi-access) supports many (more than two) routers, but have no broadcast capability. Frame Relay and X.25 are example of NBMA.

Note 2: The **CIR** (Committed Information Rate) is the committed rate (in bits per second) at which the ingress access interface trunk interfaces, and egress access interface of a Frame Relay network transfer information to the destination Frame Relay end system under normal conditions.

Incorrect Answers:

A: Does not apply.

B: Sliding windows does not apply.

C: The bandwidth can and should be configured separately on each subinterface.

Reference: Configuration Notes for the Enhanced Implementation of EIGRP

<http://www.cisco.com/warp/public/103/12.html>

QUESTION NO: 15

Router TK2 is configured for EIGRP as shown below:

```
router eigrp 100
```

```
network 10.0.0.0
```

```
eigrp stub
```

Based on this configuration, which types of routes will be advertised with the EIGRP configuration as shown? (Choose two)

- A. Static
- B. Receive only
- C. Summary
- D. Stub
- E. Connected
- F. dynamic

Answer: C, E

Explanation:

A router that is configured as a stub with the **eigrp stub** command shares connected and summary routing information with all neighbor routers by default.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a0080087026.html

QUESTION NO: 16

The TestKing EIGRP network is utilizing the EIGRP stub routing feature. Which of the following are key concepts that apply when configuring the EIGRP stub routing feature in a hub and spoke network? (Select three)

- A. A hub router prevents routes from being advertised to the remote router.
- B. Only remote routers are configured as stubs.
- C. Stub routers are not queried for routes.
- D. Spoke routers connected to hub routers answer the route queries for the stub router.
- E. A stub router should have only EIGRP hub routers as neighbors.
- F. EIGRP stub routing should be used on hub routers only.

Answer: B, C, E

Explanation:

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.

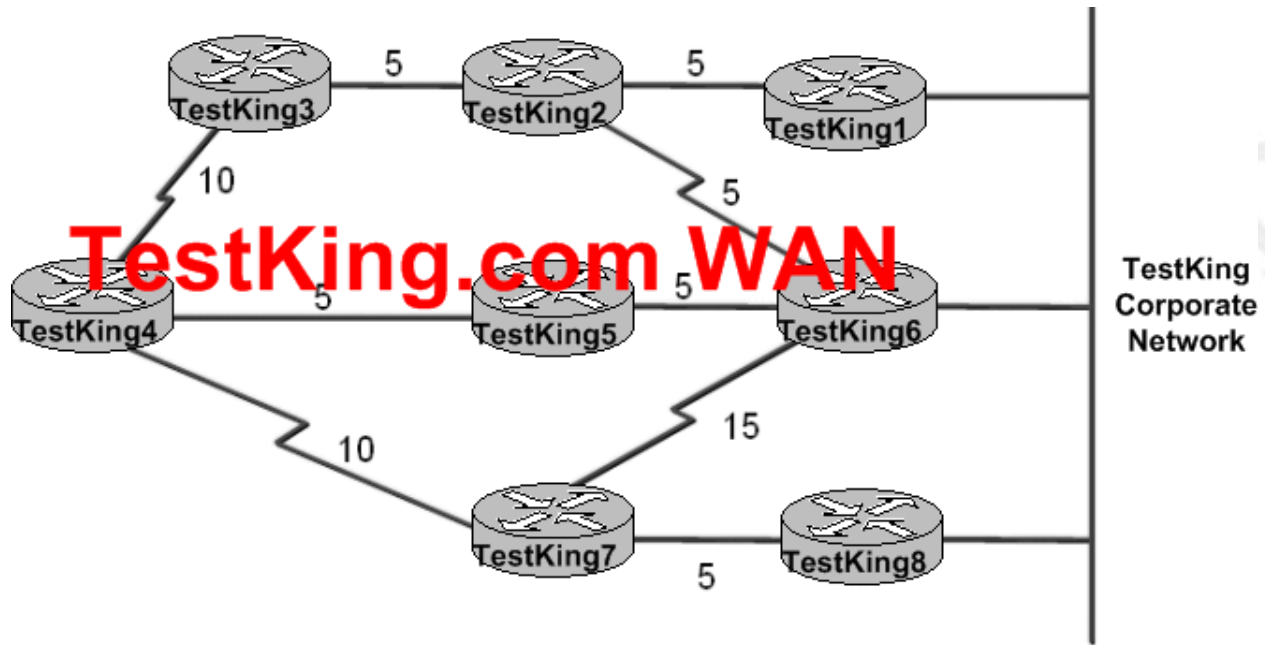
This feature should only be used on stub routers. A stub router is defined as a router connected to the network core or distribution layer through which core transit traffic should not flow. A stub router should not have any EIGRP neighbors other than distribution routers. Ignoring this restriction will cause undesirable behavior.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a0080087026.html

QUESTION NO: 17**Exhibit**

Leading the way in IT testing and certification tools, www.testking.com



Traffic from TestKing4 to the TestKing Corporate Network is distributed between the links with unequal costs in the EIGRP network by configuring the variance command on TestKing4 to 2. How many paths will participate in the load sharing?

- A. 1
- B. 2
- C. 3
- D. 4
- E. 5
- F. 6

Answer: B

Explanation:

EIGRP automatically load balances across links of equal cost. Whether the traffic is sent on a predestination or round-robin basis depends on the internal switching within the router. It is possible to configure EIGRP to load balance across unequal-cost paths using the variance command.

The variance command allows the administrator to identify the metric scope for including additional paths by the use of a multiplier parameter. The command structure follows:

```
Router(config-router)#variance multiplier
```

The multiplier argument is the metric value used for load balancing. It can be a value from 1 to 128. The default is 1, which means equal-cost load balancing.

Reference: CCNP Self-Study CCNP BSCI Exam certification guide p.485

QUESTION NO: 18

Which configuration command is used to enable EIGRP unequal-cost path load balancing?

- A. maximum-paths
- B. distance
- C. metric
- D. variance
- E. default-metric

Answer: D

Explanation:

EIGRP automatically load balances across links of equal cost. Whether the traffic is sent on a predestination or round-robin basis depends on the internal switching within the router. It is possible to configure EIGRP to load balance across unequal-cost paths using the variance command.

The variance command allows the administrator to identify the metric scope for including additional paths by the use of a multiplier parameter. The command structure follows:


```
Router(config-router)#variance multiplier
```

The multiplier argument is the metric value used for load balancing. It can be a value from 1 to 128. The default is 1, which means equal-cost load balancing.

Reference: CCNP Self-Study CCNP BSCI Exam certification guide p.485

QUESTION NO: 19

The TestKing WAN is displayed below:



Examine the EIGRP network diagram. All paths should be used for load balancing between TestKing4 and the TestKing Corporate Network. Which value should be assigned to the variance command to accomplish this?

- A. 1
- B. 2
- C. 2.5
- D. 3
- E. 5
- F. 6

Answer: D

Explanation:

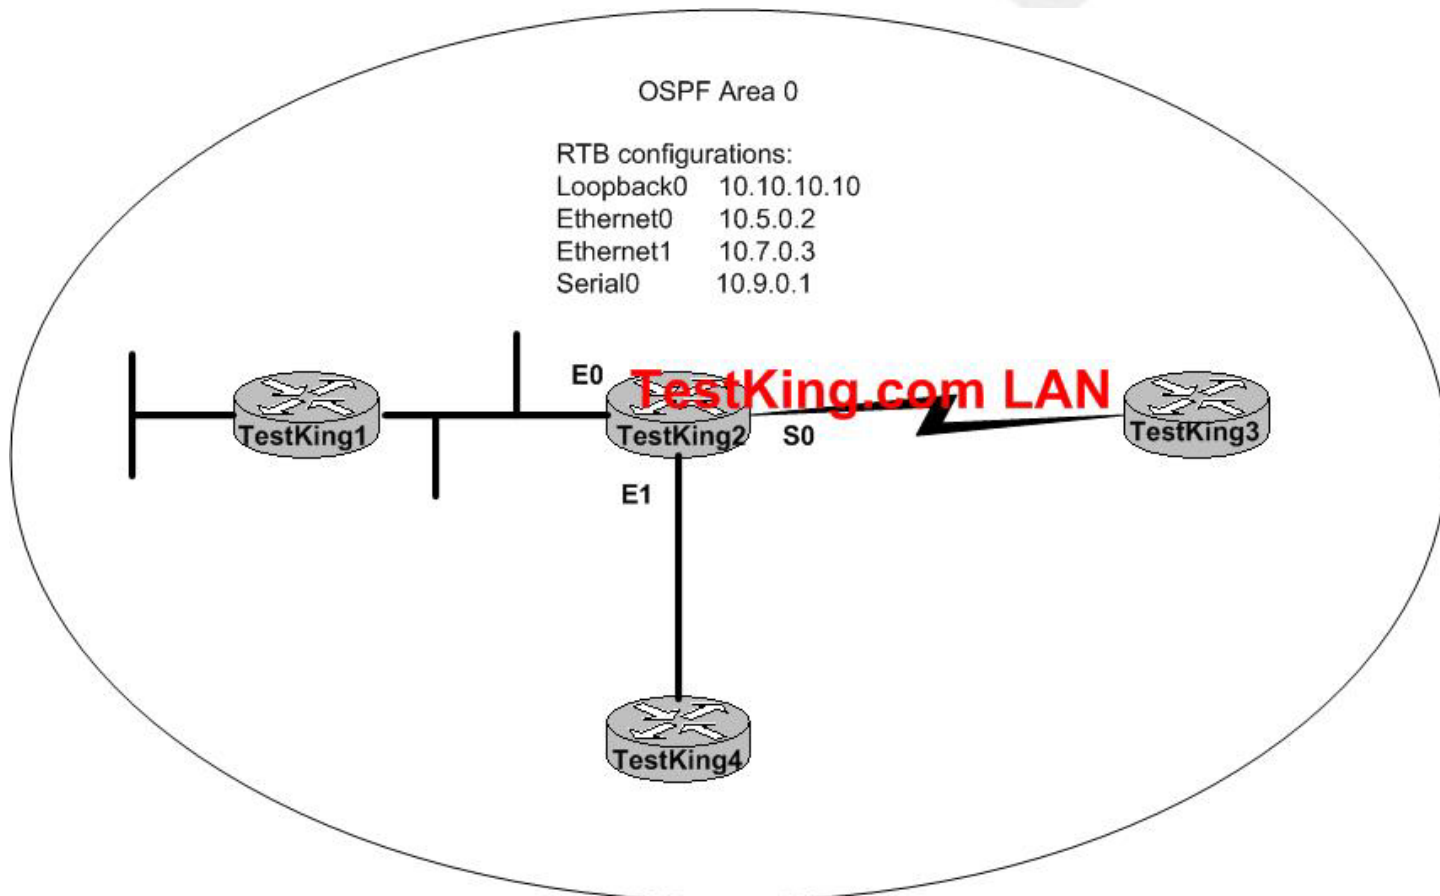
Every routing protocol supports equal cost path load balancing. In addition to that, IGRP and EIGRP also support unequal cost path load balancing. Use the variance command to instruct the router to include routes with a metric less than n times the minimum metric route for that destination, where n is the number specified by the variance command. The variable n can take a value between 1 and 128, with the default being 1, which means equal cost load balancing.

Traffic is also distributed among the links with unequal costs, proportionately, with respect to the metric. In this example, the worst route is 2.5 times worse than the best route. Since we can not use 2.5 for the variance (the value must be an integer) we must specify a minimum value of 3.

Section 2: Given an addressing scheme and other laboratory parameters, identify the steps to configure a single-area OSPF environment and verify proper operation (within described guidelines) of your routers (8 questions)

QUESTION NO: 1

The TestKing OSPF backbone consists of four routers as shown below:



Assuming that the OSPF router ID was not explicitly defined on router TK2, what would its router ID default to?

- A. 10.5.0.2
- B. 10.7.0.3
- C. 10.9.0.1

- D. 10.10.10.10
- E. 0.0.0.0
- F. None of the above

Answer: D

Explanation:

The router ID is the tie-breaker for OSPF path selection. The path selection process uses a variety of metrics to select a route. If all other metrics (accessibility, administrative weight, local preference, etc.) are equal, OSPF determines the router ID using the following priority:

1. Use the address configured by the **ospf router-id** command
2. Use the address of the loopback 0 interface
3. Use the highest IP address of any interface
4. If no interface exists, set the router-ID to 0.0.0.0

QUESTION NO: 2

The Berlin and Nuremberg OSPF networks need to be established and configured in the following manner:

- The router is named Berlin.
- The clocking is provided on the Berlin router's serial 0/0 interface.
- The secret password in the Berlin router is "TestKing".
- You should use OSPF in a single area
- Area 0 should be used for the routing protocol.
- Use 1 for the process ID of the routing protocol.
- The IP addresses and subnet masks are listed in the chart.

Berlin

FA0/0 10.60.2.1/23

S0/0 10.60.4.1/30


Lo 0 10.90.20.10/32

Nuremberg

FA0/0 10.60.6.1/24

S0/0 10.60.4.2/30

Lo 0 10.90.20.20/32



Please provide the complete configuration and also the prompt where the configuration is to be done.

Start by clicking on host that is connected to the router you want to configure.

Lab A
 Berlin
 FA0/0 10.60.2.1/23
 S0/0 10.60.4.1/30
 Lo 0 10.90.20.10/32

Secret Password: TestKing

Lab B
 Nuremberg
 FA0/0 10.60.6.1/24
 S0/0 10.60.4.2/30
 Lo 0 10.90.20.20/32

Answer:

<Click on Host2, which is connected to the Lab 1 router>


```

Berlin>enable <--- notice the prompt
Password: TestKing
Berlin# configure terminal
Berlin(config)#router ospf 1
Berlin(config-router)# network 10.60.2.1 0.0.1.255 area 0 <--- this specific mask should be use
to reflect the subnet mask of 23
Berlin(config-router)# network 10.60.4.1 0.0.0.3 area 0 <--- this specific mask should be use
to reflect the subnet mask of 30
Berlin(config-router)# network 10.90.20.10 0.0.0.0 area 0 <--- this > specific mask should be
use to reflect the subnet mask of 32
Berlin(config-router)#exit
Berlin(config)#^Z
Berlin#copy running-config startup-config

```

QUESTION NO: 3

Which command should you use to verify what networks are being routed by a given OSPF process?

- A. show ip ospf
- B. show ip route
- C. show ip protocol
- D. show ip ospf database

Answer: C

Explanation:

The information displayed by the show ip protocols command is useful in debugging routing operations. Information in the Routing Information Sources field of the show ip protocols output can help you identify a router suspected of delivering bad routing information. For OSPF routers, this command will display the routed networks.

Incorrect Answers:

A: To display general information about Open Shortest Path First (OSPF) routing processes, use the show ip ospf command in EXEC mode. This command will display the areas assigned and other useful information, but not the networks being routed.

Example:

```

TK1# show ip ospf
Routing Process "ospf 201" with ID 192.42.110.200
Supports only single TOS(TOS0) route
It is an area border and autonomous system boundary router
Redistributing External Routes from,
    igrp 200 with metric mapped to 2, includes subnets in redistribution
    rip with metric mapped to 2
    igrp 2 with metric mapped to 100
    igrp 32 with metric mapped to 1
Number of areas in this router is 3

```

Area 192.42.110.0

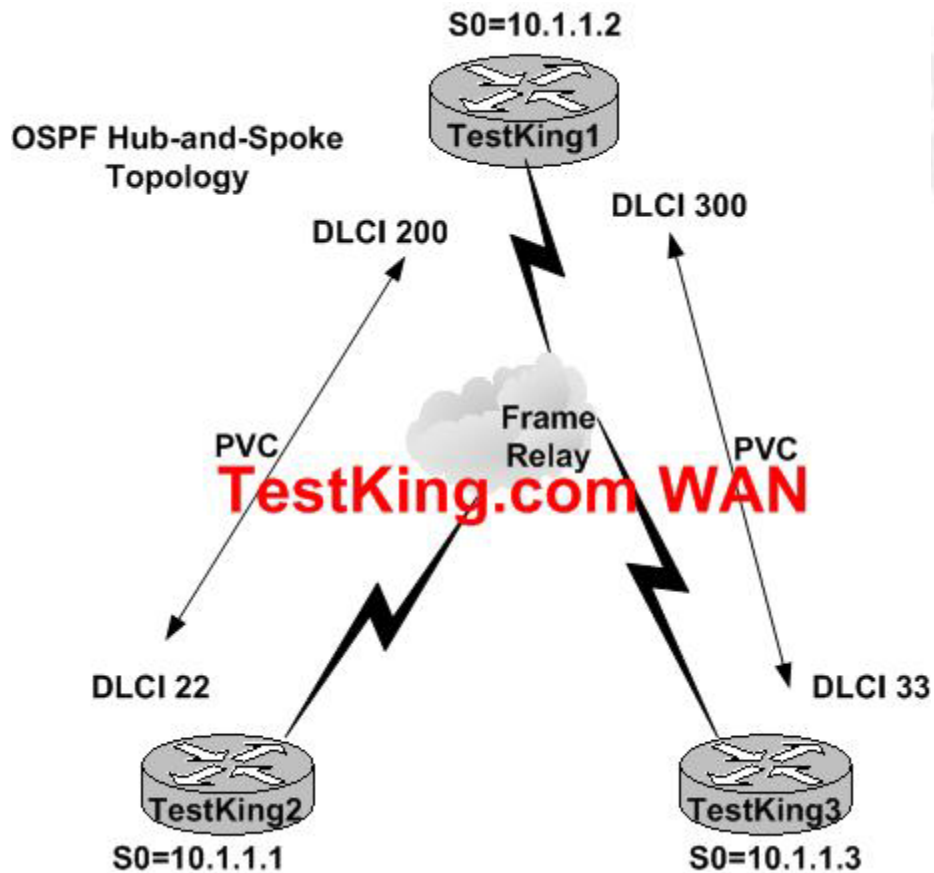
Number of interfaces in this area is 1
 Area has simple password authentication
 SPF algorithm executed 6 times

B: This will display the active routing table, but not the networks that are being routed.

D: The OSPF database does not display the networks being routed.

QUESTION NO: 4

OSPF is being configured over the TestKing frame relay network as displayed in the diagram below:



Which three of the following commands would you enter on the TestKing1 serial 0 interface, if you wanted to implement OSPF on this network? (Select three)

- A. `ip ospf network point-to-point`
- B. `ip ospf network point-to-multipoint`
- C. `frame-relay map ip 10.1.1.1 200`
- D. `frame-relay map ip 10.1.1.3 300`
- E. `frame-relay map ip 10.1.1.1 200 broadcast`

F. frame-relay map ip 10.1.1.3 300 broadcast

Answer: B, E, F

Explanation:

The **ip ospf network command**, typed under the interface configuration mode, is used to specify the OSPF network configuration and sets the network mode to **point-to-multipoint**

DLCI- Data-link connection identifier (DLCI) number.

Broadcast - Forwards broadcasts to the specified IP address. This keyword is needed on the frame relay mappings in order to transmit the OSPF information across the frame relay network.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products_feature_guide09186a0080087b42.html

QUESTION NO: 5

Router TK1 is configured as shown below:

```
interface serial0
ip address 110.1.1.1 255.255.255.0
encapsulation frame-relay
ip ospf network point-to-multipoint

router ospf 51
network 110.1.1.0 0.0.0.255 area 0
```

Based on the information above, which two of the following statements are true? (Select two)

- A. DB/BDR elections do not take place.
- B. The router is restricted to a hub and spoke topology.
- C. The area 0 NBMA cloud is configured as more than one subnet.
- D. OSPF **neighbor** statements are not necessary.
- E. The OSPF hello timers need to be adjusted manually.

Answer: A, D

Explanation:

In an OSPF Point-to-Multipoint environment, DB/BDR elections do not take place.

The **neighbor** command became somewhat obsolete with the introduction of the capability to configure other network modes for the interface, regardless of the underlying physical topology.

Reference: Building Scalable Cisco Networks (Cisco Press) page 130 and 181

Point-to-Multipoint Network:

Point-to-multipoint is a single interface that connects to multiple destinations. The underlying network treats the network as a series of point-to-point circuits. It replicates LSA packets for each circuit. OSPF traffic is sent as multicast. There is no DR or BDR election. This technology uses one IP subnet for all endpoints on the network.

By default, the network is considered to be a series of point-to-point interfaces. There is no need to specify neighbors, because the neighbors will see each other and simply become adjacent, with no need for the election of a DR or a BDR. Point-to-multipoint does not try to reduce adjacencies using a DR. Instead, it accepts the extra overhead of having a full set of adjacencies for the sake of stability. Point-to-multipoint forms an adjacency automatically along any PVC, which causes more overhead but is more resilient than NBMA.

QUESTION NO: 6

On the TestKing NBMA Frame Relay network, subinterfaces were configured. In routing OSPF over the frame relay network, what advantage does using subinterfaces provide?

- A. To conserve IP addressing space.
- B. To avoid split-horizon issues with the routing protocol.
- C. Because logical interfaces are more reliable than physical interfaces.
- D. Subinterfaces remain up when the physical interface changes to a down state.
- E. All of the above.

Answer: B

Explanation:

When configuring routers in a NBMA topology, subinterfaces are typically used. A physical interface can be split into multiple logical interfaces, called subinterfaces, with each subinterface being defined as point-to-multipoint interface. Subinterfaces originally were created to better handle issues caused by split horizon over NBMA and distance vector-based routing protocols.

Incorrect Answers:

A: Using subinterfaces creates separate IP subnets for each frame relay link, so actually more IP addresses are used this way.

C: Using logical interfaces has no impact on the stability of any network link.

D: When the physical link goes down, all subinterfaces also go down.

Reference: Building Scalable Cisco Networks (Cisco Press) page 120.

QUESTION NO: 7

You have been assigned the task of connecting two office networks together via a frame relay network, and running OSPF across this network. What kind of configuration structure would you use to accomplish this?

- A. Point-to-point over sub-interfaces.
- B. Point-to-multipoint star configuration.
- C. Point-to-multipoint using a single subnet.
- D. Point-to-multipoint nonbroadcast using a single subnet.
- E. None of the above.

Answer: A

Explanation:

Using point to point subinterfaces will eliminate the issues surrounding split horizons over NBMA networks. Configuring the network in this way is the Cisco recommended solution.

Reference: RFC1586, Guidelines for Running OSPF over Frame Relay Network.

QUESTION NO: 8

A TestKing router is configured as shown below:

```
interface serial 0
    ip address 164.67.36.1 255.255.255.224
    encapsulation frame-relay
    ip ospf network non-broadcast
!
router ospf 1
    network 164.67.36.0 0.0.0.31 area 0
    neighbor 164.67.36.2
    neighbor 164.67.36.3
```

Based on this configuration, which of the following statements are true? (Select two)

- A. There can be no DR or BDR in this configuration.
- B. This is a point-to-point configuration over Frame Relay.
- C. The network mode is nonbroadcast multiaccess (NBMA).
- D. The DR and BDR need a static list of neighbors due to non-broadcast.

Answer: C, D

Explanation:

The configuration printout shows an OSPF configuration across a frame relay network that is not configured for subinterfaces. With this setup you are required to specify neighbors so that a DR and BDR can be elected. In addition, all locations are part of the same IP subnet so the OSPF neighbors need to be explicitly assigned at the host location of an NBMA network.

Note: Had subinterfaces been used, the network would be seen as a series of logical point to point links, so the OSPF neighbors would not need to be specified.

Section 3: Given an addressing scheme and other laboratory parameters, identify the steps to configure a multiple area OSPF environment and verify proper operation (within described guidelines) of your routers (17 questions)

QUESTION NO: 1

In a network running OSPF, what is the term used to describe the administrative process of dividing a large area into smaller areas?

- A. interior areas
- B. OSPF subarea
- C. link-state protocol
- D. hierarchical routing

Answer: D

Explanation:

OSPF's ability to separate a large internetwork into multiple areas is referred to as hierarchical routing. The backbone area, area 0, is considered to be at the top of the hierarchical chain, with the other areas lying below it.

QUESTION NO: 2

Assuming that you are configuring an ABR in an OSPF area, which IOS command would you execute if your goal was to summarize the networks advertised out of the area?

- A. summary-address *address mask*
- B. area area-id range *address mask*
- C. auto-summary *address mask area area-id*
- D. network *network-number wildcard mask area area-id*

Answer: B

Explanation:

The area area-id range *address mask* command consolidates IA (intra-area) routes on an ABR. This command instructs the ABR to summarize routes for a specific area before injecting them into a different area.

Incorrect Answers:

- A: The summary-address *address mask* command consolidates external routes (inter-area) on an ASBR
- C: Auto-summarization is not useful here.
- D: The network command cannot be used for this purpose. This command is simply used to add a network to the OSPF routing process.

QUESTION NO: 3

Which two of the following characteristics are defined by the `network` command? (Select two)

- A. The OSPF area ID
- B. The OSPF router ID
- C. The OSPF process ID
- D. Which interface belongs to which OSPF area

Answer: A, D

Explanation:

The network command designates the OSPF area for an interface with the specified IP address.

Syntax: `network address wildcard-mask area area-id`

QUESTION NO: 4

Router TK1 is configured for OSPF as shown below:

```
router ospf 76
 network 172.22.23.0 0.0.0.0 area 1
 network 172.18.0.0 0.0.255.255 area 0
 area 0 range 172.18.0.0 255.255.0.0
 area 1 range 172.22.23.0 255.255.255.0
```

Which of the following statements are true regarding the above configuration? (Select three)

- A. The OSPF router ID is 76.
- B. This is an area border router.
- C. The designated router priority is 76.
- D. This router connects area 1 to the backbone area.
- E. Any router interfaces with an address of 172.18.x.x are in area 0.

Answer: B, D, E

Explanation: An interface may belong to only one area. If a router has multiple interfaces and if any of those interfaces belong to different areas, the router is considered as an area border router. The networks that follow the network command are connected to each other and Area 0 is always the backbone area. Finally, the command “network 172.18.0.0 0.0.255.255 area 0” identifies that all interfaces with IP address of 172.18.0.0 area within area 0.

Incorrect Answers:

A, C: The command router ospf 76 identifies the process ID as 76 and not the router ID nor the router priority.

Catherine Paquet and Daine Teare, “Building Scalable Cisco Networks” (Cisco Press 2001), pp 178-181.

QUESTION NO: 5

When configuring a multi-area OSPF network to summarize routes, what additional command is required by ASBR’s that is not needed by ABR’s?

- A. **area range** command
- B. **ospf summarize** command
- C. **aggregate-route** command
- D. **summary-address** command
- E. None of the above

Answer: D

Explanation:

Using this command for OSPF causes an OSPF autonomous system boundary router (ASBR) to advertise one external route as an aggregate for all redistributed routes that are covered by the address. For OSPF, this command summarizes only routes from other routing protocols that are being redistributed into OSPF.

Incorrect Answers:

A: The area range command is used only with area border routers (ABRs). It is used to consolidate or summarize routes for an area. The result is that a single summary route is advertised to other areas by the ABR.

B, C: There are no such commands for OSPF.

Reference: OSPF Commands

QUESTION NO: 6

When designing OSPF networks, maintaining stability in an area is important. Which of the following describes a reason for this?

- A. Instability causes more LSAs to be sent, requiring more CPU to recalculate routes.
- B. Convergence cannot happen until holddown timers expire, so routing loops can occur.

- C. Flooding the area topological database instances consumes excessive bandwidth.
- D. Summary link LSAs cannot be sent until all routers in the OSPF area have the same topological database.

Answer: A

Explanation:

Instability of routes (links) would force sending of LSAs (Link State Advertisements), and CPU time would be required to recalculate the routes.

Incorrect Answers:

- B: OSPF does not use hold down timers. Holddown timers are used by distance vector protocols, such as RIP.
- C: The topological database is not distributed, only link changes.
- D: Summary-link LSAs are not sent to all routers.

Summary-link LSAs originate from area border routers, and flood throughout the LSA's associated area. Each summary-LSA describes a route to a destination outside the area but within the AS.

Reference: <http://www.faqs.org/rfcs/rfc2328.html>

QUESTION NO: 7

A TestKing OSPF router is configured in the following manner:

```
router ospf 200
  network 203.42.67.0 0.0.0.255 area 7
  network 203.42.68.0 0.0.0.255 area 0
  area 7 stub no-summary
  area 7 default-cost 30
```

Which of the following statements are true regarding this configuration? (Select two)

- A. Area 7 is a totally stubby area.
- B. If the backbone becomes discontinuous, traffic can be routed through area 7.
- C. Redistribution of other routing protocols takes place at the area designated router.
- D. Area 7 non-ABR routers contain only intra-area routing information and a default route.

Answer: A, D

Explanation:

An extension to stub areas is what is called "totally stubby areas". Cisco indicates this by adding a "no-summary" keyword to the stub area configuration. A totally stubby area is one that blocks external routes and summary routes (inter-area routes) from going into the area.

Note: The **area stub** command is used to define an area as a stub area.

Syntax: `area area-id stub [no-summary]`

The **no-summary** optional parameter prevents an ABR from sending summary link advertisements into the stub area.

Reference: OSPF Design Guide

<http://www.cisco.com/warp/public/104/3.html>

QUESTION NO: 8

Router TK1 is an ABR in the TestKing OSPF network. What does an ABR connect to?

- A. Multiple OSPF areas
- B. OSPF and RIP networks
- C. Multiple designated routers
- D. Multiple OSPF autonomous systems
- E. Multiple redistributed networks.

Answer: A

Explanation:

An ABR (Area Border Router) shares an interface with at least one other OSPF area. A router that contains network/area statements for two or more different areas is an ABR.

QUESTION NO: 9

You are a systems administrator of a large multi-area OSPF network, and you've just created a new area for an upcoming remote network. Ordinarily OSPF areas are be connected to the backbone, area 0. However, circumstances dictate you to connect it to the existing area 2 at this time.

Which conditions have to be met in order to make this configuration work? (Select three)

- A. There must be a virtual link.
- B. Area 2 must be a stub area.
- C. Area 2 cannot be a stub area.
- D. Area 2 must attach directly to area 0.
- E. Network summary link LSAs must be disabled.

Answer: A, C, D

Explanation:

A: There must be a virtual link from the new area to the backbone, area 0. The virtual link provides the disconnected area a logical path to the backbone.

C: The area through which you configure the virtual link, known as a transit area (here area 2), must have full routing information. It cannot be a stub area.

D: The transit area, area 2, must attach directly to area 0.

Reference: OSPF Virtual Link

<http://www.cisco.com/warp/public/104/ospfdb7.html>

Incorrect Answers:

- B: The transit area, area 2, cannot be a stub area.
- E: This is not a requirement.

QUESTION NO: 10

Which of the following is an OSPF configuration parameter that is used on an ABR, but not on an internal router?

- A. A virtual link to area 0.
- B. OSPF summarization command.
- C. `default-cost` extension to the `area` command.
- D. `no-summary` extension to the `area stub` command.
- E. None of the above

Answer: D

Explanation:

The `no-summary` extension of the `area stub` command is used only for ABRs connected to totally stubby areas. It prevents an ABR from sending summary link advertisements into the stub area. This option is used for creating a totally stubby area.

Incorrect Answers:

- A: For a virtual link to work both ends need to be configured.
- B, C: These commands are not specific to an ABR only.

QUESTION NO: 11

A new location is being added to the Testking network and a new OSPF area is being created for it. What kind of OSPF router would you use to connect this new OSPF area to the backbone area?

- A. ABR
- B. Stub router
- C. Internal router
- D. Backbone router

Answer: A

Explanation:

ABRs are used to connect non-backbone areas to the backbone area (area 0).

Incorrect Answers:

- B: Stub routers do not apply. An area is stub, not a router and not stub routers are typically a type of ABR, but not every ABR is configured as a stub.
- C: Internal routers are only used within an area.
- D: Backbone routers sit on the perimeter of the backbone area. They have at least one interface connected to area 0. However, backbone routers do not necessarily connect to other areas.

QUESTION NO: 12

Router TK1 is an OSPF ABR that connects area 3 to the network. You wish to configure area 3 to summarize the IP address range 172.16.20.192 – 172.16.20.223. Which two of the following IOS commands would you execute to accomplish this? (Select two)

- A. `network 172.16.20.192 0.0.0.31 area 3`
- B. `area 3 range 172.16.20.192 172.16.20.223`
- C. `area 3 range 172.16.20.192 255.255.255.224`
- D. `network 172.16.20.192 255.255.255.224 area 3`

Answer: A, C

Explanation:


- A: When configuring multiple OSPF areas, make sure to associate the correct network addresses with the desired area ID. Syntax: `network address wildcard-mask area area-id`. We must use a wildcard mask (0.0.0.31) and not a network mask (255.255.255.224)
- C: We must instruct the ABR to summarize routes for a specific area before injecting them into a different area. **Syntax:** `area area-id range address mask`

Incorrect Answers:

- B: This is the wrong syntax. We should use a network mask to specify the address range.
- D: We should use a wildcard mask, not a network mask with the network command.

QUESTION NO: 13

The TestKing OSPF network is displayed in the diagram below:



Based on the information above, what kind of router is RTD in OSPF terminology?

- A. Designated router
- B. Internal router
- C. Backbone router
- D. Area border router
- E. Autonomous system boundary router
- F. None of the above.

Answer: B

Explanation:

An *area* is a set of networks and hosts within an AS that have been administratively grouped together. We recommend that you configure an area as a collection of contiguous IP subnetted

networks. Routers that are wholly contained within an area are called *internal routers*. All interfaces on internal routers are directly connected to networks within the area.

Incorrect Answers:

A: Designated routers are defined on multi-access networks. In a network consisting of point to point circuits, the DR/BDR election process does not take place.

C: Backbone routers are routers that are contained within area 0.

D: Routers that belong to more than one area are called *area border routers*.

E: Routers that exchange routing information with routers in other ASs are called *AS boundary routers*.

QUESTION NO: 14

Router TK1 is an ABR that is part of the TestKing OSPF network. TK1 has been configured with the “area range” command. What does this command provide?

- A. It provides the range of areas in the OSPF internetwork.
- B. It provides the range of areas in which this ABR participates.
- C. It provides a summary of networks to be advertised outside the area.
- D. It provides a summary of networks outside the area and inside the AS.

Answer: C

Explanation:

The area range command creates/deletes an area address range entry and optionally specifies whether to advertise the addresses. It is used for summarizing routes into an area.

Syntax: [no] area *area-id* range *address mask* [do-not-advertise]

Description: Address ranges are used to aggregate address ranges from within the area into one single advertisement sent into the adjacent areas, or to prevent the advertisement of networks.

An OSPF area is defined as a list of address ranges. Each address range consists of an address and mask pair. Area range entries are used to aggregate network information before the advertisements are flooded into the backbone. A single summary LSA is originated for each range.

QUESTION NO: 15

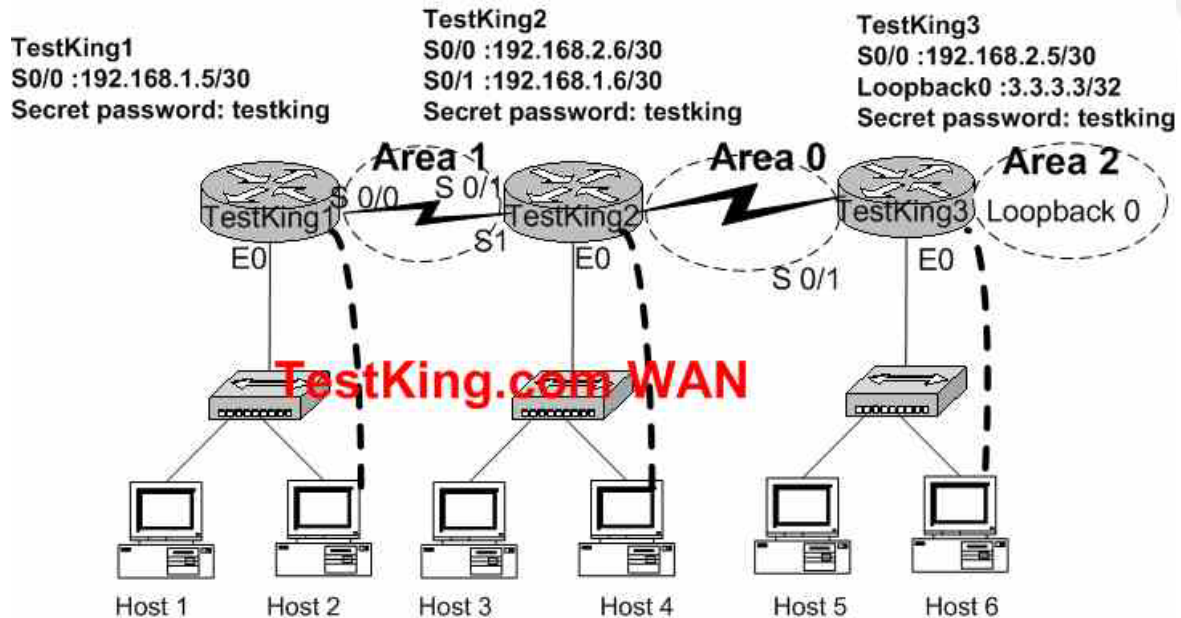
You are the network administrator at TestKing. The TestKing network consists of a single Windows 2000 Active Directory domain testking.com. Windows XP is used on all servers and client computers in the network.

Only Cisco routers are used throughout the company. The routers are named TestKing1, TestKing2, and TestKing3.

The network is characterized by:

- TestKing2 and TestKing3 are configured with OSPF.
- TestKing3's Loopback0 interface is in Area 2.
- TestKing3's S0/1 interface and TestKing2's S0/0 interface are both in Area 0.

The TestKing network is shown in the following exhibit:



You are required to configure the network as follows:

- TestKing2's S0/1 interface in Area 1.
- TestKing1's S0/0 interface in Area 1.
- TestKing1's S0/0 and TestKing2's S0/1, and only these interfaces, should be in Area 1. Use the appropriate mask!
- Area 1 should not receive any external routes.
- Area 1 should not receive any inter-area routes, except the default route.

Answer:

Click on host 2 on testking1

```
enab
confi t
router ospf 1
network 192.168.1.4 0.0.0.3 area 1
area 1 stub
```

Click on host 4 testking2


```

enab
config t
router ospf 1
network 192.168.1.4 0.0.0.3 area 1
network 192.168.2.4 0.0.0.3 area 0
area 1 stub no-summary

```

Note: No need to configure TestKing3

Variation #1

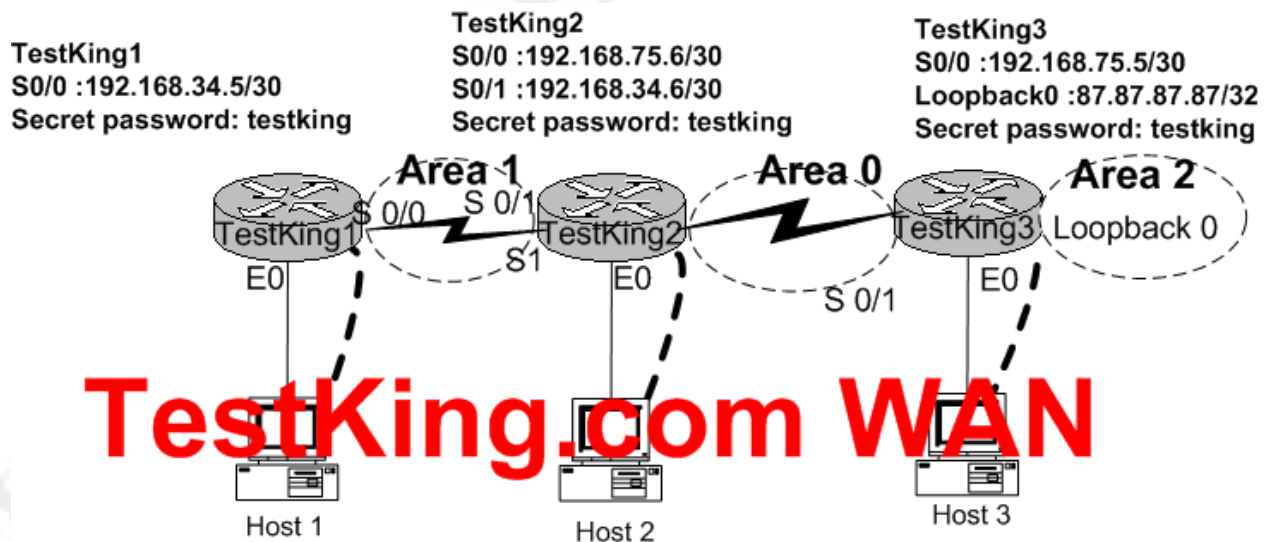
TestKing1:
S0/0: 192.168.18.5/30

TestKing2:
S0/0: 192.168.40.6/30
S0/1: 192.168.18.6/30

TestKing3:
S0/1: 192.168.40.5/30
Loopback0: 200.200.200.200/32

QUESTION NO: 16

The TestKing WAN is depicted below:



You work as a network engineer at TestKing.com. OSPF has been configured on routers TestKing1, TestKing2, and TestKing3. TestKing2's S0/0 interface and TestKing3's S0/1 interface are in Area 0. TestKing3's Loopback0 interface is in Area2.

Your task is to configure the following:

**TestKing1's S0/0 interface in Area 1.
TestKing2's S0/1 interface in Area 1.**

Use the appropriate mask such that ONLY TestKing1's S0/0 and TestKing2's S0/1 could be in Area 1.

Area 1 should not receive any external or inter-area routes (except the default route).

Answer:

Configuration for testking 1

```
Testking1#configt
Testking1(config)#int s0/0
Testking1(config-if)#ip address 192.168.34.5 255.255.255.252
Testking1(config-if)#no shutdown
Testking1(config-if)#exit
Testking1(config)#router ospf 1
Testking1(config-router)#network 192.168.34.4 0.0.0.3 area 1
Testking1(config-router)#area 1 stub
Testking1(config-router)#^Z
```

Configuration for Testking 2

```
Testking2#config
Testking2(config)#interface serial 0/0
Testking2(config-if)#ip address 192.168.75.6 255.255.255.252
Testking2(config-if)# no shutdown
Testking2#config
Testking2(config)#interface serial 0/1
```

Testking2(config-if)#ip address 192.168.34.6 255.255.255.252

Testking2(config-if)# no shutdown

Testking2(config-if)#exit

Testking2(config)#router ospf 1

Testking2(config-router)#network 192.168.75.4 0.0.0.3 area 0

Testking2(config-router)#network 192.168.34.4 0.0.0.3 area 1

Testking2(config-router)#area 1 stub no-summary

Testking2(config-router)#^z

Testking2#copy running-config startup-config

OR for testking2

Testking2#config terminal

Testking2(config)#router ospf 1

Testking2(config-router)#network 192.168.34.4 0.0.0.3 area 1

Testking2(config-router)#area 1 stub no summary

Testking2(config-router)#network 192.168.75.4 0.0.0.3 area 0

Testking2(config-router)#^z

Testking2#copy running-config startup-config

Note, variation:

TestKing1:

S0/0: 192.168.11.5/30

Secret password: testking

TestKing2:

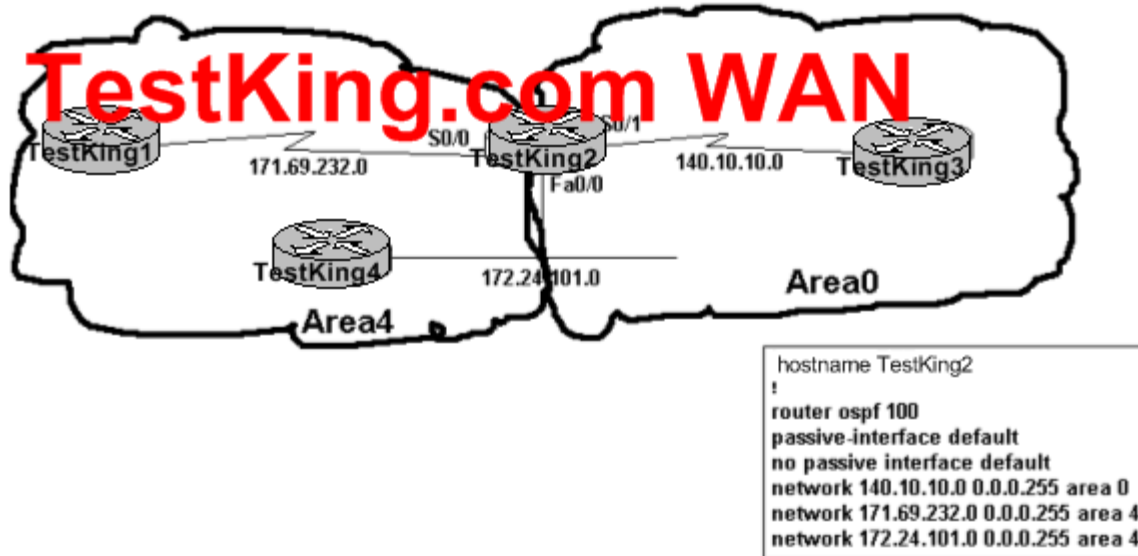
S0/0: 192.168.44.6/30

S0/1: 192.168.11.6/30

Secret password: testking

TestKing3:
 S0/0: 192.168.44.5/30
 Loopback0: 221.221.221.221
 Secret password: testking

QUESTION NO: 17



Which two statements concerning TestKing2 are true about the above configuration?
 (Choose two.)

- A. The interfaces S0/1 and Fa0/0 will not process any routing update that the router receives.
- B. The interfaces S0/1 and Fa0/0 are configured as passive interfaces.
- C. The S0/0 interface will not process any routing update that it receives.
- D. The S0/0 interface will not sending out routing updates.
- E. TestKing3 will hear about the 171.69.232.0 and 172.24.101.0 networks through OSPF update.
- F. TestKing1 and TestKing4 will hear about the 140.10.0.0 network through OSPF update.

Answer: E, F

Explanation:

Section 4: Given an addressing scheme and other laboratory parameters, identify the steps to configure Cisco routers for proper Integrated IS-IS operation (10 questions)

QUESTION NO: 1

In order to configure ISIS to route IP traffic, the following command was entered into router TK1:

```
ip router isis
```

In which IOS configuration mode was this command typed in?

- A. Line configuration mode
- B. Router configuration mode
- C. Global configuration mode
- D. Interface configuration mode
- E. Root configuration mode

Answer: D

Explanation:

To configure an IS-IS routing process for IP on an interface, use the **ip router isis** interface configuration command.

Note: To enable IS-IS, perform the following tasks starting in global configuration mode:

Step 1: router isis

Enable IS-IS routing and specify an IS-IS process for IP, which places you in router configuration mode.

Step 2: net *network-entity-title*

Configure NETs for the routing process; you can specify a name for a NET as well as an address.

Step 3: interface *type number*

Enter interface configuration mode.

Step 4: ip router isis [*tag*]

Specify the interfaces that should be actively routing IS-IS.

Incorrect Answers:

- A, B; C: The **ip router isis** cannot be used in either line, router or Global configuration mode.
- E: This mode does not exist.

QUESTION NO: 2

The TestKing IS-IS network has been configured with mesh groups. Which of the following describe an advantage of utilizing mesh groups?

Leading the way in IT testing and certification tools, www.testking.com

- A. Mesh groups optimize LSP flooding.
- B. Mesh groups keep the routers more secure.
- C. Mesh groups help the routers form adjacencies.
- D. Mesh group speed the flow of data across WAN links.
- E. None of the above.

Answer: A

Explanation: The mesh group feature is a mechanism to reduce flooding of LSPs in nonbroadcast multi-access (NBMA) networks with highly meshed, point to point topologies.

Reference: RFC 2973, IS-IS Mesh Groups

QUESTION NO: 3

When configuring IS-IS, why is it important to manually configure the link costs for each of the associated links within the network?

- A. Because by default, all links have a cost of 10 regardless of the bandwidth.
- B. It is not important because the default provides for optimal routing.
- C. Because there is no default link cost. The link cost must be configured on each interface.
- D. Because by default, all LAN interfaces have a cost of 10 and all WAN interfaces have a cost of 50.

Answer: A

Explanation:

The original IS-IS specification defines four different types of metrics: cost, delay, expense, and error. The Cisco implementation uses cost only. All links use the metric of 10 by default.

The original IS-IS specification defines four different types of metrics. Cost, being the default metric, is supported by all routers. Delay, expense, and error are optional metrics. The delay metric measures transit delay, the expense metric measures the monetary cost of link utilization, and the error metric measures the residual error probability associated with a link.

The Cisco implementation uses cost only. If the optional metrics were implemented, there would be a link-state database for each metric and SPF would be run for each link-state database.

IS-IS uses a single required default metric with a maximum path value of 1024. The metric is arbitrary and typically is assigned by a network administrator. Any single link can have a maximum value of 64, and path links are calculated by summing link values. Maximum metric values were set at these levels to provide the granularity to support various link types while at the same time ensuring that the shortest-path algorithm used for route computation will be reasonably efficient. IS-IS also defines three optional metrics (costs): delay, expense, and error.

Reference: Introduction to Intermediate System-to-Intermediate System Protocol, http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm

Incorrect Answers:

- B: While some routing protocols calculate the link metric automatically based on bandwidth (OSPF) or bandwidth/delay (Enhanced Interior Gateway Routing Protocol [EIGRP]), there is no automatic calculation for IS-IS.
- C: The default link cost is 10.
- D: There is no such distinction between LAN and WAN interfaces.

QUESTION NO: 4

Which of the following best describe the addresses that are used at the network layer of the OSI model?

- A. Internet Protocol address
- B. Media Access Control address
- C. Packet Layer Protocol address
- D. Network Service Access Point address
- E. Authority and Format Identifier address

Answer: D

Explanation:

The term "network address" is used to refer to the Network Service Access Point (NSAP) at which the OSI Network Service is made available to a Network Service user by the Network Service provider.

Incorrect Answers:

- A: IP is a protocol that is based on the OSI model.
- B: MAC addresses operate at a lower level in the OSI model.
- C, E: Do not apply.

Reference: RFC 941, Addendum to the Network Service Definition Covering Network Layer Addressing

QUESTION NO: 5

The OSI IS-IS NSAP address 47.040C.0061.040C.0056.0D12.00 is applied to the Ethernet interface of router TK1. What is the area ID of the address?

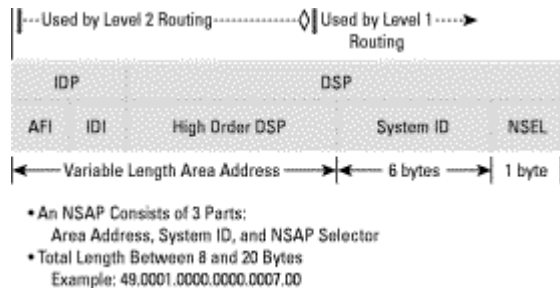
- A. 00
- B. 47
- C. 47.040C
- D. 47.040C.0061
- E. 040C.0056.0D12
- F. None of the above

Answer: D

Explanation:

An NSAP address (figure 7) has two major parts: the initial domain part (IDP) and the domain specific part (DSP) (Figure 7). The IDP consists of a 1-byte authority and format identifier (AFI) and a variable-length initial domain identifier (IDI), and the DSP is a string of digits identifying a particular transport implementation of a specified AFI authority. Everything to the left of the system ID can be thought of as the area address of a network node.

Figure 7 The NSAP address



Incorrect Answers:

A: 00 is the SEL.

B, C: 47 and 47.040C is just a part of the Domain

E: 040C.0056.0D12 is the SystemID.

Reference:

http://www.cisco.com/en/US/tech/tk365/tk381/technologies_white_paper09186a00800a3e6f.shtml

QUESTION NO: 6

Which of the answer choices below represents a valid private, locally administered NSAP address?

- A. 39.0f01.0002.0000.0c00.1111.00
- B. 48.0f01.0002.0000.0c00.1111.00
- C. 49.0004.30ac.0000.3090.c7df.00
- D. 52.0f01.0002.0000.0c00.1111.00

Answer: C

Explanation:

The private NSAP addresses have AFI beginning with 49, as these are locally administered addresses..

AESA Network Service Access Point (NSAP) ATM Addresses

There are 3 types of private ATM addresses:

- **NSAP encoding format for E.164 addresses** - The authority and format identifier (AFI) is 45. These addresses are used in establishing ISDN calls by public networks, and they are normally used in public telephony.
- **Data Country Code (DCC) AESA** - The AFI is 39. These addresses are to be used in public networks. For example, the initial domain identifier (IDI) value 0x84.0f identifies the United States.
- **International Code Designator (ICD) AESA** - The AFI is 47. These addresses are used in private organizations, and the ICD field indicates the code set or organization. Cisco uses by default ICD addresses.

Incorrect Answers:

A: Addresses beginning with 39 is administered by the ISO for ISO Data Country Code.

B, D: 48 and 52 are unassigned AFI addresses.

Reference:

http://www.cisco.com/en/US/tech/tk39/tk49/technologies_tech_note09186a00800c9761.shtml

QUESTION NO: 7

You are a guest lecturer at the TestKing Academy and are preparing a lesson on the Cisco IS-IS. Which of the following points can you include in your discussion of the Cisco IS-IS NSAP address System IDs? (Select three)

- A. System IDs can vary in size within a domain.
- B. The System ID identifies a node in an IS-IS network.
- C. The System ID must be unique within a Level-1 area.
- D. The System ID must be unique within a Level-2 area.
- E. The System ID must be the MAC address of the router.

Answer: B, C, D**Explanation:**

B: Each system ID within an area must be unique. It is used to identify a IS-IS node.

C: All Level 1 routers and hosts in an area must have an NSAP with the same area address.

D: Level 2 routers advertise their own area addresses (NSAP) to the other Level 2 routers in the backbone.

Incorrect Answers:

A: All ISs and ESs in a routing domain must have system IDs of the same length. Furthermore, Cisco implements a fixed length of 6 bytes for the system ID.

E: There are several techniques for creating unique system IDs

* Start numbering 1, 2, 3, 4, and so on.

* Use Media Access Control (MAC) addresses.

* Convert and use the loopback IP address: 192.168.11.1 --> 192.168.011.001--> 1921.6801.1001.

Reference: Introduction to Intermediate System-to-Intermediate System Protocol

http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm

QUESTION NO: 8

When comparing and contrasting the NET and NSAP used in IS-IS routing, which of the following statements is true?

- A. A NET is an NSAP address with the N-selector byte set to 00.
- B. Network Entity Titles do not have to start and stop on byte boundaries.
- C. The System ID field of the NSAP address does not uniquely identify a node.
- D. To identify a Domain, an NSAP address must be used because a NET can only identify an Area.
- E. A NET is a special version of an NSAP address restricted to 8 bytes for the Area Address, System ID and the N-Selector byte.

Answer: A

Explanation:

NETs and NSAPs are ISO addresses. The differences between the NET and NSAP addresses are subtle. The NET address is the address of the host, where the value in the NSEL field is set to 0x00. Therefore, there is no upper-layer protocol identified within the host. With no application identified within the end host, the packet can be routed to the destination, but it cannot be handed off to a process after it has been delivered. However, routers do not have upper-layer protocols to identify because they are transitory ISs. Therefore, the NSAP of the router is referred to as a NET because the NSEL field is set to 00. The NSAP is the full ISO address. It not only defines the area and destination host within the area, but also specifies where to send the incoming packet after it has reached the host. The NSEL field at the end of the ISO address specifies the upper-layer protocol and is similar to the Protocol field of the IP header.

Network Entity Title (NET) In IS-IS, this is the ISO address of the system, but not to the process destination within the system. The NET describes both the area and system ID of a system in the IS-IS network but excludes the NSEL, which is set to 0x00. If the NSEL identifies the process within the system, the ISO address is called the “NSAP address.”

Network Service Access point (NSAP) In IS-IS, this describes a service at the network layer to which the packet is to be directed. The NSAP is the NET address with the NSEL field set to a positive value, a value other than 0x00.

Network entity title (NET)

The NET describes both the area and system ID of a system in the IS-IS network but excludes the NSEL, which defines the NSAP address of the system.

Network service access point (NSAP)

Describes a service at the network layer to which the packet is to be directed. The NSAP is the NET address with the SEL field set to a value other than 0x00.

An ISO Address

IDP		DSP		
AFI (1 octet)	IDI	High Order DSP	System ID (1-8 octets)	NSEL (1 octet)
AREA			ID	SEL

QUESTION NO: 9

What are the basic configuration steps needed to enable IS-IS?

- Configure the **net system-id** command under **router isis** and enable IS-IS on each interface with the **ip router isis** command.
- Configure the **network net-id** command(s) under **router isis** and enable IS-IS on each interface with the **ip router isis** command.
- Configure the **network net-id** command(s) and the **is-type level-1-2** command under **router isis**.
- Configure the **net system-id** and the **network net-id** commands under **router isis**.
- Configure the **net system.-id** and the **network net-id** commands under **router isis** and enable IS-IS on each interface with the **ip router isis** command.

Answer: A

Explanation:

In order to enable IS-IS for IP on a Cisco router and have it exchange routing information with other IS-IS enabled routers, you must perform the following two tasks:

- Enable the IS-IS process and assign area
- Enable IS-IS for IP routing on an interface

The sample configuration below configures an IS-IS router with the following parameters:

- Area 49.0001
- Level 1 (L1) and Level 2 (L2) routers (this is the default unless otherwise specified)
- No optional parameters
- Running IS-IS for IP only
- Loopback interfaces (loopbacks are advertised by IS-IS, not IS-IS enabled)

Router TK1

```
!
interface Loopback0
ip address 172.16.1.1 255.255.255.255

!--- Creates loopback interface and assigns
!--- IP address to interface Loopback0.

!
interface Ethernet0
```

```

ip address 172.16.12.1 255.255.255.0
ip router isis

!--- Assigns IP address to interface Ethernet0
!--- and enables IS-IS for IP on the interface.

!
router isis
passive-interface Loopback0
net 49.0001.1720.1600.1001.00
!

!--- Enables the IS-IS process on the router,
!--- makes loopback interface passive
!--- (does not send IS-IS packets on interface),
!--- and assigns area and system ID to router.

```

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080093f38.shtml

QUESTION NO: 10

An administrator has decided to configure Integrated IS-IS on the network.

What must the administrator do to enable the interfaces to distribute IP information using IS-IS?

- A. The networks configured on each interface must be associated to the IS-IS routing protocol using the **network** router configuration command.
- B. Each interface must be enabled to support IS-IS with the **ip router isis** interface configuration command.
- C. All configured IP networks will automatically be enabled when IS-IS is configured with the **router isis** global configuration command.
- D. Each network must be identified with the **net** interface configuration command.

Answer: B

Explanation:

The preparation for configuring any routing protocol requires a thorough understanding of the network topology and a coherent addressing scheme. When you have these, the basic configuration of the Integrated IS-IS is as follows:

642-891

- 1) Enable the routing process Integrated IS-IS with the router isis command.
- 2) Configure the Network Entity Title (NET) address, thus assigning the area with the net network-address router subcommand.
- 3) Enable Integrated IS-IS for IP on the relevant interfaces with the IP router isis interface subcommand.

Reference: Cisco Press 642-801 p.397

Section 4: Identify the steps to select and configure the different ways to control routing update traffic (22 questions)

QUESTION NO: 1

You have a router that's been configured with multiple IP routing protocols, and you're interested in checking out if and how inbound and outbound routing updates are being filtered. Which one of the following commands would list these filters?

- A. show ip
- B. show ip route
- C. show ip protocols
- D. show ip interface
- E. show protocol filters

Answer: C

Explanation:

The command **show ip protocols** command is used to display the parameters and current state of the active routing protocol process. The information is presented on a routing protocol basis and includes applied inbound and outbound filters.

Incorrect Answers:

- A: Show ip is an incomplete command.
- B: The show ip route command displays active routes, not information on filters.
- D: The show ip interface command lists a summary of an interface's IP information and status. However, it does not list the filters applied on a routing protocol basis.
- E: This is an invalid command.

QUESTION NO: 2

You have a low speed serial connection that is being used only as a backup to a frame relay network. Which routing method could you use on your internal routers to minimize the bandwidth wasted for frivolous routing updates?

- A. Use a distance vector routing protocol.
- B. Use private IP addresses.
- C. Use dial-on-demand routing.
- D. Use route summarization.
- E. Use a routing protocol that tolerates route flapping.

Answer: D

Explanation:

An advantage to using route summarization in a large complex network is that it can isolate topology changes from other routers. That is, if a specific link in the domain were flapping (going down and up rapidly), the summary route would not change, so no router external to the domain would need to keep modifying its routing table due to this flapping activity.

Reference: Building Scalable Cisco Networks (Cisco Press) page 76

QUESTION NO: 3

You want to decrease the amount of routing query traffic across your EIGRP network. Which of the following is an effective means of controlling EIGRP query traffic?

- A. Route summarization.
- B. Configuring route filters.
- C. Using a hierarchical addressing scheme.
- D. Establishing separate autonomous systems.
- E. None of the above.

Answer: A

Explanation:

The best solution to control EIGRP queries is to reduce the range of queries. The most effective method to restrict the range of queries is the establishment of route summarization boundaries. Since route summarization decreases the total number of routing table entries, fewer routing queries are needed.

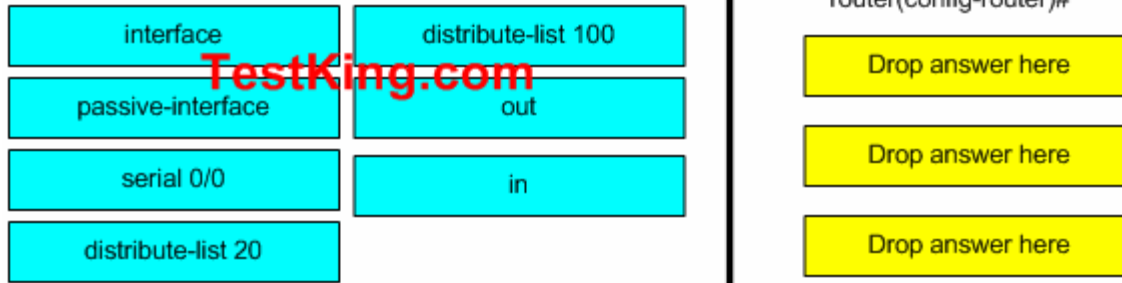
QUESTION NO: 4

You are a systems administrator of one of the TestKing EIGRP networks. Your goal is to suppress TestKing's router from sending updates out of its Serial 0/0 interface without compromising its ability to form and maintain neighbor adjacencies on the same interface.


The router has already been configured as follows:

```
Router(config)#access-list 20 deny any
Router(config)#access-list 100 permit ip any any
Router(config)#router eigrp 1
```

Drag the correct answers on the left to the correct command line space on the right, to finish configuring the router.



Answer:



Explanation:

We do not want to allow the routing updates along serial 0/0 to go out. This distribute list will deny all IP networks from being advertised out the interface.

QUESTION NO: 5

Your OSPF router has one serial interface, and one Ethernet LAN interface. The sub-interface is configured in the following manner:

```
interface serial 0.122 point-to-point
 ip address 192.168.1.1 255.255.255.0
 encapsulation frame-relay
 frame-relay interface-dlci 122
```

Your want to allow hosts on your LAN to send and receive data, but you don't want routing traffic to go through that interface. Which of the following commands should you use to complete your configuration?

- A. interface serial 0.122 point-to-point
passive-interface ethernet 0
- B. interface ethernet 0
ip address 192.168.12.1 255.255.255.0

```

    passive-interface
C. router ospf 172
    area 1 nssa
    network 192.168.1.0 0.0.0.255 area 0
    network 192.168.12.0 0.0.0.255 area 1
D. router ospf 172
    passive-interface ethernet 0
    network 192.168.1.0 0.0.0.255 area 0
    network 192.168.12.0 0.0.0.255 area 1

```

Answer: D

Explanation:

We use the passive-interface command to configure the ethernet interface to be passive. The passive-interface router configuration command is used to disable sending routing updates on an interface.

Incorrect Answers:

- A: We are not configuring the serial interface. Furthermore, the passive-interface command is a router configuration command, not an interface configuration command.
- B: We cannot use the passive-interface command like this.
- C: We should configure the Ethernet interface as passive, not the area as a not-so-stubby area (NSSA).

QUESTION NO: 6

You have a named RouterTestK running EIGRP that has already been configured with the following:

```

RouterTestK(config)#access-list 30 deny any
RouterTestK(config)#access-list 40 permit ip any
RouterTestK(config)#router eigrp 1

```

To complete your configuration you want to configure the serial 0/0 interface to prevent routing updates from going out that interface while still allowing the formation and upkeep of neighbor adjacencies on that interface. Drag the correct command phrase from the bottom and place it in the right command line sequence above.

RouterTestK(config-router)#

Drop keyword here

Drop keyword here

Drop keyword here

Select from these

in

distribute-list 40

distribute-list 30

out

serial 0/0

interface

passive interface

Answer:

RouterTestK(config-router)#

istribute-list 30

out

serial 0/0

Select from these

in

istribute-list 40

interface

passive interface

Explanation:

We use distribute list 30, which stops IP traffic, on outward traffic on the serial 0/0 interface. RouterTK continues receiving routing updates from its neighbor, but the distribute-list prevents routes from being advertised out of serial 0. Furthermore, neighbor adjacencies are allowed to be formed between RouterTK and its neighbor on serial 0/0.

Incorrect Answers:

passive interface serial 0/0:

On EIGRP, passive interface causes the router to stop sending and receiving hello packets. This will prevent the interface from maintaining neighbor adjacencies.

distribute-list 40 out serial 0/0:

We must stop outgoing traffic, not allow it.

distribute-list 30 in serial 0-0:


We must stop outgoing traffic, not incoming traffic.

Reference: How Does the Passive Interface Feature Work in EIGRP?

<http://www.cisco.com/warp/public/103/16.html>

QUESTION NO: 7

The TestKing network is displayed below:



You need to configure a static route so that users on the corporate network can reach the 172.27.6.0/24 LAN. Based on the diagram above, which of the following is the correct way to do this?

- A. TESTKING1 (config) #ip route 172.27.6.0 255.255.255.0 172.17.8.2
- B. TESTKING2 (config) #ip route 172.27.6.0 255.255.0.0 172.17.8.2
- C. TESTKING2 (config) #ip route 172.27.6.0 255.255.255.0 172.17.8.2
- D. TESTKING1 (config) #ip route 172.27.6.0 255.255.0.0 172.17.8.1
- E. TESTKING2 (config) #ip route 172.27.6.0 255.255.255.0 172.17.8.1

Answer: E

Explanation:

Only choice E provides the correct network subnet mask and next hop IP address.

IP route *prefix mask {address|Interface} [distance] [tag tag] [permanent]*

Prefix 172.27.6.0 mask 255.255.255.0 address 172.17.8.2

Address – The IP address of the next hop router that can be used to reach that network.

Reference: Building Scalable Cisco Networks (Cisco Press) page 464

QUESTION NO: 8

The TestKing network consists of a main office and a single remote office. You need to configure the remote office in using the following guidelines:

- The main office must learn all of its routes from the regional office.
- The regional office must not learn routes from the main office.
- The most scalable solution should be used.

Which of the following answer choices best describe what should be done to satisfy these requirements?

- A. Configure static routes pointing to the network behind the central office router
- B. Configure a default route pointing to the networks behind the central office router
- C. Make the interface that is connected to the central office a passive interface to block incoming updates
- D. Enable route update filtering on the interface that is connected to the central office to block incoming updates
- E. None of the above will work

Answer: D

Explanation:

The only viable solution here is to enable the main office to dynamically learn about all of the regional office routes, while still not receiving any routes is by enabling incoming route filters.

Incorrect Answers:

- A, B: Configuring static routes or default routes on the regional office will not enable the main office to learn about any of the regional office routes.
- C: A passive interface would prevent all route updates to the main office. Passive interfaces do not filter incoming routing updates, so the regional office could still learn routes from the main office.

Note: Configuring the main office using a passive interface would satisfy the conditions here, but this question asks us how to configure the regional office, not the main office.

QUESTION NO: 9

If a router already has a route to the 172.27.0.0/16 network in the routing table, which of the following commands would cause EIGRP to flag that network as the candidate default route?

- A. `ip default-network 172.27.0.0`
- B. `ip route 0.0.0.0 0.0.0.0 172.27.0.0`
- C. `ip default-network 0.0.0.0`
- D. `default-information originate`
- E. `ip classless`

Answer: A

Explanation:


Unlike the `ip default-gateway` command, you can use `ip default-network` when `ip routing` is enabled on the Cisco router. When you configure `ip default-network` the router considers routes to that network for installation as the gateway of last resort on the router.

Reference:

http://www.cisco.com/en/US/tech/tk365/tk554/technologies_tech_note09186a0080094374.shtml

QUESTION NO: 10

Two TestKing routers are connected together as shown in the diagram below.



The networks in the above diagram are all configured with RIP. On router TK1, the following configuration change was made:

```
TK1(config)#ip default-network Network TK3
```

Which router will receive a default route as a result of this change?

- A. Router TK1 only-
- B. Router TK2 only.
- C. Both routers.
- D. Neither router.

Answer: C

Explanation:

The **ip default-network** command is used as a method of distributing default route information to other routers. When running RIP, you can create the default route by using the **ip default-network** command. If the router has a directly connected interface onto the network specified in the **ip default-network** command, RIP will generate (or source) a default route to its RIP neighbor routers.

Unlike the **ip default-gateway** command, you can use **ip default-network** when **ip routing** is enabled on the Cisco router. When you configure **ip default-network** the router considers routes to that network for installation as the gateway of last resort on the router. In this example, it will create a default network for TK1, since it knows about the route to TK3, and TK1 will advertise this default route to TK2.

Reference: <http://www.cisco.com/warp/public/105/default.html#ipnetwork>

QUESTION NO: 11

Router TK1 is configured as follows:

```
interface serial 0
    ip address 185.64.1.1 255.255.255.0
!
interface ethernet 0
```

Leading the way in IT testing and certification tools, www.testking.com

```

ip address 15.10.10.1 255.255.255.0
!
router rip
  network 15.10.10.0
!
ip route 0.0.0.0 0.0.0.0 serial0

```

What result would the above commands accomplish?

- A. RIP updates are sent and received on interface serial0 of the router.
- B. A default route is sent to neighbors on interface serial0 of the router.
- C. A default route is sent to neighbors on interface ethernet0 of the router.
- D. RIP updates are sent and received on interfaces serial0 and ethernet0 of the router.
- E. None of the above.

Answer: E

Explanation:

Since RIP is configured for only the network on the Ethernet segment, RIP updates will only be sent and received on this network, not on the serial interface. In addition, static routes are not sent via RIP, unless the “redistribute static” command is configured. This includes default static routes.

Note: Had the configuration command “redistribute static” been placed under the RIP routing process, then choice C would have been correct.

QUESTION NO: 12

Which of the following commands ensures that permanently created static route entries are dynamically added into the routing process?

- A. inject static
- B. inject permanent
- C. redistribute all
- D. redistribute static

Answer: D

Explanation:

If you define a static route to an interface that is not one of the networks defined in a network command, no dynamic routing protocols will advertise the route unless a redistribute static command is specified for these protocols.

Incorrect Answers:

A: The inject static is used to configure legacy DECnet systems. It is not a valid IP routing command.

B: There is no such command.

C: The command is not valid.

Reference: Cisco, Configuring IP Routing Protocol-Independent Features

QUESTION NO: 13

Which IOS command below would be useful if you had to create a static list of the other routers in your NBMA cloud?

- A. network
- B. neighbor
- C. ip route
- D. router ospf

Answer: B

Explanation:

The **neighbor** command is used to statically map the IP address of the neighbor. This is useful in multi-point networks where the neighbors may not be able to be found dynamically due to the nature of the NBMA network.

QUESTION NO: 14

The TESTKING router had the following configuration line added:

TESTKING(config)#ip route 172.27.6.0 255.255.255.0 s0/0

What information can be gathered from this configuration change? (Select two)

- A. This is a route to a public network.
- B. There is only one path to this network from TESTKING.
- C. This is a route to interface s0/0 on the next hop router.
- D. Packets destined for this network are sent via interface s0/0 on TESTKING.
- E. Packets destined for this network enter router TESTKING through interface s0/0.

Answer: B, D

This command places a static route to the 172.27.6.0 network. All traffic destined to this network from this router will be sent across the s0/0 interface.

Incorrect Answers:

A: The 172.27.6.0 IP address range is part of the private IP address space as specified by RFC 1918.

C: The serial 0/0 indicates the interface to use going out, not the next hop.

E: Packets destined for this network will exit through serial 0/0, not enter through it.

Reference: Building Scalable Cisco Networks (Cisco Press) page 464

QUESTION NO: 15

The TestKing2 router is running RIP version 2 and is configured as shown below:

```
TestKing2 (config) #router rip
TestKing2 (config-router) #version 2
TestKing2 (config-router) #no auto-summary
```

Based on this information, which of the following is true?

- A. Subnets are summarized at the network boundary.
- B. Subnets are advertised across network boundaries.
- C. Subnet mask information is not passed in the routing updates.
- D. Subnets are made discontinuous.
- E. None of the above

Answer: B

Explanation:

By default, RIP version 2 summarizes subnets across network boundaries. To restore the default behavior of automatic summarization of subnet routes into network-level routes, use the **auto-summary** router configuration command. To disable this feature and transmit subnet routing information across classful network boundaries, use the **no** form of this command.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1826/products_command_summary_chapter09186a00800d9c56.html

QUESTION NO: 16

What would you find in an EIGRP routing table if route summarization is configured on the routers Serial0 interface, and it summarizes routes learned from its Ethernet0 interface?

- A. A summary route pointing to the Null0 interface.
- B. A summary route pointing to the Serial0 interface.
- C. A summary route pointing to the Ethernet0 interface.
- D. A summary route pointing to the Loopback0 interface.

Answer: A

Explanation:

With EIGRP, when summarization is configured on any interface, the IOS immediately creates and installs the summarized route pointing to null 0 in the routing table.

QUESTION NO: 17

With regard to controlling routing updates with the use of route filtering, which of the following statement is true?

- A. Only inbound routes can be filtered.
- B. Only outbound routes can be filtered.
- C. Routes to be filtered are selected by using distribute lists.
- D. Routes to be filtered are selected using only extended access lists.
- E. Routes to be filtered are selected by using the distribute-group command.

Answer: C**Explanation:**

The syntax for the distribute-list in/out command is:

distribute-list *access-list-number* in/out [*interface-name*]

where *access-list-number* is the standard IP access-list against which the contents of the incoming or outgoing routing update are matched. The [*interface-name*] argument is optional and specifies the interface on which the update is expected. It is important to note that the access-list referred to in *access-list-number* is applied to the contents of the update, not to the source or destination of the routing update packets. The router decides whether or not to include the contents in its routing table based on the access-lists. For example:

```
access-list 1 permit 1.0.0.0 0.255.255.255
router rip
distribute-list 1 in
```

```
!--- The distribute-list command is given
!--- under the router configuration mode.
```

Incorrect Answers:

- A, B: Distribute lists can be used to filter traffic based on incoming as well as outgoing routes.
- D: Standard access-lists are used with distribute lists, not extended access lists.
- E: This is an invalid command. The correct syntax is “distribute-list”

QUESTION NO: 18

Router TK1 has been configured to filter routes. Which of the following are reasons to control routing updates via route filtering? (Choose three)

- A. To hide certain networks from the rest of the organization.
- B. For easier implementation.
- C. To control network overhead on the wire.
- D. For simple security.
- E. To prevent adjacencies from forming.

Answer: A, C, D

Explanation:

Route filtering is used to filter out routing updates from other parts of the network, making certain networks unreachable from other parts of the organization. This can be done to simply decrease the amount of traffic overhead associated with routing updates over parts of the network, or for security reasons.

Incorrect Answers:

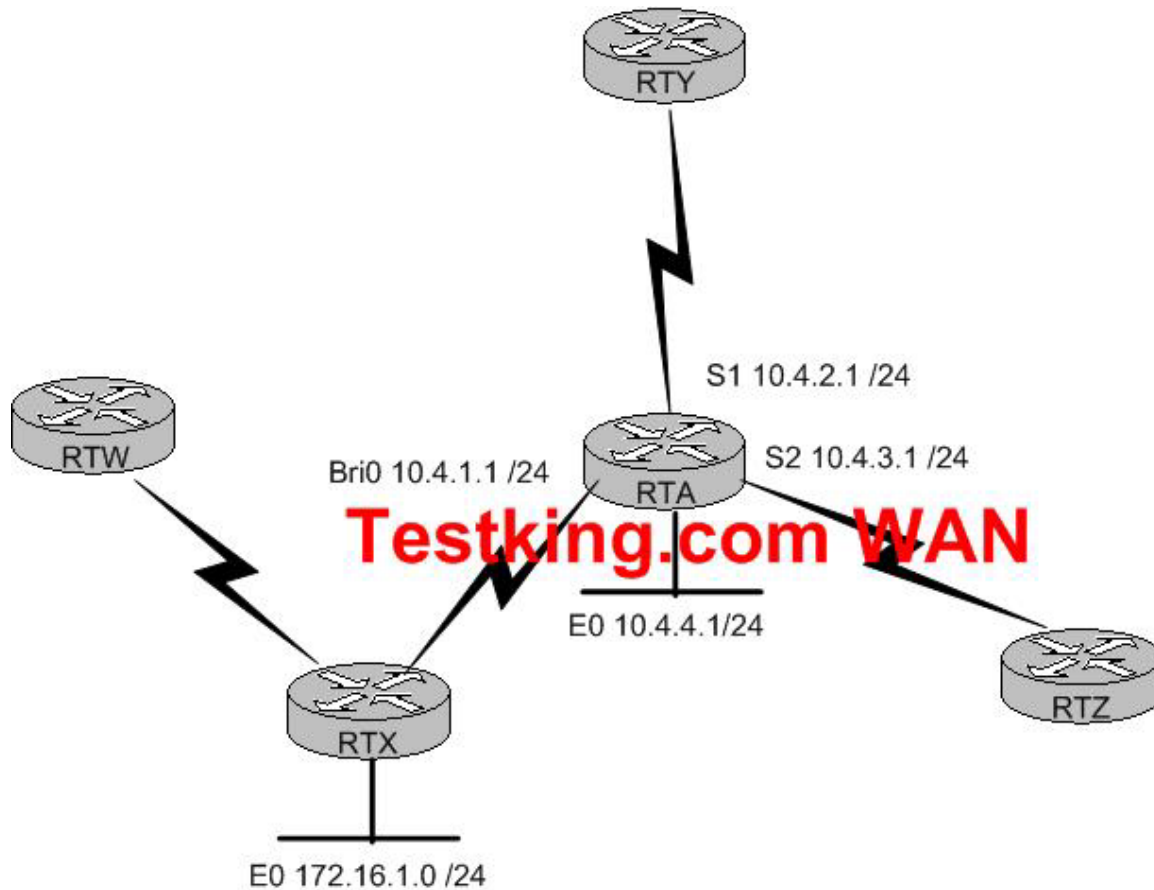
B: Route filtering requires a more complex router configuration, and can also lead to difficult troubleshooting.

E: Route filtering will only block certain routes from incoming or outgoing, but the routing protocol adjacencies will still be formed between neighboring routers.

QUESTION NO: 19

You are the network engineer at TestKing. The router topology for the TestKing network is shown in the following graphic:

642-891



Router RTA is configured as follows:

```
RTA(config)#router rip
RTA(config-router)#network 10.0.0.0
RTA(config-router)#distribute-list 44 in interface BRI0
RTA(config-router)#exit
RTA(config)#access-list 44 deny 172.16.1.0 0.0.0.255
RTA(config)#access-list 44 permit any
```

What are the effects of this RIP configuration on router RTA? (Choose two)

- A. No routing updates will be sent from interface BRI0 on router RTA to router RTX.
- B. Router RTA will not advertise the 10.0.0.0 network to router RTX.
- C. The route network 172.16.1.0 will not be entered into the routing table on router RTA.
- D. User traffic from the 172.16.1.0 network is denied by access-list 44.
- E. The routing table on router RTA will be updated with the route to router RTW.

Answer: C, E

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

Answer C is correct as the distribute list calls ACL 4 which denies network 172.16.1.0. Answer E is correct because RTW networks are not being blocked.

Incorrect Answers:

A, B: Only incoming routes are filtered, not outgoing routes. All known routing entries will be advertised from router A.

D: This is incorrect because this ACL is being called by a distribute-list which only affects route table updates, not user traffic.

QUESTION NO: 20

Router TK1 is configured as follows:

```
interface serial 0
  ip policy route-map force
    ip route-cache policy
route-map force permit 10
  match ip address 1
  set ip next-hop 172.20.16.5
```

```
access-list 1 permit 10.0.0.5 0.0.0.0
```

Based on the information above, which two of the following statements correctly describe what happens when a packet is received on serial 0 of TK1? (Choose two)

- A. If a packet is destined to 10.0.0.5, it is a candidate for fast-switch policy routing.
- B. If a packet was sourced from 10.0.0.5, it is a candidate for fast-switch policy routing.
- C. The route-map policy is incomplete because it does not account for traffic that does not match 10.0.0.5.
- D. If a packet was sourced from 10.0.0.5, it will be sent to 172.20.16.5, regardless of the desired destination IP address.

Answer: B, D

Explanation:

Policy routing is used to override the normal routing behavior of a router. In this example, the router is configured to send all traffic that matches access-list 1 to the router with the next hop IP address of 172.20.16.5.

“ip policy route-map map-tag”

map-tag is the name of the route map to use for policy routing. This must match a map tag specified by a **route-map** command.

- When policy routing is configured, turn on the fast switching with this interface command:
`ip route-cache policy`
- Fast-switched policy routing supports all of the **match** commands and most of the **set** commands, except for the following restrictions:
 - The **set ip default** command is not supported.
 - The **set interface** command is supported only over point-to-point links, unless a route-cache entry exists using the same interface specified in the **set interface** command in the route map. In addition, when process switching, the routing table is consulted to determine a path to the destination. During fast switching, the software does not make this check because fast switching is a cache of the process switch lookup. Instead, if the packet matches, the software blindly forwards the packet to the specified interface. This is a similar situation to the one described in reference to load balancing earlier.

QUESTION NO: 21

Which external EIGRP administrative distance should be used in order to allow external EIGRP routes to be more preferred over IS-IS routes but less preferred over internal EIGRP routes?

- A. 89
- B. 114
- C. 119
- D. 121
- E. 131
- F. 141

Answer: B

Explanation:

To meet the conditions of this task the Administrative Distance should be more than 90 and less than 115. B is our only valid option.

The following table displays the default AD for all routing protocols:

Default Administrative Distances	
Route Source	Default Distance
Connected interface	0
Static route	1
Enhanced IGRP summary route	5
External BGP	20

Internal Enhanced IGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
EGP	140
EIGRP external route	170
Internal BGP	200
Unknown	255

QUESTION NO: 22

With the exception of EIGRP and BGP, which type of administrative distance is changed by the distance router configuration command?

- A. External
- B. Internal
- C. Local
- D. Default

Answer: B

Explanation:

To ensure that the optimal path is chosen, it is sometimes necessary to change the administrative distance to make it less favorable. The command structure is protocol-dependent, in that EIGRP requires a separate command. The following command syntax is used for EIGRP:

```
Router(config)#distance eigrp internal-distance external-distance
```

Internal-distance is administrative distance for EIGRP internal router. These are routes learned from another entity within the same autonomous system.

Section 5: Identify the steps to configure router redistribution in a network (22 questions)

QUESTION NO: 1

A TestKing router was configured with BGP as shown below:

```
router bgp 6500
  redistribute static

ip route 164.20.0.0 255.255.0.0 null 0
```

Based on this information, is this a recommended router configuration?

- A. Yes. It allows BGP to advertise the 164.20.0.0 /16 network.
- B. Yes. It results in all traffic for all subnets of 164.20.0.0 being dropped at this router.
- C. No. Cisco prefers that you use the `aggregate-address` command to distribute IGP routes into BGP.
- D. Yes. Cisco prefers this method of redistributing IGP routes into BGP over using the `network` command.

Answer: C

Explanation:

Redistribution of static routes configured to the null 0 interface into BGP is done to advertise aggregate routes rather than specific routes from the IP table. However, Cisco recommends the use of the `aggregate-address` command instead.

QUESTION NO: 2

You are a network analyst at TestKing and are currently in the process of analyzing two autonomous systems, each running a different routing protocol, but connected between each other with redundant paths. The junior administrator wants to prevent routing loops between the two autonomous systems and asks you if there's an IOS feature that could help him. How would you answer?

- A. Route filtering.
- B. Passive interfaces.
- C. Static redistribution.
- D. Two-way redistribution.

Answer: A

Explanation:

Multiple autonomous systems or routing domains can share route information through the redistribution process. Proper implementation of redistribution requires route filters to prevent feedback loops from forming. It is strongly recommended that redistribution between multiple ASs and multiple routing protocols be accompanied by route filters.

Reference: CCNP #640-503 Building Scalable Cisco Networks (Cisco Press), More EIGRP Scalability Rules

QUESTION NO: 3

Routing loops are becoming a problem on a particular network. What can an administrator do to reduce the occurrence of routing loops on a redistributed network?

- A. Use multiple default gateways for redundancy.
- B. Use one-way redistribution for greater stability.
- C. Use two-way redistribution for greater stability.
- D. Use overlapped routing protocols for redundancy.
- E. None of the above.

Answer: B

Explanation:

One-way redistribution would help avoiding the routing loops problem.

Incorrect Answers:

- A: Multiple default gateways would increase the risk of routing loops.
- C: Two-way redistribution would increase the risk of routing loops.
- D: Using several routing protocols would increase the risk of routing loops.

QUESTION NO: 4

You are in the process of redistributing EIGRP into another routing protocol on your network. Which of the following commands would you execute if your intention was to use altered administrative distance (AD) parameters for redistributing EIGRP into the second routing protocol?

- A. `default-metric eigrp metric`
- B. `distance eigrp administrative-weight`
- C. `distance eigrp internal-distance external-distance`
- D. `distance eigrp external-distance internal-distance`

Answer: C

Explanation:

The distance eigrp command is used to allow the use of two administrative distances---internal and external---that could be a better route to a node.

Syntax: `distance eigrp internal-distance external-distance`

Note: Use the **distance eigrp** command if another protocol is known to be able to provide a better route to a node than was actually learned via external Enhanced IGRP or if some internal routes should really be preferred by Enhanced IGRP.

Incorrect Answers:

A: Incorrect usage of the default-metric command. Furthermore, this command would not be of use here.

B, D: This is the incorrect use of the **distance eigrp** command.

QUESTION NO: 5

Which three of the following IOS commands could an administrator use to verify route redistribution? (Select three)

- A. debug
- B. traceroute
- C. show summary
- D. show ip route
- E. ipconfig

Answer: A, B, D

Explanation:

A: The debug command can be used to debug redistribution.

B: We can verify connectivity, and the presence of a route, with the traceroute command.

D: We can verify that the routes have been redistributed with the show ip route command. The routes will be shown.

Incorrect Answers

C: The show summary command displays a summary of relationships among owners, content rules, and services. It is not of help in this scenario.

E: This is a command used by PC hosts to verify that the IP stack is working properly.

QUESTION NO: 6

A TestKing OSPF router has been configured in the following manner:

```
router ospf 1
  redistribute eigrp 1 metric 33 subnets
```

Based on this information, what function does the 33 parameter in the `redistribute` command serve?

- A. It specifies the metric cost to be applied to the redistributed routes.
- B. It specifies the administrative distance on the redistributed routes.
- C. It specifies the metric limit to 33 subnets in each OSPF route advertisement.
- D. It specifies the process-id for the pseudo process that injects the EIGRP routes into OSPF.

Answer: A

Explanation: It specifies the metric cost to be applied to the redistributed routes. In this example, a metric of 33 will be applied to all of the EIGRP routes when redistributed into OSPF. The “subnets” keyword specifies that subnet mask information is to be preserved during this process.

Reference: Redistributing Routing Protocols

<http://www.cisco.com/warp/public/105/redist.html#examples>

QUESTION NO: 7

If you had a static route configured on a router, and that static route had to be advertised to other routers in the network; which one of the following statements would be true?

- A. The router automatically advertises static routes to RIP routers.
- B. You should configure redistribution using the **redistribute** command.
- C. You should enable static advertisements using the **static routes advertise** command.
- D. You should include the static route in a distribution list using the **distribute-list** command.

Answer: B.

Explanation:

If you want a router to advertise a static route in a routing protocol, you will need to redistribute it into a dynamic routing protocol.

Incorrect Answers:

A: Static routes are not automatically redistributed when they are configured. Static routes must be redistributed manually.

C: There is no static route advertise command.

D: This is used for filtering inbound and outbound routes.

Reference: Catherine Paquet and Diane Teare, “Building Scalable Cisco Networks” (Cisco Press 2001), p 465.

QUESTION NO: 8

On the subject of redistributing EIGRP between other routing protocols, which three of statements below are correct? (Select three)

- A. IPX RIP redistribution with IPX EIGRP is enabled by default.
- B. AppleTalk EIGRP and RTMP redistribution is enabled by default.
- C. EIGRP for IPX automatically redistributes route information with Novell RIP.
- D. Redistribution between EIGRP and IGRP is always enabled by default regardless of the AS number used.

Answer: A, B, C

Explanation:

A: IPX RIP redistribution with Enhanced IGRP is enabled by default.

B: AppleTalk RTMP redistribution is enabled by default.

C: EIGRP automatically redistributes route information with Novell RIP.

Note: Using a routing protocol to advertise routes that are learned by some other means, such as by another routing protocol, static routes, or directly connected routes, is called redistribution.

Incorrect Answers:

D: Redistribution between EIGRP and IGRP in the same autonomous system is automatically done. Manual configuration is required if different autonomous systems are used.

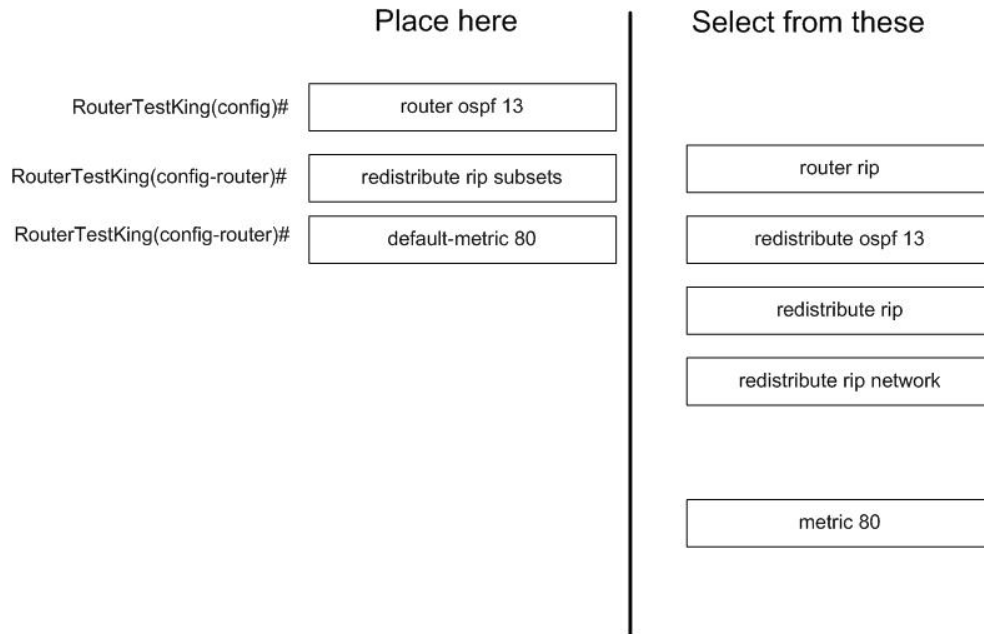
QUESTION NO: 9

You are the network administrator of an antiquated RIP network that's become too inefficient for its size, so you want to make the transition to OSPF. Your goal is to redistribute RIP routes to OSPF. Your metric will be 80, and you are using the private IP network of 10.0.0.0.

Choose the correct command phrases on the left and drag them onto the correct command sequence on the left.

	Place here	Select from these
RouterTestKing(config)#	<input type="text" value="Place here"/>	<input type="text" value="router ospf 13"/>
RouterTestKing(config-router)#	<input type="text" value="Place here"/>	<input type="text" value="router rip"/>
RouterTestKing(config-router)#	<input type="text" value="Place here"/>	<input type="text" value="redistribute ospf 13"/>
		<input type="text" value="redistribute rip"/>
		<input type="text" value="redistribute rip network"/>
		<input type="text" value="redistribute rip subsets"/>
		<input type="text" value="metric 80"/>
		<input type="text" value="default-metric 80"/>

www.testking.com

Answer:**Explanation:**

Step 1: `router ospf 13`

We are distributing into ospf. 13 denotes the AS (autonomous system) in use (which we only can assume be 13).

Step 2: `redistribute rip subnets`

The RIP routes are redistributed into OSPF. The subnets keyword tells OSPF to redistribute all subnet routes. Without the subnets keyword, only networks that are not subnetted will be redistributed by OSPF.

Step 3: `default-metric 80`


Finally the default metric is set.

Reference: RIP and OSPF Redistribution

<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/cs001.htm>

QUESTION NO: 10

The TestKing network is displayed in the diagram below:



Assuming that TK2 is configured for mutual redistribution; which of the routes below should be present on Router TK3's routing table? (Select four)

- A. 172.16.211.0/24
- B. 172.16.213.0/24
- C. 172.16.214.0/24
- D. 172.16.215.0/24
- E. 172.16.210.64/26
- F. 172.16.212.16/28

Answer: A, B, C, D

Explanation:

- A: This is a classful route.
- B: This route is redistributed from OSPF into IGRP.
- C: This route is learned through IGRP.
- D: This is a directly connected route.

Incorrect Answers:

E, F: These are subnetted routes. IGRP is a classful routing protocol. These routes will not be redistributed into IGRP.

Reference: <http://www.cisco.com/warp/public/105/52.html>

Redistributing Between Classful and Classless Protocols: EIGRP or OSPF into RIP or IGRP

QUESTION NO: 11

In order to enable a RIP network to communicate with an OSPF network, redistribution is configured on the TestKing network. What kind of router would be used to redistribute RIP into an OSPF network?

- A. ABR
- B. ASBR
- C. Internal router
- D. Backbone router

Answer: B

Explanation:

External route summarization is specific to external routes that are injected into OSPF via redistribution. Only ASBRs can summarize external routes. These types of routes cannot be summarized by any other router type.

QUESTION NO: 12

One of the TestKing routers is configured to redistribute the interior IP networks into BGP. What are two potential consequences that can occur when dynamically learned routes from IGP protocols get redistributed into BGP? (Select two)

- A. Routing loops can occur.
- B. The IGP routing table is reduced.
- C. External IGP learned routes might not necessarily have originated in this AS.
- D. Route processing is done using process switching instead of cache switching.

Answer: A, C

Explanation:

It is fairly common to redistribute IGP routes (such as Enhanced IGRP, IGRP, IS-IS, OSPF, and RIP routes) into BGP. However, precautions should be made when doing so. Some of your IGP routes might have been learned from BGP (C), so you need to use access lists to prevent the redistribution of routes back into BGP, or else routing loops can occur (A). In addition, when these IGRP routes are redistributed, all of the networks in the IGP routing table will be redistributed, including those that were learned externally.

Reference: Using the Border Gateway Protocol for Interdomain Routing
<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/icsbgp4.htm>

QUESTION NO: 13

What's the preferred method of route redistribution, when two routes of different protocols get exchanged? (Select two)

- A. Use one way route redistribution when there is one path.
- B. Use one way route distribution when there are multiple paths.
- C. Use static routes when there are multiple paths.
- D. Use two way route distribution when there is one path.
- E. Use two way route redistribution where there are multiple paths.
- F. Use static routes when there is one path.

Answer: B, C

Explanation:

- B. One way redistribution- To avoid routing loops and problems with varying convergence times, allow routes to be exchanged in only one direction, not both directions. In the other direction, you should consider a default route.
- C. When you want to prevent routing loops – Many companies have large enough networks that redundant paths are prominent. In some cases, for example, when a path to the same destination is learned from two different routing protocols, you may want to filter the propagation of one of the paths.

Reference: Building Scalable Cisco Networks (Cisco Press) page 472

QUESTION NO: 14

You are the administrator of a network and are planning on configuring route redistribution. The network you plan on working on is running different routing protocols, and the routers are connected together with redundant links. If your goal was to allow traffic to take the best route, what would you do first?

- A. Define the default metric.
- B. Identify the boundary router(s).
- C. Determine which routing protocol is running in the core router(s).
- D. Determine which routing protocol is running in the edge router(s).
- E. None of the above

Answer: A

Explanation:

Because different protocols use different metrics for finding the best path to a destination, the first step in configuring mutual redistribution is to define the default metric that will be used as a baseline standard.

QUESTION NO: 15

You have a network with multiple routing protocols running in different Autonomous Systems, redistributed together with two-way redistribution. What could you do to prevent routing loops?

- A. Manually configuring the static routes.
- B. Manually configuring the default gateway.
- C. Manually configuring the administrative k-value.
- D. Manually configuring the administrative distance.

Answer: D**Explanation:**

If you manually configure the administrative distance, you can easily decide which protocol the router will choose for each route, thus preventing the possibility of routing loops.

Incorrect Answers:

A: Manually configuring static routes could cause additional routing issues, since static routes are preferred over the dynamically learned routes they could override the information provided by the redistribution.

B: Manual routing configurations should be avoided when used with dynamically learned routes that are being redistributed.

C: These values are used only by IGRP and EIGRP, so they may not apply at all when redistributed into other routing protocols.

QUESTION NO: 16

Router TK1 is configured as shown below:

```
router igrp 100
  network 197.135.20.0
  network 197.135.24.0
  network 197.135.27.0
  redistribute rip
  default-metric 10 100 255 1 1500
  distance 140 0.0.0.0 255.255.255.255 9
```

```
access-list 9 permit 197.135.20.0
access-list 9 permit 197.135.24.0
access-list 9 permit 197.135.27.0
```

Which of the statements below correctly describe the configuration above? (Select two)

- A. Networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 are allowed into the routing table.
- B. The RIP learned routes to networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 will be assigned an administrative distance of 140.
- C. The IGRP learned routes to networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 will be assigned an administrative distance of 140.
- D. Changing the administrative distance to a number larger than the default value makes networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 unreachable.

Answer: A, C

Explanation:

We are redistributing RIP into IGRP. The **redistribute rip** command specifies that routes learned via RIP will be advertised in the IGRP updates. All of the routes specified in the network statements will be allowed into the routing table. No route filtering is taking place with this command. The access lists are used only to define which networks will have their administrative distances changed when redistributed. In this example, the 3 networks in access list 9 will have their AD changed to 140 when redistributed.

We examine the following command:

```
distance 140 0.0.0.0 255.255.255.255 9
```

The **140** defines the administrative distance that specified routes will be assigned.

The **0.0.0.0 255.255.255.255** defines the source address of the router supplying the routing information, in this case any router.

The **9** defines the access-list to be used to filter incoming routing updates to determine which will have their administrative distance changed.

And one of the access-list statements:

access-list 9 permit 197.135.27.0

9 is the access-list number.

permit allows all networks that match the address to be permitted, in this case to have their administrative distance changed.

197.135.27.0 A network to be permitted, in this case to have its administrative distance changed.

Incorrect Answers:

B: RIP routes are redistributed into IGRP; not vice versa. However, it is not clear that RIP even knows about this route, since the RIP configuration was omitted. We can, however, be certain that IGRP knows about it since these networks were added to the IGRP routing process.

D: In this case we are changing the AD to 140 for the specified networks. Although the AD has been changed to a value that is higher than the default, they are still reachable.

QUESTION NO: 17

You are a network administrator at TestKing and you're in the process of migrating their networks from RIP to OSPF (you're using the private address of 10.0.0.0 for routing). You want to redistribute the RIP routes into OSPF and assign a metric of 60 to them. Drag the commands on the left onto the corresponding boxes on the right.

router ospf 1
 router rip
 redistribute ospf 1
 redistribute rip
 redistribute rip networks
 redistribute rip subnets
 metric 60
 default-metric 60

Router(config)# Drop answer here
 Router(config-router)# Drop answer here
 Router(config-router)# Drop answer here

Answer:

router rip
 redistribute ospf 1
 redistribute rip
 redistribute rip networks
 metric 60

Router(config)# router ospf 1
 Router(config-router)# redistribute rip subnets
 Router(config-router)# default-metric 60

QUESTION NO: 18

A TestKing router is configured for route redistribution as shown below:

```
ip route 30.0.0.0 255.0.0.0 172.16.1.2
ip route 192.168.1.0 255.255.255.0 172.10.1.2

router eigrp 100
  redistribute ospf 100
```

Leading the way in IT testing and certification tools, www.testking.com

```

network 172.19.1.0

router ospf 108
 redistribute static
 redistribute eigrp 100
 network 172.16.1.0 00.0.25 area 0
 distribute-list 5 out static

access-list 5 permit 30.0.0.0 0.255.255.255

```

In the above configuration, what is the function of the command line `distribute-list 5 out static`?

- A. It denies the route to 30.0.0.0 via OSPF
- B. It denies the route to 30.0.0.0 via EIGRP
- C. It propagates the route to 30.0.0.0 via OSPF
- D. It propagates the route to 30.0.0.0 via EIGRP

Answer: C

Explanation:

The “`distribute-list 5 out static`” command filters routes learned from static entries by using access list 5, before those routes are passed to the OSPF process. In this example, the static route to the 30.0.0.0/8 network matches access list 5, so it is permitted to be redistributed into the OSPF protocol.

Incorrect Answers:

- A, B: This access-list permits, not denies, routes.
- D: The route is propagated via OSPF, not via EIGRP.

QUESTION NO: 19

You are configuring redistribution to advertise EIGRP routes into OSPF on a boundary router. Given the configuration:

```

router ospf 1
 redistribute eigrp 1 metric 25 subnets

```

Which is the function of the subnets in the redistribute command?

- A. It specifies subnetted routes should be advertised into OSPF.
- B. It specifies subnetted routes should be advertised out of OSPF.
- C. It specified routes that will be summarized on the 25-bit boundary.
- D. It specifies a limit of 25 subnets for each OSPF route advertisement.
- E. None of the above.

Answer: A

Explanation:

The subnets keyword is used to specify that subnet mask information used by classless routing protocols should be preserved when redistributed into the other routing protocol. In this case, the EIGRP routes, including subnet masks, will be redistributed into OSPF.

Incorrect Answers:

B: Here the routes are being placed into the OSPF routing process, not advertised out of.

C, D: In this configuration, the value of “25” means that the metric will be 25 when placed into OSPF. It has nothing to do with the bit boundary or the number of subnets that can be advertised.

QUESTION NO: 20

The TestKing network is in the process of changing the routing protocol from EIGRP to OSPF, as well as changing the IP network to the private 10.0.0.0/8 network. To do this, router TK1 has been configured as shown below:

```
router ospf 100
 redistribute eigrp 100 metric 100 metric-type 1
 network 172.16.0 0.0.0.255.255
```

All other configurations use the default values.

You want to successfully redistribute all networks and subnets on the TestKing network. What can you do to accomplish this goal? (Choose two)

- A. Change the OSPF process-id number from 100 to 1 in the **router ospf** command.
- B. Configure the **redistribute** command under **router eigrp 1** instead.
- C. Change the EIGRP AS number from 100 to 1 in the **redistribute** command.
- D. Add the **subnets** option to the **redistribute** command.
- E. Add the **network 10.0.0.0 0.255.255.255** command under **router ospf 100**.
- F. Change the metric to an EIGRP compatible metric value for Bandwidth, Delay, Reliability, Load, MTU, (such as 64 1000 100 1 1500) in the **redistribute** command.

Answer: D, E

Explanation:

The use of the “subnets” keyword will ensure that the subnet mask information will be redistributed into the OSPF process. Since both EIGRP and OSPF support the use of VLSM, this is recommended.

Since the network is being migrated to OSPF, it would be best to begin advertising the network via the OSPF process, as shown in answer choice E.

QUESTION NO: 21

You are using multiple protocols in different Autonomous Systems (AS). You need to redistribute between the systems. You are using two-way redistribution. Which action should help you avoid routing loop issues?

- A. Manually configuring the static routes.
- B. Manually configuring the default gateway.
- C. Manually configuring the administrative K-factor.
- D. Manually configuring the administrative distance.


Answer: D

Explanation:

If you manually configure the administrative distance, you can manually control which router will choose which route thus preventing the possibility of routing loops.

QUESTION NO: 22

Network topology exhibit



TestKing.com recently completed a merger with Foo Inc. The two companies have been using separate routing protocols on their companies corporate networks, and an immediate solution is required for the two companies to begin sharing data. A boundary router, TestKing2, has been established to perform mutual redistribution of route information between the two networks. Configure route redistribution from EIGRP into IS-IS and from IS-IS into EIGRP on the boundary router per the following requirements.

- Seed metric for EIGRP must have the following characteristics:
Bandwidth=64 Kbps

Leading the way in IT testing and certification tools, www.testking.com

- Delay=100
- Reliability=255
- Load=1
- MTU=1500

Seed metric for IS-IS must be set to 50.

Only redistribute Level-1 IS-IS routes into EIGRP.

Redistribute EIGRP routes into IS-IS as Level-1 routes.

Name: TestKing1
S0/0: 10.10.10.5/30
Loopback1: 1.1.1.1/32
Secret Password: testking

Name: TestKing2
S0/0: 192.168.1.6/30
Loopback1: 2.2.2.2/32
Secret Password: testking

Name: TestKing3
S0/0: 192.168.1.5/30
Loopback1: 3.3.3.3/32
Secret Password: testking

Answer:

Testking2> en

Testking2# config t

Testking2(config-router)# router isis

Testking2(config-router)# redistribute eigrp 100 level-1 metric 50

Testking2(config-router)# exit

Testking2(config)# router eigrp 100

Testking2(config-router)# redistribute isis level-1 metric 64 10 255 1 1500

Testking2(config-router)# redistribute connected

Testking2(config-router)# ^z

Testking2# copy running-config start up-config.

Alternative #1:

- Seed metric for EIGRP must have the following characteristics:
Bandwidth=64 Kbps
- Delay=2 ms
- Reliability=255
- Load=1
- MTU=1500

Name: TestKing1
S0/0: 10.197.197.5/30
Loopback1: 1.1.1.1/32
Secret Password: testking

Name: TestKing2
S0/0: 192.168.191.6/30
Loopback1: 2.2.2.2/32
Secret Password: testking

Name: TestKing3
S0/0: 192.168.191.5/30
Loopback1: 3.3.3.3/32
Secret Password: testking

Section 6: Identify the steps to configure policy-based routing using route maps (13 questions)

QUESTION NO: 1

Your network appears to be experiencing some issues relating to policy based routing. Which troubleshooting command could you use to view the current route maps configured on a specific interface?

- A. show interface
- B. show route-map
- C. show ip policy
- D. show ip route map
- E. None of the above

Answer: C

Explanation:

The **show ip policy** command is used to display which route map is associated with which interface.

Incorrect Answers:

- A: The show interface command shows detailed information about the Cisco router/switch. However, the output does not include information on configured route maps.
- B: The show route-map command displays the contents of all route maps or the specified route map.
- D: This is an invalid command.

QUESTION NO: 2

A router is configured for policy based routing as shown below:

```
interface serial 0
  ip policy route-map demo
route-map demo permit 10
  match ip address 4
  set interface serial2 serial3
```

```
access list 4 permit 10.3.3.2 0.0.0.0
```

Based on the information above, which of the following statements is true?

- A. If the packet is sourced from 10.3.3.2, it is a candidate for fast-switched policy routing.
- B. If the packet is destined for 10.3.3.2, it will be routed out interface serial 2 and interface serial 3 in a load-sharing fashion.
- C. If the packet is sourced from 10.3.3.2, it will be routed out interface serial 2, unless it is not up, in which case it will be routed out interface serial 3.
- D. If the packet is sourced from 10.3.3.2, the flow will be routed out interface serial 2 and interface serial 3 in a load-sharing fashion for the duration of the flow.

Answer: C

Explanation:

The set command specifies the list of interfaces which may be used to forward the traffic. If the first interface goes down, the second interface is being used. In this example, packets sourced from 10.3.3.2 (standard access list matches based on the source, not the destination) will be sent out the serial 2 interface. These packets will be sent to the serial 3 interface only if the serial 2 interface is unreachable (down).

QUESTION NO: 3

When configuring policy based routing using the “route-map” command, what is the final command you have to enter to complete the configuration?

- A. Set
- B. Match
- C. Map-list
- D. IP policy

Answer: A

Explanation:

Route maps are complex access lists: A collection of route-map statements that have the same route-map name are considered one route-map. Route maps are configured by using match commands to match the criteria for the action, followed by a set command to define the action to take. To configure route maps, perform the following:

Step 1: RouterTestKing(config)# route-map map-tag [permit | deny] [sequence-number]

First we define the conditions for policy routing.

Step2: RouterTestKing(config-route-map)#match { conditions}

Then we define the conditions to match

Step 3: RouterTestKing(config-route-map)# set { actions}

Finally we define the action to be taken on a match.

QUESTION NO: 4

If you're using policy-based routing, what could you do to prevent packets with no match in the route map from being returned to the normal forwarding process?

- A. Set the next-hop metric to 255 for packets without a match.
- B. Use a `set` statement to route packets to the `null0` interface.
- C. Use a "deny any" statement as the last statement in the route map.
- D. Use a logical "OR" in the `match` statement to send packets to the `null0` interface.

Answer: B

Explanation:

If it is desired not to revert to normal forwarding and to drop a packet that does not match the specified criteria, then a `set` statement to route the packets to interface `null 0` should be specified as the last entry in the route-map.

Incorrect Answers:

A: Setting a metric of 255 will not necessarily make the route unreachable, since this is a valid metric for OSPF, IGRP, and EIGRP.

C: The "deny all" statement is already implicitly enabled at the end of each access list that is used in a route map. As opposed to regular access lists, however, traffic that does not match the permit statements are routed using the normal forwarding process.

D: It is not possible to use a logical OR statement. Doing so would mean that all traffic could be routed to `null 0`.

QUESTION NO: 5

At the end of every route map, there's a command to implicitly "deny any" What is the end result of this rule?

- A. Packets that reach the end of the route map are discarded.
- B. Packets are forwarded to the null interface for special handling.
- C. Packets that reach the end of the route map are routed in a normal fashion.
- D. Packets that reach the end of the route map are returned in the originating interface.


Answer: C

Explanation:

The implicit deny any in every route map means that packets not meeting any of the matching criteria are routed in the normal fashion. This is in contrast to regular access lists, where packets that fail to match any of the permit statements are discarded.

QUESTION NO: 6

The TestKing network is displayed below:



You need to implement a route map on the E0 interface of router TESTKING1 (name it ISPA). Which command would be issued first to do this?

- A. TESTKING1 (config) #interface e0
TESTKING1 (config-if) ip route-map ISPA
- B. TESTKING1 (config) #interface e0
TESTKING1 (config-if) #ip policy route-map ISPA
- C. TESTKING1 (config) #interface e0
TESTKING1 (config-if) #policy route-map ISPA
- D. TESTKING1 (Config) #interface e0
TESTKING1 (Config-if) #policy route map ISPA

Answer: B

Explanation:

The correct syntax to enable an interface for policy based routing is “ip policy route-map *name*.” One interface can only have one route-map tag, but you can have multiple route map entries with different sequence numbers. These entries are evaluated in sequence number order until the first match. If there is no match, packets will be routed as usual.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps4324/products_configuration_guide_chapter09186a008019d0dd.html

QUESTION NO: 7

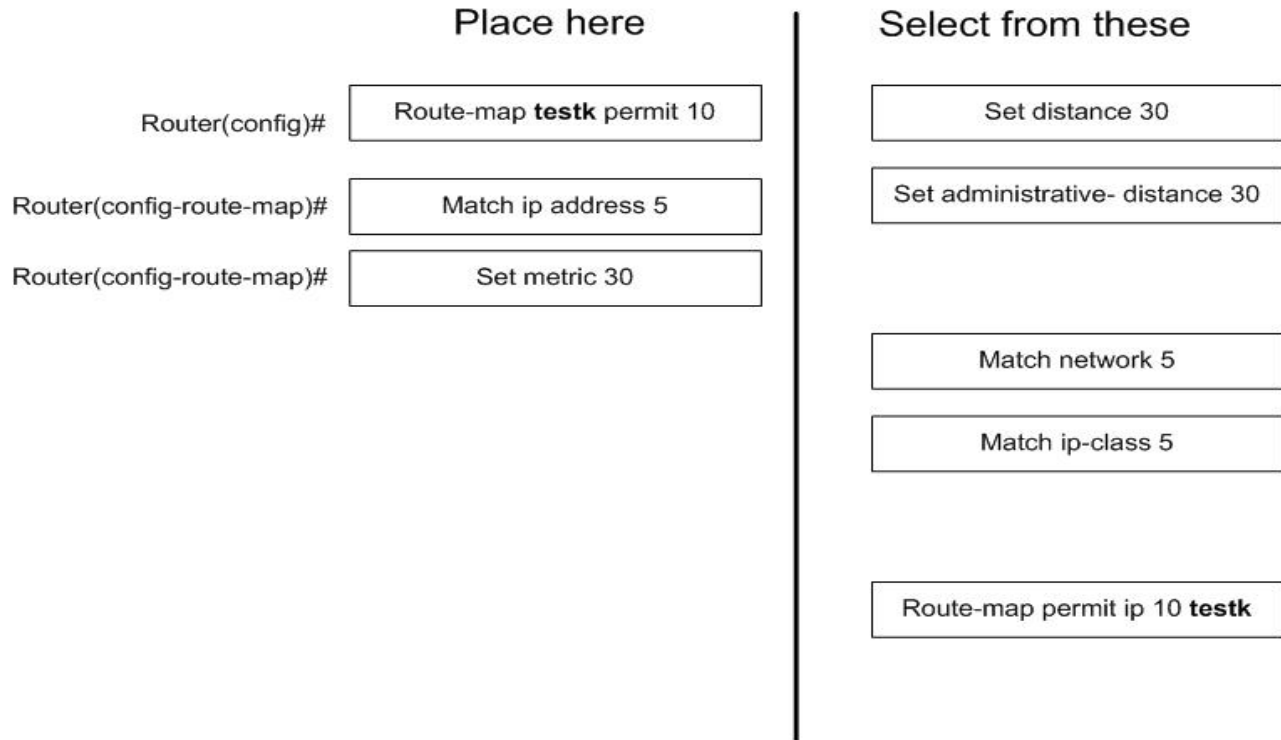
You have the following access list statement configured on your router:

```
access-list 5 permit 176.234.5.0
```

Your goal is to configure a route map to modify the metric to 30 for the network in the access list above. Drag the proper commands on the right side to its proper command sequence on the left side. (Hint: you aren't required to apply the route map yet.)

	Place here	Select from these
Router(config)#	<i>Place here</i>	Set distance 30
Router(config-route-map)#	<i>Place here</i>	Set administrative- distance 30
Router(config-route-map)#	<i>Place here</i>	Set metric 30
		Match network 5
		Match ip-class 5
		Match ip address 5
		Route-map permit ip 10 testk
		Route-map testk permit 10

Answer:

**Explanation:**

Step 1: First we must enter Route-Map Configuration mode. We issue the following command **route-map testk permit 10**

Testk is the tag, we permit (the alternative is to deny), and we use the sequence number 10.

Step 2: We define the match condition.

Match ip address 5

We match the IP address to the IP address of access-list 5, namely 176.234.5.0.

Step 3: We apply the set statement.

set metric 30

We should simply change the metric.

Note: Syntax route-map

route-map *map-tag* [[**permit** | **deny**] | [*sequence-number*]]

Each route map will consist of a list of match and set configuration. The match will specify a match criteria and set specifies a set action if the criteria enforced by the match command are met.

The related commands for **match** are: match as-path, match community, match clns, match interface, match ip address, match ip next-hop, match ip route-source, match metric, match route-type, match tag.

The related commands for **set** are: set as-path, set clns, set automatic-tag, set community, set interface, set default interface, set ip default next-hop, set level, set local-preference, set metric, set metric-type, set next-hop, set origin, set tag, set weight.

Incorrect Answers:

Route-map permit ip 10 testk

The command has the wrong syntax. The tag (here: word), must be immediately after the keyword route-map.

Match ip-class 5 and Match ip-class 5

There are no such commands.

Set distance 30 and Set administrative- distance 30

There are no such commands.

QUESTION NO: 8

The TestKing router is running on Cisco IOS version 11.2. On a router running Cisco IOS 11.2 or later, what is the default switching mode for forwarding packets that coincide with the established routing policy?

- A. fast
- B. slow
- C. NetFast
- D. process

Answer: A

Explanation:

Starting with IOS 11.2, the default switching mode is fast. Prior to this, it was process switching. The following describes these two different methods.

Process Switching

In process switching the first packet is copied to the system buffer. The router looks up the Layer 3 network address in the routing table and initializes the fast-switch cache. The frame is rewritten with the destination address and sent to the outgoing interface that services that destination. Subsequent packets for that destination are sent by the same switching path.

Fast Switching

When packets are fast switched, the first packet is copied to packet memory and the destination network or host is found in the fast-switching cache. The frame is rewritten and sent to the outgoing interface that services the destination. Subsequent packets for the same destination use the same switching path.

QUESTION NO: 9

You are the network engineer at TestKing. You want to configure a route map that will modify the metric for the network in the following access list:

Leading the way in IT testing and certification tools, www.testking.com

access-list 1 permit 192.168.1.0

Drag the commands in the proper order to configure the route map. You are not required to apply the route map at this time.

set distance 50	RouterTestKing(config)#	Drop answer here
set administrative-distance 50	Router(config-route-map)#	Drop answer here
set metric 50	RouterTestKing(config-route-map)#	Drop answer here
set network 1		
match network 1		
match ip-class 1		
match ip address 1		
route-map permit 10 word		
route-map word permit 10		

Answer:

set distance 50	RouterTestKing(config)#	route-map word permit 10
set administrative-distance 50	RouterTestKing(config-route-map)#	match ip address 1
set network 1	RouterTestKing(config-route-map)#	set metric 50
match network 1		
match ip-class 1		
route-map permit 10 word		

Explanation:

route-map word permit 10

match ip address 1

set metric 50

The **route-map** command is shown here:

```
Router(config)#route-map map-tag [{permit | deny} sequence-number]
```

The **match** commands used in policy-based routing are summarized in Table the table. These **match** commands are used to determine whether the packet is to be policy-routed, as opposed to being forwarded simply by destination. If it is to be policy-routed, the packet is sent down a different path, typically one less traveled.

The set Commands for Policy-Based Routing with Route Maps

The **set** commands used in policy-based routing are summarized in Table 18-4. These **set** commands are used after the **match** criteria has been satisfied. Whereas the **match** parameter determines whether the packet will be policy-routed, the **set** parameter determines how the packet is to be policy-routed.

Table 18-3 The match Commands used in Policy Based Routing

Command	Description
match ip address [<i>access-list-number</i> <i>name</i>] [... <i>access-list-number</i> <i>name</i>]	This states the number or name of a standard or extended access list that will be used to examine incoming packets. A standard IP access list is used to match criteria for the source address of the packet. An extended IP access list is used to specify criteria based on source and destination, application, protocol type, TOS, and precedence. If multiple access lists are specified, matching any one will result in a match.

QUESTION NO: 10

By default, what actions will a router take when the configuration command “route-map cisco” is applied? (Choose two)

- A. Any matches are considered to be permitted
- B. Any matches are considered to be denied
- C. The default sequence number will be 1.
- D. The default sequence number will be 5.
- E. The default sequence number will be 10.

Answer: A, E

Explanation:


By default, when a route-map is applied the initial sequence number that is enabled by the router is 10. In addition, all matching traffic will be permitted by default.

QUESTION NO: 11

```

hostname TestkingB
!
interface serial 0/0
ip policy route-map SNOW
!
route-map SNOW permit 5
match ip address 1
set ip next-hop 172.16.5.2
!
route-map SNOW deny 10
match ip address 2
set ip next-hop 172.16.4.2
!
access-list 1 permit 172.16.6.5 0.0.0.0
access-list 2 permit 172.16.6.10 0.0.0.0

```



Refer to the diagram above. Policy-based routing is enabled for the incoming traffic on TestkingB.


Given the above configuration, how is the packet that comes from 172.16.6.10 routed?

- A. The packet is denied access and is dropped.
- B. The packet is routed through the normal destination-based routing process.
- C. The packet is sent to TestkingC.
- D. The packet is sent back to TestkingA.

Answer: B

QUESTION NO: 12

The TestKing WAN is displayed in the diagram below:



Which configuration will force traffic to use the NY link to reach the 192.168.10.0/24 and 192.168.11.0/24 networks?

- A. TestkingC(config)# router bgp 100
TestkingC(config-router)# neighbor 192.168.5.2 route-map LOCAL in
TestkingC(config)# route-map LOCAL permit 10
TestkingC(config)# set local-preference 300
- B. TestkingB(config)# router bgp 100
TestkingB(config-router)# neighbor 172.16.20.2 route-map LOCAL in
TestkingB(config)# route-map LOCAL permit 10
TestkingB(config)# set local-preference 300
- C. TestkingC(config)# router bgp 100
TestkingC(config-router)# neighbor 192.168.5.2 route-map LOCAL out
TestkingC(config)# route-map LOCAL permit 10
TestkingC(config)# set local-preference 50
- D. TestkingB(config)# router bgp 100
TestkingB(config-router)# neighbor 172.16.20.2 route-map LOCAL out
TestkingB(config)# route-map LOCAL permit 10
TestkingB(config)# set local-preference 300

Answer: A

QUESTION NO: 13

```
router bgp 50001
network 10.0.0.0
network 172.16.0.0
neighbor 10.1.1.1 remote-as 5003
neighbor 10.1.1.1 route-map test out
! Testking.com
access-list 1 permit 10.0.0.0
access-list 2 permit any
!
route-map test permit 10
match ip address 1
set metric 200
!
route-map test permit 20
match ip address 2
```

Examine the above configuration.
What does the route map named test accomplish?

- A. Marks all prefixes advertised to the 10.1.1.1 neighbor with a MED of 200.
- B. Marks all prefixes received from the 10.1.1.1 neighbor with a MED of 200.
- C. Marks the 10.0.0.0/8 prefix advertised to the 10.1.1.1 neighbor with a MED of 200.
- D. Marks the 10.0.0.0/8 prefix received from the 10.1.1.1 neighbor with a MED of 200.
- E. Permits only the 10.0.0.0/8 prefix to be advertised to the 10.1.1.1 neighbor.
- F. Permits only the 10.0.0.0/8 prefix to be received from the 10.1.1.1 neighbor.

Answer: C

Section 7: Given a set of network requirements, identify the steps to configure a BGP environment and verify proper operation (within described guidelines) of your routers (25 questions)

QUESTION NO: 1

Your network has multiple connections to the Internet via multiple ISPs. Which of the following are common practices for configuring the network to receive BGP routes from these Internet Service Providers?

- A. Accept full routes from the ISPs.
- B. Accept only IGP routes from the ISPs
- C. Accept an external route from the ISPs.
- D. Accept only redistributed routes from the ISPs.
- E. All of the above

Answer: A

Explanation:

The configuration of the multiple connections to the ISPs can be classified depending on the routes that are provided to the AS from the ISPs. Three common ways of the configuring the connections are:

- All ISPs pass only default routes to the AS.
- All ISPs pass default routes, and selected specific routes (for example, from customers with who the AS exchanges a lot of traffic) to the AS.
- All ISPs pass all routes to the AS (A).

QUESTION NO: 2

Regarding the use of BGP prefix lists, which one of the statements below is true?

- A. They start at 10 by default.
- B. They automatically increment by ten by default.
- C. They are displayed using the `show ip prefix-list` command.
- D. They are displayed using the `show bgp prefix-list` command.
- E. None of the above.

Answer: C

Explanation:

The **show ip prefix-list** command is used to display information about a prefix list or prefix list entries.

Note: Filtering by prefix list involves matching the prefixes of routes with those listed in the prefix list, similar to using access lists.

Incorrect Answers:

A, B: By default, the entries of a prefix list will have sequence values of 5, 10, 15 etc. They start at 5 and increment by 5.

D: This is an invalid command.

QUESTION NO: 3

BGP is being used on the TestKing network for Internet routing. In EBGP, which of the following configuration lines advertises the subnet 154.2.1.0 255.255.255.0 to EBGP neighbors?

- A. Router (config-router)#network 154.2.1.0
- B. Router (config-router)#network 164.2.1.0
- C. Router (config-router)#network-advertise 154.2.1.0
- D. Router (config-router)#network 154.2.1.0 mask 255.255.255.0

Answer: D

Explanation:

When configuring BGP, we specify the mask keyword also along with the network command. If the network mask is omitted, the IP network is taken at the network class boundary.

QUESTION NO: 4

Routers A, B, and C are all part of the same autonomous system and are configured for BGP routing. The relevant configurations for RTR A and RTR B are shown below:

```

Hostname RTR A
!
interface serial 0/0
ip address 140.140.13.2 255.255.255.252
description Connection to ISP A
!
interface 10/0
ip address 10.10.10.2 255.255.0
description Connection to RTR B
!
interface f0/1
ip address 172.16.30.1 255.255.255.0

```


```

description Connection to RTR C

router bgp 5201
neighbor 10.10.10.3 remote-as 5201
neighbor 140.140.13.2 remote-as 5203

Hostname RTR B
!
interface 10/0
ip address 10.10.10.3 255.255.255.0
description Connection to RTR A
!
interface 10/1
ip address 10.10.20.2 255.255.255.0
description Connection to RTR C
!
router bgp 5201
neighbor 10.10.10.2 remote-as 5201
neighbor 10.10.20.1 remote-as 5201

```



Based on the information provided, what will RTR A do with the updates it receives from ISP A?

- A. Send the update to RTR B and C.
- B. Send the update to RTR B only.
- C. Ignore and discard the packet.
- D. Update its tables and discard the packet.

Answer: C

Explanation:

When specifying the remote external AS in the BGP peer statement, the IP address of the neighbor router should be used as the peer. In this example, the remote peer specified in RTR A is 140.140.13.2, but as we can see from the configuration file this is the IP address of its own serial interface. The correct statement should have been “neighbor 140.140.13.1 remote-as 5203” under the BGP process.

Because of this, the EBGP peering session will not be established and all BGP updates from ISP A will be ignored.

QUESTION NO: 5

Your network has 2 separate T1 connections with two different providers. TO support this, BGP is being configured. The required information regarding the ISPs is shown below:

Your network: 164.67.36.0/24
 Your AS number: 300
 AS number of ISP1: 1005
 AS number of ISP2: 1010

Based on the information provided, which of the following commands would you enter to advertise your network to ISP1 and ISP2 via BGP?

- A. router bgp 1005
 network 164.67.36.0 mask 255.255.255.0
 neighbor 15.1.1.1 remote-as 1005
 neighbor 25.1.1.1 remote-as 1010
- B. router bgp 300
 network 164.67.36.0 mask 255.255.255.0
 neighbor 15.1.1.1 remote-as 1005
 neighbor 25.1.1.1 remote-as 1010
- C. router bgp 300
 network 164.67.36.0
 neighbor 15.1.1.1 remote-as 1005
 neighbor 25.1.1.1 remote-as 1010
- D. router bgp 1010
 network 164.67.36.0
 neighbor 15.1.1.1 remote-as 1005
 neighbor 25.1.1.1 remote-as 1010

Answer: B

Explanation:

Step 1: `router bgp 300`

The **router bgp** command is used to activate the BGP protocol and identify the local autonomous system.

Step 2: `network 164.67.36.0 mask 255.255.255.0`

The **network** command controls which networks are originated by this router.

Syntax: `network network-number network-mask`

Step 3:

```
neighbor 15.1.1.1 remote-as 1005
neighbor 25.1.1.1 remote-as 1010
```

The **neighbor remote-as** command to identify a peer router with which the local router will establish a session.

Incorrect Answers:

- A, D: We must specify the local autonomous system in the router bgp command. We must use AS 300, not AS 1005 or AS 1010
- C: Both the network number and the network mask must be specified with the network command.

QUESTION NO: 6

The TestKing Internet router is configured as follows::

```
router bgp 64000
network 17.0.0.0
neighbor 178.5.1.1 remote-as 64000
neighbor 197.4.1.2 remote-as 64100
```

Based on this router configuration, which of the following statements are true? (Select two)

- A. The router with IP address 178.5.1.1 runs IBGP with RouterTK.
- B. The router with IP address 178.5.1.1 runs EBGP with RouterTK.
- C. The router with IP address 197.4.1.2 runs EBGP with RouterTK.
- D. The router with IP address 178.5.1.1 runs as a community member with RouterTK.
- E. The router with IP address 197.4.1.2 runs as a peer group member with RouterTK.

Answer: A, C

Explanation:

The **router bgp** command is used to activate the BGP protocol and identify the local autonomous system.

The **neighbor** command activates a BGP session with another router using either IBGP or EBGP.

Syntax: `neighbor { ip-address|peer-group-name} remote-as autonomous-system`

The value placed in the autonomous system field of the neighbor command determines whether the communication with the neighbor is an EBGP or an IBGP session.

A: If the autonomous system field configured in the **router bgp** command is identical to the field in the **neighbor remote-as** command, then BGP will initiate an internal session (IBGP).

Here both the local AS and the neighbor 178.5.1.1 are in the AS 64000.

C: If the field values are different, then BGP will initiate an external session (EBGP). The network 197.4.1.2 has a different AS number.

QUESTION NO: 7

You wish to permit all BGP prefixes between /10 and /18 for the 207.0.0.0 network. How should a BGP prefix list be configured to do this?

- A. `ip prefix-list 207.0.0.0/8 ge 10 le 18`
- B. `ip prefix-list 207.0.0.0/8 ge 18 le 10`
- C. `ip prefix-list 207.0.0.0/24 ge 10 le 18`
- D. `ip prefix-list 207.0.0.0/24 ge 18 le 10`
- E. None of the above

Answer: A

Explanation:

With prefix lists, the “ge” means “greater than or equal to” while the le means “less than or equal to.” Here, we want to specify all prefixes greater than or equal to 10 and less than or equal to 18 as specified in choice A. In addition, we must specify an 8 bit network mask, not a 24 bit network mask.

QUESTION NO: 8

Routers TK1 and TK2 are configured for BGP as shown below:

```
RTR TK1
router bgp 200
neighbor 183.215.22.1 remote-as 200
neighbor 183.215.22.1 update-source loopback 1
RTR TK2
router bgp 200
neighbor 147.229.1.1 remote-as 200
```

These two routers currently have an active BGP peering session between them and they are able to pass routes to each other. Based on these configuration files, which of the following are true? (Select three)

- A. RTR TK1 and RTR TK2 are running IBGP inside AS 200
- B. The IP address of RTR TK1's Loopback 1 interface is 147.229.1.1.
- C. The IP address of RTR TK1's Loopback 1 interface is 183.215.22.1.
- D. RTR TK1 and RTR TK2 are running EBGP between the autonomous systems.
- E. RTR TK1 has forced BGP to use the loopback IP address as the source in the TCP neighbor connection.

Answer: A, B, E

Explanation:

- A: The two statements that define a remote AS (autonomous) use the same AS number. Both routers must belong to the same AS and there IBGP is used.
- B: The IP address of the TK1 loopback address is 147.229.1.1 as specified in the neighbor statement of the TK2 router: **neighbor 147.229.1.1 remote-as 200**
- E: The "**neighbor 183.215.22.1 update-source loopback 1**" command issued at TK1 makes TK1 use this loopback for the BGP connection to TK2. See Note 2 below.

Note 1: Syntax of neighbor command:

```
Router (config-router) #neighbor
{ip-address | peer-group-name} remote-as as-number
```

This command specifies a BGP neighbor.

Note 2: For iBGP, you might want to allow your BGP connections to stay up regardless of which interface is used to reach a neighbor. To enable this configuration, you first configure a *loopback* interface and assign it an IP address (neighbor 183.215.22.1 update-source loopback 1). Next, configure the BGP update source to be the loopback interface (we have to assume this step – it is not indicated by the exhibit). Finally, configure your neighbor to use the address on the loopback interface (neighbor 147.229.1.1 remote-as 200).

Incorrect Answers:

- C: 183.215.22.1 is the IP address of TK2, not TK1.
- D: There is only one AS and the routers belong to that AS. Within an AS IBGP is used, not EBGP.

QUESTION NO: 9

The TestKing BGP router is configured as shown below:

```
router bgp 64000
```

```

neighbor 172.16.1.1 remote-as 64000
neighbor 10.1.1.2 remote-as 64550
network 200.52.1.192 mask 255.255.255.224
no synchronization
aggregate-address 200.52.1.0 255.255.255.0

```

The BGP peers 172.16.1.1 and 10.1.1.2 in the above configuration are active, as well as the interface with IP address 200.52.1.192. Based on this information, which statement below is true about your configuration?

- A. Router TestK has an IBGP connection with neighbor 10.1.1.2.
- B. Router TestK has an EBGP connection with the neighbor 172.16.1.1.
- C. Router TestK advertises only a route 200.52.1.0 255.255.255.0 in BGP.
- D. Router TestK advertises only a route 200.52.1.192 255.255.255.224 in BGP.
- E. Router TestK advertises both the routes to 200.52.1.0 255.255.255.0 and 200.52.1.192 255.255.255.224 in BGP.

Answer: E

Explanation:

The aggregate route, 200.52.1.0 255.255.255.0, and the more specific route, 200.52.1.192 255.255.255.224, will both be advertised. To only advertise the aggregate route the **summary-only** option of the **aggregate-address** command must be used.

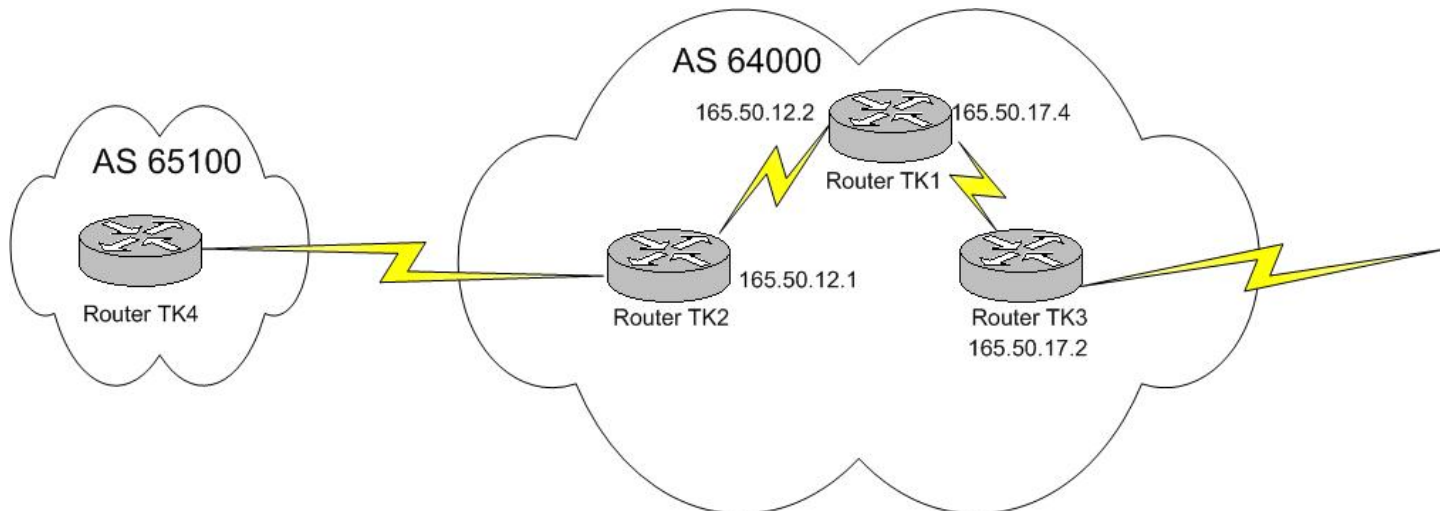
Note: The **aggregate-address** command is used to create an aggregate, or summary, entry in the BGP table.

Incorrect Answers:

- A: The neighbor 10.1.1.2 belongs to another autonomous system and is a EBGP neighbor, not an IBGP neighbor.
- B: Router A and neighbor 172.16.1.1 belong to the same autonomous system and is an IBGP neighbor, not an EBGP neighbor.
- C: The **summary-only** option of the **aggregate-address** command is used to only advertise the summary and not the specific routes.
- D: The aggregate route, 200.52.1.0 255.255.255.0, will be advertised as well.

QUESTION NO: 10

The TestKing network consists of two separate autonomous systems as shown below:



You need to configure Router TK2 as a BGP route reflector and Router TK1 as the client. Assuming that Router TK3 isn't running BGP, which two of the commands below would you enter on TK2 to satisfy your goals? (Select two)

- A. `neighbor 165.50.12.1 remote-as 65100`
- B. `neighbor 165.50.12.2 remote-as 64000`
- C. `neighbor 165.50.12.1 route-reflector-client`
- D. `neighbor 165.50.12.2 route reflector-client`

Answer: B, D

Explanation:

B: RouterTK2(config-router)# **neighbor 165.50.12.2 remote-as 64000**

We configure router TK1 (165.50.12.2) as a neighbor in AS 64000.

D: RouterTK2(config-router)# **neighbor 165.50.12.2 route-reflector-client**

Configures the router TK2 as a BGP route reflector and configures the specified neighbor TK1 (165.50.12.2) as its client.

Incorrect Answers:

A: We must specify router TK1 as neighbor, not TK2 itself (165.50.12.1). Furthermore, we should use the local AS (64000), not the remote AS 65100.

C: We must specify router TK1 as route reflector client, not TK2 itself (165.50.12.1).

QUESTION NO: 11

The TestKing network consists of a series of routers that are all configured for IBGP. Which one of the following IBGP characteristics is true?

- A. The IBGP routers must always be fully meshed.
- B. The IBGP routers can be in a different AS.

- C. The IBGP routers must be directly connected.
- D. The IBGP routers do not need to be directly connected.
- E. None of the above are true.

Answer: D

Explanation:

The IBGP routers do not have to be directly connected. The remote IBGP peers need only be reachable via a TCP connection. For example, if the network is also running an interior routing protocol such as EIGRP or OSPF, the remote IBGP router could be many hops away, as long as it is reachable via the IGP that is being used.

Incorrect Answers:

- A: Using route reflectors or confederations a full mesh topology is not necessary.
- B: The IBGP routers must be placed in the same AS. Peers that are in different autonomous systems are using EBGP, not IBGP.
- C: The IBGP routers do not have to be directly connected.

QUESTION NO: 12

Which of the following statements regarding BGP peer groups are true? (Select two)

- A. Peer members inherit all options of the peer group.
- B. Peer groups can be used to simplify BGP configurations.
- C. Peer groups are optional non-transitive attributes for BGP.
- D. A peer group allows options that affect outbound updates to be overridden.
- E. A common name should be used on all routers because this information is passed between neighbors.

Answer: A, B

Explanation:

A BGP peer group is a group of BGP neighbors with the same update policies.

- A: By default members of the peer group inherit all of the configuration options of the peer group.

Note: Members can also be configured to override these options

- B: Peer groups are normally used to simplify router configurations when many neighbors have the same policy.

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

Incorrect Answers:

- C: The only Optional non-transitive attribute in BGP is MED.
- D: Does not apply.

E: Using a common router name provides no BGP benefits at all.

Reference: BGP Peer Groups, <http://www.cisco.com/warp/public/459/29.html>

QUESTION NO: 13

A BGP router is configured as shown below:

```
interface ethernet 0
ip address 10.10.10.1 255.255.0.0
!
int serial 0
ip address 172.16.1.1 255.255.255.252
!
router bgp 65001
neighbor 192.168.1.1 remote-as 65002
```

Based on the above configuration, which of the following BGP statements would inject the 10.10.0.0/16 prefix into the BGP routing table?

- A. network 10.0.0.0
- B. network 10.10.0.0 mask 255.255.0.0
- C. network 10.10.10.1 mask 255.255.255.255
- D. network 10.10.10.0 mask 255.255.255.0
- E. network 10.0.0.0 mask 255.255.0.0

Answer: B

Explanation:

The /16 mask is equal to 255.255.0.0, so answer choice B matches the address and the mask. To specify the route as classless, the mask keyword should be included or the network will be summarized at the network boundary.

QUESTION NO: 14

Why would a systems administrator enter in the network command when configuring BGP?

- A. Local routes matching the network command are filtered from the BGP routing table.
- B. Local routes matching the network command can be installed into the BGP routing table.
- C. Routes matching the network command will be filtered from BGP routing updates.
- D. External routes matching the network command will be installed into the BGP routing table.

E. None of the above.

Answer: B

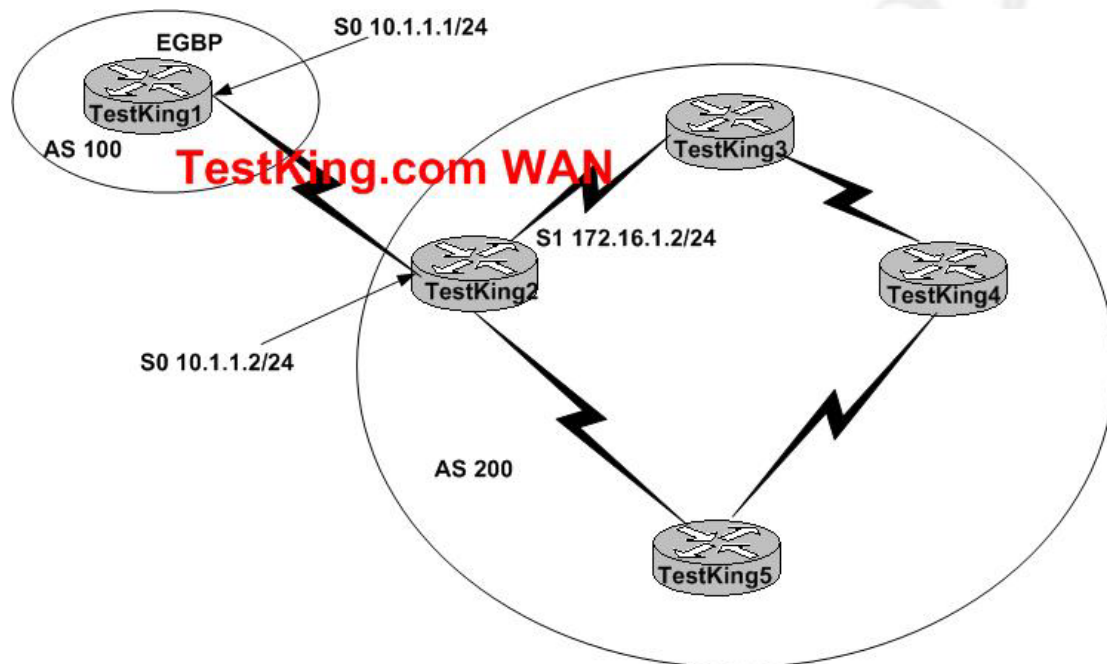
Explanation:

Use the **network** router configuration command to permit BGP to advertise a network if it is present in the IP routing table.

Reference: Building Scalable Cisco Networks (Cisco Press) page 342

QUESTION NO: 15

The TestKing network is depicted in the following diagram:



Which of the following command sets would you use if you want TESTKING1 and TestKing2 to exchange BGP routing information? (Select two)

- A. TESTKING2 (config) #router bgp 200
TESTKING2 (config-router) #neighbor 10.1.1.1 remote-as 100
- B. TESTKING1 (config) #router bgp 100
TESTKING1 (config-router) #exit
TESTKING1 (config) #interface S0
TESTKING1 (config-if) #neighbor 10.1.1.2 remote-as 200
- C. TESTKING1 (config) #router bgp 100
TESTKING1 (config-router) #neighbor 10.1.1.2 remote-as 200
- D. TESTKING2 (config) #router bgp 100

Leading the way in IT testing and certification tools, www.testking.com


```

TESTKING2 (config-router) exit
TESTKING2 (config) #interface S0
TESTKING2 (config-if) #neighbor remote-as 100
E. TESTKING1 (config) #router bgp 100
TESTKING1 (config-router) #network 10.0.0.0
TESTKING1 (config-router) #neighbor 10.1.1.2 remote-as 100
F. TESTKING2 (config) #router bgp 200
TESTKING2 (config-router) #network 10.0.0.0
TESTKING2 (config-router) #network 10.1.1.1 remote-as 200

```

Answer: A, C

Explanation:

Choices A and C give the correct syntax for configuring EBGP peering sessions.

Incorrect Answers:

B, D: The BGP peering configurations are placed under the BGP routing process, not in interface configuration mode.

E, F: The network 10.0.0.0 command is not required here, since this is an EBGP peering session using the directly connected interface. Since each neighbor shares the 10.0.0.0 network they each already know how to reach this network.

QUESTION NO: 16

Router TK-1 is configured for BGP routing as shown below:

```

router bgp 65300
network 27.0.0.0
neighbor 192.23.1.1 remote-as 65300

```

From the perspective of router TK-1, what kind of router is the router with IP address 192.23.1.1?

- A. A peer router running IBGP
- B. A peer router running EBGP
- C. A community member running IBGP
- D. A peer group member running IBGP
- E. A peer group member running EBGP

Answer: A

Explanation:

Both the local and remote router is configured with the same autonomous system number so they are peer routers running IBGP.

QUESTION NO: 17

When you're configuring BGP on a Cisco router, what is true of the command "network"?

- A. Local routes matching the `network` command are filtered from the BGP routing table.
- B. Local routes matching the `network` command can be installed into the BGP routing table.
- C. Sending and receiving BGP updates is controlled by using a number of different filtering methods.
- D. The route to a neighbor autonomous system must have the correct MED applied to be installed into BGP routing table.
- E. None of the above.

Answer: B

Explanation:

The `network` command allows BGP to advertise an IGP route if it is already in the IP table. A matching route must exist in the routing table before the network is announced. The `network` command is used to permit BGP to advertise a network if it is present in the IP routing table.

QUESTION NO: 18

Router TK1 needs to be configured to advertise a specific network. Which of the following commands would you use if you wanted to advertise the subnet 154.2.1.0 255.255.255.0 to the EBGp neighbors on your subnet?

- A. Router (config-router)#network 154.2.1.0
- B. Router (config-router)#network 164.2.1.0
- C. Router (config-router)#network-advertise 154.2.1.0
- D. Router (config-router)#network 154.2.1.0 mask 255.255.255.0
- E. None of the above

Answer: D

Explanation:

The `network` command is used to specify the networks to be advertised by the Border Gateway Protocol (BGP) and multiprotocol BGP routing processes.

Syntax: `network network-number [mask network-mask] [route-map map-name]`

Mask and route-map are optional. If the `mask` keyword is configured, then an exact match must exist in the routing table.

Incorrect Answers:

- A: If we do not specify the subnet mask then additional networks are allowed to be advertised.
The classful subnet mask of 154.2.1.0 is 255.255.0.0 – a Class B network.
- B: This is using the incorrect IP address, as well as a missing subnet mask.
- C: The **network-advertise** is an invalid command.

QUESTION NO: 19

You are the administrator of a company with BGP connections to multiple ISP's. How could you configure BGP to make it favor one particular ISP for outbound traffic?

- A. Configure weight
- B. Enable route reflector
- C. Create a distribute list
- D. Enable the Longer Autonomous System path option.
- E. All of the above.

Answer: A

Explanation:

If the router learns about more than one route to the same destination, the route with the highest weight will be preferred. Weight is a Cisco BGP parameter that is local to the router. When terminating multiple ISP connections into the same router, weight can be used to affect which path is chosen for outbound traffic.

Incorrect Answers:

- B: A route reflector cannot be used to influence outbound traffic. A route reflector modifies the BGP split horizon rule by allowing the router configured as the route reflector to propagate routes learned by IBGP to other IBGP peers. This saves on the number of BGP TCP sessions that must be maintained, and also reduces the BGP routing traffic.
- C: Distribute lists restrict the routing information that the router learns or advertises. By itself a distribute list cannot make routes from one ISP be preferred to routes from another ISP.
- D: This choice describes ASD path pre-pending, which would be used to influence the path that incoming traffic takes, not outgoing.

QUESTION NO: 20

What are two solutions to overcome the full mesh requirement with iBGP? (Choose two)

- A. BGP groups
- B. BGP local preference
- C. Route reflector
- D. Confederation
- E. Aggregate-address

Answer: C, D

Explanation:

In general, all IBGP peers must be configured to be fully meshed. If they are not, then all of the IBGP routers will not have the updated information from the external BGP routers. There are two ways to overcome the scalability issues of a full IBGP mesh: route reflectors and confederations. With route reflectors, internal BGP routers peer only with the route reflector. With confederations, the AS is broken up into smaller, more manageable sub autonomous systems.

QUESTION NO: 21

An ISP is running a large IBPG network with 25 routers. The full mesh topology that is currently in place is inefficiently using up bandwidth from all of the BGP traffic.

What can the administrator configure to reduce the number of BGP neighbor relationships within the AS?

- A. Route reflectors
- B. Route maps
- C. Route redistribution
- D. Peer groups
- E. Aggregate addresses

Answer: A

Explanation:

In general, all IBGP peers must be configured to be fully meshed. If they are not, then all of the IBGP routers will not have the updated information from the external BGP routers. There are two ways to overcome the scalability issues of a full IBGP mesh: route reflectors and confederations. With route reflectors, internal BGP routers peer only with the route reflector, and then the route reflectors connect with each other. This can considerably reduce the number of IBGP sessions. Another solution to the scalability problem of IBGP is the use of confederations. With confederations, the AS is broken up into smaller, more manageable sub autonomous systems.

QUESTION NO: 22

Arrange the BGP attributes on the left in order or priority.

Your boss at TestKing.com asks you to place the BGP attributes in the correct order used for determining a route.

BGP Attributes, select from these

Place here

Originate route

Place first here

AS_Path

Place second here

weight

Place third here

Local preference

Place fourth here

MED

Place fifth here

Answer:

Your boss at TestKing.com asks you to place the BGP attributes in the correct order used for determining a route.

BGP Attributes, select from these

Place here

TestKing.com

weight
Local preference
Originate route
AS_Path
MED

Explanation:

How the Best Path Algorithm Works:

BGP assigns the first valid path as the current best path. It then compares the best path with the next path in list, until it reaches the end of the list of valid paths. The following is a list of rules used to determine the best path.

1. Prefer the path with the highest WEIGHT.

Note: WEIGHT is a Cisco-specific parameter, local to the router on which it's configured.

2. Prefer the path with the highest LOCAL_PREF. Note the following:
 - o Path without LOCAL_PREF is considered as having the value set with the bgp default local-preference command, or 100 by default.
3. Prefer the path that was locally originated via a network or aggregate BGP subcommand, or through redistribution from an IGP. Local paths sourced by network or redistribute commands are preferred over local aggregates sourced by the aggregate-address command.
4. Prefer the path with the shortest AS_PATH. Note the following:
 - o The AS_CONFED_SEQUENCE and AS_CONFED_SET are not included in the AS_PATH length.
5. Prefer the path with the lowest origin type: IGP is lower than EGP, and EGP is lower than INCOMPLETE.
6. Prefer the path with the lowest multi-exit discriminator (MED). Note the following:
 - o This comparison is only done if the first (neighboring) AS is the same in the two paths; any confederation sub-ASs are ignored. In other words, MEDs are compared only if the first AS in the AS_SEQUENCE is the same for multiple paths. Any preceding AS_CONFED_SEQUENCE is ignored.
 - o If bgp always-compare-med is enabled, MEDs are compared for all paths. This option needs to be enabled over the entire AS, otherwise routing loops can occur.

- If `bgp bestpath med-confed` is enabled, MEDs are compared for all paths that consist only of `AS_CONFED_SEQUENCE` (paths originated within the local confederation).
- Paths received from a neighbor with a MED of 4,294,967,295 will have the MED changed to 4,294,967,294 before insertion into the BGP table.
- Paths received with no MED are assigned a MED of 0, unless `bgp bestpath missing-as-worst` is enabled, in which case they are assigned a MED of 4,294,967,294.
- The `bgp deterministic med` command can also influence this step as demonstrated in the *How BGP Routers Use the Multi-Exit Discriminator for Best Path Selection*.

7. Prefer external (eBGP) over internal (iBGP) paths. If `bestpath` is selected, go to Step 9 (multipath).

Note: Paths containing `AS_CONFED_SEQUENCE` and `AS_CONFED_SET` are local to the confederation, and therefore treated as internal paths. There is no distinction between Confederation External and Confederation Internal.

8. Prefer the path with the lowest IGP metric to the BGP next hop. Continue, even if `bestpath` is already selected.
9. Check if multiple paths need to be installed in the routing table for BGP Multipath. Continue, if `bestpath` is not selected yet.
10. When both paths are external, prefer the path that was received first (the oldest one). This step minimizes route-flap, since a newer path will not displace an older one, even if it would be the preferred route based on the next decision criteria (Steps 11, 12, and 13).

Skip this step if any of the following is true:

- The `bgp best path compare-routerid` command is enabled.

Note: This command was introduced in Cisco IOS® Software Releases 12.0.11S, 12.0.11SC, 12.0.11S3, 12.1.3, 12.1.3AA, 12.1.3.T, and 12.1.3.E.

- The router ID is the same for multiple paths, since the routes were received from the same router.
- There is no current best path. An example of losing the current best path occurs when the neighbor offering the path goes down.

11. Prefer the route coming from the BGP router with the lowest router ID. The router ID is the highest IP address on the router, with preference given to loopback addresses. It can also be set manually using the `bgp router-id` command.

Note: If a path contains route-reflector (RR) attributes, the originator ID is substituted for the router ID in the path selection process.

12. If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length. This will only be present in BGP route-reflector environments. It allows clients to peer with RRs or clients in other clusters. In this scenario, the client must be aware of the RR-specific BGP attribute.

13. Prefer the path coming from the lowest neighbor address. This is the IP address used in the BGP neighbor configuration, and corresponds to the remote peer used in the TCP connection with the local router.

QUESTION NO: 23

The TestKing BGP routing table consists of the following network routes:

TestKing.com
192.168.12.0/24
192.168.13.0/24
192.168.14.0/24
192.168.15.0/24

What is the correct command to summarize these prefixes into a single summary prefix of 192.168.12.0/22 while also allowing for the advertisement of the more specific prefixes?

- A. network 192.168.12.0 mask 255.255.252.0
- B. network 192.168.12.0 mask 0.0.3.255
- C. network 192.168.12.0
- D. aggregate-address 192.168.12.0 255.255.252.0
- E. aggregate-address 192.168.12.0 255.255.252.0 summary-only
- F. aggregate-address 192.168.12.0 255.255.252.0 as-set

Answer: D

Explanation:

To summarize BGP prefixes into one aggregated route, use the “aggregate-address” command. When used alone, this will advertise the aggregate route, along with the individual specific routing entries. To advertise only the aggregated route, use the “summary-only” keyword, as specified in choice E.

QUESTION NO: 24

The TestKing BGP network is displayed below:



TestKing3 is the hub router and TestKing1 and TestKing2 are the spokes. There are no virtual circuits between the spoke locations. Each router is in a separate AS. What is needed to successfully route traffic to the 11.11.11.0/24 network from TestKing1?

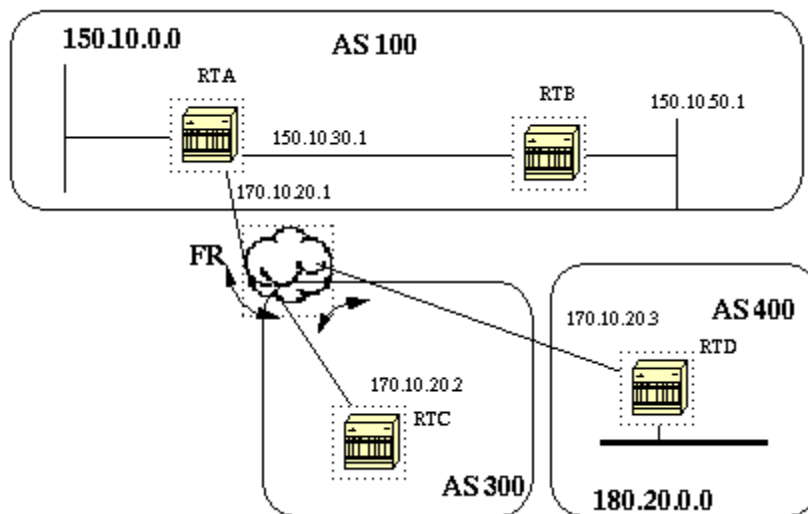
- A. The **neighbor 10.10.10.1 next-hop-self** command on TestKing1.
- B. The **neighbor 10.10.10.1 next-hop-self** command on TestKing2.
- C. The **neighbor 10.10.10.1 next-hop-self** command on TestKing3.
- D. Nothing is required. This is the default behavior on this topology

Answer: C

Explanation:

The following example illustrates the issue:

BGP Next Hop (NBMA)



If the common media as you see in the shaded area above is a frame relay or any NBMA cloud then the exact behavior will occur as if we were connected via Ethernet. RTC will advertise 180.20.0.0 to RTA with a next hop of 170.10.20.3.

The problem is that RTA does not have a direct PVC to RTD, and cannot reach the next hop. In this case routing will fail.

In order to remedy this situation a command called **next-hop-self** is created.

The next-hop-self Command

Because of certain situations with the next hop as we saw in the previous example, a command called **next-hop-self** is created. The syntax is:

neighbor {ip-address|peer-group-name} next-hop-self

The **next-hop-self** command allows us to force BGP to use a specified IP address as the next hop rather than letting the protocol choose the next hop.

In the previous example, the following configuration solves our problem:

RTC#

router bgp 300

neighbor 170.10.20.1 remote-as 100

neighbor 170.10.20.1 next-hop-self

RTC advertises 180.20.0.0 with a next hop = 170.10.20.2

Reference: <http://www.cisco.com/warp/public/459/bgp-toc.html#bgpnexthop>

QUESTION NO: 25

Exhibit

TestKing1#show ip bgp summary

BGP table version is 8, main routing table version 8

4 network entries (8/12 paths) using 832 bytes of memory

5 BGP path attribute entries using 576 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

2 received paths for inbound soft reconfiguration

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.1	4	50001	80	81	8	0	0	04:15:10	2
10.2.2.2	4	50002	79	81	0	0	0	00:00:15	Active
10.3.3.3	4	50003	80	82	0	0	0	02:00:00	Idle

Based on the show ip bgp summary output in the exhibit, which two statements are true?
Select two?

- A. The BGP session to the 10.1.1.1 neighbor is established.
- B. The BGP session to the 10.2.2.2 neighbor is established.
- C. The BGP session to the 10.3.3.3 neighbor is established.
- D. The router is attempting to establish a BGP peering to the 10.1.1.1 neighbor.

- E. The BGP session to the 10.3.3.3 neighbor is established, but the router has not received any BGP routing updates from the 10.3.3.3 neighbor.
- F. The router is attempting to establish BGP peering with the 10.2.2.2 neighbor.

Answer: A, F

Not C: How can neighbor 10.3.3.3 be established if it is in Idle state?

Not D: Neighbor 10.1.1.1 has a state prefix received of 2. Therefore it is an established neighbor.

Section 8: Identify the steps to configure a router for Network Address Translation with overload, static translations, and route maps. (4 questions)

QUESTION NO: 1

You want to hide some of your internal IP subnets from outside networks. By what means can you conceal the details of your IP addressing scheme to the outside world?

- A. Subnetting
- B. Supernetting
- C. Challenge Handshake Protocol
- D. Usernames and passwords

Answer: B

Explanation:

In large internetworks, hundreds or even thousands of networks can exist. In these environments, it is often not desirable for routers to maintain all these routes in their routing table. Route summarization (also called *route aggregation* or *supernetting*) can reduce the number of routes that a router must maintain because it is a method of representing a series of network numbers in a single summary address. By advertising only one large supernet to the outside world, the details of your IP network scheme can remain hidden.

QUESTION NO: 2

The TestKing network will be using a new ISP for their Internet connection. The new Internet provider will be allocating a new registered class C IP address subnet to use. In order to not have to change out the entire internal IP addressing scheme on your network, you plan to use NAT to translate all of the internal IP address to the new IP addresses that are being assigned to TestKing. You have been assigned the task of making changes so that the following requirements are met:

- Maintain the current IP addressing scheme
- Configure TestKing's router for network address translation (NAT) so all the internal computers use the same external IP address
- Configure a static translation so TestKing employees can access their email from the internet.

Privileged mode password: testking

The IP Addresses are shown below:

Name: TestKingNAT

SO 192.168.15.1/24

E0 10.100.5.1/24


E-mail Server's External Address

192.168.15.5/24

E-mail Server's Internal Address

10.100.5.5/24

To configure the router click on a host icon that is connected to a router by a serial console cable.

**Answer:**

TestKingNAT#Conf t

TestKingNAT(config)#Access-list 5 permit 10.100.5.0 0.0.0.255

TestKingNAT(config)#Ip nat pool lan 192.168.15.1 192.168.15.1 netmask 255.255.255.0

TestKingNAT(config)#Ip nat inside source list 5 pool lan overload

TestKingNAT(config)#Ip nat inside source static 10.100.5.5 192.168.15.5

TestKingNAT(config-if)#Int S0

TestKingNAT(config-if)#Ip nat outside

TestKingNAT(config-if)#Int E0

TestKingNAT(config-if)#Ip nat inside


TestKingNAT(config-if)#exit

TestKingNAT(config)#exit

TestKingNAT#copy run start

QUESTION NO: 3

The TestKing network is using RIP and IGRP for routing as shown below:



You have been assigned the task of configuring the routers to meet the following requirements:

- In the unlikely event that a link goes down, you have to ensure a backup link exists
- The proper metrics must be distributed through the IGRP portion of the network

Based on this information, which of the following should you do? (Select two)

- Apply the `distance` command to Main
- Apply the `distance` command to TK2
- Apply the `distribute-list` command to Main
- Apply the `distribute-list` command to TK2

Answer: B, C

Explanation:

B: Most routing protocols, including IGRP and RIP, use the distance command to modify the administrative distance. We need to change the metric of the RIP routes learned by TK2.

C: To control the advertising and processing of routes in routing updates the **distribute-list** command is used.


Incorrect Answers:

A: We change the administrative distance of RIP routes on TK2.

D: We should apply the `distribute-list` command to Main, not to TK2.

QUESTION NO: 4

Network topology exhibit



TestKing.com is changing ISPs. As a result, they will need to install an Email server and utilize network address translation (NAT) for the internal network. TestKing.com does not want to change the IP Addresses on all of the internal routers and servers. The new ISP, Foo Services, will allocate a registered class C address for TestKing to use. The internal IP Address scheme will remain the same. Configure the router to provide NAT so that all internal TestKing PCs will use the single external IP Address assigned to the router interface. Configure a static translation so that TestKing.com's Email server will be accessible from the Internet.

Name: TestKingNAT

S0: 192.168.17.1/24

E0: 10.0.216.1/24

Secret Password: testking

Answer:

```
Testking2> en
```

```
Testking2# config t
```

```
Testkingnat(config)#access-list 2 deny host 10.0.216.5
```

```
Testkingnat(config)#access-list 2 permit 10.0.216.0 0.0.0.255
```

```
Testkingnat(config)#ip nat pool nat 192.168.17.1 192.168.17.1 netmask 255.255.255.0
```

```
Testkingnat(config)#ip nat inside source list 2 pool nat overload
```

```
Testkingnat(config)#ip nat inside source static 10.0.216.5 192.168.17.5
```

```
Testkingnat(config)#interface ethernet 0
```

```
Testkingnat(config-if)#ip address 10.0.216.1 255.255.255.0
```

```
Testkingnat(config-if)#ip nat inside
```

```
Testkingnat(config-if)#interface serial 0
```

```
Testkingnat(config-if)#ip address 192.168.17.1 255.255.255.0
```

```
Testkingnat(config-if)#ip nat outside
```

```
Testkingnat(config-if)# ^ z
```

```
Testkingnat#copy running-config startup-config
```

Note, variation:**Name:** TestKingNAT**S0:** 192.168.212.1/24**E0:** 10.0.243.1/24**Secret Password:** testking**Web Server's External Address****192.168.212.5/24****Web Server's Internal Address****10.0.243.5/24**

Topic 3: Design (41 questions)

Section 1, Describe the three-layer hierarchical design model and explain the function of each layer: Access, Distribution and Core (4 questions)

QUESTION NO: 1

One of the virtues of a Link-State routing protocol is its hierarchical design. What advantages does this equate to in network design? (Select two)

- A. It allows link-state protocols to support VLSM.
- B. It allows them to support address summarization.
- C. It reduces the need to flood LSAs to all devices in the internetwork.
- D. Routers are no longer required to keep track of the topology of the entire autonomous system.

Answer: C, D

Explanation:

Hierarchical design can limit the requirement to notify all devices. The use of areas restricts the flooding to the logical boundary of the area rather than to all devices in the OSPF domain. Another advantage of hierarchical network design is that the entire network is divided into multiple smaller sections, so that much of the routing information is kept within a small area. In a non-hierarchical network, each router is required to maintain routing information for all other routers within the entire network.

Incorrect Answers:

A: Although link state protocols do indeed support VLSM, it is not due to a hierarchical network design. For example, RIP version 2 supports VLSM, but it is not hierarchical.

B: Although address summarization can indeed be configured on routers configured in a hierarchical design, it is an independent function and is not related to the fact that a tiered network design is being used.

QUESTION NO: 2

In regards to the three layered hierarchical network design, which of the following are responsibilities of the distribution layer? (Choose 2)

- A. Reliable transport structure
- B. Route redistribution

- C. Optimized transport structure
- D. Address aggregation
- E. Unauthorized entry access control lists

Answer: B, D

Explanation:

The distribution layer is responsible for what enters and exits a network, including the consolidation of traffic from multiple subnets into a core connection. Route redistribution and address aggregation happen at this layer.

Incorrect Answers:

- A, C: This is a function of the core layer.
- E. This is a function of the access layer.

QUESTION NO: 3

Which of the following are advantages of implementing a hierarchical IP addressing scheme when designing a large network? (Select two)

- A. Smaller routing tables
- B. Efficient address allocation
- C. Translation of private addresses
- D. Support for link-state routing protocols


Answer: A, B

Explanation:

Hierarchical IP addressing uses the addresses more conservatively, by planning for the correct use of subnet masks and route summarization. A hierarchical network design facilitates the use of IP routing summarization, which will reduce the size of the routing tables.

QUESTION NO: 4

The TestKing frame relay network is shown in the diagram below:



Based on the information given in this diagram, which of the following are true??

- A. The network is using a two-layer full-mesh hierarchical design.
- B. The network is using a two-layer hub-and-spoke hierarchical design.
- C. To improve scalability, route summarization at Routers TK1, TK2, TK3, TK4, TK5, and TK6 should be performed.
- D. Routers A and B are the distribution layer routers.
- E. Routers TK1, TK2, TK3, TK4, TK5 and TK6 are the access layer routers.
- F. To improve security, packet filtering that uses ACLs at Routers A and B must be implemented.

Answer: C

Routers TK1, TK2, TK3, TK4, TK5 and TK6 are the distribution layer routers. Route summarization should occur at the distribution layer.

Incorrect Answers:

- A: The network shown here is using a 3 level tiered approach: Core, Distribution, and Access.
- B: Three layers are used, not two.
- D: Routers A and B are the Core routers here.
- E: These routers make up the Distribution layer.

Section 2: Given specific requirements, choose the correct routing protocol to meet the requirements (14 questions)

QUESTION NO: 1

You are trying to determine the best routing protocol to use for the large TestKing network. Which routing protocols should you avoid when deploying a large network? (Select two)

- A. IGRP
- B. OSPF
- C. EIGRP
- D. RIP v.2
- E. RIP v.1

Answer: D, E

Explanation:

Both RIP version 1 and version 2 should only be used in relatively small networks. A large network would be congested by RIP broadcasts. These frequent broadcasts contain the entire routing table, and could saturate links.

QUESTION NO: 2

What are some of the features that are important for designing large, scalable networks? (Select three)

- A. A tiered network design model.
- B. Sufficient memory on the router.
- C. Multiple EIGRP autonomous systems.
- D. Good address space allocation scheme.

Answer: A, B, D

Explanation:

- A: A tiered network design model such as Core, Distribution, and Access is also needed for large networks.
- B: Sufficient capacity of the routers, in particular the memory, is required for large networks.
- D: Good allocation of address space is required- Each region should have a unique address space so route summarization is possible.

Incorrect Answers:

C: Configuring multiple EIGRP autonomous systems can add extra overhead and it provides no additional benefits.

QUESTION NO: 3

Two of the following routing protocols require a tiered, hierarchical topology. Which ones are they? (Select two)

- A. IS-IS
- B. IGRP
- C. OSPF
- D. EIGRP
- E. RIP v2

Answer: A, C

Explanation:

IS-IS and OSPF both require a hierarchical topology. The use of multiple areas and a core backbone area means that they are hierarchical by nature.

QUESTION NO: 4

What has become the leading Interior Routing Protocol (IGP) in use by Internet service providers?

- A. IS-IS
- B. OSPF
- C. EIGRP
- D. RIPv2
- E. BGP4

Answer: A

Explanation:

In recent years, the IS-IS routing protocol has become increasingly popular, with widespread use among Service Providers. It is a link state protocol, which enables very fast convergence with large scalability. It is also a very flexible protocol and has been extended to incorporate leading edge features such as MPLS Traffic Engineering. It is also chosen because it is an IETF standard based protocol, rather than a proprietary protocol such as IGRP or EIGRP.

Reference:

http://www.cisco.com/en/US/tech/tk365/tk381/technologies_white_paper09186a00800a3e6f.shtml

QUESTION NO: 5

Two of the following routing protocols support load balancing over unequal cost links. Which ones are they? (Select two)

- A. IGRP
- B. OSPF
- C. EIGRP
- D. RIP v2
- E. RIP v1

Answer: A, C

Explanation:

IGRP and EIGRP support unequal cost path load balancing, which is known as variance. OSPF, RIP v1 and RIP v2 do not support this.

Reference: <http://www.cisco.com/warp/public/103/19.html>

QUESTION NO: 6

You are required to choose a routing protocol that supports variable length subnet masks (VLSM) for your network. Which of the following would meet this requirement? (Choose three)

- A. IS-IS
- B. IGRP
- C. OSPF
- D. EIGRP
- E. RIP v1

Answer: A, C, D

Explanation:

IS-IS, EIGRP, OSPF, RIP V2 and static routes all support VLSM.

Incorrect Answers:

B: IGRP is the predecessor to EIGRP and does not support VLSM.

E: RIP version 1 is a distance vector routing protocol that does not support VLSM.

QUESTION NO: 7

Cisco routers automatically perform route summarization for various protocols. Which three of the following routing protocols are they? (Select three)

- A. IS-IS
- B. IGRP
- C. OSPF
- D. EIGRP
- E. RIP v.1

Answer: B, D, E

Explanation:

Sending route summaries – Routing information advertised out an interface is automatically summarized at major (classful) network address boundaries by RIP, IGRP, and EIGRP. Specifically, this autonomous summarization occurs for those routes whose classful network address differs from the major network address of the interface to which the advertisement is being sent.

Reference: Building Scalable Cisco Networks (Cisco Press) page 79

QUESTION NO: 8

In some situations, static routing is preferred over dynamic routing. In which situations would it be better to use a static routing protocol? (Select two)

- A. Medium to large networks with redundant paths.
- B. Networks with a single entry point.
- C. Low maintenance routing is required.
- D. Highly adaptable networks.
- E. High degree of control in path selection is required.

Answer: B, E

Explanation:

Networks with a single entry point only have one possible path to take for all networks, so it generally makes the most sense to simply add a single static default route out of that link. With static routes, there is not routing protocol overhead. Static routes may also be preferred when a high degree of control is required, since the entries are static and can not be overridden by a dynamic protocol.

Incorrect Answers:

- A: Larger networks with multiple links are best suited for dynamic routing protocols, since there would be too much administrative overhead with the use of static routes.
- C: Static routes require the most administrative maintenance, since everything is done manually.
- D: Statically routed networks are not adaptable, since static routes never change, regardless of any changes to the network topology. Only dynamic protocols can adapt to changing network conditions.

QUESTION NO: 9

One routing protocol considers different path attributes based on preconfigured network policies. Which of the protocols below is it?

- A. EIGRP
- B. OSPF
- C. RIPv2
- D. BGP4
- E. IS-IS
- F. IGRP

Answer: D

Explanation:

BGP uses numerous attributes. They are: origin, AS_path, next hop, MED, local preference, atomic aggregate, aggregator, community, originator ID, cluster list, and weight. Some of these values can be adjusted to give the administrator flexibility in creating a network routing policy. BGP differs in other routing protocols in that it is used for external use, so the attributes of paths are considered, rather than just the individual routes.

QUESTION NO: 10

You are the network administrator at TestKing. You are in the planning stages of upgrading the TestKing network, and need to decide which protocol to use. You need to choose a protocol that will meet the following goals:

- **Supports classless routing**
- **Supports VLSM**
- **Does not rely on TCP/IP to exchange routing information**
- **Performs auto-summarization by default**
- **Allows for manual route summarization**

Which routing protocol should you implement on the TestKing network?

- A. BGP
- B. EIGRP
- C. OSPF
- D. IS-IS
- E. IGRP
- F. RIPv2

Answer: B

Explanation:

EIGRP is a classless protocol that supports VLSM and performs auto-summarization of subnets by default. The subnet routes are summarized into a single network number aggregate. In addition, IP-EIGRP will allow aggregation on any bit boundary in an IP address and can be configured on a network interface for additional granularity. EIGRP also supports automatic summarization of network addresses at major network borders. There is an automatic redistribution mechanism used so IGRP routes are imported into EIGRP and vice versa. Since the metrics for both protocols are directly translatable, they are easily comparable as if they were routes that originated in their own AS.

QUESTION NO: 11

Which routing protocol: (1) is intended to support large routing domains consisting of combinations of many media types, (2) may be used as an IGP, and (3) supports multiple routed protocols in an integrated manner rather than “Ships in the Night” method?

- A. EIGRP
- B. OSPF
- C. IS-IS
- D. BGP
- E. None of the above

Answer: A

Explanation:

EIGRP is a very scalable Interior Routing Protocol (IGP) that supports large networks. In addition to support for IP, EIGRP can be used for routing IPX and Appletalk networks.

Incorrect Answers:

B, C: Although OSPF and IS-IS are link state protocols that are used as IGPs and are highly scalable, they do not provide support for routed networks such as those used by Novell and Apple.

D: BGP only supports IP, and is generally used only for external routing.

QUESTION NO: 12

Match the correct routing protocol on the left to its characteristics on the right side. Note that not all answer choices will be used.

OSPF	BGP	Uses NET to identify the router	Place here
RIPv1	MPLS	Maintains feasible successor(s)	Place here
IGRP	EIGRP	Uses path attributes	Place here
ISIS	RIPv	Supports special areas like NSSA	Place here

Answer

		Uses NET to identify the router	ISIS
RIPv1	MPLS	Maintains feasible successor(s)	EIGRP
IGRP		Uses path attributes	BGP
	RIPv	Supports special areas like NSSA	OSPF

QUESTION NO: 13

ODR is being used as the routing protocol for a new TestKing network. Which network topology is best suited for ODR?

- A. highly redundant
- B. fully meshed
- C. dual self-healing ring
- D. hub and spoke
- E. NBMA

Answer: D

Explanation:

On Demand Routing (ODR) is a mechanism that provides minimum-overhead IP routing for stub sites. The overhead of a general dynamic routing protocol is avoided, without incurring the configuration and management overhead of using static routing.

A stub router is the peripheral router in a hub and spoke network topology. Stub routers commonly have a WAN connection to the hub router and a small number of LAN network segments (stub networks) that are connected directly to the stub router. To provide full connectivity, the hub routers can be statically configured to know that a particular stub network is reachable via a specified access router. However, if there are multiple hub routers, many stub networks, or asynchronous connections between hubs and spokes, the overhead required to statically configure knowledge of the stub networks on the hub routers becomes too great.

ODR simplifies installation of IP stub networks in which the hub routers dynamically maintain routes to the stub networks. This is accomplished without requiring the configuration of an IP routing protocol at the stub routers. With ODR, the stub advertises IP prefixes corresponding to the IP networks that are configured on its directly connected interfaces. Because ODR advertises IP prefixes, rather than IP network numbers, ODR is able to carry Variable Length Subnet Mask (VLSM) information. Once ODR is enabled on a hub router, the router begins installing stub network routes in the IP forwarding table. The hub router can also be configured to redistribute these routes into any configured dynamic IP routing protocols. IP does not need to be configured on the stub router. With ODR, a router is automatically considered to be a stub when no IP routing protocols have been configured on it.

QUESTION NO: 14

Your company, TestKing Inc, has a central office and a remote branch office. Each site has a separate autonomous system (AS) as well as a separate Internet connection. In this network, what is the best method to route Internet traffic.

- A. IGP routing
- B. BGP routing
- C. Configure `ip route prefix mask`
- D. Configure `ip default-gateway ip address`
- E. Configure IS-IS

Answer: B

Explanation:

Border Gateway Protocol (BGP) is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP). In order to take full advantage of multiple Internet connections, as well as multiple Internet providers, BGP is the best choice to ensure full redundancy.

Reference: http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

Incorrect Answers:

A: Customer networks usually employ an Interior Gateway Protocol (IGP) such as RIP or OSPF for the exchange of routing information within their networks.

C, D: Using statically assigned default routes and default gateways is a simplistic approach that works well in many situations, but will not provide for automatic failover and does not provide the routing options that BGP provides.

E: IS-IS is used exclusively as an IGP.

Section 3: Identify the correct IP addressing scheme, including features of IPv6 (8 questions)

QUESTION NO: 1

The TestKing network is in the process of migrating the IP address scheme to use IPv6. Which of the following address types are associated with IPv6? (Select three)

- A. Unicast
- B. Private
- C. Broadcast
- D. Public
- E. Multicast
- F. Anycast

Answer: A, E, F

Explanation:

IP version 6 introduced the concept of anycasts. The three IP address types used for IPv6 are:

- **Unicast** - An IPv6 unicast address is an identifier for a single interface, on a single node. A packet that is sent to a unicast address is delivered to the interface identified by that address.
- **Anycast** - An anycast address is an address that is assigned to a set of interfaces that typically belong to different nodes. A packet sent to an anycast address is delivered to the closest interface—as defined by the routing protocols in use—identified by the anycast address.
- **Multicast** - An IPv6 multicast address is an IPv6 address that has a prefix of FF00::/8 (11111111). An IPv6 multicast address is an identifier for a set of interfaces that typically belong to different nodes.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide_chapter09186a0080110dd2.html#99899

QUESTION NO: 2

What number is a valid representation for the 200F:0000:AB00:0000:0000:0000/56 IPv6 prefix?

- A. 200F:0:0:AB/56
- B. 200F:0:AB00::/56
- C. 200F::AB00/56
- D. 200F:AB/56

Answer: B**Explanation:**

Due to the method of allocating certain styles of IPv6 addresses, it will be common for addresses to contain long strings of zero bits. In order to make writing addresses containing zero bits easier a special syntax is available to compress the zeros. The use of "::" indicates multiple groups of 16-bits of zeros. The "::" can only appear once in an address. The "::" can also be used to compress the leading and/or trailing zeros in an address.

For example the following addresses:

1080:0:0:0:8:800:200C:417A a unicast address
 FF01:0:0:0:0:0:0:43 a multicast address
 0:0:0:0:0:0:0:1 the loopback address
 0:0:0:0:0:0:0:0 the unspecified addresses

may be represented as:

1080::8:800:200C:417A a unicast address
 FF01::43 a multicast address
 ::1 the loopback address
 :: the unspecified addresses

In our example, the trailing zero's can be compressed into ::, giving us 200F:0:AB00:: as the address.

QUESTION NO: 3

You have the binary IP address, {11000000.10100100.11000000.00000001 }. What class of IP address is it, and what is a characteristic of that class?

- A. It is a Class B public address.
- B. It is a Class C public address.
- C. It is a Class B private address
- D. It is a Class C private address.
- E. It is a Class D experimental address.

Answer: B**Explanation:**

11000000.10100100.11000000.00000001 = **192.164.192.1** = A public Class C address

Incorrect Answers:

- A, C: Class B addresses begin with a "10" in the first octet.
- D. Class C private (RFC 1918) addresses are in the IP range 192.168.0.0-192.168.255.255

Reference: Building Scalable Cisco Networks (Cisco Press) page 66

QUESTION NO: 4

You are the network administrator at TestKing. TestKing is assigned a network of 200.10.5.0/24. You want to create 5 subnets from the assigned address.

Which subnet mask will give you 5 subnets and the most possible hosts per subnet?

- A. /21
- B. /23
- C. /26
- D. /27
- E. /29

Answer: D

Explanation:

We need 5 subnets from the given address 200.10.5.0/24. It means we will do 3 bits of subnetting which will provide us 8 subnets as $2^N = 2^3 = 8$ subnets where n=number of subnet bits. So, we will have a total of $24+3 = 27$ network bits.

QUESTION NO: 5

You are the network engineer at TestKing. TestKing has a branch office with five devices that require public IP addresses.

Which of the following subnet masks will provides the most efficient allocation of addresses?

- A. 255.255.255.128
- B. 255.255.255.248
- C. 255.255.255.252
- D. 255.255.255.254

Answer: B

Explanation:

Requirement = 5 public IP addresses to connect 5 devices
3-bits of subnetting in the host portion will provide $2^3 = 8$ addresses. After we subtract 2 of these addresses, for the network and broadcast IP address, we are still left with 6 useable IP addresses for each network.

QUESTION NO: 6

Which three are characteristics of IPv6? Select three.

- A. An IPv6 address is 128 bits long.
- B. An IPv6 header is 20 bits long.
- C. An IPv6 header contains the next header field.
- D. An IPv6 header contains the next protocol field.
- E. IPv6 routers send RA messages.
- F. An IPv6 header contains the header checksum field.

Answer: A, C, E

Explanation:

In IPv6, the 128 bit IP address is separated to two: network prefix, which identifies network, and interface ID, which identifies a node (interface). Interface ID is configured by the node on its own, and prefix is notified by the network (usually router). These two are combined to form an IPv6 address. IPv6 packet headers contain many of the fields found in IPv4 packet headers; some of these fields have been modified from IPv4. The next header field is the next extension header to examine. This was previously the protocol field in IPv4.

RA stands for router advertisements, used by IPv6 routers for the purpose of automatic address configuration.

QUESTION NO: 7

Which address type does the IPv6 address FF05:0:0:0:0:0:2 specify?

- A. unspecified
- B. aggregable global unicast
- C. link local
- D. site local unicast
- E. multicast

Answer: E

Explanation:

All IPv6 multicast addresses start with the first 8 bits of the address set to 1. Thus all multicast addresses start with the hexadecimal notation FF (1111 1111). The multicast range is as follows:

FF00::/8

FFFF::/8

The second octet, following to the first octet of FF, identifies both the scope and the lifetime of the multicast address. In this way, IPv6 has millions of group multicast addresses to use in current and emerging technologies.

Reference: CCNP Self-Study CCNP BSCI Exam Certification Guide p.112

QUESTION NO: 8

Which command must be globally enabled on a Cisco router to support IPv6?

- A. ip routing ipv6
- B. ipv6 unicast routing
- C. ipv6 routing
- D. ip classless
- E. ipv6 cef

Answer: B

Section 4: Describe the concepts relating to route summarization and apply them to hypothetical scenarios (15 questions)

QUESTION NO: 1

Route aggregation is a way of reducing the number of entries in a routing table. Three of the routing protocols below support automatic route summarization. Which ones are they? (Select three)

- A. IS-IS
- B. EIGRP
- C. OSPF
- D. IGRP
- E. RIP v.1

Answer: B, D, E

Explanation:

B: IGRP automatically summarize routes on the class network boundary.

D: EIGRP support both manual and automatic route summarization.

E: RIPv1 automatically summarize routes on the class network boundary.

Incorrect Answers:

A, C: IS-IS and OSPF only support manual route summarization.

QUESTION NO: 2

The TestKing network consists of the following IP subnets:

Network Number
192.168.31.0/24
192.168.32.0/24
192.168.33.0/24
192.168.34.0/24
192.168.35.0/24
192.168.36.0/24
192.168.37.0/24
192.168.38.0/24
192.168.39.0/24
192.168.40.0/24

How many CIDR blocks will be needed to summarize these subnets in the TestKing network?

- A. Two
- B. Three
- C. Four
- D. Five
- E. None of the above

Answer: B

Explanation:

To properly summarize 10 contiguous networks, they will have to be broken into three separate advertisements: one will consist of 8 subnets (the middle eight) and the other two will consist of the first and last subnets. The summarizations are as follows:

192.168.31.0/24

192.168.32.0/21

192.168.40.0/24

Reference: Building Scalable Cisco Networks (Cisco Press) page 85

QUESTION NO: 3

You are in the process of manually summarizing some routes within your network. Which of the addresses below would be included in the manual route summarization of 153.25.200.0/21?

- A. 153.25.198.0
- B. 153.25.206.0
- C. 153.25.208.0
- D. 153.25.224.0

Answer: B

Explanation:

The 21 leftmost bits must match. In particular we should study the 3rd octet.

Decimal	1 st Octet	2 nd Octet	3 rd Octet	4 th Octet
153.25.200.0	10011001	00011001	11001000	00000000
153.25.198.0	10011001	00011001	11000110	00000000
153.25.206.0	10011001	00011001	11001110	00000000
153.25.208.0	10011001	00011001	11010000	00000000
153.25.224.0	10011001	00011001	11100000	00000000

QUESTION NO: 4

Which of the following statements about route summarization are true? (Select two)

- A. Private addresses cannot be summarized.
- B. Summarization is not compatible with VLSM.
- C. RIP v.1 automatically summarizes routes on network class boundaries.
- D. Classless routing protocols support summarization on any address bit boundary.

Answer: C, D

Explanation:

C: RIP V.1 is only able to summarize routes on network class boundaries.

D: Classless routing protocols support summarization on any address bit boundary.

Incorrect Answers:

A: Private address ranges can indeed be summarized.

B: Summarization is indeed compatible with VLSM.

QUESTION NO: 5

Some of the terms below are synonymous with route summarization? Which ones are they? (Select two)

- A. supernetting
- B. route aggregation
- C. address translation
- D. Classful inter-domain routing

Answer: A, B

Explanation:

Route summarization, also called route aggregation or supernetting, can reduce the number of routes that a router must maintain because it is a method of representing a series of network numbers in a single summary address.

QUESTION NO: 6

Your network is using the Class B private IP address scheme, as defined in RFC 1918. If you wanted to summarize this entire private address space into one single supernet, which prefix should you use?

- A. /8
- B. /12
- C. /16

D. /24

Answer: B

Explanation:

In Class B, the private address is 172.16.0.0 to 172.31.255.255

Four bits are required to take care of the range 16-31. So, from the default subnet mask of 16, four more taken out and so, it becomes 12.

QUESTION NO: 7

In an EIGRP network using VLSM, where in the network does the process of route summarization occur?

- A. Manually on any router interface.
- B. Only at classless network boundaries.
- C. Only at classful network boundaries.
- D. Dynamically at the supernet boundary.

Answer: A

Explanation:

EIGRP automatically summarizes at the network boundary by default, but it can also be configured manually on any of the routed interfaces.

Incorrect Answers:

B, C: Since the word “only” is used, these two choices are incorrect. EIGRP can be configured to summarize automatically at the network boundary, or manually.

D: EIGRP provides no means for summarizing routes dynamically.

QUESTION NO: 8

Which one of the following addresses can be included in the route summarization of 172.17.200.0/21?

- A. 172.17.198.0
- B. 172.17.206.0
- C. 172.17.217.0
- D. 172.17.224.0

Answer: B

Explanation:

Here, we list the network addresses in binary form and see which subnet address match the 21 leftmost bits of the route summarization.

Decimal	1 st Octet	2 nd Octet	3 rd Octet	4 th Octet
172.17.200.0	10101100	00010001	11001000	00000000
172.17.198.0	10101100	00010001	11000110	00000000
172.17.206.0	10101100	00010001	11001110	00000000 Match!
172.17.217.0	10101100	00010001	11011001	00000000
172.17.224.0	10101100	00010001	11100000	00000000

QUESTION NO: 9

The TestKing network contains the following 4 IP subnets:

172.18.129.0/24
 172.18.130.0/24
 172.18.132.0/24
 172.18.133.0/24

You want to summarize these networks into a single supernet. Which one of the following route summarizations is capable of summarizing every one of the networks in the above exhibit?

- A. 172.18.128.0/21
- B. 172.18.128.0/22
- C. 172.18.130.0/22
- D. 172.18.132.0/20

Answer: A

Explanation:

To illustrate, we list the network addresses in binary form and see how many leftmost bits match.

Decimal	1 st Octet	2 nd Octet	3 rd Octet	4 th Octet
172.18.129.0	10101100	00010010	10000001	00000000
172.18.130.0	10101100	00010010	10000010	00000000
172.18.132.0	10101100	00010010	10000100	00000000
172.18.133.0	10101100	00010010	10000101	00000000
172.18.128.0	10101100	00010010	10000000	00000000

We see that the 21 leftmost bits match and that 172.18.128.0/21 can summarize the four networks.

Incorrect Answers:

B, C: Only the 21 leftmost bits match, not 22.

D: This is an illegal summarization. It is host address, not a network address.

QUESTION NO: 10

The TestKing network is using the 172.19.100.0/24 IP address range for its entire network and wishes to summarize these routes. How many Class C IP subnets could be summarized in the 172.19.100.0/22 supernet?

- A. 2
- B. 4
- C. 8
- D. 16
- E. 255
- F. None of the above

Answer: B

Explanation:

As the summarized route is /22 the first 22 bits must be the same. A Class C subnet mask has 24 bits so the lower 2 bits in the 3rd octet can be changed. There are 4 possible combinations (00, 01, 10, 11) so 172.19.100.0/22 can summarize 4 class C addresses.

Reference: Source: Cisco Press, CCNP Self-Study Building Scalable Cisco Internetworks (BSCI). Chapter 2 Extending IP Addresses, Calculating VLSMs, page 73

QUESTION NO: 11

You are administrating a network with subnets ranging from 172.16.94.0/24 – 172.16.118.0/24. A summary statement with the network 172.16.96.0 and the mask of 255.255.240.0 have summarized some of your subnets.

Which subnet is the last subnet to be included in the summary network?

Answer: 172.16.111.0

Explanation:

We can verify this answer as follows:

111 (3rd octet) decimal is 01101111. This is within the summarization. However, 172.16.112.0 would not be ok: 112 decimal is 01110000 (see below).

Decimal	1 st Octet	2 nd Octet	3 rd Octet	4 th Octet	
255.255.240.0	11111111	11111111	11110000	00000000	subnet mask
172.16.96.0	10101100	10101100	01100000	00000000	summarized network
172.16.111.0	10101100	10101100	01101111	00000000	Included in summary element

172.16.112.0 10101100 10101100 01110000 00000000 Not included in summary element

Therefore, 172.16.111.0 is the last IP subnet that can be included in this summary route.

QUESTION NO: 12

What strategy can a network administrator use to minimize the effect of routing table updates on internal routers when a single WAN interface frequently changes its state from up to down?

- A. Use a distance vector routing protocol.
- B. Use private IP addresses.
- C. Use dial-on-demand routing.
- D. Use route summarization.
- E. Use a routing protocol that tolerates route flapping.
- F. All of the above.

Answer: D

Explanation:

When using route summarization, multiple IP networks are advertised to downstream routers as a single route. As long as any of the IP subnets included within the summarized route are up and active, the summary route will be used, so any single link problem will not have to be propagated to the rest of the network.

QUESTION NO: 13

The TestKing EIGRP network is shown below:



```
TestKing1 # show ip route 2.0.0.0
Routing entry for 2.0.0.0/8, 4 known subnets
  Attached (2 connections)
  Variably subnetted with 2 masks
  Redistributing via eigrp 2000

C       2.1.3.0/24 is directly connected, Serial2
D       2.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1
D       2.1.0.0/8 is a summary, 00:23:20, Null0
C       2.1.1.0/24 is directly connected, Serial1
```

Router TestKing1 was configured in the following manner:

```
TestKing1(config)# router eigrp 2000
TestKing1(config-router)# network 2.0.0.0
TestKing1(config-router)# exit
TestKing1(config)# interface serial 2
TestKing1(config-if)# ip summary-address eigrp 2000 2.1.0.0 255.255.0.0
```

Based on the above information, how many routes will TestKing2 receive from router TestKing1?

- A. None
- B. One
- C. Two
- D. Three
- E. Four

Answer: B

Explanation:

Router TestKing1 was configured to summarize the network routes when advertising to the neighbor TestKing2. Because of this, router TestKing2 will only receive the 2.1.0.0/16 EIGRP route from router TestKing1.

QUESTION NO: 14

Which of the following statements are true regarding EIGRP manual summarization? (Select two)

Leading the way in IT testing and certification tools, www.testking.com

- A. Manual summarization is configured on a per interface basis.
- B. Manual summaries can be configured with the classful mask only.
- C. When manual summarization is configured, auto-summarization is automatically disabled by default.
- D. The summary address is assigned an administrative distance of 10 by default.
- E. The summary address is entered into the routing table and is shown to be sourced from the Null0 interface.

Answer: A, E

Explanation:

By default, EIGRP routers summarize routes at the classful boundary. When routes are manually summarized, they are done so while in interface configuration mode. After the manual summarization is done, the logical null 0 interface is created, and the summary network is installed into the routing table via this interface.

QUESTION NO: 15

The TK1 and TK2 networks are consolidating into one large network. To accommodate this, the IP network was separated in the middle by a different major network boundary, creating a discontinuous network. With EIGRP being used as the routing protocol, what can be done to provide full connectivity across these non-contiguous networks? (Select two)

- A. Use private addresses on LAN links.
- B. Use static route redistribution.
- C. Use the **no auto-summary** command.
- D. Use route summarization.
- E. All of the above

Answer: B, C

Explanation:

Redistributing static route information could be used to advertise the routes across the network, since these static routes will not be automatically summarized.

By default, EIGRP summarizes IP subnets at the network boundary. However, in order to support the use of non-contiguous networks, the “no auto-summary” command can be used to disable the feature.

Incorrect Answers:

- A: Using private or public addresses on any of the links will not correct this issue.
- D: Route summarization could result in duplicate network advertisements. We want to prevent this by disabling route summarization.

Topic 4: Troubleshooting (50 questions)

Section 1: Identify the steps to verify OSPF operation in a single area (7 questions)

QUESTION NO: 1

You are troubleshooting an OSPF problem and you need to view the OSPF neighbor information. Which of the following commands should you enter to examine neighbor adjacencies? (Select two)

- A. `show ip ospf database`
- B. `show ip ospf neighbors`
- C. `show ip ospf protocols`
- D. `show ip ospf interfaces`

Answer: B, D

Explanation:

B: Using the **show ip ospf neighbor** command, you can observe the neighbor data structure. This command displays OSPF-related neighbor information. The Interface field shows the interface on which the OSPF neighbor has formed adjacency.

Sample:

```
RouterTK2#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.45.1	1	FULL/DR	00:00:36	10.0.0.1	Ethernet0

D: The **show ip ospf interface** command displays area ID and adjacency information

Incorrect Answers:

- A: The “show ip ospf database” command displays the link-state database.
- C: This is an invalid command.

Reference: <http://www.cisco.com/warp/public/104/16.html>

QUESTION NO: 2

You are a network troubleshooter and you’ve been contracted to examine an OSPF network. You feel that the CPU utilization of your routers is too high and want to figure out how many SPF calculations have occurred. Which one of the commands below could you use?

- A. `show ip ospf`

- B. show ip route
- C. show ip ospf interface
- D. show ip ospf protocols

Answer: A

Explanation:

The **show ip ospf** command displays summary information regarding the global OSPF configuration. The output includes the number of times the Shortest Path First (SPF) algorithm has been run.

Sample output:

```
routerTestK#show ip ospf
OSPF is running, process id: 1, router id: 10.1.2.136
Number of areas: 1, normal: 1, stub: 0
Area: 1.2.3.4
Number of interfaces in this area is 1
Type of authentication none
SPF algorithm has run 3 times
SPF interval 5 seconds
```

Incorrect Answers:

- B: The **show ip route** command displays IP routing table entries.
- C: The **show ip interface** command displays information about one or more interfaces.
- D: There is no such command.

QUESTION NO: 3

Which of the following IOS commands will allow you to view when a topological database purges out of date routes?

- A. show ip ospf
- B. show ip route
- C. show ip ospf interface
- D. show ip ospf protocols
- E. show ip ospf neighbor
- F. None of the above

Answer: C

Explanation:

The OSPF route dead timer is used to purge the outdated routes. The dead timer information can be obtained via the “show ip ospf interface” command.

Example:

```
RouterTK1# show ip ospf interface ethernet 0
Ethernet0 is up, line protocol is up
```

```

Internet Address 10.10.10.1/24, Area 0
Process ID 1, Router ID 192.168.45.1, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State BDR, Priority 1
Designated Router (ID) 172.16.10.1, Interface address 10.10.10.2
Backup Designated router (ID) 192.168.45.1, Interface address 10.10.10.1
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 2, maximum is 2
Last flood scan time is 0 msec, maximum is 4 msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 172.16.10.1 (Designated Router)
Suppress hello for 0 neighbor(s)

```

Timer Intervals

These are the values of the OSPF timers:

- Hello—Interval time in seconds that a router sends an OSPF hello packet. On broadcast and point-to-point links, the default is 10 seconds. On NBMA, the default is 30 seconds.
- Dead—Time in seconds to wait before declaring a neighbor dead. By default, the dead timer interval is four times the hello timer interval.
- Wait—Timer interval that causes the interface to exit out of the wait period and select a DR on the network. This timer is always equal to the dead timer interval.
- Retransmit—Time to wait before retransmitting a database description (DBD) packet when it has not been acknowledged.
- Hello Due In—An OSPF hello packet is sent on this interface after this time. In this example, a hello is sent three seconds from the time the show ip ospf interface is issued.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094056.shtml#timer

QUESTION NO: 4

You are troubleshooting a convergence problem on an OSPF network. Which of the following factors can affect the OSPF convergence time? (Select three)

- Hold-down timers.
- The size of the network.
- The number of routing areas.
- Maximum allowed hop count.
- The route calculation algorithm.

Answer: B, C, E

Explanation:

B: The size of the network affects the size of the topology table on the routers. A larger topology table will correspond to a higher convergence time.

- C: Areas are introduced to put a boundary on the total number of link-state updates. The number of areas affects the convergence time.
- E: The route calculation algorithm used for OSPF affects the convergence time.

Incorrect Answers:

- A: OSPF does not use hold-down timers. Instead topology changes are flooded immediately. Distance vector routing protocols use hold-down timers.
- D: OSPF does not use hop count. Distance vector routing protocols use hop counts.

QUESTION NO: 5

Which IOS command could you use to see which networks are being routed by an OSPF process?

- A. show ospf
- B. show ip route
- C. show ip protocols
- D. show ip ospf database
- E. show ip ospf protocols

Answer: C

Explanation:

The **show ip protocols** command displays current routing protocols. It displays the parameters and current state of the active routing protocol process. The output includes a list of the networks routing for individual ospf processes.

Sample output:

```
RtTestKing# show ip protocols
Routing Protocol is "ospf 200"
Sending updates every 0 seconds
Invalid after 0 seconds, hold down 0, flushed after 0
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Redistributing: ospf 200
Routing for Networks:
172.6.31.5/32
Routing Information Sources:
Gateway Distance Last Update
Distance: (default is 110)
```


Incorrect Answers:

- A: The **show ospf** command displays summary information regarding the global OSPF configuration.
- B: The **show ip route** command displays the IP routing table.
- D: The **show ip ospf database** command displays the contents of the topological database maintained by the router. The command also shows the router ID and the OSPF process ID. However, the output does not include the networks routing for individual ospf processes.

E: This is an invalid command.

QUESTION NO: 6

The TestKing OSPF network is displayed in the following diagram.



In the TestKing network, if the link to the designated router were to go down, what will Router TK1 and Router TK3 do?

- A. They send a flash update with poison reverse.
- B. They perform the designated router election process.
- C. They re-broadcast their routing tables to all other neighboring routers.
- D. They send a query to neighboring routers for other routers to the failed link.
- E. None of the above

Answer: E

Explanation:

A Designated Router (DR) router is elected by all other routers on the same LAN to represent all the routers. Each network has one DR. Normally, in this case the BDR would take over as the DR. In this case, the router that was already acting as the BDR (either TK1 or TK3) would take over as the DR. A new BDR election would then take place.

Incorrect Answers

- A: IGRP and EIGRP use poison reverse. Poison reverse is not used by OSPF.
- B: The router that was already acting as the BDR would take over as the DR.
- C, D: This is not a necessary step in the OSPF process.

QUESTION NO: 7

When issuing the “show ip ospf neighbor” command on router TK1, a neighbor is reported as a DROTHER. Which of the following is true regarding this OSPF neighbor?

- A. The neighbor is down and currently not forwarding traffic.
- B. It is not in full state and is still exchange LSAs.
- C. It is the backup DR:
- D. It is neither the DR nor BDR.
- E. It is acting as both the DR and the BDR

Answer: D

Explanation:

When viewing the OSPF neighbor relationships, a neighbor can be one of three different values:

DR - This neighbor is the DR

BDR – This neighbor is the backup DR

DROTHER – This neighbor is neither the DR nor the BDR.

Section 2: Identify verification methods which ensure proper operation of Integrated IS-IS on Cisco routers (4 questions)

QUESTION NO: 1

You're troubleshooting a router that's running Integrated IS-IS and you suspect something abnormal with the neighboring routers. Which IOS command would you enter to display the adjacencies?

- A. show isis route
- B. show clns route
- C. show isis database
- D. show clns neighbors
- E. None of the above

Answer: D

Explanation:

The **show clns neighbors** command displays ES and IS neighbors. The output includes adjacency information.

Example:

```
routerTestK# show clns neighbors
System Id      SNPA          Interface State Holdtime Type Protocol
0000.0000.0007 aa00.0400.6408 Ethernet0 Init 277      IS    ES-IS
0000.0C00.0C35 0000.0c00.0c36 Ethernet1 Up 91       L1    IS-IS
0800.2B16.24EA aa00.0400.2d05 Ethernet0 Up 29       L1L2  IS-IS
0800.2B14.060E aa00.0400.9205 Ethernet0 Up 1698     ES    ES-IS
0000.0C00.3E51 *HDLC*       Serial1    Up 28       L2    IS-IS
0000.0C00.62E6 0000.0c00.62e7 Ethernet1 Up 22       L1    IS-IS
0A00.0400.2D05 aa00.0400.2d05 Ethernet0 Init 24       IS    ES-IS
```

Incorrect Answers

- A: The **show isis routes** command is used to display the IS-IS Level 1 forwarding table for IS-IS learned routes.
- B: The **show clns route** command is used to display all of the destinations to which this router knows how to route packets.
- C: The **show isis database** command is used to display the Intermediate System-to-Intermediate System (IS-IS) link state database.

QUESTION NO: 2

The IS-IS routing table for TestKing1 is displayed in the diagram below:

```

TestKing1#show ip route isis
10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
i L2  10.200.200.14/32 [115/30] via 10.1.0.2, Serial1/0
i L1  10.200.200.13/32 [115/20] via 10.1.1.3, Ethernet0/0
i L1  10.1.3.0/24 [115/20] via 10.1.1.3, Ethernet0/0
i L2  10.1.2.0/23 [115/20] via 10.1.0.2, Serial1/0
i su  10.1.0.0/23 [115/10] via 0.0.0.0, Null0

```

Regarding the command output on TestKing1 in the exhibit, which statement is true?

- A. The 10.1.0.0/23 route is an IS-IS external route
- B. The TestKing1 IS-IS router is an ASBR.
- C. The 10.1.0.0/23 route is a suppressed route.
- D. The 10.1.0.0/23 route is a summary route.
- E. The TestKing1 IS-IS router is an ASBR that belongs to multiple IS-IS areas.
- F. The TestKing1 IS-IS router is performing route aggregation and is suppressing the more specific 10.1.0.0/23 prefix.

Answer: D

Explanation:

The “su” in the routing table denotes an IS-IS summary null route. This route is automatically created within the router when IS-IS route summarization is configured.

QUESTION NO: 3

The “show isis route” command was issued on router TestKing3 as shown below:

```

TestKing3#show isis route
IS-IS Level-1 Routing Table
System Id      Next-Hop      Interface      SNPA          Metric  State
TestKing3     TestKing1     Et0/0         aabb.cc00.3300  10     Up    L2-IS
TestKing1     --

Default route out of area - (via 1 L2-attached IS)
System Id      Next-Hop      Interface      SNPA          Metric  State
TestKing1     TestKing1     Et0/0         aabb.cc00.3300  10     Up

```

Based on the information above, which statement is true?

- A. TestKing1 is the exit point out of the area for TestKing3
- B. TestKing1 is a level-1 only IS-IS router.
- C. TestKing1 has been configured with a non-default IS-IS metric.
- D. TestKing3 is a level-2 only IS-IS router.

- E. TestKing3 has been configured with a non-default IS-IS metric.
- F. TestKing3 routing table should contain L2 entries.
- G. None of the above

Answer: A

Explanation:

In the IS-IS routing table of router TestKing3, the default router is displayed as TestKing1, which is a Level 2 router. Level 2 routers are used for routing inter-area traffic, so TestKing3 will use this router as the exit point for all traffic destined to another area.

Incorrect Answers:

B: TestKing1 is a Level 2 router.

C, E: The default IS-IS metric for all links is 10, which is the value shown here.

D, F: TestKing3 is a level 1 router.

QUESTION NO: 4

The routing table of TestKing1 is displayed below:

TestKing1# show ip route

Gateway of last resort is not set

140.140.0.0 is subnetted (mask is 255.255.255.0), 3 subnets

```
C      140.140.64.0 255.255.255.0 is possibly down, routing via 0.0.0.0, Ethernet 0
i L2   140.140.67.0 [115/20] via 140.140.64.240, 0:00:12, Ethernet 0
i L2   140.140.66.0 [115/20] via 140.140.64.240, 0:00:12, Ethernet 0
```

Based on this information, how was network 140.140.66.0 discovered?

- A. It is an IGRP route that has been redistributed in from an external protocol.
- B. It is an IS-IS route via inter-area routing.
- C. It is an IS-IS route via intra-area routing.
- D. It is an IBGP route.
- E. It is a static route.

Answer: B

Explanation:

Level-2 routers are used for routing traffic between areas (inter-area). Based on the routing table of TestKing1, the 140.140.66.0/24 network is known via a Level 2 IS-IS router.

Section 3: Identify the steps to verify OSPF operation in multiple areas (8 questions)

QUESTION NO: 1

You have a multi-area OSPF network and you're concerned because one of the sites is having connectivity problem to resources in a different area. Which IOS privileged mode command would you enter to confirm that your network: A) has a path to its ABR, B) has a path to its ASBR, and C) the SPF calculation is functional?

- A. show ip protocols
- B. show running-config
- C. show ip ospf neighbor
- D. show ip ospf border-routers

Answer: D

Explanation:

The **show ip ospf border-routers** command displays the internal OSPF routing table entries to an area border router (ABR) and autonomous system boundary router (ASBR). The SPF No in the output is the internal number of SPF calculation that installs this route.

Example:

```
RouterTestKing# show ip ospf border-routers
```

```
OSPF Process 109 internal Routing Table
```

Destination No	Next Hop	Cost	Type	Rte Type	Area	SPF
160.89.97.53	144.144.1.53	10	ABR	INTRA	0.0.0.3	3
160.89.103.51	160.89.96.51	10	ABR	INTRA	0.0.0.3	3
160.89.103.52	160.89.96.51	20	ASBR	INTER	0.0.0.3	3
160.89.103.52	144.144.1.53	22	ASBR	INTER	0.0.0.3	3

Incorrect Answers:

- A: The **show ip protocols** command only displays routing protocol parameters and current timer values.
- B: The **show running-config** command displays the currently used configuration mode. The required information will not be displayed.
- C: The **show ip ospf neighbor** command displays OSPF-neighbor information on a per-interface basis. It does not include ABR, ASBR or SPF information.

QUESTION NO: 2

You have logged onto the console of a router running OSPF, and for your analysis log you need information on: the filters, default metric, maximum paths, and the number of areas configured. Which of the IOS commands below should you use?

- A. show ip protocol
- B. show ip route
- C. show ip ospf interface
- D. show ip ospf
- E. Show ip routing

Answer: A

Explanation:

The **show ip protocols** command, displays parameters about timers, filters, metrics, network, and other information for the entire router.

Reference: Building Scalable Cisco Networks (Cisco Press) page 133

QUESTION NO: 3

Which IOS command tells you how many times the OSPF shortest path first algorithm was executed?

- A. show ip protocol
- B. show ip ospf interface
- C. show ip ospf
- D. show ip ospf database
- E. None of the above

Answer: C

Explanation:

The following table describes the output of the “show ip ospf” command and their meanings:

Table 52 show ip ospf Field Descriptions	
Field	Description
Routing process "ospf 201" with ID 10.0.0.1	Process ID and OSPF router ID.
Supports...	Number of types of service supported (Type 0 only).
SPF schedule delay	Delay time of SPF calculations.
Minimum LSA interval	Minimum interval between link-state

	advertisements.
LSA group pacing timer	Configured LSA group pacing timer (in seconds).
Interface flood pacing timer	Configured LSA flood pacing timer (in milliseconds).
Retransmission pacing timer	Configured LSA retransmission pacing timer (in milliseconds).
Number of...	Number and type of link-state advertisements that have been received.
Number of external LSA	Number of external link-state advertisements.
Number of opaque AS LSA	Number of opaque link-state advertisements.
Number of DCbitless external and opaque AS LSA	Number of demand circuit external and opaque link-state advertisements.
Number of DoNotAge external and opaque AS LSA	Number of do not age external and opaque link-state advertisements.
Number of areas in this router is	Number of areas configured for the router.
External flood list length	External flood list length.

Reference: http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/iprrp_r/ip2_s3g.htm#wp1036469

QUESTION NO: 4

An OSPF link can be in multiple states at any given moment (ie. Exstart, exchange, full). Which two IOS commands let you view the state of the link? (Select two)

- A. show ip ospf
- B. show ip protocols
- C. show ip ospf neighbor
- D. show ip ospf interface

Answer: C, D

Explanation:

The link state exstart is an OSPF link state (see note below). We need retrieve OSPF link state information.

C: The output of the **show ip ospf neighbor** command is used To display OSPF-neighbor information on a per-interface basis. It includes link state information.

D: The **show ip ospf interface** command is used to display OSPF-related interface information for a particular interface. This includes the link state of the specified interface.

Note: exstart state: After two OSPF neighboring routers establish bi-directional communication and complete DR/BDR election (on multi-access networks), the routers transition to the exstart state.

Incorrect Answers:

- A: The **show ip ospf** command is used to display general information about OSPF routing processes. However, it does not include any link state information.
- B: The command “show ip protocols” displays the parameters and current state of the active routing protocol process. It does not show any link state information.

QUESTION NO: 5

While performing a routine inspection of your OSPF network you come to the realization that the Designated Router (DR) may be overloaded, and you’re considering changing the DR. What are three ways you could manipulate the election of the DR? (Select three)

- A. Use of the `priority` command.
- B. Use of the `router-id` command.
- C. Assignment of the loopback address.
- D. Assigning a lower IP address
- E. By adding additional memory to the desired router.

Answer: A, B, C

Explanation:

- A: The OSPF router priority is only used on multi-access networks such as LAN's. This establishes whether the router is eligible to become the Designated Router (DR) for the LAN. A priority of zero means that the router is not eligible to become DR.
- B: The router ID is used to break ties during the DR and BDR election processes if the priority values are equal.
- C: The loopback address is used in DR election. The highest IP address on the router is the router ID. If a loopback address is configured, then it is the router ID.

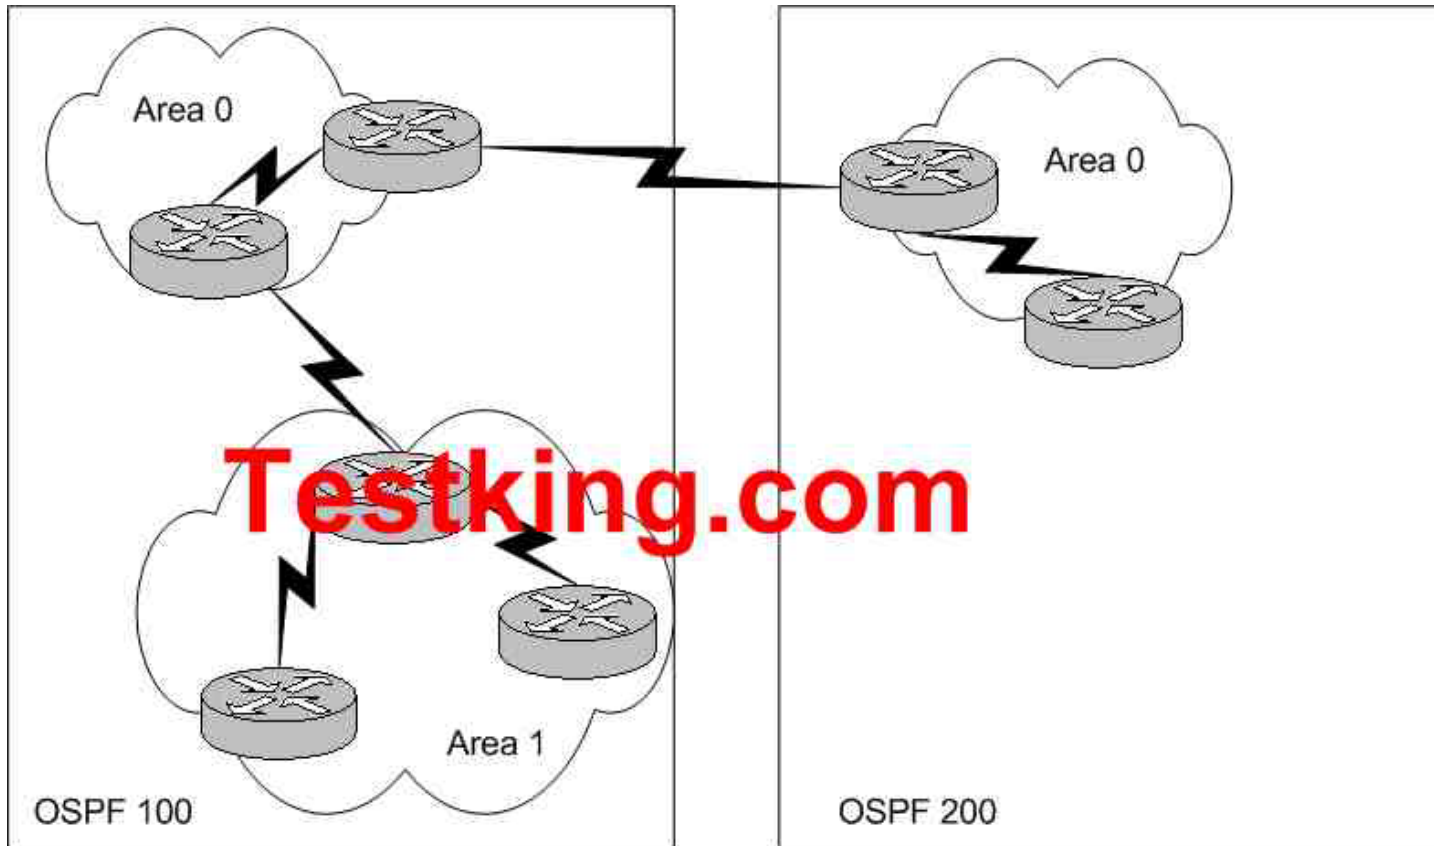
Incorrect Answers:

- D: The highest IP address is used as the router ID, not the lowest.
- E: This would not make it a more likely DR candidate, did the DR election process does not take into consideration the amount of memory.

QUESTION NO: 6

You are the network engineer at TestKing. The TestKing network topology is shown in the following graphic:

Leading the way in IT testing and certification tools, www.testking.com



OSPF area 1 has been configured to receive only the intra-area routes and a default route from the area 1 ABR. During routine maintenance, you issue the “show ip route” command on area 1 internal routers and notice that O and O IA routes appear, but no O E2 routes appear.

What is the probable cause of this problem?

- A. At the ABR for area 1, the **area 1 stub** command is missing the **no-summary** option.
- B. All the routers in area 1 are missing the **area 1 stub no summary** option.
- C. At the internal routers in area 1, the **area 1 stub** command is missing.
- D. At the ABR for area 1, the **area 1 stub** command is missing.
- E. The virtual link that transits area 1 is configured incorrectly.
- F. Area 1 should be configured as a not-so-stubby-area with the **area 1 nssa** command.

Answer: A

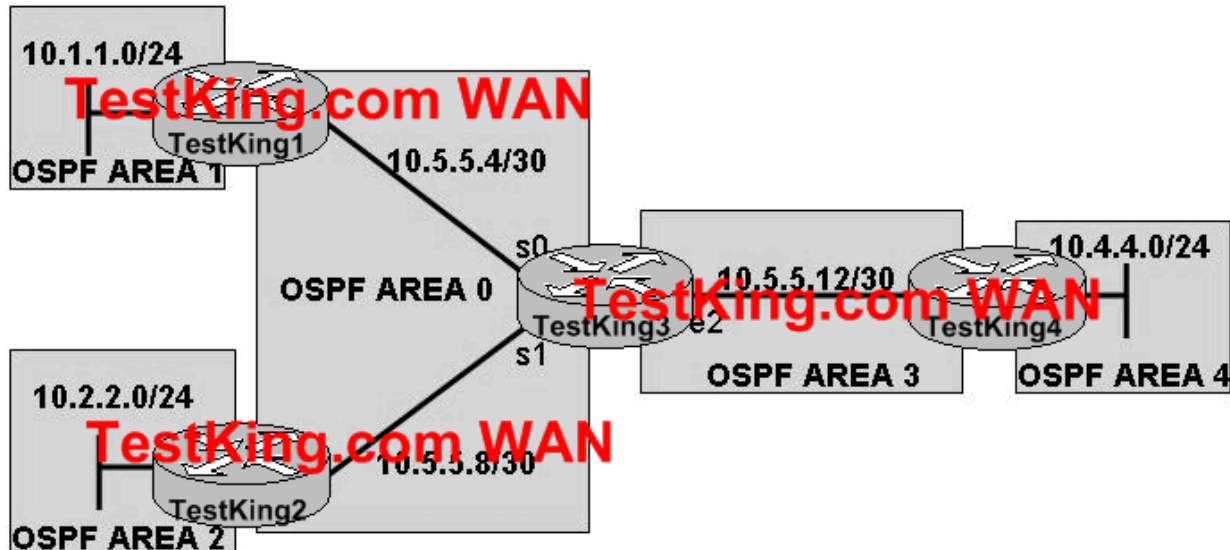
Explanation:

An autonomous system boundary router (ASBR) advertises external destinations throughout the OSPF autonomous system. In many cases, external link states make up a large percentage of the link states in the databases of every router. A stub area is an area in which you don't allow

advertisements of external routes, thus reducing the size of the database even more. Instead, a default summary route (0.0.0.0) is inserted into the stub area in order to reach these external routes. If you have no external routes in your network, then you have no need to define stub areas.

QUESTION NO: 7

The complex TestKing OSPF network topology is displayed in the following diagram:



The routing table of TestKing3 contains all the proper OSPF routes except the 10.4.4.0/24 OSPF route. Which show command on TestKing3 will be most useful to determine the cause of this problem and why?

- Perform **show is ospf virtual-link** to examine if the virtual link between TestKing3 and TestKing4 is configured and in the up state.
- Perform **show run** to examine if area 3 is configured as a stubby area.
- Perform **show run** to examine if area 3 is configured as a nssa area.
- Perform **show ip ospf neighbor** to examine which router is elected as the DR on the e2 interface.
- Perform **show ip ospf neighbor** to examine the OSPF database entries.

Answer: A

Explanation:

All OSPF areas need to be directly connected to the backbone (area 0) in order to ensure proper OSPF operation. If an area is not able to be directly connected with area 0, a virtual link needs to be configured in order to logically connect that area to the backbone network. In this case, a virtual link needs to be configured between routers TestKing3 and TestKing4. This will

logically extend the OSPF area 3 into area 4, connecting it into the backbone. In this example, if the route to 10.4.4.0/24 is missing, a misconfigured or broken virtual link is the most likely cause of the problem.

QUESTION NO: 8

Which command would display OSPF parameters such as filters, default metric, maximum paths, and number of areas configured on a router?

- A. show ip protocol
- B. show ip route
- C. show ip ospf interface
- D. show ip ospf
- E. show ip interface
- F. None of the above

Answer: A

Explanation:

The “show ip protocol” command displays values about routing timers and network information associated with the entire router. This includes, the AS number associated with the routing process, number of areas configured on the router, the metric, and the maximum paths.

Section 4: Describe the scalability problems associated with internal BGP (4 questions)

QUESTION NO: 1

While troubleshooting a BGP problem on the TestKing network you notice that the 10.10.10.0/24 prefix is not being injected into the local BGP table of TK1. The relevant configuration of router TK1 is shown below:

```
router bgp 65001
network 10.0.0.0
neighbor 172.16.1.1 remote-as 65002
no auto-summary
```

Routing table information:

```
show ip route | include 10
O 10.10.10.0/24 [110/11] via 192.168.1.1, 2d00h, Ethernet0/0
```

Why doesn't the local BGP table have the prefix?

- A. This route is not a BGP learned route.
- B. The **network** command is wrong.
- C. The 172.16.1.1 neighbor is down.
- D. The prefix 10.10.10.0/24 is not a connected route.

Answer: B

Explanation:

The correct syntax should be “network 10.10.10.0 mask 255.255.255.0” under the BGP routing process. Without the correct subnet mask specified, the route will not get injected into the BGP routing table, even if it is learned via an IGP. In this case, the route is known via OSPF.

QUESTION NO: 2

What would happen if a full mesh of BGP sessions were configured within a single large autonomous system? (Select two)

- A. Many UDP sessions will be created.
- B. More memory and CPU will be consumed.
- C. This configuration is not permitted by default.
- D. A significant amount of bandwidth for BGP updates and retransmissions can be used.

- E. Permanent Virtual Circuits (PVCs) must be created to link the fully meshed BGP sessions.

Answer: B, D

Explanation:

The number of TCP connections that are required for a full mesh grows exponentially. The formula for the total number of BGP sessions is $N(N-1)/2$. For example, for 10 IBGP routers to become fully meshed 45 peering sessions will need to be configured.

B: More memory and CPU resources are required on the routers to support the full mesh BGP.

D: There will be many TCP sessions and a significant amount of bandwidth will be required for BGP traffic.

Incorrect Answers:

A: Many TCP, not UDP, sessions will be created.

C: It is permitted.

E: PVCs are not required.

QUESTION NO: 3

An ISP is running a large IBGP network with 25 routers. The full mesh topology that is currently in place is inefficient on bandwidth and BGP traffic.

What can the administrator configure to reduce the number of BGP neighbor relationships within the AS?

- A. route reflector
- B. route maps
- C. route redistribution
- D. peer groups
- E. aggregate addresses

Answer: A

Explanation:

Route reflectors can be used to overcome this scalability issues with configuring a full IBGP mesh within a network. With route reflectors, the clients need to only peer with one single router (the route reflector). In this way, fewer peering sessions need to be configured and maintained within the IBGP network.

QUESTION NO: 4

To establish a full mesh IBGP between ten routers, how many IBGP sessions are required?

- A. 10
- B. 45

- C. 50
- D. 90
- E. 100

Answer: B

Explanation:

For full mesh $X*(X-1)/2$ connections are needed. So, $10(10-1)/2 = (10*9)/2=45$.

Section 5: Interpret the output of various show and debug commands to determine the cause of route selection errors and configuration problems (20 questions)

QUESTION NO: 1

While verifying the OSPF operation on your network, you issue the following command:

```
show ip ospf interface
```

What information parameters will you get after you execute this command? (Select two)

- A. The router ID
- B. The summary link counts
- C. The neighbor adjacencies
- D. The link-state update interval

Answer: A, C

Explanation:

A: The router ID is displayed.

C: Neighbor adjacencies are included in the output of this command.

D: OSPF use the hello interval as update interval. The hello interval is displayed by this command.

QUESTION NO: 2

You issue the “show ip route” command on router TK1 and receive the following output:

```
S    62.99.153.0/24 [1/0] via 209.177.64.130
    172.209.12.0/32 is subnetted, 1 subnets
D EX  172.209.12.1
    [170/2590720] via 209.179.2.114, 06:47:28,
Serial0/0/0.1239
    62.113.17.0/24 is variably subnetted, 2 subnets, 2 masks
D    62.113.17.0/29 [90/30208] via 62.113.20.10, 07:35:24,
ATMO/1/0.130
S    62.113.17.0/24 [1/0] via 62.113.1.25
D EX 99.3.215.0/24
    [170/27316] via 209.180.96.45, 09:52:10,
FastEthernet11/0/0
```

```

        [170/27316] via 209.180.96.44, 09:52:10,
FastEthernet11/0/0
D    25.248.17.0/24
        [90/1512111] via 209.179.66.25, 10:33:13,
Serial0/0/0.1400001
        [90/1512111] via 209.179.66.41, 10:33:13,
Serial0/0/0.1402001
        62.113.1.0/24 is variably subnetted, 12 subnets, 2 masks
D    62.113.1.227/32
        [90/24823552] via 209.180.96.45, 07:35:24,
FastEthernet1/0/0
        [90/24823552] via 209.180.96.44, 07:35:24,
FastEthernet1/0/0
S*   0.0.0.0/0 [1/0] via 209.180.96.14

```

One of the routing entries in router TK1 was learned via a static route. Which one is it?

- A. 99.3.215.0
- B. 62.99.153.0
- C. 172.209.12.1
- D. 62.113.1.227

Answer: B

Explanation:

The **S** denotes a static router (see below).

```
S    62.99.153.0/24 [1/0] via 209.177.64.130
```

Incorrect Answers:

A, C: The **D EX** denotes a route learned through EIGRP external.

D: The **D** denotes a route learned through EIGRP.

QUESTION NO: 3

While verifying BGP operation on the TestKing router, you issue the “show ip bgp” command as shown below:

```

routerTK>show ip bgp
BGP table version is 1046033, local router ID is 198.32.162.100
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
* > 143.16.0.0    128.214.63.2    0     400     0     200 1
*    143.16.0.0    192.208.10.5    0     300     0     300 1

```

Leading the way in IT testing and certification tools, www.testking.com

```
* 143.16.0.0 143.16.63.5 0 100 0 200 1
* 143.16.0.0 203.250.13.41 0 100 0 500 1
```

From the information above, which path will the network 143.16.0.0 prefer to take to exit the AS?

- A. 128.214.63.2
- B. 192.208.10.5
- C. 128.213.63.5
- D. 203.250.13.41
- E. All of the above will be used in a round robin fashion.

Answer: A

Explanation:

Local preference (LocPref) is a well-known discretionary attribute that provides an indication to routers in the AS about which path is preferred to exit the AS. A path with a higher local preference is more preferred. In this scenario the following entry has the highest local preference value of 400.

```
Network      Next Hop  Metric  LocPrf  Weight  Path
* > 128.213.0.0 128.214.63.2 0 400 0 200 1
```

The preferred exit path of the AS is therefore 128.214.63.2, as noted by the ">" which refers to the best path for this destination.

QUESTION NO: 4

Which IOS command would you use to find out which networks are routed by a particular OSPF process?

- A. show ospf
- B. show ip route
- C. show ip protocols
- D. show ip ospf database
- E. None of the above

Answer: C

Explanation:

The **show ip protocols** command display current routing protocols. It displays the parameters and current state of the active routing protocol process. The output includes a list of the networks routing for individual ospf processes.

Sample output:

```
RtTestKing# show ip protocols
```



```

Routing Protocol is "ospf 200"
Sending updates every 0 seconds
Invalid after 0 seconds, hold down 0, flushed after 0
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Redistributing: ospf 200
Routing for Networks:
172.6.31.5/32
Routing Information Sources:
Gateway Distance Last Update
Distance: (default is 110)

```

Incorrect Answers:

- A: The **show ospf** command displays summary information regarding the global OSPF configuration.
- B: The **show ip route** command displays the IP routing table.
- D: The **show ip ospf database** command displays the contents of the topological database maintained by the router. The command also shows the router ID and the OSPF process ID. However, the output does not include the networks routing for individual ospf processes.

QUESTION NO: 5

The routing table of router TK is shown below:

```

TK#show ip route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate
       default

<Some output deleted>

Gateway of last resort is 30.64.0.2 to network 0.0.0.0

   30.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
O IA  30.2.0.0/16 [110/74] via 30.64.0.2, 00:09:13, Ethernet0
C     30.1.3.0/24 is directly connected, Serial0
O IA  30.3.0.0/16 [110/148] via 30.64.0.2, 00:05:22, Ethernet0
C     30.1.2.0/24 is directly connected, Serial1

```

Based on the output above, which of the following statements are true regarding the routing table? (Select two)

- A. The area is a stub area.
- B. The area is totally stubby.

- C. Network 30 is using VLSM
- D. The routing table is for an ABR

Answer: A, C

Explanation:

Incorrect Answers:

A: This is a stub area.

C: The exhibit states that the 30.0.0.0/8 subnet is variably subnetted and the networks used are /16 and /24, so VLSM is used

Incorrect Answers:

B: The area is stub, not totally stubby.

D: This is not an ABR (Area border router) since there are no external routes, only internal and directly connected.

QUESTION NO: 6

While troubleshooting an OSPF routing issue you issue the “show ip ospf interface” command. Which of the following parameters are provided from this? (Select two)

- A. Router ID
- B. Summary link counts
- C. Neighbor adjacencies
- D. Link-state update interval

Answer: A, C

Explanation:

The **show ip ospf interface** command verifies that interfaces have been configured in the intended areas. If no loopback address is specified, the interface with the highest address is the taken router ID. It also gives the timer intervals, including the hello interval, and shows the neighbor adjacencies.

Reference: Building Scalable Cisco Networks (Cisco Press) page 134.

QUESTION NO: 7

Which command should you use to obtain information on OSPF link state advertisements? In particular: which LSA’s have been sent, which LSA’s have been received, and the time when the last LSA was received?

- A. show ip ospf database
- B. show ip ospf neighbors
- C. show ip ospf protocols

- D. show ip ospf interfaces
- E. None of the above

Answer: A

Explanation:

The **show ip ospf database** command is used to display lists of information related to the OSPF topological, the link state, database for a specific router.

Sample output:

```
R_TestKing#show ip ospf database
OSPF Router with ID (192.168.0.12) (Process ID 1)
Router Link States (Area 0)
Link ID      ADV Router  Age   Seq#           Checksum      Link count
192.168.0.10 192.1680.10 817   0x80000003    0xFF56        1
192.168.0.11 192.1680.11 817   0x80000003    0xFD55        1
192.168.0.12 192.168.0.12 816   0x80000003    0xFB54        1
192.168.0.13 192.168.0.13 816   0x80000003    0xF953        1
192.168.0.14 192.168.0.14 817   0x80000003    0xD990        1
Net Link States (Area 0)
Link ID      ADV Router  Age   Seq#           Checksum
192.168.0.14 192.168.0.14 812   0x80000002    0x4AC8
```

Incorrect Answers:

B: The **show ip ospf neighbor** is used to display OSPF-neighbor information on a per-interface basis. It does show the required information though.

Sample output:

```
Neighbor ID Pri State Dead Time Address Interface
192.168.0.13 1 2WAY/DROTHER 00:00:31 192.168.0.13 Ethernet0
192.168.0.14 1 FULL/BDR 00:00:38 192.168.0.14 Ethernet0
```

C: There is no such command.

D: The **show ip ospf interface** command is used to display OSPF-related interface information. It displays the circuit name and state, IP address, network mask, broadcast address, redundancy, Internet Control Message Protocol (ICMP) settings, and RIP settings. However, it does not display LSAs.

Sample output:

```
RouterTK# show ip interfaces
IP Interface Summary:
Circuit Name:      VLAN2      State:      active
IP Address: 172.16.1.200  Network Mask: 255.255.0.0
Broadcast Address: 172.16.255.255 Redundancy: disabled
ICMP Redirect: enabled ICMP Unreachable: enabled
RIP: enabled
```

QUESTION NO: 8

You have a console connection on a router running RIP. Which IOS command would you enter if you wanted to view the RIP routing transactions that are occurring?

- A. show ip rip database
- B. show ip route
- C. show ip protocols rip
- D. debug ip rip
- E. debug ip routing
- F. None of the above

Answer: D

Explanation:

debug ip rip

Use the **debug ip rip** EXEC command to display information on RIP routing transactions in real time. The **no** form of this command disables debugging output.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_command_reference_chapter09186a008007ff66.html

QUESTION NO: 9

While troubleshooting a problem on the TestKing network, you issue the “show ip protocol” as shown below:

TK1#show ip protocol

Outgoing update filter list for all interface is

Incoming update filter list for all interface is

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 101, ospf 101 (internal, external 1, external 2)

Automatic network summarization is in effect

Routing for Networks:

192.168.1.0

Routing Information Sources:

Gateway Distance Last Update

192.168.1.34 90 00:05:21

192.168.1.40 90 00:05:21

192.168.1.18 90 00:05:21

Distance: internal 90 external 170

Leading the way in IT testing and certification tools, www.testking.com

Routing Protocol is "ospf 101"
 Sending updates every 0 seconds
 Invalid after 0 seconds, hold down 0, flushed after 0
 Outgoing update filter list for all interfaces is
 Incoming update filter list for all interfaces is
 Redistributing: eigrp 101, ospf 101
 Routing for Networks:
 10.1.1.1/32
 Routing Information Sources:
 Gateway Distance Last Update
 172.16.11.100 110 00:00:16

Based on this information, which of the following statements are true for Router TK1? (Choose three)

- A. Redistribution has been configured for EIGRP into OSPF.
- B. Redistribution has been configured for OSPF into EIGRP.
- C. EIGRP auto-summary has been disabled.
- D. The OSPF process-id and the EIGRP autonomous system number can not be identical.
- E. EIGRP has been configured for un-equal cost paths load balancing.
- F. Router TK1 has received routing updates from three EIGRP neighbors and from one OSPF neighbor.

Answer: A, B, F

Explanation:

A: Under the EIGRP configuration the line "Redistributing: eigrp 101, ospf 101" is present.

B: Under the OSPF configuration the line "Redistributing: eigrp 101, ospf 101" is present.

F: This is correct because of the three neighbor updates listed under the Gateway Distance Last Update section of the EIGRP configuration

Incorrect Answers:

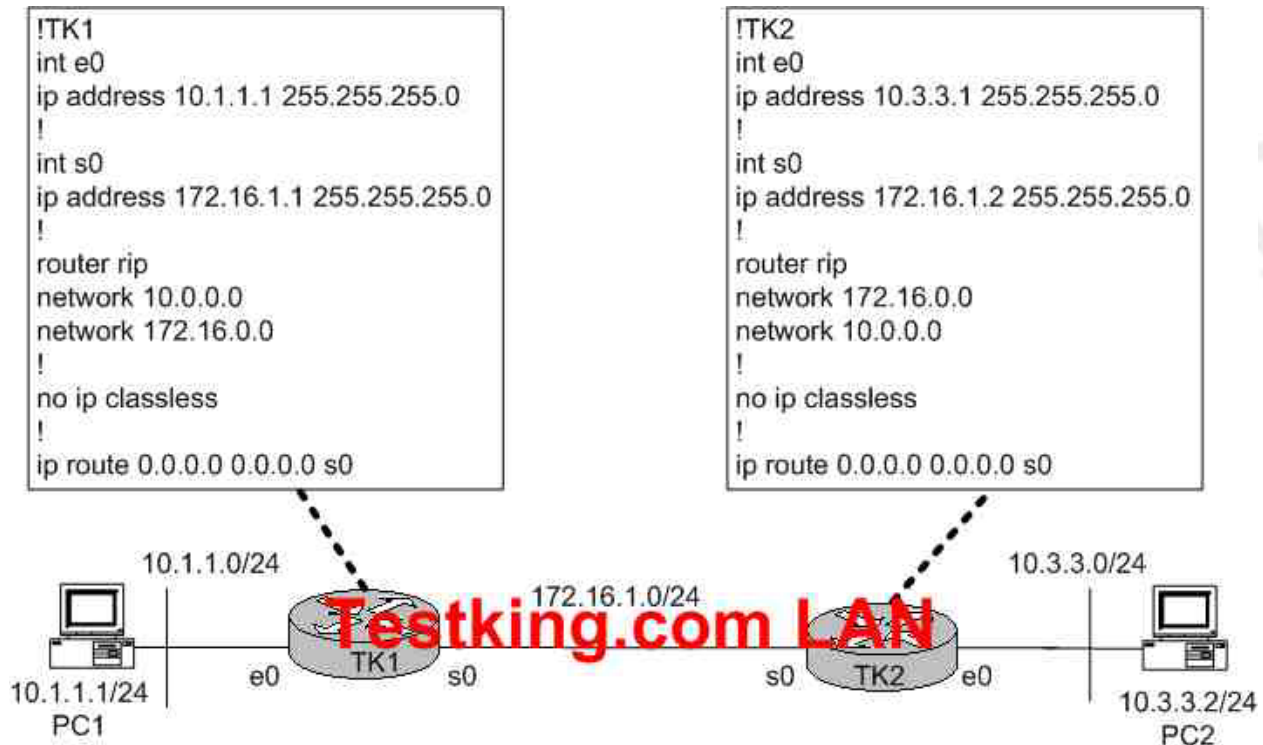
C: The output explicitly states that automatic network summarization is in effect.

D: This is not true. In this router, they are both configured for AS 101

E: is incorrect because under the EIGRP configuration the EIGRP maximum metric variance is 1, which is the default, specifying equal cost load balancing.

QUESTION NO: 10

The TestKing network is displayed below, along with the relevant router configurations for TK1 and TK2.



PC1 is unable to ping PC2 successfully.

Given the above diagram and configurations, which solution would fix this problem?

- A. Enable **IP classless** on TK1 and TK2.
- B. To support discontinuous subnet, enable RIPv2 on TK1 and TK2 then disable **auto-summary**.
- C. Enable both RIPv1 and RIPv2 on TK1 and TK2.
- D. Enable **no auto-summary** in RIP router configuration mode on TK1 and TK2.

Answer: B

Explanation:

With RIP version 1 configured, both routers will summarize the LAN subnets and each will advertise the 10.0.0.0/8 route to each other. If we enable use a routing protocol that supports VLSM, such as RIPv2, and then disable the auto-summarization feature, then router TK1 will advertise the 10.1.1.0/24 subnet to TK2, and TK2 will advertise the 10.3.3.0/24 subnet to TK1.

Incorrect Answers:

- A: This will not solve the problem, since each router has a 10s network on each of their LANs.
- C: This alone will not solve the problem, since we also need to disable the automatic summarization of network routes.
- D: This is not an option for RIP version 1 networks, since it is a classful routing protocol.

QUESTION NO: 11

The IP BGP table of router P1R3 is displayed in the following diagram:

```

P1R3#show ip bgp
BGP table version is 16, local router ID is 10.200.200.13
Status codes: s suppressed, d daaped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i192.168.11.0	10.200.200.12	0	100	101	64998 65222 65223 i
* i	10.200.200.11	0	200	0	64998 65222 65223 i
*>i192.168.12.0	10.200.200.12	0	100	101	64998 65222 65223 i
* i	10.200.200.11	0	200	0	64998 65222 65223 i

Based on the P1R3 output shown above, which statement is true?

- A. The best path to reach the 192.168.11.0 prefix is via 10.200.200.11.
- B. The best path to reach the 192.168.11.0 prefix is via 10.200.200.12.
- C. The best path to reach the 192.168.11.0 prefix is via both 10.200.200.11 and 10.200.200.12; BGP will automatically load balance between the two.
- D. The 192.168.11.0 and 192.168.12.0 prefixes were learned via EBGP from the 10.200.200.11 and 10.200.200.12 EBGP neighbors.
- E. None of the above

Answer: B

Explanation:

The best path to any given destination is noted by the ">" in the IP BGP table. In this case, the best path to 192.168.11.0 is via next hop 10.200.200.12 due to the fact that the weight is higher (101) than the path via the alternative next hop. Weight is a Cisco proprietary method for path determination and the weight value is used above all other values. Within a router, the path with the highest weight will be preferred.

QUESTION NO: 12

Which IOS command could you use to verify if a BGP router is or isn't a route reflector?

- A. show bgp neighbor
- B. show running-config
- C. show route-reflector
- D. show route-reflector-client
- E. None of the above

Answer: A

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

The **show bgp neighbor** command indicates if a particular neighbor is a route reflector client.

Note: Route reflectors modify the BGP split horizon rule by allowing the router configured as the route reflector to propagate routes learned by IBGP to other IBGP peers.

Incorrect Answers:

B: Displaying the current configuration would not be useful in finding information in the route reflector status.

C, D: There are no such commands.

QUESTION NO: 13

Which IOS command would you enter if you wanted to view a list of IBGP and EBGP neighbor relationships that are configured?

- A. show ip bgp
- B. show ip bgp paths
- C. show ip bgp peers
- D. show ip bgp summary
- E. show ip bgp protocols

Answer: D

Explanation:

The **show ip bgp summary** command displays the status of all BGP connections. Neighbors with corresponding AS values will be listed; both interior and external.

Incorrect Answers:

A: The **show ip bgp** command displays routes in the BGP routing table, not the neighbors.

B: The **show ip bgp paths** command is used to display all the BGP paths in the database. However, it does not list the neighbors.

C, E: There is no such command.

Reference: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_r/1rprt1/1rbgp.htm

QUESTION NO: 14

The 192.168.0.0/16 network is not being propagated throughout the network via BGP as expected. Observe the BGP configuration commands from the advertising router shown below.

```
Router bgp 65111
neighbor 172.16.1.1 remote-as 65111
neighbor 172.16.2.1 remote-as 65112
network 192.168.0.0
```



```
network 10.0.0.0
!
ip route 192.168.0.0 255.255.0.0 null0
```

What is the reason the 192.168.0.0/16 route is not being advertised?

- A. The **network 192.168.0.0** statement is missing **mask 255.255.0.0**
- B. The **network 192.168.0.0** statement is missing **mask 0.0.255.255**
- C. The **network 10.0.0.0** statement is missing **mask 255.0.0.0**
- D. The **network 10.0.0.0** statement is missing **mask 0.255.255.255**
- E. The **auto-summary** configuration is missing

Answer: A

Explanation:

The network 192.168.0.0 statement is missing mask 255.255.0.0. Without the mask command used in a network statement, the route may not get properly injected into the BGP routing process.

QUESTION NO: 15

The TestKing network is redistributing routing information from OSPF and EIGRP.

Which three steps are most helpful in verifying proper route redistribution? (Select three)

- A. On the routers not performing the route redistribution, use the **show ip route** command to see if the redistributed routes show up.
- B. On the ASBR router performing the route redistribution, use the **show ip protocol** command to verify the redistribution configurations
- C. On the ASBR router performing the route redistribution, use the **show ip route** command to verify that the proper routes from each routing protocol are there.
- D. On the routers not performing the route redistribution, use the **show ip protocols** command to verify the routing information sources.
- E. On the routers not performing the route redistribution, use the **debug ip routing** command to verify the routing updates from the ASBR.

Answer: A, B, C

Explanation:

In order to verify proper route redistribution, use the “show ip route” command on all routers within the network, as well as the ASBR, to verify that the routes are properly being advertised to all routers. In addition, issuing the “show ip protocol” can be used on the router performing the redistribution to verify that routes are being redistributed into each other.

Incorrect Answers:

D: Issuing this command on a non-redistribution router will not tell us where and how the routes are originating from. This command will only be useful on the redistributing routers.

E: This command can not be used to verify the redistributed routes.

QUESTION NO: 16

Exhibit:

```
TestKingA# show ip route isis
10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
I L2 10.200.200.14/32 [115/30] via 10.1.0.2, Serial 1/0
I L1 10.200.200.13/32 [115/20] via 10.1.1.3, Ethernet 0/0
I L1 10.1.3.0/24 [115/20] via 10.1.1.3, Ethernet 0/0
I L2 10.1.2.0/23 [115/20] via 10.1.0.2, Serial 1/0
I su 10.1.0.0/23 [115/10] via 0.0.0.0, Null0
```

Based on the show ip route isis output on Router TestKingA, which statement is true?

- A. The “I su” 10.1.0.0/23 route is an IS-IS external route.
- B. The TestKingA IS-IS router is an ASBR.
- C. The “I su” 10.1.0.0/23 route is a suppressed route.
- D. The “I su” 10.1.0.0/23 route is a summary route.
- E. The R1 IS-IS router is an ABR that belongs to multiple IS-IS areas.
- F. The R1 IS-IS route is performing route aggregation and is suppressing the more 10.1.0.0/23 prefix.

Answer: D

Explanation:

When viewing the “show IP route ISIS” output, the “su” entry represents a summarized route. This is the direct result of the “summary-address” router configuration command. When creating the summary route. The Cisco IOS automatically creates the summary route and point it to interface Null 0.

QUESTION NO: 17

Exhibit:

```
TESTKINGA# show ip eigrp topology
IP-EIGRP Topology Table for process 200
```

Codes:P - Passive, A - Active, U- Update, Q - Query, R - Reply, r - Reply status

```
P 192.168.1.64/28, 1 successors, FD is 281600
    Via Connected, Ethernet0
P 192.168.1.32/28, 1 successors, FD is 40512000
```

```

Via Connected, Serial1
P 192.168.1.48/28, 1 successors, FD is 40537600
  Via 192.168.1.66 (40537600/40512000), Ethernet0
  Via 192.168.1.77 (41024000/40512000), Serial0
  Via 192.168.1.33 (41024000/40512000), Serial1
P 192.168.1.16/28, 1 successors, FD is 40512000
  Via Connected, Serial0

```

Based on the above show ip eigrp topology output, which three statements are true? (Choose three.)

- A. TESTKINGA is in AS 200
- B. TESTKINGA will balance between three paths to reach the 192.168.1.48/28 prefix, because all three paths have the same AD of 40512000.
- C. The best path for TESTKINGA to reach the 192.168.1.48/28 prefix is via 192.168.1.66.
- D. 40512000 is the AD to reach the 192.168.1.48/28 prefix.
- E. All of the routes are in the passive mode because these routes are in the hold-down state.
- F. All the routes are in the passive mode, because TESTKINGA is in the query process for those routes.

Answer: A, C, D

Explanation:

It can be determined that AS 200 is used, from the fact that the IS-IS process ID is labeled as 200. The best path to reach the network 192.168.1.48/28 is the first one displayed in the routing table. This can be further demonstrated by the fact that the metric is less than the alternative route, via serial 0. Finally, the AD can be found by viewing the second number within the parentheses, which in this case is 40512000.

QUESTION NO: 18

Exhibit

```

TestKing1#sh ip ospf neighbor
Neighbor ID   Pri  State           Dead Time   Address      Interface
10.200.200.13  1    FULL/BDR        00:00:33    10.1.1.3     Ethernet0/0

TestKing3#sh ip ospf neighbor
Neighbor ID   Pri  State           Dead Time   Address      Interface
172.31.1.1    2    FULL/DR         00:00:31    10.1.1.1     Ethernet0/0

```

The TestKing1 and TestKing3 routers are OSPF neighbors over the Ethernet 0/0 connection. Based on the show ip ospf neighbor output from the TestKing1 and TestKing3 routers, which statement is true?

Leading the way in IT testing and certification tools, www.testking.com

- A. TestKing1 is the DR because it has a higher OSPF router priority.
- B. TestKing1 is the DR because it has a lower OSPF router ID.
- C. TestKing3 is the DR because it has a higher OSPF router priority.
- D. TestKing3 is the DR because it has a lower OSPF router ID.
- E. Both TestKing1 and TestKing3 are using the default OSPF router priority.

Answer: A

QUESTION NO: 19

Exhibit, Network topology



TestKing3 is redistributing the EIGRP routes into OSPF. What will the EIGRP routes appear in the routing table of TestKing1?

- A. O
- B. O IA
- C. O E2
- D. D
- E. D EX

Answer: C

Explanation:

O E1 or O E2. The routes in this LSA are external to the autonomous system. They can be configured to have one of two values. E1 will include the internal cost to the ASBR added to the external cost reported by the ASBR. E2 does not compute the internal cost – it just reports the external cost to the remote destination.

QUESTION NO: 20

Which three types of OSPF route entries can be found in the routing table of an internal OSPF router within an OSPF not so stubby area? (Select three)

- A. O
- B. O*IA
- C. O*OA
- D. O E1
- E. O E2
- F. O N1

Answer: A, B, F

Explanation:

The various route types used by OSPF are as follows:

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

An OSPF NSSA will receive inter-area, external type 2, OSPF routes.

Section 6: Identify the steps to verify Enhanced IGRP operation (7 questions)

QUESTION NO: 1

EIGRP uses five generic packet types (hello, updates, queries, replies, acknowledgements). If you wished to view the statistics for these packets, which IOS command should you use?

- A. debug eigrp packets
- B. show ip eigrp traffic
- C. show ip eigrp topology
- D. show ip eigrp neighbors

Answer: B

Explanation:

The **show ip eigrp traffic** command displays the number of Enhanced IGRP (EIGRP) packets sent and received.

Example:

The following is sample output from the show ip eigrp traffic command:

```
Router# show ip eigrp traffic
```

```
IP-EIGRP Traffic Statistics for process 77
```

```
Hellos sent/received: 218/205
```

```
Updates sent/received: 7/23
```

```
Queries sent/received: 2/0
```

```
Replies sent/received: 0/2
```

```
Acks sent/received: 21/14
```

Reference: http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_command_reference_chapter09186a00800ca5a9.html#wp1018815

QUESTION NO: 2

While troubleshooting a routing problem on the TestKing EIGRP network you discover that one of the routers is failing to establish adjacencies with its neighbor. What is a likely cause of this problem between neighbors? (Select two)

- A. The K-values do not match.
- B. The hold times do not match.

- C. The hello times do not match.
- D. The AS numbers do not match.

Answer: A, D

Explanation:

Peer relationships and adjacencies between routers will not be formed between EIGRP routers if the neighbor resides in a different autonomous system or if the metric-calculation mechanism (K values) is misaligned for that link.

Incorrect Answers:

B, C: It is possible for two routers to become EIGRP neighbors even though the hello and hold timers do not match.

QUESTION NO: 3

While troubleshooting an EIGRP routing issue, you are seeing a high number of SIA (stuck in active) routes. Which of the following are causes of a route becoming SIA? (Select two)

- A. Some query or reply packets are lost between the routers.
- B. The neighboring router stops receiving ACK packets from this router.
- C. The neighboring router starts receiving route updates from this router.
- D. A failure causes traffic on a link between two neighboring routers to flow in only one direction (unidirectional link).

Answer: A, D

Explanation:

The acknowledgement does not reach the destination or they are too delayed. This is normally due to too many routing topology changes, or a router with insufficient memory.

Note: In some circumstances, it takes a very long time for a query to be answered. So long, in fact, that the router that issued the query gives up and clears its connection to the router that isn't answering, effectively restarting the neighbor session. This is known as a stuck in active (SIA) route. The most basic SIA routes occur when it simply takes too long for a query to reach the other end of the network and for a reply to travel back.

Incorrect Answers:

B: Ack packets don't reply to Query, only Reply do.

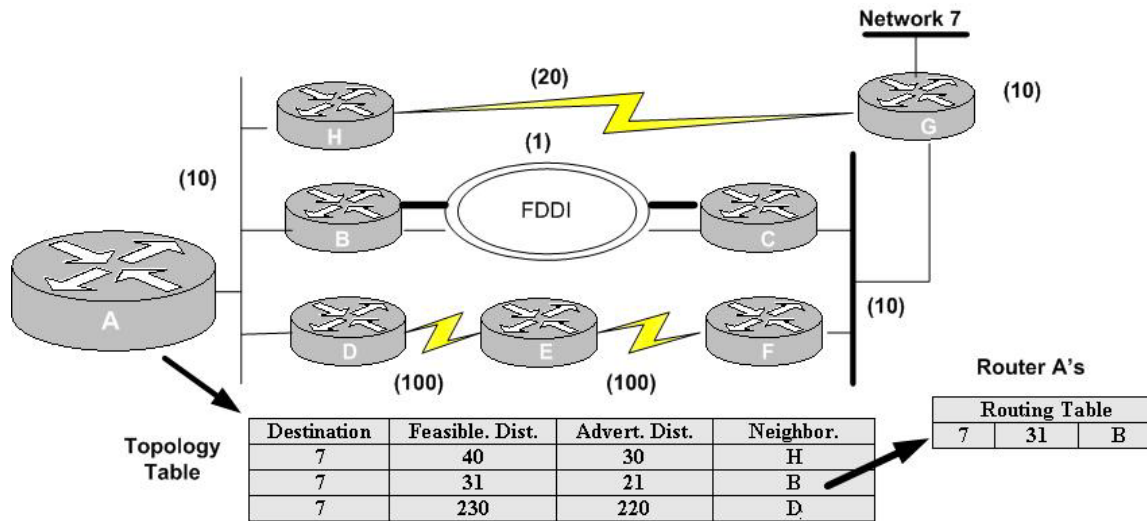
C: Does not apply to SIA. This is the normal operation of EIGRP.

Reference: <http://www.cisco.com/warp/public/103/eigrp3.html>

QUESTION NO: 4

The TestKing network is displayed in the diagram below:

Leading the way in IT testing and certification tools, www.testking.com



If the FDDI interface in the EIGRP network above were to fail or shut down, which router(s) will become Router A's next-hop to Network 7?

- A. B
- B. D only
- C. H only
- D. D and H

Answer: C

Explanation:

Router H will be the successor, and that route will be placed in the Routing table.

Router A detects the link failure between Router B and network 7. It checks the topology table for a successor. It finds that H is the successor since the advertised distance for H (30) is less than the feasible distance for B (31).

However, there is no next best route – no feasible successor. The candidate route through D has an advertised distance (220) that is higher than the feasible distance of the successor route (40).

QUESTION NO: 5

The TestKing network consists of a hub and spoke topology with a main router supporting about 20 regional offices. A point-to-point Frame Relay WAN connects the regional offices to the main office, and EIGRP is deployed as the routing protocol. The committed information rate (CIR) for each of the Frame Relay PVC's is different, and the bandwidth command IS NOT configured on any of the interfaces or subinterfaces. You want to ensure that EIGRP routes everything properly. How should you configure the network?

- A. Convert each Frame Relay PVC to point-to-multipoint connection

- B. Manually configure the bandwidth of the major interface to the lowest CIR x 24
- C. Manually configure the bandwidth of the major interface to the highest CIR x 24
- D. Manually configure the bandwidth of each of these PVCs to equal to their respective CIR.

Answer: D

Explanation:


Although this will require some administrative effort, the only way to ensure that EIGRP properly considers the actual bandwidth to use in the routing decision, each link should be set to the CIR.

Incorrect Answers:

- A: This is not a Cisco recommended solution.
- B: This will force all PVC to run at a low speed.
- C: This would give too high a bandwidth.

QUESTION NO: 6

The TestKing EIGRP network is displayed in the diagram below:



All routers are running EIGRP. Based on this, what will router TK3 do if the link between TK3 and TK1 were to go down? (Select two)

- A. It elects a new designated router.
- B. It sends a flash update with poison reverse.
- C. It checks its topology table for an alternate route.
- D. It re-broadcasts its routing table to all other neighboring routers.
- E. It sends a query to neighboring routers for other routes to the failed link.

Answer: C, E

Explanation:

The steps of convergence in an EIGRP network are as follows:

1. Router TK3 detects the link failure between TK1 and TK3. It checks the topology table for a feasible successor, but it doesn't find a qualifying alternate route and enters in an active convergence state. (C)
2. TK3 sends a Query out all interfaces for other routes to the failed link (E). The neighboring routers acknowledge the query.
3. The reply from TK4 indicates no other route to the failed link.
4. TK2's reply contains a route to the failed link, although it has a higher feasible distance.
5. Router TK3 accepts the new path and metric information, places it in the topology table, and creates an entry for the routing table.
6. TK3 sends an update about the new route out all interfaces.

QUESTION NO: 7

The EIGRP topology table for router TestKing1 is displayed below:

```
TestKing1#show ip eigrp topology
IP: EIGRP Topology Table for process 200
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status

P 192.168.1.64/28, 1 successors, FD is 281600
   via Connected, Ethernet0
P 192.168.1.32/28, 1 successors, FD is 40512000
   via Connected, Serial1
P 192.168.1.48/28, 1 successors, FD is 40537600
   via 192.168.1.66 (40537600/40512000), Ethernet0
   via 192.168.1.17 (41024000/40512000), Serial0
   via 192.168.1.33 (41024000/40512000), Serial1
P 192.168.1.16/28, 1 successors, FD is 40512000
   via Connected, Serial0
```

Regarding the command output on TestKing1 in the exhibit, which statements are true?
(Select three)

- A. TestKing1 is in AS 200
- B. TestKing1 will load balance between three paths to reach the 192.168.1.48/28 prefix, because all three paths have the same AD of 40512000.
- C. The best path for TestKing1 to reach the 192.168.1.48/28 prefix is via 192.168.1.66.
- D. 40512000 is the advertised metric via 192.168.1.66 to reach the 192.168.1.48/28 prefix.
- E. All the routes are in the passive mode because these routes are in the hold-down state.
- F. All the routes are in the passive mode, because TestKing1 is in the query process for those routes.

Answer: A, C, D

Explanation:

The TestKing1 router resides in AS 200, as displayed by the “IP EIGRP topology for process 200” output.

Regarding the numbers specified in the parenthesis, the first number is the EIGRP metric that represents the cost to the destination. The second number is the EIGRP metric that this peer advertised.

Based on this, the best path to the 192.168.1.48/28 destination is via 192.168.1.66, because the metric is less than the alternatives.

Topic 5: Mixed Questions (19 Questions)

QUESTION NO: 1

Which of the following OSPF routes are supported by the Not-So-Stubby-Area (NSSA) type? (Select three) (801)

- A. O
- B. O N2 **
- C. O* N2 **
- D. O E1
- E. O E2

Answer: A, B, C

Explanation:

The various route types used by OSPF are as follows:

O – OSPF

IA - OSPF inter area

N1 - OSPF NSSA external type 1

N2 - OSPF NSSA external type 2

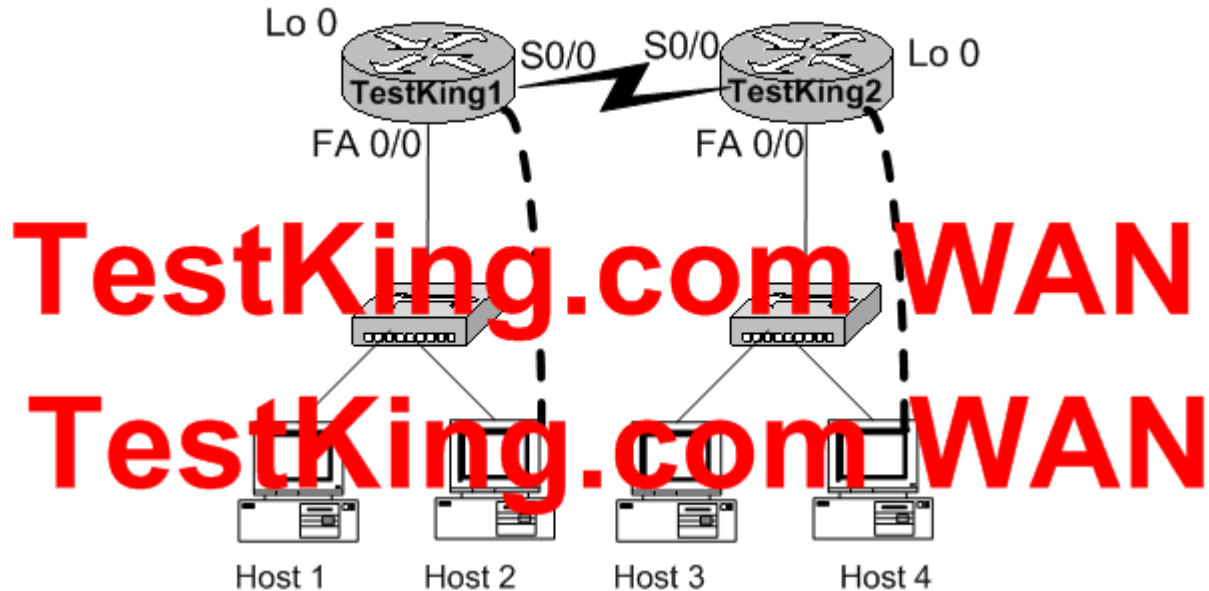
E1 - OSPF external type 1

E2 - OSPF external type 2

An OSPF NSSA will receive inter-area, external type 2, OSPF routes.

QUESTION NO: 2

Network topology exhibit



TestKing.com is planning to provide network connectivity for two its sites: TestKing1 and TestKing2. Each site is to have one LAN. You have been asked to configure the first router at the TestKing1 location. The TestKing1 router has been completely configured except for the routing protocol. Configure OSPF as the routing in a single area to allow a host on the LAN of the TestKing1 router to communicate with a host on the LAN of the TestKing2 router.

Due to the fact that adjacent subnets are in use or are planned for future growth on other routers in the AS, make sure you use specific subnet information in your configuration. Please note, the TestKing2 router will be installed at a later time. The TestKing1 router has been configured with the following specifications:

The router is named TestKing1.

The clocking is provided on the TestKing1's serial 0/0 interface.

The secret password on the TestKing1 router is 'testking'.

Name: TestKing1

FA0/0 172.10.143.1/24

S0/0 172.20.44.1/30

Lo 0 172.30.10.197/32

Name: TestKing2

FA0/0 172.10.144.1/24

S0/0 172.20.44.2/30

Lo 0 172.30.10.198/32

Secret password: testking

Answer:

```
Click on Host 2
en
testking
config t
router ospf 1
network 172.10.143.0 0.0.0.255 area 0
network 172.20.44.0 0.0.0.3 area 0
network 17230.10.197 0.0.0.0 area 0
exit
ctrl z
copy run start
```

QUESTION NO: 3

Exhibit, Network Topology



Exhibit, Configuration

```
hostname TestKing2
!
interface serial0
 ip policy route-map TEST
!
access-list 1 permit 192.168.10.0 0.0.0.255
access-list 101 permit tcp host 192.168.10.5 eq telnet any
!
route-map TEST permit 5
 match ip address 1
 set ip next-hop 192.168.16.254
!
route-map TEST permit 15
 match ip address 101
 set ip next-hop 192.168.17.254
```

You work as a network technician at TestKing.com. Study the two exhibits. Policy-based routing is configured on TestKing2 to direct all traffic coming from sources with address prefixes 192.168.10.0/24 to the next hop 192.168.16.254. The exception should be made for packets originating from the Telnet port on host 192.168.10.5 to be forwarded to the next hop 192.168.17.254. The debug ip packet command executed on TestKing3 shows that the policy routing does not work correctly.

What could be the problem?

- A. The **telnet** keyword in the access-list 101 is associated with the wrong port.
- B. The standard access list 1 should be replaced with an extended access-list.
- C. The **ip local route-map** should be configured on TestKing1 Serial0 interface.
- D. The route map statements are in the wrong order.

Answer: D

QUESTION NO: 4

Which three statements are true when configuring redistribution for OSPF? Select three.

- A. The default metric is 10.
- B. The default metric is 20.
- C. The default metric type is 2.
- D. The default metric type is 1.
- E. Subnets do not redistribute by default.
- F. Subnets redistribute by default.

Answer: B, C, E

QUESTION NO: 5

Exhibit

```

router ospf 5
 network 10.0.0.0 0.0.255.255 area 0
router eigrp 100
 network 172.10.0.0
 redistribute ospf 5

```

You work as a network technician at TestKing.com. You are troubleshooting a redistribution of OSPF routes into EIGRP. Study the exhibit carefully.

Leading the way in IT testing and certification tools, www.testking.com


Which statement is true?

- A. Redistributed routes will have an external type of 1 and a metric of 1.
- B. Redistributed routes will have an external type of 2 and a metric of 20
- C. Redistributed routes will maintain their original OSPF routing metric.
- D. Redistributed routes will have a default metric of 0 and will be treated as reachable and advertised.
- E. Redistributed routes will have a default metric of 0 but will be treated as unreachable and not advertised.

Answer: E

QUESTION NO: 6

Exhibit, Network Topology



Exhibit, Router Configuration

TestKing2# show ip route

Gateway of last resort is not set

2.0.0.0/24 is subnetted, 1 subnets
 C 2.2.2.0 is directly connected, Ethernet0/0
 C 3.0.0.0/8 is directly connected, Serial1/0
 O E2 200.1.1.0/24 [110/20] via 2.2.2.2, 00:16:17, Ethernet0/0
 O E1 200.2.2.0/24 [110/104] via 2.2.2.2, 00:00:41, Ethernet0/0
 131.108.0.0/24 is subnetted, 2 subnets
 O 131.108.2.0 [110/74] via 2.2.2.2, 00:16:17, Ethernet0/0
 O IA 131.108.1.0 [110/84] via 2.2.2.2, 00:16:17, Ethernet0/0

Which command should be added to TestKingB under router bgp 100 to allow only the external OSPF routes to be redistributed to TestKingC?

- A. redistribute ospf 1
- B. redistribute ospf 1 match external 1
- C. redistribute ospf 1 match external 2
- D. redistribute ospf 1 match external 1 external 2

Answer: D

QUESTION NO: 7

Which command should you use to verify what networks are being routed by a given OSPF process?

- A. show ospf
- B. show ip route
- C. show ip protocols
- D. show ip ospf database

Answer: C

QUESTION NO: 8

Part of the configuration file and debug output for one of the TestKing routers is displayed in the following exhibit:

```

!
interface Ethernet0/0
 ip address 100.100.100.1 255.255.255.0
 ip policy route-map blah
!
interface Serial1/0
 ip address 10.10.10.1 255.255.255.0
!
interface Serial2/0
 ip address 20.20.20.1 255.255.255.0
!
ip classless
no ip http server
!
access-list 100 permit ip host 100.100.100.3 host 200.200.200.4
!
route-map blah permit 10
 match ip address 100
 set ip default next-hop 10.10.10.2
!
end
!
Routers debug ip policy

Policy routing debugging is on
*Dec 4 12:50:57.363: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4, len 100, policy match
*Dec 4 12:50:57.363: IP: route map blah, item 10, permit
*Dec 4 12:50:57.363: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4 (Serial2/0), len 100,
policy rejected -- normal forwarding
*Dec 4 12:50:57.431: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4, len 100, policy match
*Dec 4 12:50:57.431: IP: route map blah, item 10, permit
*Dec 4 12:50:57.431: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4 (Serial2/0), len 100,
policy rejected -- normal forwarding
*Dec 4 12:50:57.491: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4, len 100, policy match
*Dec 4 12:50:57.491: IP: route map blah, item 10, permit
*Dec 4 12:50:57.491: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4 (Serial2/0), len 100,
policy rejected -- normal forwarding

```

Study the exhibit carefully.

Which statement is true regarding policy-based routing?

- A. IP packets matching the access list are being forwarded to 10.10.10.2 because of the **set ip default next-hop** command.
- B. IP packets not matching the access list are being forwarded to 10.10.10.2 because of the **set ip default next-hop** command.
- C. IP packets matching the access list are not being forwarded to 10.10.10.2 because a path to 200.200.200.4 exists in the routing table.
- D. IP packets matching the access list are not being forwarded to 10.10.10.2 because a path to 200.200.200.4 does not exist in the routing table.
- E. IP packets matching the access list are being forwarded to 10.10.10.2 because OSPF has not been configured on the Serial1/0 interface.

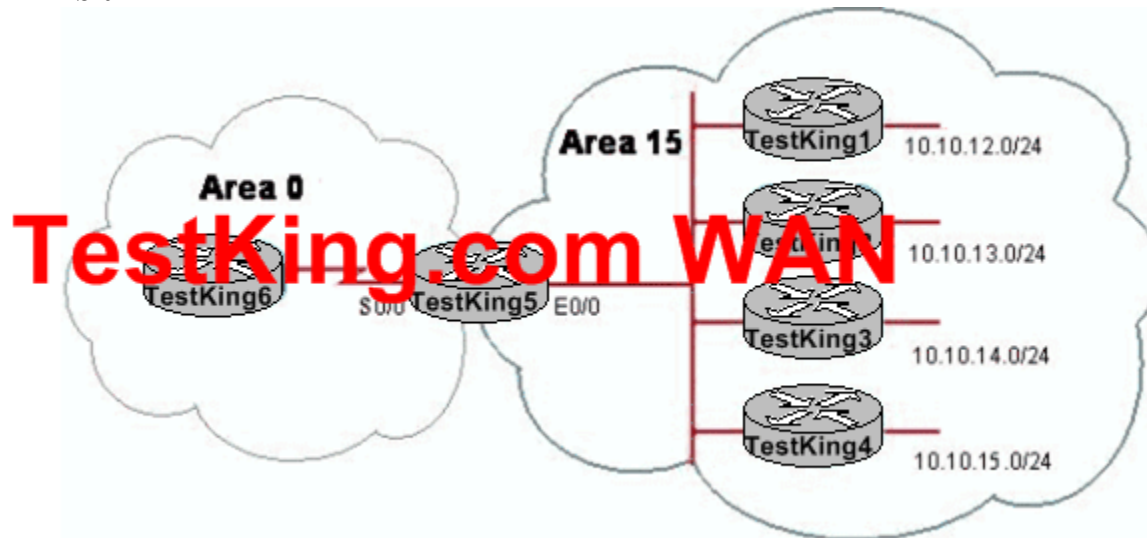
Answer: D

Explanation:

In the example above, the packets sourced from 100.100.100.3 destined to 200.200.200.4 are correctly matching the policy map as specified in access-list 100, but a route to 200.200.200.4 does not exist in the routing table, so the policy is rejected and normal forwarding will occur.

QUESTION NO: 9

Exhibit



You work as a network administrator at TestKing.com. Study the network topology exhibit carefully.

Which command will summarize area 15?

- A. on the ABR E0/0 interface: **area 15 range 10.10.8.0 255.255.248.0**
- B. on the ABR S0/0 interface: **area 15 range 10.10.8.0 255.255.252.0**
- C. on the ABR in the OSPF config-router mode: **area 15 range 10.10.12.0 255.255.252.0**
- D. on all routers in area 15 in the OSPF config-router mode: **area 15 range 10.10.12.0 255.255.248.0**
- E. on the ABR S0/0 interface: **summary-address 10.10.8.0 255.255.248.0**
- F. on the ABR in the OSPF config-router mode: **summary-address 10.10.12.0 255.255.252.0**

Answer: C

QUESTION NO: 10

Exhibit

```

!
access-list 100 permit ip host 100.100.100.3 host 200.200.200.4
!
TestKing.com
route-map dft permit 10
match ip address 100
set ip default next-hop 10.10.10.2
!

```

You work as a network administrator at TestKing.com. Study the network topology exhibit carefully. Which statement describes the result of using the set ip default next-hop command?

- A. If a packet received by the router matches the access list, it will always be forwarded to 10.10.10.2
- B. If a packet received by the router matches the access list, it will be forwarded to 10.10.10.2 only if a path to 200.200.200.0 is present in the routing table.
- C. If a packet received by the router matches the access list, it will be forwarded to 10.10.10.2 only if a path to 200.200.200.0 is not present in the routing table.
- D. If a packet received by the router matches the access list, it will be forwarded to 200.200.200.4 only if a path to 100.100.100.0 is present in the routing table.
- E. If a packet received by the router matches the access list, it will be forwarded to 200.200.200.4 only if a path to 100.100.100.0 is not present in the routing table.

Answer: C

QUESTION NO: 11

When an IPv6 enabled host boots, it sends a router solicitation (RS) message. An IPv6 router responds with a router advertisement (RA). Which two items are contained in the RA? Select two.

- A. IPv6 address for the host
- B. lifetime of the prefix
- C. prefixes for the link
- D. keepalive timers
- E. request for the local host IP address
- F. any route advertisement it has received

Answer: B, C

QUESTION NO: 12

What action does an EIGRP router take when it cannot find a feasible successor for a network?

- A. It examines the routing and neighbor tables for the next best path.
- B. It transitions from passive to active state for that network and queries its neighbors.
- C. It examines the topology for a next best path.
- D. It transitions from active to passive state for the next network and queries its neighbors.

Answer: B

QUESTION NO: 13

What does a spoke router require to be enabled to ensure proper functionality of on-demand routing?

- A. dynamic routing
- B. static routing
- C. CDP
- D. broadcasting

Answer: C

QUESTION NO: 14

Which three IPv6 notations represent the same address? Select three.

- A. 2031:0000:130F:0000:0000:09C0:876A:130B
- B. 2031::130F:9C0:876A:130B
- C. 2031:0:130F::9C0:876A:130B
- D. 2031::130F:0::9C0:876A:130B
- E. 2031:0:130F:0:0:09C0:876A:130B
- F. 2031:0:130F:::9C0:876A:130B

Answer: A, C, E

QUESTION NO: 15

Network topology exhibit



```
TestKing1# show run
hostname TestKing1
.
ip route 10.0.0.0 255.0.0.0 interface fa0/0
ip route 10.0.0.0 255.0.0.0 172.17.20.3
```

You work as a network technician at TestKing.com. Study the exhibit carefully. When performing a show run on router TestKing1, you notice that two static entries were configured for the 10.0.0.0 network. Which path will be selected by the router as the best path?

- A. Per-destination load sharing will be implemented.
- B. The path through TestKing2 will be used.
- C. The path through TestKing3 will be used.
- D. Per-packet load sharing will be implemented.

Answer: B

QUESTION NO: 16

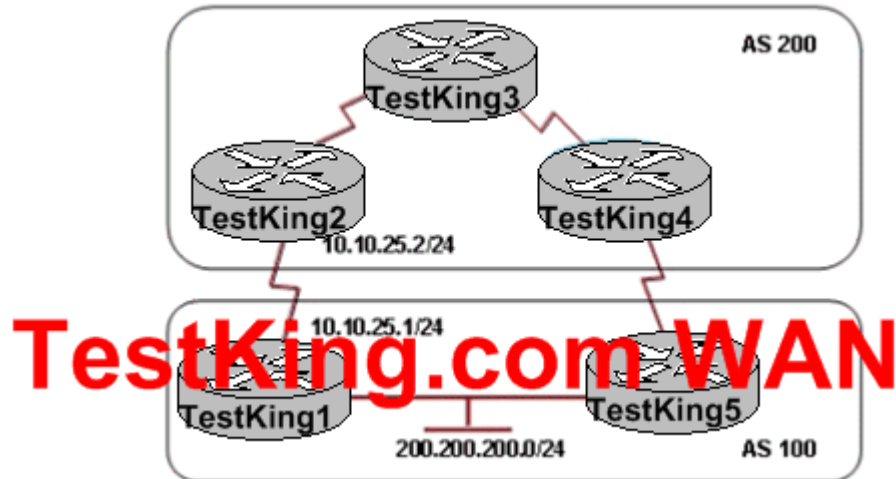
Given the NSAP, 39.0100.0102.0001.0c00.1211.00, which portion is interpreted by IS-IS as the area?

- A. 39
- B. 39.0100
- C. 39.0100.0102
- D. 0001
- E. 0001.0c00
- F. 0001.0c00.1211

Answer: C

QUESTION NO: 17

Exhibit, Network topology



Exhibit, Router configuration

```
TestKing1#show running-config
TestKing.com
router bgp 100
network 200.200.200.0 mask 255.255.255.0
neighbor 10.10.25.2 remote-as 200
neighbor 10.10.25.2 route-map BEST out
!
access-list 101 permit ip host 200.200.200.0 host 255.255.255.0
!
route-map BEST permit 10
match ip address 101
set metric 150
!
route-map BEST permit 20
```

You work as an administrator at TestKing.com. Study the exhibits carefully. BGP is configured on all routers, synchronization is turned off, and one of the default attributes have been changed except the MED attribute on TestKing5. Which path is preferred by TestKing2 to reach the network 200.200.200.0/24?

- A. TestKing2-TestKing1 because it is the shortest path.
- B. TestKing2-TestKing1 because it has a higher metric.
- C. TestKing3-TestKing4-TestKing5 because it has a lower administrative distance.
- D. TestKing3-TestKing4-TestKing5 because it has a lower metric.

Answer: D

Leading the way in IT testing and certification tools, www.testking.com

QUESTION NO: 18

Which three statements are true concerning redistributed routes when a default metric is not configured? Select three.

- A. RIP, IGRP, and EIGRP assign a default metric of 0 (infinity) to redistributed routes and will advertise these routes accordingly.
- B. RIP, IGRP, and EIGRP assign a default metric 0 (infinity) to redistributed routes and will only advertise these routes if a valid seed metric is configured.
- C. IS-IS assigns a default metric of 0 to redistributed routes.
- D. IS-IS assigns a default metric of 10 to redistributed routes.
- E. OSPF assigns a default metric of 1 for routes from all protocols except BGP, which gets a metric of 20.
- F. OSPF assigns a default metric of 20 for routes from all protocols except BGP, which gets a metric of 1.

Answer: B, C, F

QUESTION NO: 19

In IS-IS routing, which level is used to route between different areas within the same domain?

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

Answer: C

Topic 6: Extra practice Questions (200 questions)

For students that want extra practice. These questions are not strictly necessary to pass the exam. These questions should be used with the purpose to reinforce exam concepts.

QUESTION NO: 1

You are the network administrator at TestKing. You have configured multiple IP routing protocols on a single router on the TestKing network.

Which command lists the filters applied to routing updates on a routing protocol basis?

- A. show ip
- B. show ip route
- C. show ip protocols
- D. show ip interface

Answer: C

Explanation:

The show ip protocols command will display the IP routing protocols configured on the router and will also show what each routing process is redistributing. In addition, it will list the redistribution filters applied to interfaces. Specifically, the output will show:

Routing protocol and process ID

Update frequency

Hold down timers

Incoming and outgoing filters

Default distribution metric

Redistribution parameters

Chapter: 1

QUESTION NO: 2

Using route summarization, which two of these networks fall into the 174.69.16.0/20 range?

- A. 174.69.33.0/24
- B. 174.69.31.0/24
- C. 174.69.17.0/24
- D. 174.69.32.0/24

Answer: B, C

Explanation:

The valid 24-bit subnets from the address 174.69.16.0/20 are:

174.69.16.0/24 174.69.24.0/24
 174.69.17.0/24 174.69.25.0/24
 174.69.18.0/24 174.69.26.0/24
 174.69.19.0/24 174.69.27.0/24
 174.69.20.0/24 174.69.28.0/24
 174.69.21.0/24 174.69.29.0/24
 174.69.22.0/24 174.69.30.0/24
 174.69.23.0/24 174.69.31.0/24

Chapter: 1

QUESTION NO: 3

Which of the following statement is true when a static route is configured on a router and that static route is advertised throughout the network?

- A. The router automatically advertises static routes to all routers
- B. You should configure redistribution using the redistribute static command
- C. You should enable static advertisements using the advertise static route command
- D. You should include the static route in a distribution list and specify which interface to use when redistributing the route.

Answer: B

Explanation:

To redistribute static routes that have been created on the local router to other routers in the network, use the redistribute static command.

Chapter: 1

QUESTION NO: 4

You are the network administrator at TestKing. You are configuring redistribution to advertise OSPF routes into EIGRP on a boundary router on the TestKing network. You specify a seed metric with the default-metric command.

What is the format of the metric being specified?

- A. hop-count
- B. hop-count ticks
- C. bandwidth delay hop-count load
- D. load delay hop-count reliability mtu
- E. bandwidth delay reliability load mtu

Answer: E

Explanation:

When redistributing static routes or other protocols within EIGRP, metrics can be set for these routes using the default-metric command. The range of values for each parameter is listed below:

bandwidth - 0 to 4,294,967,295 in Kbps

delay - 0 to 4,294,967,295 in 10-microsecond units

reliability - 0 to 255 with 255 being the most reliable

load - 0 to 255 with 255 being a saturated link

MTU - 0 to 4,294,967,295

Chapter: 1

QUESTION NO: 5

Which switching mode is enabled by default on a router running Cisco IOS 11.2 or later, to forward packets that match the established policy routing?

- A. fast
- B. wire-speed
- C. NetFast
- D. packet

Answer: A

Explanation:

Fast switching is the default switching mechanism on all Cisco router platforms. It is accomplished by maintaining a cache of recently switched destinations therefore reducing the number of full route table lookups. It also allows the information required for MAC header rewrites to be stored in cache rather than being recalculated.

Chapter: 1

QUESTION NO: 6

You are the network administrator at TestKing. Router TK1 is configured as follows:

```
router igrp 300
network 192.168.20.0
network 192.168.24.0
network 192.168.27.0
redistribute rip
default-metric 10 100 255 1 1500
distance 140 0.0.0.0 255.255.255.255 9
access-list 9 permit 192.168.20.0
```

```
access-list 9 permit 192.168.24.0
access-list 9 permit 192.168.27.0
```

Which of the following statements are true? (Choose all that apply.)

- A. Networks 192.168.20.0, 192.168.24.0, and 192.168.27.0 are allowed into the routing table
- B. The RIP learned routes to networks 192.168.20.0, 192.168.24.0, and 192.168.27.0, will be assigned an administrative distance of 140
- C. The IGRP learned routes to networks 192.168.20.0, 192.168.24.0, and 192.168.27.0, will be assigned an administrative distance of 140
- D. Changing the administrative distance to a number larger than the default value makes networks 192.168.20.0, 192.168.24.0, and 192.168.27.0 unreachable

Answer: A, B

Explanation:

The networks listed in under the IGRP section are advertised to other routers on the network and installed into the routing table. The routes that are redistributed from RIP from those same networks are assigned an administrative distance of 140 because of the distance command listed above. The distance command is used to define an administrative distance for routes learned from other routing protocols. The last argument (9) specifies that access-list 9 be used to permit/deny networks.

Chapter: 1

QUESTION NO: 7

Which command could you use to verify proper operation of multiple routing protocols that are sharing routes?

- A. ping
- B. show ip route
- C. show cdp neighbor
- D. show ip ospf neighbor

Answer: B

Explanation:

This is the only command listed that can show any information regarding the state of routes or routing protocols. Answer-A will not show any information regarding route selection or the route to a target address. Ping can be used to verify connectivity to another IP address.

Chapter: 1

QUESTION NO: 8

Which Cisco IOS command can be used to display the route maps configured on an interface?

- A. show interface
- B. show route-map
- C. show ip policy
- D. show ip route map

Answer: B

Explanation:

This command will display all route-maps that are configured. If you specify a route-map as an argument, then only that route-map is displayed. See the sample output below:

```
Router# show route-map
route-map new, permit, sequence 10
Match clauses:
tag 1 2
Set clauses:
metric 5
route-map new2, permit, sequence 20
Match clauses:
tag 3 4
Set clauses:
metric 6
Chapter: 1
```

QUESTION NO: 9

What happens due to the implicit deny at the end of a route-map?

- A. Packets that reach the end of the route map are discarded
- B. Packets are forwarded to the null interface for special handling
- C. Packets that reach the end of the route map are routed in normal fashion
- D. Packets that reach the end of the route map are returned to the originating interface

Answer: C

Explanation:

The use of route maps for policy-based routing is a little different than other application of route maps. When used for policy-based routing, if a packet does not match the criteria specified in the route map or a matched route map statement specifies deny, then the packet is not dropped. It is sent to the routing process and routed normally, by destination, as if it had never encountered a

route map. If your intention is to drop packets that do not match the criteria, it is necessary to use the set command to route packets to the null interface as the last entry in the route map.

Source: Self-Study CCNP BSCI Exam Certification Guide Third Edition P.674

Topic: Understanding Policy-Based Routing

QUESTION NO: 10

You are the network administrator at TestKing. You want to redistribute and advertise EIGRP routes into OSPF on a boundary router. The router has the following configuration:

```
router ospf 1
 redistribute eigrp 1 metric 25 subnets
```

What does the 25 parameter in the redistribute command specify?

- A. It specifies the seed cost to be applied to the redistributed routes
- B. It specifies the administrative distance on the redistributed routes
- C. It specifies the metric limit of 25 subnets in each OSPF route advertisement
- D. It specifies the process-id for the pseudo process that injects the EIGRP routes into OSPF

Answer: A

Explanation:

The metric {value} command specifies the seed metric for use in redistributed routes.

Reference: Building Scalable Cisco Networks (Cisco Press) page 456

QUESTION NO: 11

You are the network administrator at TestKing. A router on the TestKing network has one serial interface and one Ethernet interface. Given the serial interface to a WAN configuration:

```
interface serial 0/122 point-to-point
 ip address 192.168.1.2 255.255.255.0
 encapsulation frame-relay
 frame-relay interface-dlci 122
```

Which command prevents routing protocol information from being sent on the Ethernet interface?

- A. interface serial 0.122 point-to-point
passive-interface ethernet 0
- B. interface Ethernet 0
ip address 192.168.12.1 255.255.255.0
passive interface
- C. router ospf 102
area 1 ospf
network 192.168.1.0 0.0.0.255 area 0
network 192.168.12.0 0.0.0.255 area 1
- D. router ospf 102
passive-interface Ethernet 0
network 192.168.1.0 0.0.0.255 area 0
network 192.168.12.0 0.0.0.255 area 1

Answer: D

Explanation:

When a passive interface is defined for any routing process, then updates are not sent on the specified interface by that routing process. Passive interfaces must be defined for each routing protocol (process). The passive interface command is not a valid interface configuration command.

Chapter: 1

QUESTION NO: 12

Which of the following commands would produce output that can be used to verify route redistribution? (Choose all that apply.)

- A. debug
- B. traceroute
- C. show tech-support
- D. show ip route

Answer: A, B, D

Explanation:

Debug can be used to view routing protocol information exchanged between routers. Traceroute can be used to determine the path an IP packet will take when traversing the network. The show ip route command will display all known routes and indicate the source of the route (Static, OSPF, RIP, etc)

Chapter: 1

QUESTION NO: 13

Which command forces manually entered route entries are injected into the routing process?

- A. inject static
- B. inject permanent
- C. redistribute all
- D. redistribute static

Answer: D

Explanation:

The redistribute static command is used to inject static routes into the routing protocol's route table and subsequent updates.

Chapter: 1

QUESTION NO: 14

You are the network administrator at TestKing. The TestKing network includes Router TK1. Router TK1 is configured as follows:

```
interface serial 0
ip address 10.1.1.1 255.255.255.0
encapsulation frame-relay
ip ospf network point-to-multipoint
router ospf 7
network 10.1.1.0 0.0.0.255 area 0
```

Which of the following statements are true? (Choose all that apply.)

- A. DR/BDR elections do not take place
- B. Neighbor statements are required
- C. Communication between neighbors is broadcast to 255.255.255.255
- D. The area 0 NBMA cloud is configured as more than one subnet

Answer: A, B

Explanation:

When configuring OSPF in a point-to-multipoint environment, DR/BDR elections do not take place. Neighbor statements must be statically defined due to the NBMA architecture. The point-to-multipoint environment removes the assumption that there is a full mesh and communication between neighbors is done via unicast.

Chapter: 1

QUESTION NO: 15

When configuring a router to participate in an OSPF area, what is the default priority used in DR/BDR elections?

- A. 0
- B. 1
- C. 16
- D. 255

Answer: B

Explanation:

The ip ospf priority command can be used to administer which router becomes the DR. This number ranges from 0-255 and defaults to 1. A router configured with a priority of 0 can never be elected DR.

Chapter: 1

QUESTION NO: 16

Which two are benefits of using OSPF over RIP as a routing protocol in a large network? (Choose all that apply.)

- A. OSPF has fewer tables to manage
- B. OSPF is a simpler protocol than RIP
- C. OSPF has virtually no reachability limits
- D. OSPF uses a metric that is based on bandwidth to select a path through a network

Answer: C, D

Explanation:

OSPF has virtually no limits with regard to scalability in large networks because of its hierarchal design. RIP uses a hop count limit (15 hops) to prevent routing loops. It is possible in a very large network to outgrow this limitation with RIP. RIP also uses hop count as its metric for selecting the best route. Cisco's implantation of OSPF calculates link cost based on bandwidth ($10^8 / \text{Interface Bandwidth}$) to determine path selection across the network.

Chapter: 1

QUESTION NO: 17

You are a technician at TestKing. You want to assign an OSPF router ID of 172.16.20.127.

Which series of commands should you use?

- A. ospf loopback 0
ip address 172.16.20.127 255.255.255.0
- B. router loopback 0
ip address 172.16.20.127 255.255.255.0
- C. interface loopback 0
ip address 172.16.20.127 255.255.255.0
- D. ospf interface loopback 0
ip address 172.16.20.127 255.255.255.0

Answer: C

Explanation:

In Cisco's OSPF implementation, the Loopback interface address is used as the Router ID. If the Loopback interface is not configured with an IP address, the highest IP address configured on any router interface is used.

Chapter: 1

QUESTION NO: 18

What is used to determine which router that will become the DR in an OSPF network?

- A. the lowest router ID
- B. the highest priority value
- C. the first router to attach to the network
- D. a router that is connected to more than one OSPF area and designated ASBR

Answer: B

Explanation:

The router with the highest priority value is elected as the DR. The second highest priority value becomes the BDR. The ip ospf priority command can be used to administer which router becomes the DR. This number ranges from 0-255 and defaults to 1. A router configured with a priority of 0 can never be elected DR.

Chapter: 1

QUESTION NO: 19

Before an running OSPF can route traffic to another OSPF neighbor, what state must the router be in to route traffic?

- A. full state

- B. INIT state
- C. 2wy
- D. forwarding state

Answer: A

Explanation:

In order to route traffic to an OSPF neighbor router, the adjacency must be established before any traffic can be passed. The adjacency is not established until DR/BDR elections are completed and link-state information is exchanged (full routing information).

Chapter: 1

QUESTION NO: 20

In OSPF, what is defined using the network command? (Choose all that apply.)

- A. the OSPF area ID
- B. the OSPF router ID
- C. the OSPF process ID
- D. which interface is in which OSPF area

Answer: A, D

Explanation:

In OSPF, the network area command defines the interfaces on which OSPF runs and the area ID for those interfaces. The syntax for the command is below: network [address] [wildcard-mask] area [area-id]

Chapter: 1

QUESTION NO: 21

How does OSPF simulate a broadcast environment in an NBMA point-to-multipoint configuration for routed traffic?

- A. by creating adjacencies with each endpoint
- B. by sending replicated traffic to each neighbor
- C. by using the 224.0.0.5 multicast address on serial links
- D. by separating out each endpoint using the hello protocol

Answer: A

Explanation:

In a point-to-multipoint configuration, neighbors must be statically defined and communication are done via unicast instead of multicast.

Chapter: 1

QUESTION NO: 22

Which command can be used to verify when out-of-date routes will be removed from the topological database?

- A. show ip ospf
- B. show ip route
- C. show ip ospf interface
- D. show ip ospf topo-database

Answer: A

Explanation:

Displays the Link State Update Interval and the Link State Age Interval and when an update is due. Each route is flooded throughout the area via an LSA. Each LSA has an age field that is incremented while it is contained in the database or as it gets flooded throughout the area. When an LSA reaches a Maxage it gets flushed from the database if that LSA is not on any neighbors retransmission list.

Router# show ip ospf

Routing Process "ospf 201" with ID 192.42.110.200

Supports only single TOS(TOS0) route

It is an area border and autonomous system boundary router

Summary Link update interval is 0:30:00 and the update due in 0:16:26

External Link update interval is 0:30:00 and the update due in 0:16:27

Redistributing External Routes from,

igrp 200 with metric mapped to 2, includes subnets in redistribution

rip with metric mapped to 2

igrp 2 with metric mapped to 100

igrp 32 with metric mapped to 1

Number of areas in this router is 3

Area 192.42.110.0

Number of interfaces in this area is 1

Area has simple password authentication

SPF algorithm executed 6 times

Area ranges are

Link State Update Interval is 0:30:00 and due in 0:16:55

Link State Age Interval is 0:20:00 and due in 0:06:55

Chapter: 1

QUESTION NO: 23

At a minimum, which two configuration commands are required to configure OSPF on a single internal router? (Choose all that apply.)

- A. network
- B. neighbor
- C. router ospf dr 1
- D. router ospf

Answer: A, D

Explanation:

OSPF is enabled on a router by specifying an OSPF process ID and defining the network, interfaces used, and area-id that will be included in the OSPF process. The network command defines the interfaces on which OSPF runs and the area ID for those interfaces.

Chapter: 1

QUESTION NO: 24

You are a trainee technician at TestKing. Your instructor shows you the following router configuration:

```
interface serial 0
ip address 172.14.12.1 255.255.255.224
encapsulation frame-relay
ip ospf network non-broadcast
!
router ospf
network 172.14.12.0 31.255.255.255
neighbor 172.14.12.2
neighbor 172.14.12.3
```

Your instructor wants to know which of the following statements are true.

What would your reply be? (Choose all that apply.)

- A. DR/BDR elections are not held
- B. This is a point-to-multipoint configuration
- C. The network type is non-broadcast multi-access (NBMA)
- D. The DR and BDR require a static list of neighbors

Answer: C, D

Explanation:

Because the network type is defined as non-broadcast, DR/BDR election take place based on statically defined neighbors. In addition, communication between neighbors is done via unicast instead of multicast. If this were a point-to-multipoint configuration, the ip ospf network point-to-multipoint command would have been used on the serial interface.

Chapter: 1

QUESTION NO: 25

You are a technician at TestKing. TestKing has an OSPF network. Your newly appointed TestKing trainee wants to know what is used to send link-state information to all other routers within an OSPF area.

What would your reply be?

- A. LSA - router link, type 1
- B. LSA - network link, type2
- C. LSA - external link, type 5
- D. NSA - network summary link, type 3

Answer: B

Explanation:

Network Link Advertisements are sent during the adjacency process to inform the neighbor of its network links. When a link changes state or a new link added on an existing router, the router that owns the link generates a new LSA.

Chapter: 1

QUESTION NO: 26

Which of the following features require the subnet mask to be carried within OSPF routing protocol updates? (Choose all that apply.)

- A. VLSM
- B. NBMA
- C. summarization
- D. SPF route calculation

Answer: A, C

Explanation:

OSPF like all classless routing protocols, carry the network number and mask in its updates. This is required to provide support for VLSM, route summarization, and super netting.

Chapter: 1

QUESTION NO: 27

If an OSPF router has interfaces connected in two or more areas, what kind of router is it considered?

- A. ABR
- B. ASBR
- C. MAR
- D. backbone router

Answer: A

Explanation:

An ABR is a router that has multiple interfaces with at least two interfaces in two different OSPF areas. An ASBR is a router with at least one interface connected to an external network or AS.

Chapter: 1

QUESTION NO: 28

Which is true of an OSPF area with too many routers?

- A. The second BDR cannot keep all the LSA information up to date
- B. Convergence time can be slower
- C. A second backbone area must be created to split the traffic into two areas.
- D. Route processing time is decreased because the information is dispersed among all routers in the area

Answer: B

Explanation:

With too many routers will take longer to converge. Answer should be B. Also not possible to have two backbone areas in ospf. Only one is allowed.

QUESTION NO: 29

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know what is used to connect a new OSPF area to area 0

What would your reply be?

- A. external router

- B. DR
- C. BDR
- D. backbone router

Answer: D

Explanation:

Any router that will have any interface connected to an OSPF backbone area is considered a backbone router.

Chapter: 1

QUESTION NO: 30

You are the network administrator at TestKing. You are using OSPF as your IGP throughout the TestKing network. You want to connect the network to the outside world or to a different routing protocol.

Which kind of OSPF router must you configure?

- A. ABR
- B. BDR
- C. ASBR
- D. neighbor border router
- E. backbone router

Answer: C

Explanation:

ASBR's are used to connect two separate autonomous systems together. The role of the ASBR is to exchange routing information between the two routing processes.

Chapter: 1

QUESTION NO: 31

In order to summarize routes, which configuration requirement does an ASBR have that an ABR does not?

- A. area range command
- B. ospf summarize command
- C. aggregate-route command
- D. summary-address command

Answer: D

Explanation:

The router subordinate command summary-address is used on ASBR's to consolidate external routes. It can be used in combination with the stub areas or used stand-alone

Chapter: 1

QUESTION NO: 32

Which command should you use to verify which process is responsible for routing which network?

- A. show ospf
- B. show ip route
- C. show ip protocols
- D. show ip ospf database

Answer: D

Explanation:

Sample output is shown below. Note the router ID and process ID in the first line. Each link ID is

representative of a network route.

router#show ip ospf database

OSPF Router with id(190.20.239.66) (Process ID 300)

Displaying Router Link States(Area 0.0.0.0)

Link ID ADV Router Age Seq# Checksum Link count

155.187.21.6 155.187.21.6 1731 0x80002CFB 0x69BC 8

155.187.21.5 155.187.21.5 1112 0x800009D2 0xA2B8 5

155.187.1.2 155.187.1.2 1662 0x80000A98 0x4CB6 9

155.187.1.1 155.187.1.1 1115 0x800009B6 0x5F2C 1

155.187.1.5 155.187.1.5 1691 0x80002BC 0x2A1A 5

155.187.65.6 155.187.65.6 1395 0x80001947 0xEEE1 4

155.187.241.5 155.187.241.5 1161 0x8000007C 0x7C70 1

155.187.27.6 155.187.27.6 1723 0x80000548 0x8641 4

155.187.70.6 155.187.70.6 1485 0x80000B97 0xEB84 6

Chapter: 1

QUESTION NO: 33

You are the network administrator at TestKing. TestKing has an address range of 172.16.20.192 to 172.16.20.223. You want to configure the area 3 border router for network summarization

Which configuration command must you use? Select two.

- A. summarize 172.16.20.192 0.0.0.31 area 3
- B. area 3 range 172.16.20.192 172.16.20.223
- B. area 3 range 172.16.20.192 255.255.255.224
- C. D. network 172.16.20.192 255.255.255.224 area 3

Answer: A, C

Explanation:

The area range command is used on ABR's to summarize and advertise routes. A network statement must also be defined, but answer D does not use a wildcard mask and therefore is an invalid command.

Chapter: 1

QUESTION NO: 34

You are a technician at TestKing. TestKing has an OSPF network. Your newly appointed TestKing trainee wants to know why VLSM is used in an OSPF network.

What would your reply be? (Choose all that apply.)

- A. to allow for address summarization
- B. to allow use of the all zero's subnet
- C. to make efficient use of available addresses
- D. it is required for a point-to-multipoint nonbroadcast network

Answer: A, C

Explanation:

VLSM is utilized in OSPF for address summarization and it also allows for more efficient use of networks due to support for classless boundaries.

Chapter: 1

QUESTION NO: 35

What does an ABR connect in an OSPF network?

- A. multiple OSPF areas
- B. OSPF and RIP networks
- C. multiple designated routers
- D. multiple autonomous systems

Answer: A

Explanation:

An ABR is a router that has multiple interfaces with at least two interfaces in two different OSPF areas. It is used to connect different OSPF areas.

Chapter: 1

QUESTION NO: 36

Which OSPF router is responsible for flooding an of OSPF area with type 2 link LSAs?

- A. DR
- B. ABR
- C. BDR
- D. ASBR

Answer: A

Explanation:

The DR is responsible for flooding the network with network LSA's when a change occurs within the OSPF area.

Chapter: 1

QUESTION NO: 37

Which statement regarding route summarization within OSPF is true?

- A. Summarization must be performed by every router within an OSPF area.
- B. Summarization prevents type 1 link LSAs from being propagated into the backbone area 0.
- C. Route summarization can be performed at any point in the network where enough contiguous addresses are present
- D. Route summarization reduces the amount of bandwidth, CPU, and memory resources consumed by the OSPF process.

Answer: D

Explanation:

Route summarization reduces the amount of bandwidth, CPU, and memory resources consumed by the OSPF process. Each additional network requires an individual entry in the routing table and must be propagated throughout the network and added to each router's table.

Chapter: 1

QUESTION NO: 38

Which of the following Cisco IOS commands that can be used to view neighbor adjacencies?

(Choose all that apply.)

- A. show ip ospf database
- B. show ip ospf neighbor
- C. show ip ospf protocols
- D. show ip ospf interface ethernet 0

Answer: B, D

Explanation:

Both of these commands display neighbor adjacency information.

```
Router# show ip ospf interface ethernet 0
```

```
Ethernet 0 is up, line protocol is up
```

```
Internet Address 131.119.254.202, Mask 255.255.255.0, Area 0.0.0.0
```

```
AS 201, Router ID 192.77.99.1, Network Type BROADCAST, Cost: 10
```

```
Transmit Delay is 1 sec, State OTHER, Priority 1
```

```
Designated Router id 131.119.254.10, Interface address 131.119.254.10
```

```
Backup Designated router id 131.119.254.28, Interface addr 131.119.254.28
```

```
Timer intervals configured, Hello 10, Dead 60, Wait 40, Retransmit 5
```

```
Hello due in 0:00:05
```

```
Neighbor Count is 8, Adjacent neighbor count is 2
```

```
Adjacent with neighbor 131.119.254.28 (Backup Designated Router)
```

```
Adjacent with neighbor 131.119.254.10 (Designated Router)
```

```
Router# show ip ospf neighbor
```

```
ID Pri State Dead Time Address Interface
```

```
199.199.199.137 1 FULL/DR 0:00:31 160.89.80.37 Ethernet0
```

```
192.31.48.1 1 FULL/DROTHER 0:00:33 192.31.48.1 Fddi0
```

```
192.31.48.200 1 FULL/DROTHER 0:00:33 192.31.48.200 Fddi0
```

```
199.199.199.137 5 FULL/DR 0:00:33 192.31.48.189 Fddi0
```

```
Chapter: 1
```

QUESTION NO: 39

What term is used to describe a BGP autonomous system connected to two different BGP autonomous systems for increased reliability?

- A. multi-exit
- B. multisource
- C. multihomed
- D. multi-neighbor

Answer: C

Explanation:

Multihoming refers to a single network or AS as having more than one connection to another network or AS to improve reliability and/or performance.

Chapter: 1

QUESTION NO: 40

You are the network administrator at TestKing. TestKing has subscribed to multiple ISPs. You use BGP to connect to multiple ISPs. You want to force outbound Internet traffic to one ISP unless there is a link failure.

Which tool would you use?

- A. configure weight
- B. enable route reflector
- C. create a distribute list
- D. enable the Longer Autonomous System path option

Answer: A

Explanation:

The weight attribute is a Cisco proprietary attribute used by BGP in path selection. This allows as administrator to “prefer” one path over one or more paths to the same destination. In the event the preferred route or link fails, the secondary route will automatically be used by outbound traffic.

Chapter: 1

QUESTION NO: 41

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know which command displays both the configured iBGP and eBGP neighbors.

What would your reply be?

- A. show bgp neighbors
- B. show ip bgp paths
- C. show ip bgp peers
- D. show ip bgp summary

Answer: D

Explanation:

This is the only valid command that is listed that will show any information about BGP neighbors. The show ip bgp neighbors will show detailed information about each neighbor but the syntax listed is incorrect. The show ip bgp summary will show the status of all configured BGP connections.

Chapter: 1

QUESTION NO: 42

Which method makes it possible to receive BGP routes from multiple ISPs?

- A. accept only IGP routes from the ISPs
- B. accept an external route from the ISPs
- C. accept only default routes from the ISPs
- D. accept only redistributed routes from the ISPs

Answer: B

Explanation:

This is a major difference between iBGP and eBGP. A BGP router will never forward a path learned from an iBGP peer to another iBGP peer. However, eBGP peers always forward routes learned from one eBGP peer to both eBGP and iBGP peers. Thus receiving routes from 2 upstream ISP's (eBGP peers) is allowed.

Chapter: 1

QUESTION NO: 43

With regard to BGP updates, which of the following statements is true?

- A. A BGP router will forward a learned path from an iBGP peer to another iBGP peer.
- B. A eBGP peer will never forward a learned path to an iBGP peer
- C. BGP Route reflectors will propagate a route learned from an iBGP peer to eBGP peers
- D. If a BGP route was learned via an update from an iBGP peer, it will propagate this information to iBGP and eBGP peers.
- E. If a BGP route was learned via an update from an eBGP peer, it will propagate this information to iBGP and eBGP peers.

Answer: E

Explanation:

This is a major difference between iBGP and eBGP. A BGP router will never forward a path learned from an iBGP peer to another iBGP peer. However, eBGP peers always forward routes learned from one eBGP peer to both eBGP and iBGP peers. Thus receiving routes from 2 upstream ISP's (eBGP peers) is allowed.

Chapter: 1

QUESTION NO: 44

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know why it is necessary to redistribute or advertise IGP routes into BGP.

What would your reply be?

- A. so BGP can propagate this information to other IGP neighbors
- B. so BGP can propagate this information to other iBGP neighbors
- C. so BGP can propagate this information to other OSPF neighbors
- D. so BGP can propagate this information to other eBGP neighbors

Answer: B

Explanation:

iBGP will advertise these redistributed routes to all configured iBGP neighbors on the network. This will provide the entire AS with the routing information required to reach networks that are routed with different routing protocols.

Chapter: 1

QUESTION NO: 45

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know why an iBGP router must be peered with all iBGP routers within an AS.

What would your reply be? (Choose two.)

- A. iBGP routes are not propagated to other eBGP peers
- B. iBGP routes that a router originates are propagated to other iBGP peers
- C. iBGP routes are propagated to other iBGP speakers in the AS that are not peers
- D. iBGP routes that are learned from an eBGP neighbor are propagated to only eBGP peers

Answer: B, D

Explanation:

A BGP router will never forward a path learned from an iBGP peer to another iBGP peer. So, in order for all routers to know about the routes originated on any router in the network, they must

maintain the full mesh because another iBGP peer will never send the update to its iBGP neighbors.

Chapter: 1

QUESTION NO: 46

You are the network administrator at TestKing. You have limited router memory. Under these conditions, what does Cisco suggest as the best way to connect to multiple ISPs using BGP?

- A. receive only default routes
- B. receive only external BGP routes
- C. receive only internal BGP routes
- D. receive only redistributed routes

Answer: A

Explanation:

When working with limited router resources, use default routes instead of BGP routes from the ISP. Your internal AS decides which ISP to use and sends the traffic to the appropriate ISP.

Chapter: 1

QUESTION NO: 47

You are a trainee technician at TestKing. Your instructor shows you the following router configuration:

```
router bgp 6500
 redistribute static
 ip route 164.20.0.0 255.255.0.0 null 0
```

Your instructor wants to know which of the following statements are true.

What would your reply be?

- A. It allows BGP to advertise the 164.20.0.0/16 network
- B. It results in all traffic for all subnets of 172.16.0.0 being dropped at this router
- C. Cisco prefers that you use the aggregate-address command to distribute IGP routes into BGP
- D. Cisco prefers this method of distributing IGP routes into BGP over using the network command

Answer: A

Explanation:

By using the redistribute static command, the static route will be propagated throughout the network. Because the route directs traffic that is destined for this network to null0, all traffic will be dropped at all local routers within the AS.

Chapter: 1

QUESTION NO: 48

You are a trainee technician at TestKing. Your instructor shows you the following partial information from the output of a BGP command on Router TK1:

```
Network Next Hop Metric LocPrf Weight Path
192.168.2.0 10.15.10.2 0 100 65250 65000 i
10.15.20.2 0 120 65200 65000 i
10.15.30.2 0 130 65000 i
10.15.40.2 0 140 65000 i
```

Your instructor wants to know which next-hop address Router TK1 uses to send data destined for the network 192.168.2.0.

What would your reply be? (Choose all that apply.)

- A. 10.15.30.2
- B. 10.15.40.2
- C. 10.15.10.2
- D. 10.15.20.2

Answer: B**Explanation:**

The next hop router 10.15.40.2 has the highest weight value. The weight attribute is Cisco proprietary and when multiple paths to the same destination exist, the connection with the highest weight value is used. This allows an administrator to prefer one path over one or more others to the same destination.

Chapter: 1

QUESTION NO: 49

What is the correct command to create a BGP prefix list that will permit all prefixes between /8 and /24 for the 10.0.0.0 network?

- A. ip prefix-list 10.0.0.0/8 ge 8 le 24
- B. ip prefix-list 10.0.0.0/8 ge 24 le 8

- C. ip prefix-list 10.0.0.0/24 ge 24 le 8
- D. ip prefix-list 10.0.0.0/24 ge 8 le 24

Answer: A

Explanation:

The prefix list optional syntaxes ge-value and le-value are used when you need to specify a range of the prefix that is more specific than identified in the network/len syntax. Use the following rule when specifying these values:

len < ge-value < le-value <=32.

Chapter: 1

QUESTION NO: 50

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know what can be used to advertise iBGP learned routes to other iBGP neighbors within the AS.

What would your reply be?

- A. client router
- B. EBGP peer
- C. route reflector
- D. community router

Answer: C

Explanation:

Route reflectors are the only BGP routers that can propagate iBGP routes to other iBGP peers. By configuring route reflectors, you reduce the number of neighbor peering relationships in an AS. This creates a central source for updates to the route reflector clients and eliminates the need for a fully meshed iBGP network.

Chapter: 1

QUESTION NO: 51

You are the network administrator at TestKing. The TestKing network has three configured BGP route reflectors. Each route reflector has a minimum of 2 clients. Your newly appointed TestKing trainee wants to know what action a BGP route reflector takes if it receives updates from a peer in another autonomous system.

What would your reply be?

- A. It discards the route

- B. It sends the update to all iBGP peers
- C. It sends the update only to nonclients
- D. It sends the update only to its configured route reflector clients
- E. It send the update to all routers in the autonomous system

Answer: B

Explanation:

When a route reflector receives an update from a peer in an external AS, the routes are only advertised to the reflector's peers (established BGP sessions). This would mean that the new route would be sent to the reflector's clients and/or other configured route reflectors within the AS. The reflector that receives the update would NOT send to all routers within the AS.

Chapter: 2

QUESTION NO: 52

Which methods advertises your internal networks to external ISPs via BGP? Select two.

- A. using aggregate routes
- B. disabling synchronization
- C. forcing the next-hop address
- D. defining routes via the network statement

Answer: A, D

Explanation:

BGP will advertise the network number and mask specified in the network statement unless the community attribute is changed to NO ADVERTISE or a route filter is used to block the advertisement.

Chapter: 2

QUESTION NO: 53

You are the network administrator at TestKing. Router TK1 is the headquarters router in a hub and spoke topology supporting 24 remote offices. Point-to-point Frame Relay EIGRP network is deployed between the headquarters and the remote offices. The CIR for each Frame Relay PVC is different and that there is no bandwidth command configured under either the major serial interface nor the subinterfaces on Router TK1.

What is a possible fix for the potential EIGRP packet pacing problem because of the different CIR each PVC has?

- A. convert each Frame Relay PVC to a point-to-multipoint connection

- B. manually configure the bandwidth of the major interface to the lowest CIRx24
- C. manually configure the bandwidth of the major interface to the highest CIRx24
- D. manually configure the bandwidth of each of these PVCs to equal to their respective CIR

Answer: D

Explanation:

EIGRP assumes that all serial interfaces operate at T-1 speed. By configuring a bandwidth for each subinterface, EIGRP can identify slow links (< T-1)and will not generate packets faster than the configured line speed.

Chapter: 2

QUESTION NO: 54

Which type of packet is used by EIGRP routers build a neighbor table?

- A. hello
- B. ACK
- C. LSA-Type 1
- D. query
- E. update

Answer: A

Explanation:

The hello protocol is used to establish neighbor relationships on a common network. Two routers become neighbors (establish adjacency) when they acknowledge each other's hello packets and their K values match

Chapter: 2

QUESTION NO: 55

You are a trainee technician at TestKing. Your instructor shows you the following configuration commands:

- 1) router eigrp 200
- 2) network 172.16.0.0
- 3) network 3.0.0.0

Your instructor wants to know which of the following statements are true.

What would your reply be? (Choose all that apply.)

- A. Line 1 defines EIGRP as an ip routing process in area 200
- B. The command network 172.16.0.0 causes this router to become the access point for the default network
- C. Line 2 causes all interfaces connected to network 172.16.0.0 to send and receive EIGRP updates to/from other EIGRP routers
- D. The number at the end of line 1 restricts this EIGRP routing process to only communication with other EIGRP routing process that have the same number

Answer: C, D

Explanation:

EIGRP will send/receive updates on all interfaces that are included on the specified network. The number at the end of the line is the autonomous system number and routers will only exchange protocol information with routers within the same autonomous system.

Chapter: 2

QUESTION NO: 56

What appears in a routing table after EIGRP route summarization is configured on a router's Serial0 interface summarizing routes learned from Ethernet0 interface?

- A. a summary route pointing to the Null0 interface
- B. a summary route pointing to the Serial0 interface
- C. a summary route pointing to the Ethernet0 interface
- D. a summary route pointing to the Loopback0 interface

Answer: A

Explanation:

EIGRP creates a null route that matches the summary-address network and mask entry to prevent routing loops.

Chapter: 2

QUESTION NO: 57

You are the network administrator at TestKing. An EIGRP router on the TestKing network has not established adjacency with a neighbor.

What are the possible causes for this? (Choose all that apply.)

- A. K-values do not match
- B. Hold times do not match
- C. Hello times do not match

D. AS numbers do not match

Answer: A, D

Explanation:

If the autonomous system (AS) numbers do not match, the routers will not form an adjacency. When EIGRP is enabled on an interface, the router begins sending hellos to a multicast address. The hello packet includes the configured EIGRP metric K values. The two routers become adjacent if their K values match.

Chapter: 2

QUESTION NO: 58

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know what the correct command format to configure EIGRP summary route is.

What would your reply be?

- A. ip auto-summary *as-number address mask*
- B. ip summary-address eigrp *as-number address mask*
- C. ip auto-summary eigrp *as-number address mask*
- D. ip summary-route eigrp *as-number address mask*

Answer: B

Explanation:

This is the correct syntax to summarize external routes in EIGRP.

Chapter: 2

QUESTION NO: 59

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know how bandwidth information per neighbor is determined on a multipoint Frame Relay interface.

What would your reply be?

- A. bandwidth command per neighbor
- B. the configured CIR per subinterface
- C. the configured CIR divided by the number of neighbors on that interface
- D. bandwidth of the main interface divided by the number of neighbors on that interface

Answer: B

Explanation:

The bandwidth is assumed to be 1.54Mbps on serial interfaces by EIGRP. The only way to overcome this assumption, is to configure the bandwidth on each subinterface.

Chapter: 2

QUESTION NO: 60

Which command shows the active or passive state of EIGRP routes, the number of successors, and the feasible distance to the destination?

- A. show ip route eigrp
- B. show ip eigrp traffic
- C. show ip eigrp neighbors
- D. show ip eigrp topology

Answer: D

Explanation:

If the show ip route command were issued, only the current route would appear in the routing table. The EIGRP traffic and neighbor command do not show any information about the routes in the topology database.

IP-EIGRP Topology Table for process 77

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
r - Reply status

P 172.16.90.0 255.255.255.0, 2 successors, FD is 0

via 172.16.80.28 (46251776/46226176), Ethernet0

via 172.16.81.28 (46251776/46226176), Ethernet1

via 172.16.80.31 (46277376/46251776), Ethernet0

P 172.16.81.0 255.255.255.0, 1 successors, FD is 307200

via Connected, Ethernet1

via 172.16.81.28 (307200/281600), Ethernet1

via 172.16.80.28 (307200/281600), Ethernet0

via 172.16.80.31 (332800/307200), Ethernet0

Chapter: 2

QUESTION NO: 61

What are two classless routing protocol features supported by EIGRP? (Choose all that apply.)

- A. triggered updates
- B. variable length subnet masks

- C. periodic update announcements
- D. unequal path-cost load balancing

Answer: A, B

Explanation:

Eigrp does not send out periodic updates only triggered.

QUESTION NO: 62

When point-to-point Frame Relay sub-interfaces are used on a router running EIGRP, what is the default line speed used in calculating routes associated with the WAN interfaces?

- A. 256 Kbps
- B. 1.544 Mbps
- C. set by the PVC
- D. set by the DLCI

Answer: B

Explanation:

EIGRP assumes that all serial interfaces are operating at T-1 speed. If the actual line speed is different, the interface line speed should be specified using the bandwidth command.

Chapter: 2

QUESTION NO: 63

You are the network administrator at TestKing. Router TK1 is the central router in a hub and spoke topology supporting 24 remote locations. Point-to-point Frame Relay EIGRP networks are deployed between the central router and the remote locations. There is no bandwidth command configured under either the major serial interface or the subinterfaces on the central router.

What does EIGRP perceive as the bandwidth of each Frame Relay connection?

- A. 64 kbps
- B. 128 kbps
- C. 512 kbps
- D. 1.544 Mbps

Answer: D

Explanation:

EIGRP assumes that all serial interfaces are operating at T-1 speed. If the actual line speed is different, the interface line speed should be specified using the bandwidth command.

Chapter: 2

QUESTION NO: 64

What is a use of VLSM in EIGRP?

- A. disjointed networks
- B. address cumulative
- C. address aggregation
- D. contiguous networks

Answer: A

Explanation:

Variable Length subnet masks allows the router to separate networks according to the amount of Hosts required per network. Disjointed networks would be the result of using EIGRP. EIGRP does not require Contiguous Network Assignment, as it does not require hierarchical Address.

Chapter: 2

QUESTION NO: 65

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know what could possibly cause EIGRP Stuck-In-Active routes.

What would your reply be? (Choose all that apply.)

- A. The neighboring router has a better route than the active route on this router
- B. The neighboring router starts receiving route updates from this router
- C. The neighboring router is having memory problems and cannot allocate the memory to process the query or build the reply packet
- D. Packets are being dropped because EIGRP is sending packets faster than the actual line speed of the circuit.

Answer: C, D

Explanation:

An EIGRP SIA message means that an EIGRP router hasn't received a reply to a query from one or more neighbors within the allotted time. When this happens, EIGRP clears the neighbors that didn't send a reply and logs a DUAL-3-SIA error message for the route that went active.

Chapter: 2

QUESTION NO: 66**Which EIGRP information is added to a routing table?**

- A. successor only
- B. feasible successor only
- C. successor and back up successor
- D. successor and feasible successor

Answer: A**Explanation:**

EIGRP maintains its own topology database where up to six routes for every destination can be stored. Only the best route (successor) is installed into the routing table. A backup is registered with the routing table maintenance process, but not installed into the routing table. In the event the route in the routing table fails, the routing table maintenance process calls each routing protocol process that has registered a backup route, and asks them to reinstall the route in the routing table. Then, the route with the preferred route is chosen based administrative distance.
Chapter: 2

QUESTION NO: 67**Which is the most effective technique to contain EIGRP queries?**

- A. route summarization
- B. configuring route filters
- C. using a hierarchical addressing scheme
- D. establishing separate autonomous systems

Answer: A**Explanation:**

After you determine the minimum routing requirements, you can make EIGRP more scalable. Two of the best options are the following: 1. Configure route summarization using the ip summary-address eigrp command on the outbound interface of the appropriate routers 2. Configure the remote routers as stub EIGRP routers. Summarizing routes limits the queries scope by limiting a routers knowledge of networks subnets. If a subnet goes down, queries go only as far as the routers that have knowledge of that subnet.

Source: CCNP Self-Study Second Edition P.185

Topic: Limiting the EIGRP Query Range

QUESTION NO: 68**What is the default hold time for EIGRP hellos on NBMA media?**

- A. 30 seconds
- B. 60 seconds
- C. 90 seconds
- D. 180 seconds

Answer: D**Explanation:**

EIGRP sends hello packets every 5 seconds on high bandwidth links and every 60 seconds on low bandwidth NBMA media. The default hold time is three times the hello interval or 180 seconds for NBMA 5-second hello:

broadcast media, such as Ethernet, Token Ring, and FDDI point-to-point serial links, such as PPP or HDLC leased circuits, Frame Relay point-to-point subinterfaces, and ATM point-to-point subinterfaces high bandwidth (greater than T1) multipoint circuits, such as ISDN PRI and Frame Relay

Chapter: 2

60-second hello:

multipoint circuits T1 bandwidth or slower, such as Frame Relay multipoint interfaces, ATM multipoint interfaces, ATM switched virtual circuits, and ISDN BRIs

Chapter: 2

QUESTION NO: 69**With regard to iBGP routers, which of the following statements is true?**

- A. They are level-1 routers
- B. They are level-2 routers
- C. They are in the same AS
- D. They are in a different AS

Answer: C**Explanation:**

BGP routers that are all within the same AS are considered to be internal BGP routers (iBGP).

Chapter: 2

QUESTION NO: 70

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know what the characteristics of an autonomous system in a BGP network is.

What would your reply be? (Choose all that apply.)

- A. It used only Interior Gateway Protocols (IGPs)
- B. EGPs are used to connect different autonomous systems
- C. It is a set of routers under a single technical administration
- D. It uses EGPs to route packets to other autonomous systems and IGPs to route packets within the autonomous system
- E. It uses IGPs to route packets to other autonomous systems and EGPs to route packets within the autonomous system

Answer: B, C, D

Explanation:

EGP (Exterior Gateway Protocol) is used to exchange routing information between two different autonomous systems. IGP (Interior Gateway Protocol) is used to exchange routing information within an autonomous system. Routers within the same AS are normally under a single technical administration.

Chapter: 2

QUESTION NO: 71

You are a trainee technician at TestKing. Your instructor shows you the following router configuration:

```
router bgp 65000
neighbor 172.16.1.1 remote-as 65000
neighbor 10.1.1.2 remote-as 64550
network 192.168.1.192 mask 255.255.255.224
```

Your instructor wants to know which of the following statements is true if interfaces 192.168.1.193 172.16.1.1 and 10.1.1.2 are active.

What would your reply be? (Choose all that apply.)

- A. Router A is in autonomous system 64550
- B. Router A advertises network 192.168.1.192/27
- C. Router A forms an iBGP relationship with neighbor 10.1.1.2
- D. Router A forms an eBGP relationship with neighbor 172.16.1.1

Answer: B

Explanation:

Router A is in AS 65000 as indicated on line 1.

Line 2 specifies 172.16.1.1 as a neighbor and belonging to the same AS (iBGP peer).

Line 3 specifies 10.1.1.2 as a neighbor and belonging to a different AS (eBGP peer).

Line 4 specifies that 192.168.1.192/27 be advertised

Chapter: 2

QUESTION NO: 72

In which instance is it appropriate to use BGP?

- A. If there is single connection to the Internet
- B. If you have limited understanding of route filtering
- C. If there is a low-bandwidth connection between autonomous systems
- D. If route selection to routes outside of your autonomous system is not a concern
- E. If an autonomous system allows packets to transit through it to reach other autonomous systems

Answer: E

Explanation:

This is the only reason listed that would require you to run BGP. If you are serving as a transit AS for other downstream AS's, then you must run BGP so that all paths are known into the downstream AS's.

Chapter: 2

QUESTION NO: 73

You are a trainee technician at TestKing. Your instructor shows you the following router configuration on Router TK1:

```
router bgp 65000
network 10.0.0.0
neighbor 172.17.1.1 remote-as 65000
```

Your instructor wants to know what type of relationship is neighbor 172.17.1.1 to Router TK1.

What would your reply be? (Choose all that apply.)

- A. a peer router running iBGP
- B. a peer router running eBGP

- C. a community member running iBGP
- D. a peer group member running iBGP
- E. a peer group member running eBGP

Answer: A

Explanation:

Line 1 specifies that Router A belongs to AS 65000 and the neighbor statement on line 3 indicates that the neighbor is also a member of AS 65000. Therefore, these routers are considered to be iBGP peers.

Chapter: 2

QUESTION NO: 74

Which two statements are true about BGP peering? (Choose two.)

- A. Periodic keepalives are used to verify connectivity
- B. Incremental keepalives are used to verify connectivity
- C. It provides a “best effort” connection between two BGP providers
- D. It provides a reliable connection between two BGP providers

Answer: A, D

Explanation:

BGP uses periodic keepalives to maintain connectivity. The interval can be changed to suit your needs (fast fail-over). Each neighbor sessions runs over TCP (port 179) and ensures reliable delivery of routing information.

Chapter: 2

QUESTION NO: 75

You are a trainee technician at TestKing. Your instructor shows you the following router configuration:

```
interface serial 0
ip address 172.16.1.1 255.255.255.0
!
interface ethernet 0
ip address 10.1.1.1 255.255.255.0
!
router rip
network 10.0.0.0
!
```

```
ip route 0.0.0.0 0.0.0.0 serial0  
!
```

Your instructor wants to know which of the following statements is true.

What would your reply be?

- A. RIP updates are sent and received on interface serial 0 of the router
- B. A default route is sent to neighbors on interface serial 0 of the router
- C. A default route is sent to neighbors on interface ethernet 0 of the router
- D. RIP updates are sent and received on interfaces serial 0 and ethernet 0 of the router

Answer: D

Explanation:

RIP is a true classful routing protocol. One of the problems with classful routing protocols is that periodic routing updates are sent out all active interfaces on every router. Even if the router is not running RIP, it will still receive broadcast RIP packets.

RIPv2 is a classless routing protocol and instead of using broadcast updates it sends multicast packets to a multicast address of 224.0.0.9.

Chapter: 2

QUESTION NO: 76

When using BGP policy-based routing, which two statements are true? (Choose two.)

- A. Policy routing cannot be used to modify the AS-path
- B. Policy routing can be used to alter the final destination of the IP packet
- C. Policy routing allows traffic to be directed based on the source address
- D. Policy routing can influence which router will be used as the next-hop router for a given packet

Answer: C, D

Explanation:

Answers A and B are false. With the use of route maps, BGP AS paths may be lengthened by adding fictitious AS numbers. This technique is called AS path pre-pending. The final destination of a packet cannot be altered by a routing protocol or a routing protocol policy.

Chapter: 2

QUESTION NO: 77

When should BGP synchronization be unnecessary?

- A. when only the edge routers in the AS will be running BGP
- B. when traffic from a different AS passes through an AS to a third AS
- C. when traffic from a different AS will not pass through an AS to a third AS
- D. when sending and receiving of external BGP updates is controlled by using a number of different filtering methods

Answer: C

Explanation:

All of the other items require BGP synchronization. When traffic from a different AS passes through an AS to a third AS, BGP will not advertise the route until all routers within the AS have learned of the of the route through IGP.

Chapter: 2

QUESTION NO: 78

Identify two statements regarding BGP peer groups from the choices below?

(Choose all that apply.)

- A. The peer group name is passed to other routers in the peer group during routing updates.
- B. A peer group is a group of BGP neighbors with different update policies.
- C. The peer group name is only local to the router on which it is configured.
- D. A peer group allows options that affect outbound updates to be overridden.
- E. BGP configurations can be placed on one peer group router and the configuration is applied to all members of the peer group.

Answer: C, E

Explanation:

A BGP peer group is a group of routers that share similar configurations. Every configuration line supplied to a peer group definition is applied to each peer group member. A peer group name may be specified, but it is not passed to any other router; the name is local only to the router it is configured on.

Chapter: 2

QUESTION NO: 79

With regard to BGP attributes, which of the following statements are true? (Choose all that apply.)

- A. Med is an optional attribute
- B. Origin is an optional attribute

- C. Next-hop is an optional attribute
- D. Local Preference is an optional attribute
- E. AS-Path is a well-known mandatory attribute
- F. Community is a well-known mandatory attribute

Answer: A, E

Explanation:

MED is an optional attribute

ORIGIN is a well-known mandatory attribute

NEXT_HOP is a well-known mandatory attribute

LOCAL_PREF is a well-known discretionary attribute

AS_PATH is a well-known mandatory attribute

COMMUNITY is an optional attribute

Chapter: 2

Cisco Press BSCI Third edition, Pages 526-527

QUESTION NO: 80

With regard to the network command on a BGP router, which of the following statements is true?

- A. The local route matching the network command can be learned dynamically
- B. The local route matching the network command are blocked from the BGP routing table
- C. The route to a neighbor autonomous system must have the correct MED applied to be installed into BGP's routing table
- D. The specified network is identified as a transit AS and traffic must pass through this AS to reach its final destination.

Answer: A

Explanation:

The network command only specifies which networks are to be advertised by BGP. This can be learned from static routes, other routing protocols, or directly connected interfaces.

Chapter: 2

QUESTION NO: 81

With regard to BGP community attributes, which of the following statements are true?

- A. Communities are tagged by default in outgoing updates
- B. Communities are local to the autonomous system where specified and can only be used within that autonomous system
- C. Communities are a means of tagging routes to ensure consistent filtering

D. Communities perform summarization of blocks of contiguous network prefixes

Answer: C

Explanation:

The BGP COMMUNITY attribute is used to tag/mark routes. Once these routes are marked, route maps can be used to limit the distribution and acceptance of routes with a particular mark.

Some commonly used communities:

No-Export: The route will not be passed outside the AS

No-Advertise: The route will not be advertised to other routers

No-Export-Subconfed: Routes will not be advertised to eBGP peers (including eBGP peers in the same confederation)

Chapter: 2

QUESTION NO: 82

With regard to BGP policy-based routing, which of the following statements is true?

A. If the next-hop router goes down and no alternative path is in place, policy routing will route to null 0

B. If the next-hop router goes down and no alternative path is in place, policy routing will default to another BGP path

C. If the next-hop router goes down and no alternative path is in place, policy routing will deny all traffic to that destination

D. If the next-hop router goes down and no alternative path is in place, policy routing will default to dynamic routing decisions.

Answer: D

Explanation:

Some things should be considered before arbitrarily deciding to implement policy-based routing. Understand that any additional configurations require additional CPU, particularly when every packet characteristic must be examined. It is also wise to have a backup path in place in case the defined next-hop router goes down. If there is no alternative defined, policy-routing will default to dynamic routing decisions.

Source: Self-Study CCNP BSCI Exam Certification Guide Third Edition P.523

Topic: BGP and Policy-Based Routing

A backup path should be defined in case the defined next-hop router goes down. If there is no alternative defined, policy-based routing uses the IP routing table

Source: Self-Study CCNP BSCI Exam Certification Guide Third Edition P.675

Topic: Disadvantages of Policy-Based Routing

QUESTION NO: 83

You are a trainee technician at TestKing. Your instructor shows you the following router configuration for Router TK1:]

ROUTER TK1

```
router bgp 500
```

```
neighbor 190.225.11.1 remote-as 500
```

```
neighbor 190.225.11.1 update-source loopback 1
```

ROUTER B

```
router bgp 500
```

```
neighbor 150.212.1.1 remote-as 500
```

Your instructor wants to know which of the following statements are true.

What would your reply be? (Choose all that apply.)

- A. ROUTER A and ROUTER B are running iBGP inside as 500
- B. The IP address of ROUTER A's loopback 1 interface is 150.212.1.1
- C. The IP address of ROUTER A's loopback 1 interface is 190.225.11.1
- D. ROUTER A and ROUTER B are running eBGP between autonomous systems
- E. ROUTER A is configured to use the loopback IP address as the source in the BGP neighbor connection with neighbor 190.225.11.1

Answer: A, C, E

Explanation:

Based upon the router bgp 300 commands listed on both routers, we know that these routers belong to the same AS, and therefore are running iBGP sessions. If they were in different AS's, then the session would be considered eBGP.

Line 3 in Router A forces the local BGP session to use the Loopback 1 interface for peering. The neighbor command on Router B specifies the IP address of the remote BGP peer, which must be the Loopback 1 address if these two routers are to establish a BGP session. So, we must assume that the IP address listed is the IP address for the Loopback 1 interface on Router A.

Chapter: 2

QUESTION NO: 84

When the default-information originate always command used?

- A. It is required whenever you want to propagate a default route into a RIP autonomous system
- B. It is required whenever you want to propagate a default route into an IGRP autonomous system
- C. It is required whenever you want to propagate a default route into an OSPF autonomous system

D. It is required whenever you want to propagate a default route into an EIGRP autonomous system

Answer: C

Explanation:

This command is used to direct all other OSPF routers to place a default route into its routing table. This command is not used with IGRP/EIGRP. A similar command is used in RIP to specify some interfaces as the default route even when the local router does not have a default route in its routing table.

Chapter: 2

QUESTION NO: 85

Which routing protocol uses the Diffusing Update Algorithm (DUAL) for route calculation?

- A. BGP
- B. OSPF
- C. EIGRP
- D. RIPv2

Answer: C

Explanation:

DUAL is used by EIGRP. This allows EIGRP to achieve fast, loop-free convergence with little impact on CPU cost and overhead. DUAL takes corrective action when topology changes occur and in doing so, only involves the routers that are affected.

Chapter: 2

QUESTION NO: 86

Which fields is included in a RIP version 2 routing update packet that is not included in RIP version 1?

- A. metric
- B. next hop
- C. subnet mask
- D. autonomous system number

Answer: C

Explanation:

RIPv1 only carries the next hop address and its associated metric in its routing updates. One of the major improvements in RIPv2 is that it is a classless routing protocol. This means it carries the subnet mask along with the network number in its routing updates.

Chapter: 2

QUESTION NO: 87

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know what a router must determine in order to route data.

What would your reply be?

- A. the route age of the next-hop device
- B. the subnet mask of the source network
- C. the cost metric of the path to the destination
- D. the outbound interface of the best path to the destination

Answer: D

Explanation:

The router must know which interface that the data will be forwarded to. The other items listed are not used to determine if the router can route the data.

Chapter: 2

QUESTION NO: 88

Which routing protocol does not have a hop count limit?

- A. IGRP
- B. OSPF
- C. EIGRP
- D. RIPv1
- E. RIPv2

Answer: B

Explanation:

RIP has a maximum allowable hop count of 15 meaning a value of 16 is considered unreachable. IGRP and EIGRP have a limit of 255 (100 by default).

Chapter: 2

QUESTION NO: 89

You are the network administrator at TestKing. TestKing uses the RIPv2 routing protocol. Due to a dropped connection, your core router has not received an update from a neighboring router for four minutes and the route is marked invalid.

What is the next step for the core router?

- A. It recalculates the network topology
- B. It purges that link from its routing table
- C. It places a hold-down on the routes from that link
- D. An invalid timer is started

Answer: D

Explanation:

RIP versions 1 and 2 use timed updates to respond to changes in topology. Updates are sent every 30 seconds by all participating routers. If a neighbor does not respond for 180 seconds, the router assumes the link or neighbor is down. The router then marks the route as invalid and an invalid timer is started (180seconds). When the invalid timer expires, a hold-down time begins. During this period, the route is marked as possibly down and the metric is set and advertised with infinity. If a new route was received during the invalid or hold-down periods, the router begins advertising the new route. If not, a request is sent to neighbors for an alternate route to the destination. The last timer is the flush timer. It is activated when the invalid timer expires and last 60 seconds after the hold-down timer expires. When the flush timer expires (240 seconds), the route is finally deleted if no alternate route was found.

Chapter: 2

QUESTION NO: 90

Which routing protocol provides support for load balancing over unequal cost paths?

- A. IGRP
- B. OSPF
- C. RIPv1
- D. RIPv2

Answer: A

Explanation:

IGRP and EIGRP provide support for load balancing via unequal cost paths using the variance command. In order for RIP to perform balancing, the hop count must be the same as RIP looks at hop count only when determining the best path to a destination. IGRP and EIGRP use a composite metric to determine the cost to a remote network. The variance command controls the load balancing between the best path and the worst acceptable path.

Chapter: 2

QUESTION NO: 91

An LSA is used by which routing protocol to send topology change information across the network?

- A. RIP
- B. IGRP
- C. OSPF
- D. EIGRP

Answer: C

Explanation:

The OSPF routing protocol uses different types of LSA to send information to other OSPF routers. An OSPF router's link state database is comprised of LSA's Listed below are the five types of LSA's used by the OSPF protocol:

LSA Type 1 Router link state

LSA Type 2 Network link state

LSA Type 3 Summary link state (type 3)

LSA Type 4 Summary link state (type 4)

LSA Type 5 External link state

Chapter: 2

QUESTION NO: 92

Which of the following is a characteristic of link-state protocols?

- A. A network topology change generates an update that is broadcast to every router on the network.
- B. The periodic routing update interval is slightly different in each router on a subnet.
- C. Each router creates a routing table that only includes its directly connected neighbors.
- D. Link-state protocols send routing updates only when there is a change in the network topology.

Answer: D

Explanation:

One of the differences between distance-vector and link-state routing protocols is the way topology changes are conveyed across the network. Distance-vector protocols use periodic updates and send out complete routing tables. If a link goes down somewhere in the network,

routers running only distance-vector protocols will not know about the change until the next update.

Link-state routing protocols send updates only when there is a topology change and the change is also included within the update. In addition, only the change is sent in the update. This means that each router does not have to rebuild its entire routing table with each update. The update is either added or removed from the routing table.

Chapter: 2

QUESTION NO: 93

What is the term associated with synchronizing the routing tables after a topology change occurs?

- A. flooding
- B. broadcasting
- C. convergence
- D. summarization

Answer: C

Explanation:

Convergence is the time required for all routers on the network to agree on the network topology after a change has occurred. Simply stated, all routers synchronize the routing tables with the same information.

Chapter: 2

QUESTION NO: 94

Which of the following must be determined first by the router in order to route the data?

- A. the distance metric of the data
- B. the source address of the data
- C. which routing protocol is used by the data
- D. whether the protocol suite of the data is active

Answer: C

Explanation:

Before a routing decision can be made three major decisions must be made in the following order:

- 1) Is the logical destination address a known protocol and is the protocol active/enabled on the router?
- 2) Is the destination address in the routing table?

3) If the destination address is in the routing table, which interface will be used to forward the data?

Chapter: 2

QUESTION NO: 95

You are a technician at TestKing. TestKing has the subnet/mask of 172.29.100.0/26 set aside for small remote locations. Each location will have 5 IP devices connected to the network.

Which VLSM mask will provide the minimum number of hosts you require?

- A. /27
- B. /28
- C. /29
- D. /30
- E. /31

Answer: C

Explanation:

/29 mask is equal to 255.255.255.248 in decimal format. If 29 bits are used for the network portion of the address, then 3 bits are left for the host portion of the address. Once the reserved network and the broadcast addresses are removed, there are six usable addresses for hosts on the subnet. $2 \times 2 \times 2 = 8$ minus network and broadcast address = 6 valid host addresses (Requirement from question is 5).

Chapter: 2

QUESTION NO: 96

You are a trainee technician at TestKing. Your instructor shows you the following route summarization entry:

192.168.134.0/22

Your instructor wants to know how many class C addresses are contained in this summarization.

What would your reply be?

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

- D. 8
- E. 16

Answer: C

Explanation:

The valid class C addresses are:

- 192.168.134.0
- 192.168.135.0
- 192.168.136.0
- 192.168.137.0

Chapter: 2

QUESTION NO: 97

Which two UDP ports are NOT enabled automatically when the ip helper-address command is used on a router? (Choose all that apply.)

- A. 53 (DNS)
- B. 69 (TFTP)
- C. 515 (LPR)
- D. 161 (SNMP)
- E. 49 (TACACS)

Answer: C, D

Explanation:

By default, if an IP helper address is specified, UDP forwarding is enabled on certain ports.

Broadcast packets destined to the following port numbers are forwarded by default:

Trivial File Transfer Protocol (TFTP) (port 69)

Domain Naming System (port 53)

Time service (port 37)

NetBIOS Name Server (port 137)

NetBIOS Datagram Server (port 138)

Boot Protocol (BOOTP) client and server datagrams (ports 67 and 68)

TACACS service (port 49)

IEN-116 Name Service (port 42)

Chapter: 2

QUESTION NO: 98

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know what the benefit of hierarchical IP addressing is.

What would your reply be?

- A. support for network address translation
- B. efficient address allocation
- C. translation of private addresses
- D. support for link-state routing protocols

Answer: B

Explanation:

The only listed benefit of hierarchical IP address is efficient address allocation. Another major benefit in large enterprise networks is the reduced number of entries in the routing table. Combined with link-state routing protocols, overall CPU and memory requirements are reduced on deployed routers.

Chapter: 2

QUESTION NO: 99

What entry would be used to summarize all the address space between 172.18.129.0 and 172.18.133.255?

- A. 172.18.128.0/21
- B. 172.18.128.0/22
- C. 172.18.130.0/22
- D. 172.18.132.0/20

Answer: A

Explanation:

172.18.128.0/21 is the only network/mask combination that will cover the entire listed range.

Using this network/mask, the valid subnets would be (partial listing):

172.18.0.0 172.18.104.0 172.18.208.0
 172.18.8.0 172.18.112.0 172.18.216.0
 172.18.16.0 172.18.120.0 172.18.224.0
 172.18.24.0 172.18.128.0 172.18.232.0
 172.18.32.0 172.18.136.0 172.18.240.0
 172.18.40.0 172.18.144.0 172.18.248.0
 172.18.48.0 172.18.152.0
 172.18.56.0 172.18.160.0
 172.18.64.0 172.18.168.0
 172.18.72.0 172.18.176.0
 172.18.80.0 172.18.184.0
 172.18.88.0 172.18.192.0

172.18.96.0 172.18.200.0

Chapter: 2

QUESTION NO: 100

With regard to route filtering, which of the following statements are true?

- A. It uses the same technique as packet filtering; it's based on standard access-lists
- B. It is required when using EIGRP and OSPF at the same time
- C. Unlike packet filtering (which uses inverted subnet masks, route filtering uses normal subnet masks
- D. Route filtering is used only while performing route redistribution

Answer: A

Explanation:

Routes can be filtered so that only a subset of routes can be advertised or received. Two configuration commands are necessary to filter a route. First, a standard access list is created that specifies which networks are allowed/denied. Second, the distribute-list command is configured for the routing protocol.

```
!
access-list 2 permit 192.168.1.0 0.0.0.255
access-list 2 permit 192.168.2.0 0.0.0.255
```

```
!
router eigrp 100
distribute-list 2 in
```

```
!
Chapter: 2
```

QUESTION NO: 101

Which of these are benefits of incorporating hierarchical addressing?

- A. You can summarize multiple routes into a single route summaries, making the network more scaleable
- B. A contiguous address assignment allows the most efficient use of address
- C. Reduction in the number of routing table entries for participating routers
- D. A more efficient allocation of addresses network-wide

Answer: B, C, D

Explanation:

Incorporating hierarchical and contiguous addressing minimizes the number of wasted addresses, reduces the number of entries in routing tables (summarizing), and simplifies network design and administration.

Chapter: 3

QUESTION NO: 102

How can a single IP statement indicate many IP addresses?

- A. Source-route bridge
- B. Route summarization
- C. Helper address
- D. Default gateway
- E. Passive interface

Answer: B

Explanation:

A single IP network entry can represent multiple subnets using route summarization and VLSM. Classless routing protocols carry/distribute the mask information when managing routing tables and sending updates to topology changes.

Chapter: 3

QUESTION NO: 103

What is NOT a reason reasons makes OSPF a better and/or more efficient routing protocol over RIP?

- A. OSPF can select paths based on bandwidth
- B. Link-State protocols generally have faster convergence than Distance-Vector
- C. OSPF sends its topology database to configured neighbors on a regular basis to ensure that all routers have the same view of the network
- D. OSPF supports VLSM
- E. OSPF has no hop count limitation because it is based on path cost

Answer: C

Explanation:

RIP periodically (60 seconds) broadcasts its entire routing table to all routers in the network. RIP also employs a series of timers that must expire before alternate routes can be discovered or routes removed in the event of a link failure. These are considered drawbacks to distance-vector protocols as it increases the time for the network to convergence.

Chapter: 3

QUESTION NO: 104

With regard to EIGRP, which of the following are true?

- A. It only sends change-updates to the systems that are affected by the event
- B. It can provide client address conflict resolution
- C. It is considerably more complex than OSPF
- D. It supports Variable Length Subnet Masking
- E. It supports compression over the WAN link

Answer: A, D

Explanation:

EIGRP is considered an advanced distance-vector routing protocol. EIGRP uses the DUAL (Distributed Update Algorithm) to ensure that updates and queries are not propagated beyond affected routers. EIGRP is also a classless routing protocol meaning that the network mask information is retained and advertised by the protocol. By default, EIGRP advertises the natural classful network boundary for all EIGRP internal routes. This behavior can be modified using the no auto-summary command and EIGRP will leave the network mask unchanged.

!

```
router eigrp 93
no auto-summary
```

!

Chapter: 3

QUESTION NO: 105

Which of the statements below correctly describes the BGP regular expression ^200_[0-9]*\$?

- A. Matches AS path 200 only
- B. Matches all AS's
- C. Matches AS path that is originated in AS 200 regardless of location
- D. Matches AS path that is originated in AS 200 and AS's that are directly connected to AS 200

Answer: D

Explanation:

You can use regular expressions in the ip as-path access-list command with BGP. This allows you to setup filters to allow incoming and outgoing updates based on the BGP autonomous system paths. In the example above, the ^ starts the input string and designates "AS". The _

signifies a null string or space follows 200 and [0-9]*\$ indicates that any connected AS with a valid AS string can pass the filter. The \$ matches the end of the input string.

Chapter: 3

QUESTION NO: 106

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know what a BGP peer is.

What would your reply be?

- A. A BGP neighbor that the router has formed a TCP connection with
- B. Another BGP router that the router is currently speaking with
- C. Another router on the network running BGP
- D. A BGP neighbor with a broken TCP connection

Answer: A, B

Explanation:

If BGP cannot establish the TCP connection with its specified neighbor, the routers are not considered to be peering and thus not peers. A peering relationship only exists between routers that explicitly identify other routers as neighbors. BGP does not “auto-discover” neighbors as other routing protocols such as EIGRP.

Chapter: 3

QUESTION NO: 107

With regard to BGP, which of the following statements are true?

- A. iBGP is used between the AS
- B. iBGP is used within the AS
- C. eBGP is used between the AS
- D. eBGP is used within the AS

Answer: B, C

Explanation:

Internal BGP (iBGP) sessions are contained within the autonomous system. The active session is between routers with the same AS number.

External BGP (eBGP) sessions are connections between two different autonomous systems. The active session is between routers with different AS numbers.

Chapter: 3

QUESTION NO: 108

OSPF stub or totally stubby Areas can only exist if adjacent to which of the following?

- A. LSA
- B. LSP
- C. BDR
- D. ABR
- E. DR
- F. LSU

Answer: D

Explanation:

There are only two possible answers - BDR and ABR. The function of the BDR is not directly related to stub or stubby areas. The BDR is a hot standby for the DR and would assume the role of the DR in the event the DR failed. The purpose of a DR is to minimize the number of adjacencies formed and disseminate/receive routing updates throughout the network. An ABR has multiple area assignments and two or more interfaces in two or more areas. An ABR would be used to connect a stub or totally stubby area to the rest of the network.

Chapter: 3

QUESTION NO: 109

You are the network administrator at TestKing. You stop sending RIP broadcasts from all router interfaces that do not require it. You notice that RIP is broadcasting on Ethernet0, Serial0, and Serial1 on the router in question. You determine that Serial0 does not need to participate in RIP updates.

What command would you use to disable RIP on this interface?

- A. Configure RIP using the 'passive interface' command
- B. Configure RIP using the 'no router rip' command
- C. Configure Serial0 using the 'shutdown rip' command
- D. Configure the serial interface using the passive interface command

Answer: A

Explanation:

The passive interface would be defined in the RIP configuration as follows:

```
!  
router rip  
network 192.168.1.0
```


passive interface serial 0
!
Chapter: 3

QUESTION NO: 110**Why will an EIGRP route get stuck in the Active state?**

- A. The EIGRP router on the far end of the link has crashed
- B. The reply to a query caused by a change in a route takes too long and times out
- C. The physical link between the two AS's has gone down
- D. The network is unstable and probably has a routing loop

Answer: B**Explanation:**

EIGRP sends out queries to neighbors when it detects a change in topology. In some situations, the response to the query could be delayed longer than the router issuing the query is willing to wait. When this happens, the route in question is stuck in the active state. This means that EIGRP is manipulating the route and cannot route data to the remote network. Eventually, the router sending the query will give up and clear its connection with the neighbor. This can be caused by slow links, bad connections, dropped packets, or an overloaded router.

Chapter: 3

QUESTION NO: 111**You are a technician at TestKing. Your newly appointed TestKing trainee wants to know what the most effective method is to reduce large routing tables.****What would your reply be?**

- A. Route filters
- B. CIDR summary blocks
- C. Compression
- D. Switching

Answer: B**Explanation:**

The only valid choice is to use CIDR (Classless Interdomain Routing) summaries whenever possible. With CIDR, masks in address space are grouped together to form one update. The lack of midsize address space (between a Class C and Class B), the growth of the Internet routing tables, and the eventual exhaustion of the 32-bit IP address space all fueled the creation of CIDR.

Chapter: 3

QUESTION NO: 112

You are a technician at TestKing. You need to create filter for EIGRP network 198.30.64.0/24.

What mask should be used to create the filter?

- A. 0.0.0.255
- B. 0.0.255.255
- C. 255.255.1.0
- D. 255.255.254.0

Answer: A

Explanation:

A distribution list is used to block (filter) routing updates or queries and relies on an access-list to match the criteria. The configuration commands are listed below to allow only route information for the 198.30.60.0/24 network:

```
!  
router eigrp 1  
network 198.30.64.0  
distribute-list 2 in  
!  
access-list 2 permit 198.30.64.0 0.0.0.255
```

Chapter: 3

QUESTION NO: 113

Which of the following best describes a Type 3 or 4 LSA?

- A. Summary link entry which shows routing tables as IA for OSPF inter-area
- B. Autonomous System External type-1 which shows in routing tables as 'E1' for OSPF External-1
- C. Network link entry that shows in routing tables as '0' for OSPF
- D. Router link entry/record which show in routing tables as '0' for OSPF

Answer: A

Explanation:

Type 3 and 4 LSA's are generated by ABR's and sent to all routers within an area. These type LSA's advertise intra-area routes to the backbone area and both intra-area and inter-area routes to non-backbone areas.

Chapter: 3

QUESTION NO: 114

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know how long OSPF will wait after receiving no updates before it sends a new LSA.

What would your reply be?

- A. 10 minutes
- B. 20 minutes
- C. 30 minutes
- D. 1 hour

Answer: C

Explanation:

The fixed OSPF constant "Link State Refresh" is defined as the maximum amount of time between distinct origination of the same LSA. When the LSA age reaches this interval, the router must originate a new instance of the same LSA, keeping everything the same. The value of this constant is 30 minutes.

Chapter: 3

QUESTION NO: 115

By default, how does BGP handle IGP routes?

- A. BGP automatically redistributes all IGP routes
- B. BGP advertisements are independent of IGP route propagation
- C. BGP must wait until the IGP has propagated routing information across the autonomous system
- D. BGP can immediately advertise routes without waiting until the IGP has propagated information across the autonomous system

Answer: C

Explanation:

BGP synchronization rules dictate when traffic from a different AS passes through an AS to a third AS, BGP will not advertise the route until all routers within the AS have learned of the route through IGP.

Chapter: 3

QUESTION NO: 116

What methods does EIGRP use to determine if a previous neighbor is dead?

- A. Unicast
- B. Hold Time
- C. Multicast
- D. Broadcast

Answer: B, C

Explanation:

EIGRP multicasts hello packets to all neighbors at regular intervals (5 or 60 seconds). Every neighbor must send these hello packets. By default, if three consecutive hello packets are not received, the hold time expires and the neighbor is declared dead.

Chapter: 3

QUESTION NO: 117

Select the valid subnet number and broadcast address combinations for mask /26:

- A. Subnet 10.0.0.128, broadcast 10.0.0.192
- B. Subnet 10.0.0.64, broadcast 10.0.0.127
- C. Subnet 10.0.0.128, broadcast 10.0.0.191
- D. Subnet 10.0.0.0, broadcast 10.255.255.255

Answer: B, C

Explanation:

If 26 bits are used for the network portion of the address, then 6 bits are used in the mask portion. $2 \times 2 \times 2 \times 2 \times 2 \times 2 = 64$ addresses per subnet. The broadcast address is always the all 1's address or the last address within the subnet.

Subnet 10.0.0.0 broadcast 10.0.0.63

Subnet 10.0.0.64 broadcast 10.0.0.127

Subnet 10.0.0.128 broadcast 10.0.0.191

Subnet 10.0.0.192 broadcast 10.0.0.255

Chapter: 3

QUESTION NO: 118

Leading the way in IT testing and certification tools, www.testking.com

Policy-Based routing has which of these criteria?

- A. Can Deny based on subnet mask
- B. Can Deny based on RIP version
- C. Can Match based on the IP address
- D. Can Match based on Next Hop address

Answer: C

Explanation:

Policy-based routing relies on standard or extended access list to define the traffic matching criteria. Using standard or extended access lists, we cannot match subnet mask, RIP version, or next hop address. IP addresses can easily be matched using access lists.

Chapter: 3

QUESTION NO: 119

Which of the following BGP Path Attributes (metrics) is a proprietary enhancement?

- A. Community
- B. Origin
- C. Aggregator
- D. Weight

Answer: D

Explanation:

The weight attribute is a Cisco proprietary attribute used for path selection. This allows as administrator to “prefer” one path over one or more paths to the same destination.

Chapter: 3

QUESTION NO: 120

Which of the following are needed for an OSPF to form an adjacency?

- A. Autonomous System Number
- B. Area ID
- C. Hello and Dead intervals
- D. Stub Area Flag

Answer: C

Explanation:

In order to form an adjacency hello packets are exchanged. All routers connected to a common network must agree on all the parameters from the hello packet listed below:

Network Mask

Hello Interval

Dead Interval

Neighbor

Chapter: 3

QUESTION NO: 121

Which of the following correctly describes an OSPF Area Border Router?

- A. It is used to connect various OSPF Areas
- B. It is used only at the outside edge of an OSPF network for protocol translation
- C. It is used only internally for OSPF to EIGRP redistribution
- D. It is used for routing protocol redistribution

Answer: A

Explanation:

An ABR is defined as having multiple area assignments and two or more interfaces in two or more areas. An ABR would be used to connect a different, stub or totally stubby area to the rest of the network.

Chapter: 3

QUESTION NO: 122

You are a technician at TestKing. You are troubleshooting a problem on the TestKing network. You issue the debug ip ospf adj command. Your newly appointed TestKing trainee wants to know what the purpose of this command is.

What would your reply be?

- A. LSA type 1
- B. Elections
- C. LSA type 3 and 4
- D. OSPF neighbor relationships

Answer: D

Explanation:

The output from this command can be very useful in troubleshooting neighbor relationships. A sample output from the command is listed below:

```
Router#debug ip ospf adj
OSPF: Receive dbd from 70.70.70.70 seq 0x14B
OSPF: 2 Way Communication to neighbor 70.70.70.70
OSPF: send DBD packet to 192.16.64.2 seq 0x1797
OSPF: Receive dbd from 70.70.70.70 seq 0x1797
OSPF: NBR Negotiation Done We are the MASTER
OSPF: send DBD packet to 192.16.64.2 seq 0x1798
OSPF: Database request to 70.70.70.70
OSPF: sent LS REQ packet to 192.16.64.2, length 12
OSPF: Receive dbd from 70.70.70.70 seq 0x1798
OSPF: send DBD packet to 192.16.64.2 seq 0x1799
OSPF: Receive dbd from 70.70.70.70 seq 0x1799
OSPF: Exchange Done with neighbor 70.70.70.70
OSPF: Synchronized with neighbor 70.70.70.70, state:FULL
OSPF: Build router LSA, router ID 172.16.13.1
Chapter: 3]
```

QUESTION NO: 123

What are BGP routers with the same AS number exchange information known as?

- A. BGP
- B. eBGP
- C. IGRP
- D. iBGP

Answer: D

Explanation:

Internal BGP (iBGP) sessions are contained within the autonomous system. The active session is between routers with the same AS number.

External BGP (eBGP) sessions are connections between two different autonomous systems. The active session is between routers with different AS numbers.

Chapter: 3

QUESTION NO: 124

BGP falls into which routing protocol category below?

- A. Link-state protocol
- B. Interior gateway protocol
- C. Exterior gateway protocol
- D. Distance-Path state protocol

Answer: C

Explanation:

BGP is an EGP (Exterior Gateway Protocol used to connect and find routes to/from autonomous systems.

Chapter: 3

QUESTION NO: 125

Which of the following is untrue of routers running internal BGP?

- A. Routers are running iBGP when two BGP speakers are in different AS's
- B. They must form a TCP port 179 connection with each other
- C. They do not have to be directly connected
- D. Routers are running iBGP when two BGP speaking routers are in the same AS

Answer: A

Explanation:

Internal BGP (iBGP) sessions are contained within the autonomous system. The active session is between routers with the same AS number.

External BGP (eBGP) sessions are connections between two different autonomous systems. The active session is between routers with different AS numbers.

Chapter: 3

QUESTION NO: 126

Which of the following are true of a network design with a high number of OSPF routers in a single area?

- A. Increased problems with routing table 'black holes'
- B. Increased SPF algorithm re-calculation times
- C. Decreased OSPF update traffic overhead
- D. Increased OSPF update traffic overhead

Answer: B, D

Explanation:

SPF path re-calculation times are increased as more routers (and paths) are added to the network. In addition, as more routers and paths are added additional LSA's must be generated for each network and summary that exists within the area.

Chapter: 3

QUESTION NO: 127

Which of these values is required by EIGRP as a Seed Metric when performing Route Redistribution?

- A. Bandwidth, Delay, Reliability, Load, MTU
- B. Cost, Delay, Reliability
- C. Route-Tag
- D. Cost

Answer: A

Explanation:

When redistributing routes into IGRP or EIGRP, the bandwidth, delay, reliability, load and MTU (maximum transmission unit) must be specified.

!

```
router eigrp 91
network 10.10.108.0
redistribute static
redistribute rip
default-metric 10000 100 255 1 1500
Chapter: 3
```

QUESTION NO: 128

With regard to an OSPF internal router, which of the following statements are true?

- A. Internal routers are in more than one area
- B. At least one interface is in Area 0
- C. All interfaces are in the same Area
- D. Internal routers redistribute routing information with routers in other areas

Answer: C

Explanation:

An internal OSPF router is defined by having all its interfaces members of the same area.

Chapter: 3

QUESTION NO: 129

What is the inverted wildcard mask of /26?

- A. 63.255.255.255
- B. 0.0.0.63
- C. 192.255.255.255
- D. 0.0.0.192

Answer: B

Explanation:

The regular dotted decimal format of a /26 prefix is represented as: 255.255.255.192 or 11111111.11111111.11111111.11000000.

To get the wildcard mask, use $256-192=64$, which is the block size of the subnet mask. The wildcard mask is always one less than the block size or 63 in this case. A wildcard of zero (0) means all network bits are on. So the answer is 0.0.0.63.

Chapter: 3

QUESTION NO: 130

You are the network administrator at TestKing. You attempt to start an OSPF process on Router TK1 but you receive an 'Unable to allocate router ID' error.

What could be the cause of this problem?

- A. The OSPF area is Stubby
- B. No IP address assigned to any interface
- C. Router ID is used by another router
- D. All interfaces are shut down

Answer: B, D

Explanation:

When OSPF is enabled on a router, the loopback interface IP address is used to assign the Router ID by default. If the loopback interface is not configured, then the highest IP address of all interfaces is assigned as the Router ID. At least one interface must be active/enabled and configured with an IP address to enable OSPF.

Chapter: 3

QUESTION NO: 131

Which types of LSA's are passed into a Totally Stubby Area?

- A. Summary Type 3
- B. Summary Type 4

- C. Router Link Advertisement Type 1
- D. External Link Advertisement Type 5

Answer: C

Explanation:

Only possible answer might be C. Type 3, 4 and 5 are all blocked from going into a Totally stubby area. So only intra-area and default.

Source: BSCI Self Study Cisco Press (2003) 2nd Edition, P190

QUESTION NO: 132

With regard to Route Reflectors, which of the following are true?

- A. They increase the size of the BGP routing table
- B. They enable the use of Route Summarization
- C. They reduce the number of TCP sessions
- D. A route reflector cannot propagate iBGP routes to iBGP peers.

Answer: C

Explanation:

One of the major reasons route reflectors are deployed is to relieve the iBGP full-mesh requirement. The BGP split-horizon rule is modified by route reflectors. It is modified by allowing a route reflector to be the only router that propagates routes learned by iBGP to other iBGP peers.

Chapter: 3

QUESTION NO: 133

Among these methods of sending an IGP route into BGP, which is not recommended?

- A. Redistribute the IGP into BGP
- B. Use the network command with the 'mask' option
- C. Redistribute BGP into the IGP
- D. Use the network command without the 'mask' option

Answer: A

Explanation:

Redistribution can cause routing loops and route flapping. In addition, BGP actually pulls information that other IGP's have learned about their environment. BGP handles the translation of one protocol to another if multiple routing protocols are used in the AS.

Chapter: 3

QUESTION NO: 134

Which of the following are reasons a network administrator would choose to run multiple routing protocols at the same time on the same network?

- A. Reduce routing protocol overhead on the network
- B. Optimize the route redistribution between areas
- C. Convert from an older routing protocol to a more modern one
- D. Acquisition of a company running a different routing protocol than your own

Answer: C

Explanation:

In general, multiple routing protocols are not needed simultaneously on the same network. However, when migrating to a new routing protocol, this is a good methodology as migration configurations can be tested without affecting existing traffic.

Chapter: 3

QUESTION NO: 135

You are a technician at TestKing. TestKing has a BGP network. Your newly appointed TestKing trainee wants to know when BGP is not suitable for use.

What would your reply be?

- A. Routing policy and route selection are not important for your AS
- B. A single connection to the Internet
- C. When you need to perform load-sharing to multiple ISP's simultaneously
- D. Low bandwidth between your AS and the other AS

Answer: A, C, D

Explanation:

"B" "A single connection to the Internet" is suitable for BGP. According to this Question, "ACD" is correct, because these are unsuitable responses to this Question.

QUESTION NO: 136

Which command can be used to view the state of the link, such as exstart, exchange, or full?

- A. show ospf neighbor
- B. show ip protocols
- C. show ip ospf neighbor
- D. show ospf interface

Answer: C

Explanation:

Answer D lists the state of the DR on the interface. Answer C shows all information about all OSPF neighbor routers including the state.

```
router2#show ip ospf neighbor
```

```
Neighbor ID Pri State Dead Time Address Interface
170.170.5.1 1 INIT/- 00:00:34 170.170.1.1 Serial0
```

```
router-2#
```

Chapter: 3

QUESTION NO: 137

Which routing protocol features provide solutions to the problems associated with very large routing tables?

- A. Filtering
- B. Data Compression
- C. Route summarization
- D. Incremental updates

Answer: C, D

Explanation:

Route summarization reduces the number of entries in the routing tables. Incremental updates only send information about the network topology that has changed. Both of these features free up resources on routers, and bandwidth on the network.

Chapter: 3

QUESTION NO: 138

What is the router command to diagnose and display the entire route, including delays, given the path to a destination?

- A. routepath
- B. routetrace
- C. pathtrace
- D. traceroute

Answer: D

Explanation:

The traceroute command provides each hop and delay information about the path to a destination.

Chapter: 3

QUESTION NO: 139

You are the network administrator at TestKing. You configure a static route on router running EIGRP. Your newly appointed TestKing trainee wants to know what the default administrative distance of this static route would be.

What would your reply be?

- A. 255
- B. 5
- C. 1
- D. 90

Answer: C

Explanation:

Static routes are assigned a default administrative distance of 1.

Chapter: 3

QUESTION NO: 140

With regard to an OSPF backbone router, which of the following statements is true?

- A. It redistributes routing information with routers in other areas
- B. All backbone routers must be in more than one area
- C. At least one interface is in Area 0
- D. All interfaces are in the same Area

Answer: C

Explanation:

A backbone router has at least one interface in the backbone area.

Chapter: 3

QUESTION NO: 141

Which of these packet types does an OSPF router use at bootup to receive information?

- A. SPF
- B. LCP
- C. Flood
- D. Hello

Answer: D

Explanation:

OSPF uses hello packets to build adjacencies when OSPF is enabled on a router or when a router boots up with OSPF already configured.

Chapter: 3

QUESTION NO: 142

You are a technician at TestKing. You reboot the existing DR. This results in the BDR being promoted to DR.

With regard to this scenario, which of the following statements are true?

- A. Upon boot of the old DR, the newly elected DR will automatically demote itself back to BDR
- B. When the old DR comes back online, it will send out an LSA to override the new DR
- C. The newly elected OSPF DR will remain DR even when the old DR comes back on-line
- D. Upon boot of the old DR another election will occur

Answer: C

Explanation:

Once a DR and BDR have been elected, the DR/BDR is sent to all routers via hello packets. In this case, the existing DR is rebooted and the BDR is promoted to DR and a new BDR is elected. When the previous DR comes back on-line, it accepts the newly promoted/elected DR/BDR that is received in the hello packets.

Chapter: 3

QUESTION NO: 143

Which of the following supports multiple routed protocols?

- A. IP
- B. OSPF

- C. BGP
- D. EIGRP
- E. IGRP

Answer: D

Explanation:

One of the enhancements to EIGRP from IGRP is its support for all major Layer 3 routed protocols. EIGRP supports IP, IPX, and AppleTalk. All of the other routing protocols listed provide support for routing IP only.

Chapter: 3

QUESTION NO: 144

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know what the BGP MED attribute is used for.

What would your reply be?

- A. Setting the route maps peer ID
- B. Setting IGP synchronization
- C. Setting the route reflectors peer ID
- D. Setting a preferred return-pathway back into the originating AS

Answer: D

Explanation:

The MED attribute in BGP is used to indicate the best entry point or path to reach a particular destination to a neighboring AS. A lower MED is preferred over a higher MED. The MED attribute is also non-transitive because the AS does not pass the MEDs it learns from one AS to another.

Chapter: 3

QUESTION NO: 145

With regard to RIPv1, which of the following statements are true? (Choose all that apply.)

- A. Maximum hop count is 16
- B. Update interval is 90 seconds
- C. Update interval is 30 seconds
- D. VLSM is not supported
- E. Maximum hop count is 15

Answer: C, D, E

Explanation:

Answer A is tricky as the maximum allowable hop count is 15. If the hop count is set to 16 (infinity), the destination is considered unreachable. The RIP periodic update interval is 30 seconds.

Chapter: 3

QUESTION NO: 146

How does the command 'ip bandwidth-percent eigrp' set EIGRP maximum bandwidth allocation?

- A. On a per AS basis
- B. On a per packet basis
- C. On a per interface basis
- D. Globally for all IGRP and EIGRP AS's on the router

Answer: C

Explanation:

This command is applied to an interface to set the maximum amount of bandwidth to be used on an interface for EIGRP traffic.

Chapter: 3

QUESTION NO: 147

Which of the following are possible reasons that EIGRP might not make an adjacency to a neighbor router?

- A. Both routers are not running DUAL
- B. Hop counts do not match on both routers
- C. The feasible successor is still in the Hello state
- D. Update packets have not been sent

Answer: D

Explanation:

When EIGRP begins making adjacencies, it multicasts hello packets out an interface(s). The neighbors become adjacent when they acknowledge each others hello packets and their K values match.

Chapter: 3

QUESTION NO: 148

With regard to OSPF, which of the following statements are correct? (Choose all that apply)

- A. OSPF computes cost based on the interface's bandwidth setting
- B. OSPF is not limited to a hop count metric
- C. OSPF uses LSA packets instead of broadcasts
- D. OSPF sends the complete routing table inside of each LSA packet

Answer: A, B, C

Explanation:

OSPF does not send its complete routing table inside each LSA. An LSA is sent for each network that is included in the routing process.

Chapter: 3

QUESTION NO: 149

All BGP routers belong to which of these communities by default?

- A. no-export
- B. classless
- C. Internet
- D. no-advertised

Answer: C

Explanation:

All routers by default belong to the "Internet" community and can be used to advertise routes to all other routers. No-export tags the routes so that it will not be sent outside the AS. The No-advertise tag will prevent the route from being advertised to any other BGP router. Classless is not a valid BGP community.

Chapter: 3

QUESTION NO: 150

Which Administrative Distances listed below are correct? (Choose all that apply)

- A. BGP=20
- B. Static Route=1
- C. External IGRP=170

- D. EIGRP=120
- E. Summary EIGRP=5

Answer: A, B, E

Explanation:

Administrative distances are used to rate the reliability or trustworthiness of a route. Connected interfaces are assigned a value of 0, and static routes assigned a value of 1. These two are the most "trusted" routes, while an administrative distance of 255 indicates the route will never be used. Different routing protocols are assigned different values and can have different values for different types of routes for each protocol (internal, external and summary). Choice C is incorrect because IGRP does not distinguish between internal and external routes. EIGRP is assigned a default administrative distance of 90 not 120 as specified in choice D.

Chapter: 3

QUESTION NO: 151

With regard to RIPv1, which of the following statements are true? (Choose all that apply.)

- A. rip v.1 route updates include a subnet mask field.
- B. rip v.1 is considered a "classful" routing protocol.
- C. rip v.1 is considered a "classless" routing protocol.
- D. rip v.1 route updates do not include a subnet mask field.

Answer: B, D

Explanation:

Ok, set the "wayback" machine to the first day of networking class. On that day, we learned some interesting things about routing protocols and their behaviors. We know RIP v.1 is a "distance vector" routing protocol that relies on "hop count" for pathing decisions. We also know that RIP v.1 is a "classful" routing protocol, which means that it "observes" address class boundaries. When you advertise a network, using the "network" command, you are not given an opportunity to enter a subnet mask. RIP is going to assume the default mask for that address class (255.0.0.0 for a Class "A", 255.255.0.0 for a Class "B" and so on). RIP does NOT include subnet masks in its routing updates. Only a classless routing protocol like OSPF would need to furnish that kind of information.

```
Router(config)#router rip
Router(config-router)#network 192.168.1.10
Router(config-router)#network 192.168.2.10
Router(config-router)#network 192.168.3.10
Router(config-router)#^Z
Router#
Chapter: 4
```

QUESTION NO: 152

Which two routing protocols will most likely be used in a network with limited scalability in mind?

- A. IGRP
- B. OSPF
- C. EIGRP
- D. RIPv2
- E. RIPv1

Answer: A, E

Explanation:

This should be a relatively easy one to answer since there are only two "classful" routing protocols in the list of possible answer choices. Due to the fact that RIP v1 summarizes all known routes along classful network boundaries, using VLSM (Variable Length Subnet Mask) to extend your addressing scheme is out of the question. RIP limits network scalability thanks to its classful nature and its need to send updates every 30 seconds. Imagine how long it would take a network of 200+ routers to converge with RIP sending updates every 30 seconds.

Chapter: 4

QUESTION NO: 153

You are a technician at TestKing. TestKing has an OSPF network. Your newly appointed TestKing trainee wants to know what he benefits of running a "link-state" protocol like OSPF are.

What would your reply be? (Choose all that apply.)

- A. link-state protocols maintain a central network topology database
- B. link-state protocols react quickly to topology changes.
- C. link-state protocol updates are sent out every 30 seconds regardless of topology change
- D. link-state protocols use very little bandwidth in a stable network.

Answer: A, B, D

Explanation:

Link-state routing protocols like OSPF maintain a central network topology database that is built during convergence. They also react quickly to topology changes. If a router's interface goes down, that information is sent out at once in the form of an LSA (Link-State Advertisement), which informs the other routers of the change. Link-state routing protocols like OSPF (starting to notice a theme here?) use very little bandwidth as their routing updates are "incremental" in

nature and only advertise which part of the network has changed. This behavior is the opposite of RIP, which sends a complete copy of its routing table with every update.

Chapter: 4

QUESTION NO: 154

The BGP routing protocol maintain which two types of tables?

- A. IP
- B. BGP topology
- C. BGP attributes
- D. BGP information sent to and received from other BGP routers
- E. shared table that combines UDP and BGP route information

Answer: A, D

Explanation:

The BGP routing protocol maintains two sets of tables, an IP table and a table composed of BGP information sent to and received from other BGP routers. The information in these tables is by default maintained separately, however you can configure the router to share the information if you so desire.

Chapter: 4

QUESTION NO: 155

If you want a permanently created static route entry picked up by the routing process, what command would you use?

- A. static insert
- B. static announce perm
- C. static redistribute
- D. redistribute static

Answer: D

Explanation:

When you create a static route that is not included in the list of the networks advertised by the IOS "network" command, it will not be picked up/learned/injected into the routing protocol unless you specifically advertise it. This feat is accomplished through the use of the "redistribute static" command.

Chapter: 4

QUESTION NO: 156

BGP peers communicate via which method?

- A. RIP
- B. TCP
- C. UDP
- D. ICMP

Answer: B

Explanation:

BGP peers communicate via TCP port 179. The peers initiate a TCP session in which they exchange a series of "OPEN" messages, which ultimately forms the connection. The newly formed connections are kept active by exchanging "KEEPALIVE" messages. Once all this has been completed, the peers exchange routing table, and periodic incremental updates.

Chapter: 4

QUESTION NO: 157

What BGP is considered to be?

- A. both an internal and external routing protocol
- B. an internal routing protocol only
- C. an external routing protocol only
- D. an independent routing protocol

Answer: A

Explanation:

BGP is truly a magical thing! It runs as both an internal routing protocol (iBGP) within the AS and an external routing protocol (eBGP) between ASes. When BGP is used internally (iBGP) its routing duties do not leave the confines of the AS. iBGP learns of new routes within the AS and passes these updates to eBGP peers for distribution. eBGP learns and distributes routes between different ASes.

Chapter: 4

QUESTION NO: 158

What does BGP "ingress filtering" allows you to do?

- A. allows filtering of data packets based on protocol type
- B. allows you to decide which routes will be accepted from peers or neighbors

- C. allows you to decide which routes will be advertised to peers or neighbors
- D. allows filtering of data packets based on protocol number

Answer: B

Explanation:

AAPT has deployed "ingress filtering" on all customer, peer and provider links. This means that filters will be installed to prevent traffic sourced from customer networks not registered for routing with AAPT from entering the AAPT network and will thus prevent this traffic from entering the networks of our peers and providers via AAPT

QUESTION NO: 159

BGP communities perform what function?

- A. they group routes into a single administrative control group
- B. they delete routes based on administrative control group names
- C. they restrict routes based on network class
- D. they allow routers to filter incoming or outgoing BGP routes

Answer: D

Explanation:

BGP communities allow a router a more effective way to filter BGP route updates than would be possible using distribute lists and prefix lists. A BGP community is a tag that is applied to a route update indicating a destination route (or other common attribute) that is shared with other routes. The router can therefore make pathing decisions based on the commonality of a group of updates as opposed to individually assessing each route.

Chapter: 4

QUESTION NO: 160

You are the network administrator at TestKing. You are logged into the ASBR in Area 0 (backbone). You want to advertise the address class 192.168.1.0.

Which of the following commands will make that route available to routers outside the area?

- A. network 192.168.1.0 0.0.0.255 area 0
- B. network 192.168.1.0 255.255.255.0 area 0
- C. adv address 192.168.1.0 255.255.255.0
- D. adv address 192.168.1.0 0.0.0.255

Answer: A

Explanation:

This should have been another easy one to answer. Since the "adv address" command is not used to advertise networks, you should have been able to exclude the answers that used it. That left us with the two commands containing the "network" command. One of the network commands uses a standard subnet mask "255.255.255.0" and one uses a "reverse" or "wildcard" mask. The "network" command uses the "wildcard" mask to indicate how far into the class it should advertise. The correct answer is "network 192.168.1.0 0.0.0.255 area 0"

Chapter: 4

QUESTION NO: 161

Which of the following is NOT a type of OSPF Area?

- A. Normal
- B. Stub
- C. Totally Stub
- D. Not-So-Stubby
- E. Not-So-Normal

Answer: E

Explanation:

There are only four OSPF area types (according to Cisco) and they are stub, totally stub, not-so-stubby and normal. The answer "Not-So-Normal" is the made up one.

Chapter: 4

QUESTION NO: 162

With regard to VLSM, which of the following statements are true? (Choose all that apply.)

- A. the area number must be carried with the routing update
- B. the addresses being summarized must be classful in nature
- C. the subnet mask is carried with the routing update
- D. addresses being summarized must contain the same "high-order" bits

Answer: C, D

Explanation:

Variable Length Subnet Masking or VLSM gives the router the ability to announce a series of addresses with a single routing table entry. Pretty neat trick huh?

The way that it does this, is it looks at the addresses and ascertains which of those addresses are similar out to a designated bit pattern, also known as "high-order bits". It then looks at the subnet

mask accompanying the update and decides which series of addresses are considered matches. You could potentially advertise as many as five destination addresses with a single routing table entry.

Chapter: 4

QUESTION NO: 163

You are a technician at TestKing. You are designing a multi-area OSPF network.

What two things must you do?

- A. make sure that each area connects to the backbone (Area 0).
- B. assign each area a unique number
- C. configure Areas 1-4 as Stub Areas
- D. configure Areas 1-4 as Not-So-Stubby Areas

Answer: A, B

Explanation:

Creating a multi-area OSPF network requires a few design considerations such as, all areas must have a connection to Area 0 so that routing updates will have a central distribution point, all created areas must have a unique area number (that one is kind of obvious) to prevent routing updates from being confused by the fact that two areas are using the same number.

Chapter: 4

QUESTION NO: 164

What function does the "area range" command perform?

- A. instructs the ABR to summarize the routes for a specific area before advertising them
- B. instructs the ABR to segregate the routes for a specific area before advertising them
- C. specifies a range of areas to advertise exclusively
- D. specifies a range of areas to observe and report the routing status

Answer: A

Explanation:

The "area range" command instructs the ABR to summarize the addresses within the area before sending them out in a routing update.

Chapter: 4

QUESTION NO: 165

You are a technician at TestKing. You want to check the status of configured virtual links.

Which command should you issue?

- A. show ospf virtual-links
- B. show virtual-links ospf
- C. show ip ospf virtual-links
- D. show ip ospf v-links

Answer: C

Explanation:

The command "show ip ospf virtual-links" will show you information about configured virtual-links.

```
new-york router#show ip ospf virtual-links Interface address: 192.168.1.10 (POS 1/1/1) cost: 1,
state: P To P, transit area: 1.2.3.4 hello: 10, dead: 40, retrans: 5 nbr id: 192.168.1.20, nbr address:
192.168.1.20 nbr state: Full, nbr mode: Master, last hello: 38
```

Chapter: 4

QUESTION NO: 166

An OSPF router must be in what state in order to route traffic?

- A. mixed
- B. partial
- C. two-way
- D. full

Answer: D

Explanation:

An OSPF router must be in "full" state (down, attempt, init, 2-way, exstart, exchange, loading, or full). When a router is in "full" state, adjacencies are fully established and the router is ready to route traffic.

Chapter: 4

QUESTION NO: 167

You are the network administrator at TestKing. Your newly appointed TestKing network engineer has not quite grasped the concept of classless routing.

How would you explain the concept? (Choose all that apply.)

Leading the way in IT testing and certification tools, www.testking.com

- A. classless routing protocols carry the subnet mask with updates
- B. classful routing protocols carry the subnet mask with updates
- C. classless routing enables the use of VLSM
- D. classless routing enables the use of high-order routing

Answer: A, C

Explanation:

Classless routing is a pretty complex topic for new engineers to fully grasp. The best way to explain the basic principles is to re-enforce the fact that the classless routing is not bound by address class limitations, routing updates carry the specialized subnet mask which makes VLSM possible.

Chapter: 4

QUESTION NO: 168

The "ip bandwidth-percentage eigrp" performs what function?

- A. it dictates what the maximum bandwidth percentage that EIGRP packets are allowed to use on a single router interface.
- B. it dictates what the maximum bandwidth percentage that IP packets are allowed to use on a single router interface
- C. it dictates what the maximum bandwidth percentage that EIGRP packets are allowed to use on all router interfaces.
- D. it dictates what the maximum bandwidth percentage that packets routed via EIGRP are allowed to use on a single router interface.

Answer: A

Explanation:

The ip bandwidth-percent eigrp command is used to configure the percentage of bandwidth that may be used by EIGRP on an interface. If you issue the "ip bandwidth-percentage eigrp" command at the command line, the router will only allow 50% of the interface's bandwidth (which is the default) to be used for EIGRP.

Chapter: 4

QUESTION NO: 169

What is the cost between the next hop router and the destination?

- A. reliable distance
- B. calculated distance

- C. reasonable distance
- D. advertised distance
- E. administrative distance

Answer: D

Explanation:

The Diffusing Update Algorithm (DUAL) is the "brains" behind the EIGRP routing protocol's path decision making process. When routes to a destination network go down, the DUAL algorithm calculates new paths to make sure that traffic gets to its intended destination. One of the metrics involved in this decision making process is "advertised distance" which is equal to the cost between the next hop router and the destination network.

Chapter: 4

QUESTION NO: 170

What is the cost between the local router and the next hop router?

- A. feasible distance
- B. partial distance
- C. next-hop distance
- D. agregate distance
- E. advertised distance

Answer: C

Explanation: Only possible answer might be C.

Not A: Feasible distance is the cost of the route from the current router to the remote network.

Source BSCI Self Study Cisco Press (2003) 2nd Edition P362-363.

QUESTION NO: 171

In order to determine the "best" path to a destination network, you need to add which two path costs together? (Select two.)

- A. feasible distance
- B. partial distance
- C. advertised distance.
- D. agregate distance
- E. next hop distance

Answer: C, E

Explanation:

Two path costs would be advertised distance plus the next hop distance.
Source BSCI Self Study Cisco Press (2003) 2nd Edition P362-363.

QUESTION NO: 172

What could cause EIGRP routes to be considered "stuck-in-active"?

- A. the route being reported by the "stuck-in-active" message no longer exists.
- B. the route being reported by the "stuck-in-active" message has been activated.
- C. a query for the route generated by an EIGRP neighbor has not yet been replied to.
- D. a query for the route generated by an OSPF neighbor has not yet been replied to.

Answer: A, C

Explanation:

The "stuck-in-active" message refers to the condition in which a primary route is no longer available and no feasible successor is available. The EIGRP router has sent out a query to its neighbors and no reply has been heard back for three minutes.

Chapter: 4

QUESTION NO: 173

What are Query, Update and Reply known as?

- A. EIGRP renewable packets
- B. EIGRP reliable packets
- C. TCP/IP reliable packets
- D. QueryAll packets
- E. OSPF reliable packets

Answer: B

Explanation:

EIGRP utilizes five different types of packets:

- 1) ACK
- 2) hello
- 3) query
- 4) reply
- 5) update

These packets are instrumental in EIGRP's ability to reliably route packets from source to destination network. The "reliable packets" are query, update, and reply.

Chapter: 4

QUESTION NO: 174

You are a technician at TestKing. TestKing has an IS-IS network. Your newly appointed TestKing trainee wants to know what IS-IS stand for.

What would your reply be?

- A. intercontinental station-to-intercontinental station
- B. intermediate section-to-intermediate section
- C. intermediate station-to-intermediate station
- D. intermediate system-to-intermediate system

Answer: D

Explanation:

The intermediate system-to-intermediate system routing protocol, more commonly known as IS-IS was developed by Digital Equipment Corporation.

Chapter: 4

QUESTION NO: 175

Which of the following attributes are common to both OSPF and IS-IS?

- A. they both maintain a link-state database.
- B. they both run the SPF algorithm to determine the shortest path to destination networks
- C. they both utilize the concept of "areas" to maintain hierarchical network topologies
- D. they both use RIP as their data transfer medium
- E. they both communicate via port 80 when sending routing updates

Answer: A, B, C

Explanation:

IS-IS and OSPF share a number of common features; link-state databases, the running of the SPF algorithm, the use of "areas" among other similarities.

Chapter: 4

QUESTION NO: 176

Which of the following subnet masks is the most optimized for providing five internet addresses?

- A. 225.255.255.192
- B. 225.255.255.248
- C. 225.255.255.255
- D. 225.255.255.254

Answer: B

Explanation:

Do the math and win a prize! This should have taken you about 5 minutes to answer. The subnet mask 255.255.255.248 will give you 32 subnets with 6 hosts per subnet. Sounds pretty efficient to me. If you want a great shortcut for figuring out number of hosts per subnet, subtract 248 from 256, that leaves you with 8 right?

Subtract two from 8 (can't use the first and last addresses in the range) and you are left with 6. There are six hosts in each of the subnets. It works every time.

The other subnet masks:

255.255.255.192 = 4 subnets with 62 hosts per subnet, way too many hosts.

255.255.255.255 = Gong! I hope no one fell for this one

255.255.255.254 = 128 subnets with two hosts per subnet, won't work if you need five hosts.

(*TIP* use this for point to point serial connections)

Chapter: 4

QUESTION NO: 177

On a point-to-point network, what acknowledges each LSP that it receives.

- A. PSNP
- B. POPN
- C. BPDU
- D. LMNOP

Answer: A

Explanation:

On a point-to-point IS-IS network, an LSP(Link State Packet) is generated to announce a router's links and the status of the those links. When the LSP is received, a PSNP (Partial Sequence Number Packet) is sent to acknowledge each of the LSPs that are received.

Chapter: 4

QUESTION NO: 178

What is the IS-IS equivalent of the OSPF backbone Area 0?

- A. level-0 backbone

- B. area 0 backbone
- C. level-1 backbone
- D. level-2 backbone

Answer: D

Explanation:

In order to answer this question, we are only going to deal with Level-1 and Level-2 routers (I see the guy with his hand raised...Yes, there are Level1/2 routers in IS-IS as well). The Level-1 routers are similar to OSPF areas; they are independent of each other and require a concentration point to share routing updates.

It is in this capacity that the Level-2 backbone functions. Level-2 routers provide the common connection point through which route updates from the other Level-1 routers flow. This concept has been highly simplified, but with good reason...hint-hint.

Chapter: 4

QUESTION NO: 179

With which of the following can Level-2 IS routers establish an adjacency?

- A. Level-1 IS in the same area
- B. Level-2 IS in a different area
- C. Level-2 IS in the same area
- D. Level-1 IS in a different area

Answer: B, C

Explanation:

L2 router only establish adjacency with other L1/L2 or L2 routers - - never with a plain L1 router

QUESTION NO: 180

Which of the following are most like an OSPF Area Border Router (ABR)?

- A. Level 1 IS
- B. Level 2 IS
- C. Level1/Level2 IS
- D. Level2/Level3 IS
- E. IS-IS Media Border Router (IMBR)

Answer: C

Explanation:

IS-IS L1/L2 routers are a hybrid router type (which is the default setting when activating an IS-IS router). The L1/L2 router is most like the OSPF ABR because L1/L2 routers can connect to both the backbone (L2) and a router in a different area (L1), which as we know an ABR is also capable of doing.

Chapter: 4

QUESTION NO: 181

From the answer choices below, choose the statements that are true regarding IS-IS.

- A. L1 IS routers can be used to connect areas together
- B. L2 IS routers can be used to connect areas together
- C. IS-IS routers can either be an L1 or L2
- D. IS-IS routers can be both an L1 and L2 at the same time

Answer: B, D

Explanation:

IS-IS L1/L2 routers are a hybrid router type (which is the default setting when activating an IS-IS router). The L1/L2 router is most like the OSPF ABR because L1/L2 routers can connect to both the backbone (L2) and a router in a different area (L1).

Chapter: 4

QUESTION NO: 182

What will an OSPF ABR connect to when introduced into an existing network?

- A. one or more OSPF areas
- B. a single IS-IS area
- C. L2 IS router
- D. L1/L2 IS router

Answer: A

Explanation:

The OSPF Area Border Router's job is to sit on the border of one or more OSPF areas and provide a communication conduit between the areas and the backbone. The ABR runs the SPF algorithm and maintains the data for each area that it is connected to.

Chapter: 4

QUESTION NO: 183

Level 1/2 IS routers provide which service for IS-IS networks?

- A. intra-area routing services
- B. inter-area routing services
- C. intra-level routing services
- D. consolidation of OSPF and IS-IS routing updates

Answer: B

Explanation:

IS-IS L1/L2 routers are a hybrid router type (which is the default setting when activating an IS-IS router). The L1/L2 router is most like the OSPF ABR because L1/L2 routers can connect to both the backbone (L2) and a router in a different area (L1). Much in the same way that an OSPF ABR sits on the borders of one or more areas so does the L1/L2 router. Since the L1/L2 router straddles the area "fence", it is capable of providing inter-area routing services.

Chapter: 4

QUESTION NO: 184

OSPF can handle VLSM because it has what capability?

- A. OSPF organizes the network hierarchy when the SPF algorithm is run
- B. OSPF's nature as a link-state database automatically supports VLSM
- C. OSPF carries the subnet mask within its routing updates
- D. OSPF organizes areas and their ABRs into hierarchical groups

Answer: C

Explanation:

In order for VLSM to be deployed on a network infrastructure, you need to use a routing protocol that is capable of sending specialized subnet masks as part of the routing table updates. When you configure OSPF, one of the steps is to enter the network addresses with corresponding subnet masks.

Chapter: 4

QUESTION NO: 185

Which of the following are metrics available to IS-IS?

- A. delay
- B. error
- C. expense
- D. hop count

Answer: A, B, C

Explanation:

The IS-IS routing protocol incorporates three metrics: delay, expense, and error.

delay - measures the amount of delay on the link

expense - measures cost in resources of using the link

error - measures the number of errors occurring on the link

Chapter: 4

QUESTION NO: 186

What is the purpose of the "show isis route" command?

- A. It displays the Level-1 routing table for integrated IS-IS
- B. It displays the Level-2 routing table for integrated IS-IS
- C. It displays the routes most recently taken to a Level-1 ABR
- D. It displays the hops taken by a packet leaving the nearest Level-1 router

Answer: A

Explanation:

The "show isis route" command will show the Level-1 routing table as seen in the example below:

```
TestKing#show isis route
```

```
IS-IS Level-1 Routing Table - Version X
```

```
System Id Next-Hop SNPA Interface Metric State
```

```
BB00.0400.020C BB00.0400.020C bb00.0400.020c Serial0 10 Up
```

```
0800.2BB1.4434 0000.0000.0000 -- -- 0 Up
```

```
0800.2BB3.785B 0800.2BB3.785B bb00.0400.020c Serial0 10 Up
```

Chapter: 4

QUESTION NO: 187

You are a technician at TestKing. TestKing has an OSPF network. Your newly appointed TestKing trainee wants to know what an OSPF router does when a link-state change occurs.

What would your reply be?

- A. it shuts down and runs the SPF algorithm
- B. it multicasts the link-state change on ip address 224.0.0.6 to the DRs and BDRs via an LSU

C. it multicasts the link-state change on ip address 255.255.255.224 to the DRs and BDRs via an LSU

D. it unicasts the link-state change on ip address 255.255.255.224 to the ABRs and ASBRs via an LSU

Answer: B

Explanation:

An OSPF router reacts to link-state changes in the follow manner:

- 1) Link-state change is detected by a router
 - 2) The router multicasts an LSU (Link-state update) to DRs and BDRs on address 224.0.0.6
 - 3) The DR confirms receipt of the LSU and then floods it to other routers via address 224.0.0.5
- There is more to the story, but this gives you the basic idea.

Chapter: 4

QUESTION NO: 188

You are a technician at TestKing. TestKing has an OSPF network. Your newly appointed TestKing trainee wants to know OSPF which multi-cast addresses are used on the TestKing network.

What would your reply be? (Choose all that apply.)

- A. 224.0.0.6
- B. 224.0.0.5
- C. 225.0.0.3
- D. 255.255.255.224

Answer: A, B

Explanation:

An OSPF router reacts to link-state changes in the follow manner:

- 1) Link-state change is detected by a router
- 2) The router multicasts an LSU (Link-state update) to DRs and BDRs on address 224.0.0.6
- 3) The DR confirms receipt of the LSU and then floods it to other routers via address 224.0.0.5

These addresses are defined in RFC-1583.

224.0.0.5 OSPFIGP OSPFIGP All Routers

224.0.0.6 OSPFIGP OSPFIGP Designated Routers

Chapter: 4

QUESTION NO: 189

You are a technician at TestKing. TestKing has an OSPF network. Your newly appointed TestKing trainee wants to know why OSPF scales better than other routing protocols.

What would your reply be? (Choose all that apply.)

- A. OSPF converges faster
- B. OSPF uses a faster update protocol
- C. OSPF sends updates in a compressed format
- D. OSPF updates are incremental in nature, thus smaller in size

Answer: A, D

Explanation:

OSPF can scale better than some of its contemporaries due to a couple of factors. OSPF updates are based on topology changes, as opposed to RIP, which sends updates regardless of network topology status. When OSPF does send out an update, the update only contains elements that are different due to a linkstate change. RIP sends out a copy of its entire routing table whether anything has changed or not. OSPF is also aware of link costs, which means it won't send updates out over a slow link.

Chapter: 4

QUESTION NO: 190

You are a technician at TestKing. Your newly appointed TestKing trainee wants to know why route summarization is so desirable

What would your reply be?

- A. summarization results in smaller routing tables
- B. summarization results in larger, more complex routing tables
- C. summarization is easier on router cpu due to decreased number of routes to manage
- D. summarization makes packets travel faster due to summarization compression routines

Answer: A, C

Explanation:

Route summarization is the process of condensing a series of routing table entries into a single entry that is maintained by the router. When you summarize, you compare a series of addresses and try to determine which ones have the same "high order" bit pattern (also known as significant bits). If you have a number of addresses with similar "high order" bit patterns, you can represent them with a single routing table entry as opposed to listing them separately. An end result of route summarization is that multiple "down stream" routes can be represented by a single table entry.

Chapter: 4

QUESTION NO: 191

You are a technician at TestKing. The TestKing network is connected to the internet via more than one ISP. Your newly appointed TestKing trainee wants to know what this type of network is known as.

What would your reply be?

- A. multi-gnomed
- B. mega-homed
- C. multi-homed
- D. maxi-homed

Answer: C

Explanation:

If your company requires high availability to the internet you might consider having more than one ISP provide your internet connectivity. This arrangement is referred to as being "multi-homed". Many businesses that do a large amount of their business via "the net" utilize two or more ISP in order to provide fault tolerance should one service provider experience difficulties.

Chapter: 4

QUESTION NO: 192

Which of the following routing protocols support VLSM?

- A. EIGRP
- B. IGRP
- C. OSPF
- D. RIP version 1
- E. RIP version 2

Answer: A, C, E

Explanation:

EIGRP, RIP v.2 and OSPF are classless routing protocols, which means that they are capable of including a specialized subnet mask in routing updates. Any protocol that carries a subnet mask in updates is a candidate for supporting VLSM. IGRP and RIP v.1 are classful routing protocols, which they are excluded from VLSM membership.

Chapter: 4

QUESTION NO: 193

You are a trainee technician at TestKing. The TestKing network implements the Border Gateway Protocol (BGP). Your instructor tells you that BGP comes in two flavors. He wants to what these two flavors are.

What would your reply be? (Choose all that apply.)

- A. reliable gateway protocol (RGP)
- B. interior gateway protocol (IGP)
- C. silent gateway protocol (SGP)
- D. exterior gateway protocol (EGP)

Answer: B, D

Explanation:

BGP runs in either of two capacities, as an IGP where it maintains routing tables inside the AS and EGP in which it maintains routing tables between AS's.

Chapter: 4

QUESTION NO: 194

You are a network engineer at TestKing. Your newly appointed TestKing trainee wants to know what command she should issue to view EIGRP entries in the routing table.

What would your reply be?

- A. show ip eigrp current
- B. show ip eigrp route recent
- C. show current ip eigrp route
- D. show ip eigrp route

Answer: D

Explanation:

The "show ip eigrp route" command will show only the routing table entries that are EIGRP relevant.

Chapter: 4

QUESTION NO: 195

With regard to BGP Prefix Lists, which of the following statements are true?

- A. prefix list sequence numbers by default start at 5 and increment by 5 as more are added
- B. if a router matches a prefix list entry, processing of that list ceases
- C. if a router matches a prefix list entry, processing of that list continues till another match is made
- D. the command "no ip prefix-list sequence-number" deletes the specified prefix list from memory

Answer: A, B

Explanation:

BGP prefix lists function much in the same way as ACLs, but with a few differences. Prefix lists allow increment updating, while ACLs do not. A router will cease processing a prefix list once a match is made; ACLs get processed all the way to the end. If you do not specify a BGP Prefix List sequence-numbering scheme, the list entries will increment by 5 as you enter new statements.

Chapter: 4

QUESTION NO: 196

You are a network engineer at TestKing. Your newly appointed TestKing trainee wants to know what the term "feasible distance" means.

What would your reply be?

- A. It is the distance from the web server to the mail server
- B. It is the sum of the cost of the route from the next-hop to the next-hop
- C. It is the distance from the firewall to the ISP
- D. It is the distance from the next-hop router to the destination
- E. None of the above answers is correct

Answer: E

QUESTION NO: 197

What is the relevance of "priority" when dealing with DR/BDR elections?

- A. a priority of 1 means the router is eligible to become a DR
- B. a priority of -1 means the router is eligible to become a DR
- C. a priority of 0 means the router is ineligible to become a DR or BDR
- D. DR/BDR elections run at regular intervals regardless of network topology status.

Answer: A, C

Explanation:

DR and BDR elections are handled as follows:

- Highest priority will become the DR
- Second Highest priority will become the BDR
- Only the malfunction of a DR or BDR will cause an election.
- A priority of "1" indicates eligibility to become a DR
- A priority of "0" indicates ineligibility to become a DR or BDR

Chapter: 4

QUESTION NO: 198

You are the network administrator at TestKing. Your newly appointed TestKing trainee wants to know when the use of BGP would be appropriate.

What would your reply be? (Choose all that apply.)

- A. if the border router in your AS is behind your corporate firewall
- B. if your AS is at the edge of another routing protocol's border
- C. if your network is multi-homed to the internet
- D. if you have a need to manage the traffic entering and exiting your autonomous system
- E. if your AS is a transit area for packets destined for another ASs

Answer: C, D, E

Explanation:

When to use BGP is a tough question.

The rule of thumb for when to use BGP usage is as follows:

- 1) if your network is multi-homed to the internet
- 2) if you have a need to manage the traffic entering and exiting your autonomous system
- 3) if your AS is a transit area for packets destined for another ASs

Chapter: 4

QUESTION NO: 199

You are the network administrator at TestKing. Your newly appointed TestKing trainee wants to know when the use of BGP would NOT be appropriate.

What would your reply be? (Choose all that apply.)

- A. if your network only has a single connection to the internet
- B. if you have a low bandwidth link between ASs

- C. if the selection of routes to outside ASs is not a high priority
- D. if the AS connecting you to the internet charges by the packet
- E. if the router connecting to your external AS is running IOS version 11.2 or earlier

Answer: A, B, C

Explanation:

When not to use BGP is an even tougher question.

The rule of thumb for when NOT to use BGP usage is as follows:

- 1) if your network only has a single connection to the internet
- 2) if you have a low bandwidth link between ASs
- 3) if the selection of routes to outside ASs is not a high priority

Most of the choices for not to BGP relate to the fact that BGP can easily overwhelm an underpowered link or router. So you should only use BGP if your situation really warrants the configuration and management that comes along with it.

Chapter: 4

QUESTION NO: 200

What attribute must all BGP implementations recognize?

- A. customized
- B. synchronized
- C. well-known
- D. well-adjusted
- E. optional

Answer: C

Explanation:

A BGP "well-known" attribute is one that must be recognized by all implementations of BGP and must be included in every update message. There are other types of updates that may or may not need to be included in the update messages.

The following BGP attributes are the most commonly known:

- 1) Well-known mandatory.
- 2) Well-known discretionary.
- 3) Optional transitive.
- 4) Optional non-transitive.

Chapter: 4

Part 2, Building Cisco® Multilayer Switched Networks (BCMSN®) (411 questions)

TOPIC 1: TECHNOLOGY (205 questions)

Section 1: Describe the Enterprise Composite Model used for designing networks and explain how it addresses enterprise network needs for performance, scalability and availability (1 question)

QUESTION NO: 1

In the Enterprise Composite Network Model; what are three of the functional areas? (Choose three)


- A. Enterprise Campus
- B. Enterprise Edge
- C. Service Provider Edge
- D. Building Access
- E. Server Farm
- F. Campus Backbone
- G. Wiring Closet

Answer: A, B, C

Explanation:

Although most enterprise networks have evolved with growing IT requirements, the Cisco SAFE architecture uses a modular approach, which has two main advantages. First, it allows the architecture to address the security relationship between the various functional blocks of the network. Second, it permits designers to evaluate and implement security on a module-by-module basis, instead of attempting the complete architecture in a single phase.

The following figure illustrates the first layer of modularity in SAFE. Each block represents a functional area. The Internet service provider (ISP) module is not implemented by the enterprise, but is included to the extent that specific security features should be requested of an ISP in order to mitigate against certain attacks.



This figure illustrates the three functional areas of the Network model: The enterprise campus, enterprise edge, and the service provider edge.

Reference: http://www.cisco.com/en/US/netsol/ns340/ns394/ns171/ns128/networking_solutions_white_paper09186a008009c8b6.shtml

Section 2, Describe the physical, data-link and network layer technologies used in a switched network, and identify when to use each (10 questions)

QUESTION NO: 1

The TestKing network is upgrading all switches to be layer 3 capable. What are some of the advantages experienced with layer 3 switching (Select all that apply)?

- A. High-performance packet switching
- B. Security
- C. Flow accounting
- D. Low latency
- E. Low per-port cost
- F. Quality of service
- G. Increased Scalability
- H. Hardware-based packet forwarding

Answer: A, B, C, D, E, F, G, H

Explanation:

Traditional software-based routers are simply not fast enough to do the job. Layer 3 switching is primarily a routing solution implemented in the switch's fabric. Instead of using traditional software-based routers, savvy switch manufacturers are integrating router functionality into their switch hardware, offering faster, more secure, more reliable routing solutions

The advantages are clear. Layer 3 switching provides hardware-based routing at wire speeds, which significantly improves overall performance. Routing packets via hardware eliminates bottlenecks associated with software-based routers and delivers seamless implementation into existing networks.

Providing this boost in performance and removing LAN router bottlenecks, switched networks more efficiently utilize available bandwidth. Users get a responsive, high-speed network that is more stable and reliable while protecting the existing investment in their network infrastructure. Layer 3 switches are usually managed, allowing network managers to effortlessly configure and manage the routing process. Reduced support costs make this a true cost-effective solution with the added benefits of higher network reliability and a quicker response time

QUESTION NO: 2

Which statement correctly describes the extended system ID?

- A. It is the 2-bit number of the MSTP instance.
- B. It is the VLAN identifier value and allows for 4096 BIDs to be uniquely identified.
- C. It is the bridge MAC address which is allocated from a pool of MAC addresses that are factory assigned.
- D. It is a hex number used to measure the preference of a bridge in the spanning-tree algorithm.
- E. None of the above

Answer: B

Explanation:

Each VLAN on each network device has a unique 64-bit bridge ID consisting of a bridge priority value, an extended system ID, and an STP MAC address allocation.

Extended system IDs are VLAN IDs between 1025 and 4096. Cisco Switches support a 12-bit extended system ID field as part of the bridge ID. Chassis that support only 64 MAC addresses always use the 12-bit extended system ID. On chassis that support 1024 MAC addresses, you can enable use of the extended system ID. STP uses the VLAN ID as the extended system ID.

Through the use of the extended system ID, up to 4096 VLANS can be used instead of the 1024 that are used as the default.

QUESTION NO: 3

A new TestKing branch is being opened and you are contemplating the use of Unshielded Twisted Pair (UTP) cable for this new office. What is the maximum distance that can be used between two nodes on this UTP network?

- A. 100 meters
- B. 150 meters
- C. 100 feet
- D. 2 kilometers
- E. 300 meters
- F. None of the above

Answer: A

Explanation:

UTP cabling does not offer as high bandwidth or as good protection from interference as coaxial or fiber optic cables, but it is less expensive and easier to work with. The maximum length for an unshielded twisted pair (UTP) cable segment is 100 meters

QUESTION NO: 4

A new TestKing branch office is being installed and connected, with individual stations and servers being plugged in to the LAN switch. What kind of cable should be used to connect a router, server, or individual workstation to a switch?

- A. rollover cable
- B. crossover cable
- C. straight-through cable
- D. coax cable

Answer: C

Explanation:

To connect any end device to a switch you have to use a straight cable.

Incorrect Answers:

A: A rollover cable is used to connect to the console port of a switch or a router.

B: Crossover cables are used to connect: two computers directly together, two hubs, a hub to a switch, a cable modem to a router, or two router interfaces together. It is also used for directly connecting a PC into a router's Ethernet port.

D: Coaxial cable is typically used for DS3 interfaces. It is not normally used in switched LAN networks.

QUESTION NO: 5

Gigabit Ethernet switches are being installed throughout the TestKing network. Which of the following cable types are appropriate for Gigabit Ethernet applications? (Select two)

- A. Cat-3 UTP
- B. Cat-5 UTP
- C. RG-58 coax
- D. 50 micron MMF
- E. 62.5 micron SMF

Answer: B, D

Explanation:

The following lists the Gigabit Ethernet Cabling options, along with their respective Distance Limitations:

1000BASE-T EIA/TIA Category 5 UTP 4 100 m

1000BASE-SX Multimode fiber (MMF) with 62.5 micron core; 850 nm laser 1 275 m

MMF with 50 micron core; 850 nm laser 1 550 m

1000BASE-LX/LH MMF with 62.5 micron core; 1300 nm laser 1 550 m

Single-mode fiber (SMF) with 50 micron core;

1300 nm laser
1 550 m

B: 1000BaseT use category 5 UTP.
D: 1000BaseSX use 62.5 and 50-micron MMF

QUESTION NO: 6


You need to ensure full connectivity exists between all stations on the TestKing LAN. By what means could you provide a Layer 3 data path between two separate VLANs? (Choose two)

- A. A VLAN trunking
- B. An external router
- C. An internal processor
- D. VLAN capable bridge
- E. EtherChannel

Answer: B, C

Explanation:

The only connectivity that we want between VLANs is achieved at Layer 3 (L3) by a router. This is Inter-VLAN routing. To further simplify the diagrams, we will represent VLANs as different physical Ethernet segments, as we are not really interested in the specific bridging functions provided by the switch.



In the above diagram, the two VLANs are considered as two different Ethernet segments. Inter-VLAN traffic needs to go through the external router. If host A wants to communicate with host B, it will typically use the router as a default gateway.

In order to provide connectivity between VLANs, traffic must be routed. This can be either achieved via an external router, or an internal route processor such as the Route Switch Module (RSM) found in Cisco Catalyst 6500 switches.

QUESTION NO: 7

Which layer 3 switching method utilizes a forwarding information base (FIB)?

- A. Route caching
- B. Demand-based switching
- C. Flow-based switching
- D. Topology-based switching

Answer: D

Explanation:

Cisco Express Forwarding (CEF) is an example of a topology-based switching mechanism that uses a FIB. CEF provides a topology-based switching mechanism that switches packets at hundreds of millions of packets per second, while maintaining high-speed services.

In a non-Cisco Express Forwarding implementation, the first packet of any flow needs to be processed by the CPU. This can lead to decreased performance, particularly if many new flows are being set up. In a Cisco Express Forwarding-based switch, the forwarding table is prepopulated based on the routing table. This helps to ensure both predictability in the event of a route flap and that CPU overload will not affect performance. All Cisco Catalyst switching products support Cisco Express Forwarding today.

Cisco Express Forwarding (CEF) switching is a proprietary form of scalable switching intended to tackle the problems associated with demand caching. With CEF switching, the information which is conventionally stored in a route cache is split up over several data structures. The data structures that provide optimized lookup for efficient packet forwarding include:

- The Forwarding Information Base (FIB) table - CEF uses a FIB to make IP destination prefix-based switching decisions. The FIB is conceptually similar to a routing table or information base. It maintains a mirror image of the forwarding information contained in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and these changes are reflected in the FIB. The FIB maintains next-hop address information based on the information in the IP routing table. This table is used in this topology based switching method.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 412.

QUESTION NO: 8

Leading the way in IT testing and certification tools, www.testking.com

New access layer switches are being installed in the 3-tiered TestKing network. Which of the attributes below correctly describe the characteristics of access layer switches? (Choose all that apply.)

- A. High port density to connect to end users.
- B. Robust Layer 3 routing throughput
- C. Inter-VLAN routing
- D. Low cost
- E. Security
- F. None of the above

Answer: A, D

Explanation:

The Access Layer:

The main criteria for access devices are to provide this functionality with low-cost, high port density devices. Access layer switches should provide connections for as many end devices as possible, as fast as possible.

Incorrect Answers:

B, C: Layer 3 (Inter VLAN) routing is processor intensive and should generally be used only in larger, more expensive distribution layer switches instead of at the access layer.

E: The use of security features such as access lists should be used at the distribution layer of the network.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 21

QUESTION NO: 9

In an effort to reduce the number of broadcast traffic within the TestKing network, new Catalyst switches are being installed. Which of the following statements correctly describe Layer 2 broadcast traffic?

- A. Layer 2 broadcast traffic is blocked by Layer 3 devices.
- B. A new packet is sent each time the client requests it.
- C. Each frame uses a special address for which only interested clients listen.
- D. It is the most efficient way to send data to a small group of clients.
- E. Each packet is refreshed when requested.

Answer: A

Explanation:

LAN broadcasts do not cross routers (Layer 3 devices). By default, routers do not forward any broadcast packets, unless the "IP helper-address" command is configured.

Incorrect Answers:

B: Each broadcast is only sent once.

C: Multicast is more efficient. Broadcast reach all clients, multicast will only reach the member of the multicast group.

D: All clients on the subnet receive the broadcast traffic.

E: Broadcast traffic is not refreshed or resent. Doing so could result in a broadcast storm.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 17.

QUESTION NO: 10

The TestKing network is upgrading the network to use switches that are capable of multilayer switching. Which statement below best describes the concept of multilayer switching (MLS)?

- A. Switches that operate at the access, distribution, and core layer at the design model.
- B. An OSI Layer 1 and 2 bridging technique.
- C. A technique to provide hardware switching of Layer 3 unicast packets.
- D. A flow-based Layer 3 packet routing methodology.

Answer: C

Explanation:

Switches are layer two devices originally developed to contain broadcasts. A multilayer switch is an improvement because it contains extra processing power to consider layer 3 address information, so it effectively works at more than one layer.

Multi-Layer Switching (MLS) has become a highly desired method of accelerating routing performance through the use of dedicated Application Specific Integrated Circuits (ASICs).

Traditional routing is done through a central CPU and software. MLS offloads a significant portion of routing (packet rewrite) to hardware, and thus has also been termed switching. MLS and Layer 3 switching are equivalent terms.

Section 3: Explain the role of switches in the various modules of the Enterprise Composite Model (Campus Infrastructure, Server Farm, Enterprise Edge, Network Management) (2 questions)

QUESTION NO: 1

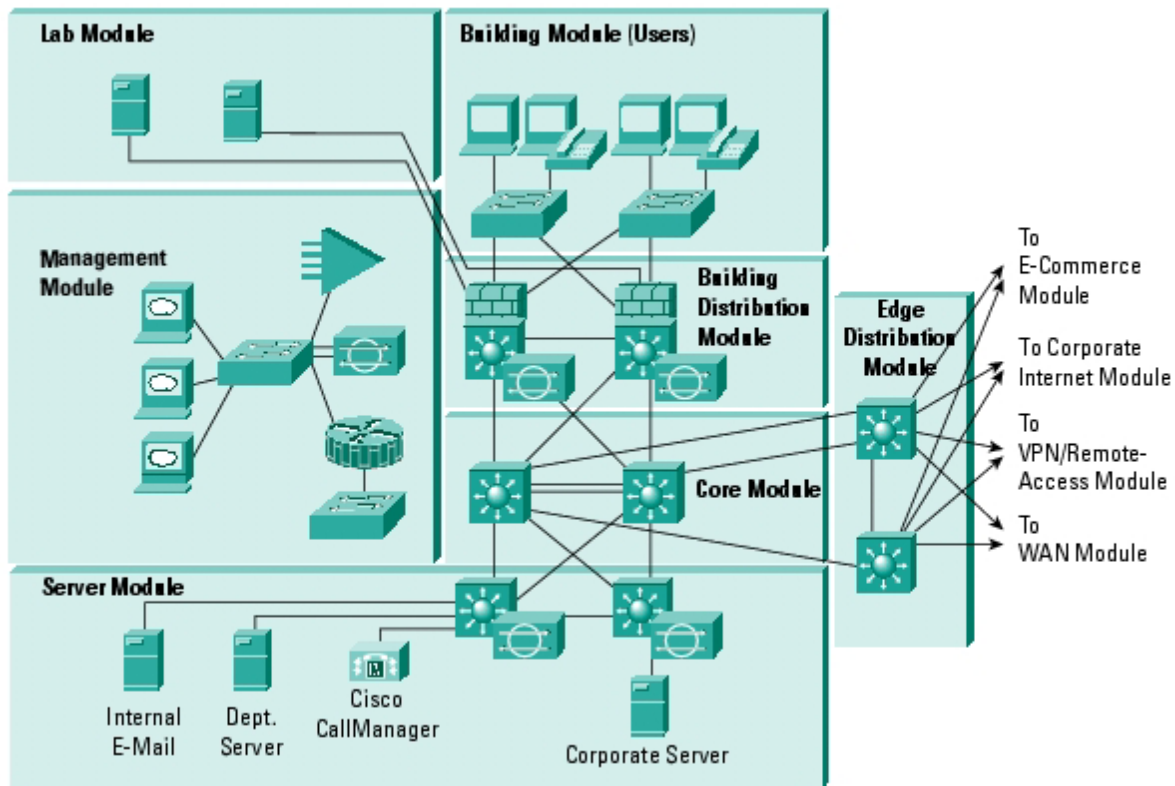
In the Enterprise Composite Model, what are the four major modules of the Campus functional area? (Choose four)

- A. Campus Infrastructure
- B. Network Management
- C. Server Farm
- D. Edge Distribution (Enterprise Edge)
- E. Access Distribution
- F. Core Layer

Answer: B, C, D, F

Explanation:

Following is a detailed analysis of all of the modules contained within the enterprise campus. The following figure shows this campus:
Enterprise Campus Detail



Management Module

The primary goal of the management module is to facilitate the secure management of all devices and hosts within the enterprise architecture.

Core Module

The core module in the network architecture is nearly identical to the core module of any other network architecture. It merely routes and switches traffic as fast as possible from one network to another.

Building Distribution Module

This module provides distribution layer services to the building switches. These include routing, quality of service (QoS), and access control. Requests for data flow into these switches and onto the core, and responses follow the identical path in reverse.

Building Access Module

This module is described as the extensive network portion that contains end-user workstations, phones, and their associated Layer 2 access points. Its primary goal is to provide services to end users.

Server Module

The server module's primary goal is to provide application services to end users and devices. Traffic flows on the server module are inspected by on-board intrusion detection within the Layer 3 switches.

Edge Distribution Module

This module aggregates the connectivity from the various elements at the edge. Traffic is filtered and routed from the edge modules and routed into the core.

Incorrect Answers:

A: This is incorrect because 'Campus Infrastructure' refers to the collective of all network equipment on the campus.

E: This is incorrect because the 'Access Distribution' area is not a defined area, it is just a combination of the already familiar terms 'access' (from the OSI access layer) and 'distribution' (from this models Edge Distribution).

QUESTION NO: 2

The TestKing network is a large campus network. Which of the following layers are typically found on this type of campus network? (Select all that apply)

- A. Access
- B. Front
- C. Distribution
- D. Back
- E. Core

Answer: A, C, E

Explanation:

An enterprise campus network can be broken down to small, medium, and large locations. In most instances large campus locations will have a three-tier design with a wiring closet component (Ethernet access layer), a distribution layer, and core layer. Small campus locations will likely have a two-tier design with wiring closet component (Ethernet access layer) and a backbone core (collapsed core and distribution layers). Medium-sized campus network designs will sometimes use a three-tier implementation or a two-tier implementation depending on the number of ports, service requirements, manageability, performance, and availability levels that are required.

Section 4: Explain the function of the Switching Database Manager [specifically Content Addressable Memory (CAM) and Ternary Content Addressable Memory (TCAM)] within a Catalyst switch (1 question)

QUESTION NO: 1

You are troubleshooting a problem between two workstations (TK1 & TK2). Workstation TK1 is unable to ping workstation TK2. They are both connected to the same switch, the same VLAN, and to they're both in the same subnets. What should you do to verify connectivity? (Select two)

- A. Verify that the router for the VLAN is operational.
- B. Check the speed and duplex settings.
- C. Check both devices for proper default gateway settings.
- D. Check to see if the MAC addresses are in the CAM table.

Answer: B, D

Explanation:

Because the two workstations are physically connected to the same switch (which isn't necessarily required to be in the same VLAN), you can rule out the possibility of a compromised physical layer connection. If the speed and duplex settings on each device are mismatched then there could indeed be connectivity issues so B is correct. If for whatever reason the MAC address for TK2 isn't on the switches CAM table then the switch won't know the whereabouts of TK2 making the ping ineffective.

Incorrect Answers:

- A: Since the two switches are the on same VLAN there is no need to check the router, so A is incorrect.
- C: Since the two switches are both in the same subnet and the same VLAN the default gateway settings wouldn't be an issue, so C is incorrect.

Section 5: Describe the features and operation of VLANs on a switched network (6 questions)

QUESTION NO: 1

The TestKing LAN switches are being configured to support the use of Dynamic VLANs. Which of the following are true of dynamic VLAN membership? (Select all that apply)

- A. VLAN membership of a user always remains the same even when he/she is moved to another location.
- B. VLAN membership of a user always changes when he/she is moved to another location.
- C. Membership can be static or dynamic.
- D. Membership can be static only.
- E. None of the above.

Answer: A, C

Explanation:

Dynamic VLAN memberships are based on the users MAC address connected to the port. If you have VTP server, a VTP database file, a VTP client switch, and a dynamic port; regardless of where your physical location is, you can still remain in the same VLAN.

Incorrect Answers:

B: This was true before the use of Dynamic VLAN membership, as VLANs were assigned to ports, not users.

D: VLAN memberships can be either static or dynamic.

QUESTION NO: 2

The TestKing LAN switches are being configured to support the use of Dynamic VLANs. What should be considered when implementing a dynamic VLAN solution? (Select two)

- A. Each switch port is assigned to a specific VLAN.
- B. Dynamic VLANs require a VLAN Membership Policy Server.
- C. Devices are in the same VLAN regardless of which port they attach to.
- D. Dynamic VLAN assignments are made through the command line interface.

Answer: B, C

Explanation:

With VLAN Membership Policy Server (VMPS), you can assign switch ports to VLANs dynamically, based on the source Media Access Control (MAC) address of the device connected to the port. When you move a host from a port on one switch in the network to a port on another switch in the network, the switch assigns the new port to the proper VLAN for that host dynamically.

Note: There are two types of VLAN port configurations: static and dynamic.

Incorrect Answers

A: In a static VLAN, the administrator assigns switch ports to the VLAN, and the association does not change until the administrator changes the port assignment. However, this is not the case of dynamic VLANs.

D: The Command Line Interface is not used for dynamic VLAN assignments.

Reference: Cisco Online, Configuring Dynamic Port VLAN Membership with VMPS

QUESTION NO: 3

What is the preferred method of filtering bridged traffic within a VLAN?

- A. Ethernet maps
- B. Router ACLs
- C. VLAN access maps
- D. IP ACLs

Answer: C

Explanation:

VLAN ACLs or VLAN maps access-control all packets (bridged and routed). You can use VLAN maps to filter traffic between devices in the same VLAN.

Each VLAN access map can consist of one or more map sequences, each sequence with a match clause and an action clause. The match clause specifies IP, IPX, or MAC ACLs for traffic filtering and the action clause specifies the action to be taken when a match occurs. When a flow matches a permit ACL entry, the associated action is taken and the flow is not checked against the remaining sequences. When a flow matches a deny ACL entry, it will be checked against the next ACL in the same sequence or the next sequence. If a flow does not match any ACL entry and at least one ACL is configured for that packet type, the packet is denied.

Reference:http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a0080160a7e.html#wp1053799

QUESTION NO: 4

You are assigning VLANs to the ports of switch TK1. What VLAN number value is an assigned to the default VLAN?

- A. VLAN 1003

- B. VLAN 1
- C. VLAN ON
- D. VLAN A
- E. VLAN 0

Answer: B

Explanation: The default VLAN is VLAN 1. Although this VLAN can be modified, it can not be deleted from the switch. The following VLANs are on by default for all Cisco Catalyst switches:

- VLAN 1 – Default VLAN
- VLAN 1002 – Default FDDI VLAN
- VLAN 1003 – Default Token Ring VLAN
- VLAN 1004 – Default FDDI Net VLAN
- VLAN 1005 – Default Token Ring Net VLAN

Incorrect Answers:

- A: This is the default Token Ring VLAN that is installed in the switch IOS. It is seldom used.
- C: ON is a VTP configuration mode, but is not a normal VLAN name.
- D: Although any VLAN can be named VLAN A, it is not created by default.
- E: Although in Cisco IOS the number 0 has significance (i.e. ethernet 0, console port 0, serial 0) in VLANs 1 is the default. VLAN 0 is an invalid VLAN and can not be used.

QUESTION NO: 5

The VLANs in switch TK1 are being modified. Which of the following are updated in TK1 every time a VLAN is modified? (Select all that apply)

- A. Configuration revision number
- B. Configuration revision flag field
- C. Configuration revision reset switch
- D. Configuration revision database
- E. None of the above.

Answer: A, D**Explanation:**

For accountability reasons, every time a VLAN is modified the revision number changes, as does the information in the configuration revision database (as that is where the VLAN information is stored).

Incorrect Answers:

B, C: The configuration revision flag field, and the configuration revision reset switch don't exist in this context.

QUESTION NO: 6

If you needed to transport traffic coming from multiple VLANs (connected between switches), and your CTO was insistent on using an open standard, which protocol would you use?

- A. 802.11B
- B. spanning-tree
- C. 802.1Q
- D. ISL
- E. VTP
- F. Q.921

Answer: C

Explanation:

The act involved in the above question is trunking. The two trunking protocols in the answer choices are: 802.1Q and ISL. ISL is Cisco proprietary and IEEE 802.1Q is based on an open standard. When non-Cisco switches are used along with Cisco switches and trunking is required, it is best to use the 802.1Q encapsulation.

Incorrect Answers:

- A: This standard is used in wireless networking and has nothing to do with VLAN switching.
- B: The Spanning Tree Protocol (STP) is used to prevent loops within a bridged network. Each VLAN runs a separate instance of the STP and this is enabled by default.
- D: This is the alternative Cisco proprietary method of trunking.
- E: VLAN Trunking Protocol (VTP) is a Layer 2 messaging protocol that manages the addition, deletion, and renaming of VLANs on a network-wide basis. It is not used to actually transport VLAN traffic.
- F: This is an ISDN signalling standard and is not related with VLAN switching.

Section 6: Describe the features of the VLAN trunking protocols including 802.1Q, ISL (emphasis on 802.1Q) and dynamic trunking protocol (17 questions)

QUESTION NO: 1

Which of the following technologies would an Internet Service Provider use to support overlapping customer VLAN ID's over transparent LAN services?

- A. 802.1q tunneling
- B. ATM
- C. SDH
- D. IP Over Optical Networking
- E. ISL

Answer: A

Explanation:

Understanding How 802.1Q Tunneling Works:

The 802.1Q tunnelling feature supports secure virtual private networks (VPNs). 802.1Q tunnelling enables service providers to keep traffic from different customers segregated in the service provider infrastructure while significantly reducing the number of VLANs required to support the VPNs. 802.1Q tunnelling allows multiple customer VLANs to be carried by a single VLAN on the Catalyst 6000 family switch without losing their unique VLAN IDs.

When you configure 802.1Q tunnelling on the Catalyst 6000 family switch, traffic to be tunnelled comes into the switch from an 802.1Q trunk port on a neighboring device and enters the switch through a port configured to support 802.1Q tunnelling (a tunnel port). When the tunnel port receives traffic from an 802.1Q trunk port, it does not strip the 802.1Q tags from the frame header but, instead, leaves the 802.1Q tags intact and puts all the received 802.1Q traffic into the VLAN assigned to the tunnel port. The VLAN assigned to the tunnel port then carries the tunnelled customer traffic to the other neighboring devices participating in the tunnel port VLAN. When the tunnelled traffic is received by an 802.1Q trunk port on a neighboring device, the 802.1Q tag is stripped and the traffic is removed from the tunnel.

Reference:http://www.cisco.com/en/US/products/hw/switches/ps700/products_configuration_guide_chapter09186a008007fa06.html

QUESTION NO: 2

If you were to configure an ISL Ethernet trunk between two Cisco switches, named TK1 and TK2, what would you have to include at the end of the link for the trunk to operate correctly? (Select two)

- A. An identical VTP mode.
- B. An identical speed/duplex.
- C. An identical trunk negotiation parameter.
- D. An identical trunk encapsulation parameter.

Answer: B, D

Explanation:

In order for a trunk to be operational, the speed and duplex settings must match at each end of the trunk, and both switches must use the same trunking encapsulation (802.1Q or ISL).

Incorrect Answers:

A: It is common for switches to have trunk links operating, while the VTP modes differ. For example, a switch configured with VTP mode server can have a trunk connected to a switch with VTP mode client.

C: This is incorrect, as there are a number of configurations that are supported where the trunk negotiation parameters differ between switches. For example, switch TK1 could have the trunk configured for “on” while switch TK2 could have the switch trunk configured for “desirable” and the trunk would be operational.

QUESTION NO: 3

Drag-and-drop the technology term on the left to the correct options column on the right (not all of the options will be used.)

LANE	embedded VLAN tag
ISL	fiber links, FDDI
802.1Q	encapsulation frames
802.10	ATM
VLAN	
VMPS	

Answer:

LANE	- ATM
ISL	- Encapsulation frames
802.1Q	- embedded VLAN tag
802.10	- Fiber links, FDDI
VLAN	
VMPS	

Explanation:

- **LANE** - LAN Emulation – An IEEE standard method for transporting VLANs over Asynchronous Transfer Mode (ATM) networks.
- **ISL** – A Cisco Proprietary encapsulation protocol for interconnection multiple switches.
- **802.1Q** – An IEEE standard method for identifying VLANs by inserting a VLAN identifier into the frame header. This process is called frame tagging.
- **802.10** – A Cisco Proprietary method of transporting VLAN information inside the standard 802.10 frame (Fiber Distributed Data Interface [FDDI]).

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 99

QUESTION NO: 4

You are the network administrator at TestKing and switch TK1 is configured as shown below:

```
Interface gigethernet 0/1
Switchport mode trunk
Switchport trunk encapsulation dot1q
Switchport trunk native vlan 5
```

If untagged frames are arriving on interface gigethernet 0/1 of TK1, which of the following statement are correct?

- Untagged frames are automatically assumed to be in VLAN 5.
- Untagged frames are defaulted to VLAN 1 traffic.
- Untagged frames are dropped because all packets are tagged when dot1q trunked.
- Untagged frames are determined on the other switch
- Untagged frames are not supported on 802.1Q trunks.

Answer: A**Explanation:**

Each physical port has a parameter called PVID. Every 802.1Q port is assigned a PVID value that is of its native VLAN ID (default is VLAN 1). All untagged frames are assigned to the LAN specified in the PVID parameter. When a tagged frame is received by a port, the tag is respected. If the frame is untagged, the value contained in the PVID is considered as a tag. All untagged frames will be assigned to the native VLAN. The native VLAN is 1 by default, but in this case the native VLAN is configured as VLAN 5 so choice A is correct.

QUESTION NO: 5

If you were to set up a VLAN trunk over a Fast Ethernet link on switch TK1, which trunk mode would you set the local port to on TK1 if you wanted it to respond to requests from its link partner (TK2) and become a trunk?

- A. Auto
- B. Negotiate
- C. Designate
- D. Nonegotiate

Answer: A

Explanation:

Only ports in desirable and auto mode will negotiate a channel (either desirable-auto or desirable-desirable). Ports in on mode will only form a functional channel with other ports in on mode (they will not negotiate a channel with ports in desirable or auto mode).

Reference: Cisco, Troubleshooting Tips

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/trbl_ja.htm

QUESTION NO: 6

Which of the following trunking modes are unable to request their ports to convert their links into trunk links? (Select all that apply)

- A. Negotiate
- B. Designate
- C. Nonegotiate
- D. Auto
- E. Manual
- F. Off

Answer: C, D

Explanation:

Auto is a trunking mode but does not actively negotiate a trunk. It requires opposite side to be trunk or desirable, and will only respond to requests from the other trunk link. No-negotiate will configure the link to be unable to dynamically become a trunk; since no requests will be sent it will not respond to requests from other trunk links from a different switch.

Incorrect Answers:

A, B, E, F: These choices are wrong because they are not valid trunking modes

QUESTION NO: 7

ISL is being configured on a TestKing switch. Which of the following choices are true regarding the ISL protocol? (Select two)

- A. It can be used between Cisco and non-Cisco switch devices.
- B. It calculates a new CRC field on top of the existing CRC field.
- C. It adds 4 bytes of protocol-specific information to the original Ethernet frame.
- D. It adds 30 bytes of protocol-specific information to the original Ethernet frame.

Answer: B, D

Explanation:

ISL adds a total of 30bytes to the Ethernet frame. A 26 byte header (10bytes identifies the VLAN ID) and a 4 byte trailer (containing a separate CRC).

Incorrect Answers:

- A: This is incorrect because ISL is Cisco proprietary and can only be used on Cisco devices. For configuring a trunk to a non-Cisco switch, 802.1Q encapsulation should be used.
- C: This is incorrect because it is contradictory to D. 30 bytes are added with ISL, not 4 bytes. This choice describes what is used in 802.1Q frames, not ISL

QUESTION NO: 8

You are the network administrator tasked with designing a switching solution for the TestKing network. Which of the following statements describing trunk links are INCORRECT? (Select all that apply)

- A. The trunk link belongs to a specific VLAN.
- B. Multiple trunk links are used to connect multiple devices.
- C. A trunk link only supports native VLAN.
- D. Trunk links use 802.10 to identify a VLAN.
- E. The native VLAN of the trunk link is the VLAN that the trunk uses if that link fails for any reason.

Answer: A, B, C, D

Explanation:

A trunk is a point-to-point link that transmits and receives traffic between switches or between switches and routers. Trunks carry the traffic of multiple VLANs and can extend VLANs across an entire network. 100BaseT and Gigabit Ethernet trunks use Cisco ISL (the default protocol) or industry-standard IEEE 802.1Q to carry traffic for multiple VLANs over a single link. Frames received from users in the administratively-defined VLANs are identified or tagged for transmission to other devices. Based on rules you define, a unique identifier (the tag) is inserted in each frame header before it is forwarded. The tag is examined and understood by each device before any broadcasts or transmission to other switches, routers, or end stations. When the frame reaches the last switch or router, the tag is removed before the frame is transmitted to the target end station.

Incorrect Answers:

E: This statement is true, as untagged frames are always used with the native VLAN. The native VLAN is VLAN 1 by default in Cisco switches.

QUESTION NO: 9

A TestKing switch port is configured as a trunk using 802.1Q encapsulation. Which three statements regarding the IEEE 802.1Q standard are true? (Select three)

- A. The packet is encapsulated with a 26 byte header and a 4 byte FCS.
- B. The IEEE 802.1Q frame format adds a 4 byte field to an Ethernet frame.
- C. The IEEE 802.1Q frame retains the original MAC destination address.
- D. The IEEE 802.1Q frame uses multicast destination of 0x01-00-0c-00-00
- E. The 802.1Q protocol uses point-to-point connectivity.
- F. The 802.1Q protocol uses point-to-multipoint connectivity.

Answer: B, C, E

Explanation:

802.1Q frames add 4 bytes to the Ethernet frame. The original MAC address is left unaltered so the destination MAC is not changed. Trunks are always defined in a point to point configuration, with two switch ports used as the endpoints.

Incorrect Answers:

- A: This describes the frame that is added to an ISL encapsulated frame, not an 802.1Q frame.
- D: The destination MAC address is not altered when trunks are configured.
- F: All trunks are always configured in a point to point fashion, there is no method available to support point to multipoint trunk configurations.

QUESTION NO: 10

Which DTP switchport mode parameter would you use to set a switch port to actively send and respond to DTP negotiation frames on switch TK1?

- A. access
- B. trunk
- C. no negotiate
- D. dynamic desirable
- E. dynamic auto

Answer: D

Explanation:

There are five DTP switchport modes, and you should be familiar with all of them.

Access: This 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.

Dynamic Desirable: The interface actively attempts 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.

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.

QUESTION NO: 11

A switch port on TK1 is being configured to support 802.1Q trunking. Which of the following are true about 802.1Q trunking? (Select one)

- A. Both switches must be in the same VTP domain.
- B. The encapsulation type of both ends of the trunk does not have to match.
- C. The native VLAN on both ends of the trunk must be VLAN 1.
- D. 802.1Q trunking can only be configured on a Layer 2 port.
- E. In 802.1Q trunking, all VLAN packets are tagged on the trunk link, except the native VLAN.

Answer: E

Explanation:

E is correct because, “frames from the native VLAN of an 802.1Q trunk are not tagged with the VLAN number.”

Incorrect Answers:

B: This is incorrect because the encapsulations types do have to match or it won't work properly. You can't use 802.1Q on one side and ISL on the other. C is incorrect because the native VLAN doesn't necessarily have to be VLAN 1.

C: By default, the native VLAN is VLAN 1 but this can be effectively changed to a different VLAN and the trunk will still be functional.

Reference: <http://www.cisco.com/warp/public/473/27.html>

QUESTION NO: 12


A TestKing switch is configured for 802.1Q trunking. What are valid characteristics of IEEE 802.1Q? (Select all that apply)

- A. Use frame tagging.
- B. None of the answers
- C. It is a method for identifying VLANs
- D. It inserts VLAN identifier into the frame header

Answer: A, C, D

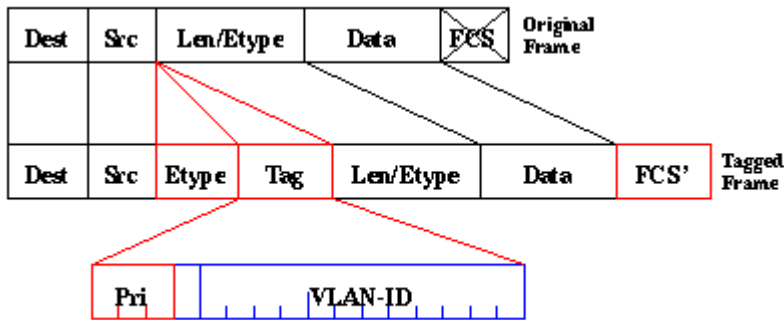
Explanation:

802.1Q uses an internal tagging mechanism. Internal means that a tag is inserted within the frame (with ISL, the frame is encapsulated instead):



Note that on an 802.1Q trunk, one VLAN is NOT tagged. This VLAN, named the native VLAN, must be configured the same on each side of the trunk. This way, we can deduce to which VLAN a frame belongs when we receive a frame with no tag.

The tagging mechanism implies a modification of the frame; the trunking device inserts a 4-byte tag and recomputes the frame check sequence (FCS):



The EtherType field identifying the 802.1Q frame is 0x8100. In addition to the 12-bit VLAN-ID, 3 bits are reserved for 802.1p priority tagging.

Also, note that inserting a tag into a frame that already has the maximum Ethernet size creates a 1522 byte frame that can be considered as a "baby giant" by the receiving equipment. The 802.3 committee is extending the maximum standard frame size to address this issue.

QUESTION NO: 13

What are the reasons as to why an administrator would deploy Dynamic Trunking Protocol (DTP) on the TestKing switched LAN? (Select all that apply)

- A. For supporting auto-negotiation of IEEE 802.1Q trunks
- B. For supporting auto-negotiation of ISL
- C. For managing trunk negotiation in 2500 router supervisor engine software R 4.2 or later
- D. For managing trunk negotiation in Catalyst supervisor engine software R 4.2 or later.
- E. None of the above.

Answer: A, B, D

Explanation:

DTP was developed for the specific purpose of supporting automatic trunk negotiation for 802.1Q and ISL trunks. It is used only with Cisco Catalyst switches.

Incorrect Answers:

DTP is supported only on Cisco Catalyst switches. It is not supported on Cisco 2500 series routers.

QUESTION NO: 14

A fast Ethernet port on switch TK1 is configured as a trunk. What is true of this trunk link?

- A. A trunk link only supports the native VLAN for a given port.

- B. A trunk link uses 802.10 to identify VLANs over an Ethernet backbone.
- C. A trunk link connects multiple devices on a single subnet to a switch port.
- D. The native VLAN of the trunk link is the VLAN to which the port will belong if that link becomes non-trunk.
- E. All of the above.

Answer: C

Explanation:

Trunks are used to connect multiple VLANs together. Individual switches configured with VLANs over the entire LAN subnet are connected to each other via a trunk port.

Incorrect Answers:

A: By default all VLANs within the range of 1 to 1000 is allowed to traverse the trunk port.

B: 802.10 is the standard used on FDDI networks and is not related to Ethernet VLAN trunks.

D: This is wrong because Native VLAN Number of the native VLAN for the trunk link (for 802.1Q trunks, the VLAN for which untagged traffic can be transmitted and received over the trunk; for ISL trunks, packets are tagged on all VLANs, including the native VLAN).

QUESTION NO: 15

You are a technician at TestKing and your newly appointed trainee asks you what the Dynamic Trunking Protocol (DTP) mode 'desirable' means. What would your reply be?

- A. The interface is put into permanent trunking mode but prevented from generating DTP frames.
- B. The interface actively attempts to convert the link to a trunk link.
- C. The interface is put into a passive mode, waiting to convert the link to a trunk link.
- D. The interface is put into permanent trunking mode and negotiates to convert the link into a trunk link.

Answer: B

Explanation:

The DTP mode of desirable configured the trunk port to actively attempt to convert the link to a trunk link.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 105

QUESTION NO: 16

What would happen to a frame if a VLAN port configured as a trunk on the Catalyst switch TK1 were to receive an untagged frame?

- A. The frame will cause an error message to be sent.

- B. The frame will be dropped.
- C. The frame will be processed as a native VLAN frame.
- D. The frame will be tagged, and then processed as a native VLAN frame.

Answer: C

Explanation:

On an IEEE 802.1Q trunk port, all transmitted and received frames are tagged except for those on the VLAN configured as the native VLAN for the port. Frames on the native VLAN are always transmitted untagged and are normally received untagged. The default native VLAN is VLAN 1.

Reference:

http://www.cisco.com/en/US/products/hw/optical/ps2006/products_module_configuration_guide_chapter09186a0080154a4a.html

QUESTION NO: 17

Switch TK1 has a trunk link configured with IEEE 802.1Q encapsulation. What is the maximum Ethernet frame size on this trunk port?

- A. 1496 Bytes
- B. 1500 Bytes
- C. 1518 Bytes
- D. 1522 Bytes
- E. 1548 Bytes

Answer: D

Explanation:

The 802.1q tag is 4 bytes; hence the resulting ethernet frame can be as large as 1522 bytes (1518 for the maximum Ethernet frame size plus the 4 byte 802.1Q tag). The minimum size of the Ethernet frame with 802.1q tagging is 68 bytes.


Reference:

http://www.cisco.com/en/US/tech/tk389/tk390/technologies_tech_note09186a0080094665.shtml

Section 7: Describe the features and operation of 802.1Q Tunneling (802.1QinQ) within a service provider network (4 questions)

QUESTION NO: 1

The TestKing Enterprise network is depicted in the following topology exhibit:



An ISP is currently providing services to the enterprise TestKing.com. TestKing.com is composed of a series of VLANs ranging from 1 to 15.

Enterprise Acme.com is currently configured to support VLANs with the range of 1 to 25 and they're requesting support from the same ISP.

Which Layer 2 technology should the ISP use to keep the traffic segmented so Enterprise TestKing.com and Enterprise Acme.com so that each can maintain their current VLAN configurations?

- A. Transparent LAN services
- B. Metro Ethernet over DWDM
- C. 802.1Q in Q
- D. EoMPLS
- E. L2TPv3

Answer: C

Explanation:

Encapsulating IEEE 802.1Q VLAN tags within 802.1Q enables service providers to use a single VLAN to support customers who have multiple VLANs. The IEEE 802.1Q-in-Q VLAN Tag Termination feature on the subinterface level preserves VLAN IDs and keeps traffic in different customer VLANs segregated.

QUESTION NO: 2

The TestKing campus is utilizing the services of a Metro Ethernet provider. What is the function of VLAN tunnelling in this Metro Ethernet environment?

- A. Renumbers their LANs.
- B. Extends their logical network topology across wide geographic networks.
- C. Provides combined wavelength routing.
- D. Translates their VLANs at the service provider edge.

Answer: D

Explanation:

An ideal scenario to support multiple customers in the service provider environment would be to have customers utilizing any range of VLAN numbers while the service provider forwards the traffic independent of those VLAN IDs. By assigning a unique VLAN to each customer, the identity of multiple VLAN IDs from the customer site will not be lost. This builds a Layer 2 VPN where traffic from different business customers is segregated inside the service provider core and is dot1q tagged with appropriate VLAN IDs. Dot1q tunneling is in essence a 1q-in-1q technique that expands the VLAN space by retagging the tagged packets entering the service provider infrastructure. This is used to translate the customer's VLAN information at the provider edge, ensuring customer traffic separation will keeping the original VLAN information intact.

Reference:

http://www.cisco.com/en/US/netsol/ns110/ns221/ns223/ns227/networking_solutions_white_paper09186a00800a1195.shtml

QUESTION NO: 3

The TestKing campus is utilizing the services of a Metro Ethernet provider, which is using 802.1Q-in-Q. What's true about this Metro 802.1Q-in-Q model? (Select all that apply.)

- A. Customer VLAN traffic is isolated from the service provide network's VLAN traffic.
- B. Quality of service can be easily implemented using the Customer's ToS and CoS.

- C. It has limited scalability in a service provider WAN.
- D. Customer traffic retains original VLAN tags.
- E. It provides efficient Layer 3 access.
- F. It can connect disparate customer networks (Frame Relay, Ethernet, ATM, etc).

Answer: A, C, D

Explanation:

Encapsulating IEEE 802.1Q VLAN tags within 802.1Q enables service providers to use a single VLAN to support customers who have multiple VLANs. The IEEE 802.1Q-in-Q VLAN Tag Termination feature on the subinterface level preserves VLAN IDs and keeps traffic in different customer VLANs segregated. This is used to translate the customer's VLAN information at the provider edge, ensuring customer traffic separation will keeping the original VLAN information intact.

Incorrect Answers:

- B: QoS services become difficult to manage and implement in metro Ethernet networks.
- E: Layer 2 access is granted, since the customer handoff from the service provider is an Ethernet port.
- F: This is used to connect Ethernet-only networks within a Metropolitan area Network, and other encapsulation types (ATM, Frame-Relay, PPP, HDLC, etc) are not supported.

QUESTION NO: 4

You are a network administrator at TestKing and the CTO has asked you to allow a customer's LAN traffic to be transmitted on a single VLAN across multiple service provider networks. Which technology would you use?

- A. Transparent LAN Services
- B. Metro network segmentation
- C. VLAN tunneling
- D. OC 192
- E. DWDM

Answer: C

Explanation:

Virtual private networks (VPNs) provide enterprise-scale connectivity on a shared infrastructure, often Ethernet-based, with the same security, prioritization, reliability, and manageability requirements of private networks. Tunneling is a feature designed for service providers who carry traffic of multiple customers across their networks and are required to maintain the VLAN and Layer 2 protocol configurations of each customer without impacting the traffic of other customers. VLAN tunneling is used to translate the customer's VLAN information at the provider edge, ensuring customer traffic separation will keeping the original VLAN information intact.

Reference:

http://www.cisco.com/en/US/products/hw/optical/ps2006/products_module_configuration_guide_chapter09186a008014f544.html#19505

www.testking.com

Section 8: Describe the operation and purpose of managed VLAN services (3 questions)**QUESTION NO: 1****What are the reasons for deploying VLANs? (Select all that apply)**

- A. to address the addition of network management through layer 3 routing protocols.
- B. to address the redundancy issues of a flat network topology
- C. to address the performance issues of a non-flat network topology
- D. to address the scalability issues of a flat network topology

Answer: A, D**Explanation:**

There are several benefits to using VLANs. In summary, VLAN architecture benefits include:

- Increased performance
- Improved manageability
- Network tuning and simplification of software configurations
- Physical topology independence
- Increased security options

Increased performance

Switched networks by nature will increase performance over shared media devices in use today, primarily by reducing the size of collision domains. Grouping users into logical networks will also increase performance by limiting broadcast traffic to users performing similar functions or within individual workgroups. Additionally, less traffic will need to be routed, and the latency added by routers will be reduced.

Improved manageability

VLANs provide an easy, flexible, less costly way to modify logical groups in changing environments. VLANs make large networks more manageable by allowing centralized configuration of devices located in physically diverse locations.

Network tuning and simplification of software configurations

VLANs will allow LAN administrators to "fine tune" their networks by logically grouping users. Software configurations can be made uniform across machines with the consolidation of a department's resources into a single subnet.

Physical topology independence

VLANs provide independence from the physical topology of the network by allowing physically diverse workgroups to be logically connected within a single broadcast domain. If the physical infrastructure is already in place, it now becomes a simple matter to add ports in new locations to existing VLANs if a department expands or relocates.

VLANs have the ability to provide additional security not available in a shared media network environment. By nature, a switched network delivers frames only to the intended recipients, and broadcast frames only to other members of the VLAN. This allows the network administrator to segment users requiring access to sensitive information into separate VLANs from the rest of the general user community regardless of physical location. In addition, monitoring of a port with a traffic analyzer will only view the traffic associated with that particular port, making discreet monitoring of network traffic more difficult.

QUESTION NO: 2

The TestKing network consists of numerous VLANs. Which two statements characterize VLANs? (Select two)

- A. All hosts in the same VLAN are in the same broadcast domain.
- B. All hosts in the same VLAN are in the same collision domain
- C. VLAN membership is based upon port membership or assigned dynamically based on MAC.
- D. It is a physical network segment.
- E. They must be created in the vlan database mode on all IOS and Cat OS based Catalyst Switches.

Answer: A, C

Explanation:

At the most basic level, a VLAN is nothing more than a broadcast domain. The only difference between a traditional broadcast domain and one defined by a VLAN is that traditionally a broadcast domain has been seen as a distinct physical entity whose boundaries consist of a router. In fact, VLANs are very similar – their boundaries are also defined by a routing device, just like any broadcast domain. However, a VLAN is a logical construct, meaning that hosts are not necessarily groups within the physical confines of a traditional broadcast domain.

Switched networks by nature will increase performance over shared media devices in use today, primarily by reducing the size of collision domains. Grouping users into logical networks will also increase performance by limiting broadcast traffic to users performing similar functions or within individual workgroups. Additionally, less traffic will need to be routed, and the latency added by routers will be reduced.

A host can only become a member of any particular VLAN via one of two methods: static VLAN membership or dynamic VLAN membership. Static VLANs are configured at the port level, meaning that when a switch port is configured for a particular VLAN, any device plugged into that port becomes part of that VLAN. Dynamic VLAN membership is performed at the MAC level. No matter where a user is plugged in, the information pertaining to the MAC address of the station is used (normally via a VMPS server) to determine which VLAN the station belongs to.

QUESTION NO: 3

The TestKing network is in need of a robust network management application. What is Cisco's flagship web based device management tool?

- A. CWOS
- B. CiscoView
- C. CWIS
- D. VLANDirector
- E. Traffic Director
- F. NetSYS
- G. Open View
- H. All of the above

Answer: B

Explanation:

CiscoView is a member of the CiscoWorks2000 family, a Web based device management application providing dynamic status, monitoring, and configuration information for the broad range of Cisco internetworking products. CiscoView displays a physical view of a device chassis, with color-coding of modules and ports for at-a-glance status. Monitoring capabilities display performance and other statistics. Configuration capabilities allow comprehensive changes to devices, given requisite security privileges are granted.

Section 9: Describe how VTP versions 1 and 2 operate including domains, modes, advertisements, and pruning (26 questions)

QUESTION NO: 1

VTP is configured on switch TK1. Which of the following features were added in VTP version 2 that were not previously supported in VTP version 1? (Select two)

- A. Supports Token Ring VLANs.
- B. Allows VLAN consistency checks.
- C. Saves VLAN configuration memory.
- D. Reduces the amount of configuration necessary.
- E. Allows active redundant links when used with spanning tree.

Answer: A, B

Explanation:

VTP Version 2 includes the following improvements: Token Ring VLAN support, TLV support, transparent mode, and Consistency checks.

Incorrect Answers:

C, D: These were not improvements added to VTP Version 2.

E: STP detects and prevents loops by logically disabling the redundant path ports so there are no active redundant links.

QUESTION NO: 2

The TestKing switches are configured to use VTP. What's true about the VLAN trunking protocol (VTP)? (Select two)

- A. VTP messages will not be forwarded over nontrunk links.
- B. VTP domain names need to be identical. However, case doesn't matter.
- C. A VTP enabled device which receives multiple advertisements will ignore advertisements with higher configuration revision numbers.
- D. A device in "transparent" VTP v.1 mode will not forward VTP messages.
- E. VTP pruning allows switches to prune VLANs that do not have any active ports associated with them.

Answer: A, D

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

VTP messages are only transmitted across trunk links.

If the receiving switch is in transparent mode, the configuration is not changed. Switches in transparent mode do not participate in VTP. If you make VTP or VLAN configuration changes on a switch in transparent mode, the changes are not propagated to the other switches in the network.

Incorrect Answers:

B: The VTP domain name is case sensitive and it must be identical with the domain name configured on the VTP server.

C: This is incorrect because if a VTP client receives an advertisement with a higher revision number, it won't ignore it. In fact, the advertisement with a higher revision level takes precedence when the switch is configured in client mode.

E: VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices. It does not prune the individual VLANs.

QUESTION NO: 3

Switch TK1 and TK2 both belong to the TestKing VTP domain. What's true about the switch operation in VTP domains? (Select all that apply)

- A. A switch can only reside in one management domain
- B. A switch is listening to VTP advertisements from their own domain only
- C. A switch is listening to VTP advertisements from multi domains
- D. A switch can reside in one or more domains
- E. VTP is no longer supported on Catalyst switches

Answer: A, B

Explanation:

A VTP domain is made up of one or more interconnected devices that share the same VTP domain name. A switch can be configured to be in only one VTP domain, and each VLAN has a name that is unique within a management domain.

Typically, you use a VTP domain to ease administrative control of your network or to account for physical boundaries within your network. However, you can set up as many or as few VTP domains as are appropriate for your administrative needs. Consider that VTP is transmitted on all trunk connections, including ISL, IEEE 802.1Q, 802.10, and LANE.

Switches can only belong to one management domain with common VLAN requirements, and they only care about the neighbors in their own domains.

Reference: CCNP Switching Exam Certification Guide: David Hucaby & Tim Boyles, Cisco Press 2001, ISBN 1-58720 000-7 page 114

QUESTION NO: 4

VTP devices in a network track the VTP revision number. What is a VTP configuration revision number?

- A. A number for identifying changes to the network switch.
- B. A number for identifying changes to the network router.
- C. A number for identifying changes to the network topology.
- D. None of the above.

Answer: C

Explanation:

The configuration revision number is a 32-bit number that indicates the level of revision for a VTP packet. Each VTP device tracks the VTP configuration revision number assigned to it, and most of the VTP packets contain the VTP configuration revision number of the sender.

This information is used to determine whether the received information is more recent than the current version. Each time you make a VLAN change in a VTP device, the configuration revision is incremented by one. In order to reset the configuration revision of a switch, change the VTP domain name and then change it back to the original name.

Incorrect Answers:

A: Not all switch configuration changes will impact the VTP revision number. Only changes made to the VLAN configuration will cause an increment in the revision number.

B: VTP revision numbers are only used on network switches configured for VTP and are not used by Cisco routers.

Reference: [Understanding and Configuring VLAN trunk protocol \(VTP\)](http://www.cisco.com/warp/public/473/21.html) Document ID: 10558

QUESTION NO: 5

Switch TK1 is configured to use the VLAN Trunking Protocol (VTP). What does TK1 advertise in its VTP domain?

- A. The VLAN ID of all known VLANs, the management domain name, and the total number of trunk links on the switch.
- B. The VLAN ID of all known VLANs, a 1-bit canonical format (CF1 Indicator), and the switch configuration revision number.
- C. The management domain name, the switch configuration revision number, the known VLANs, and their specific parameters.
- D. A 2-byte TPID with a fixed value of 0x8100 for the management domain number, the switch configuration revision number, the known VLANs, and their specific parameters.
- E. None of the above.

Answer: C

Explanation:

“Each switch participating in VTP advertises VLAN information, revision numbers, and VLAN parameters on its trunk ports to notify other switches in the management domain. VTP advertisements are sent as multicast frames. The switch intercepts frames sent to the VTP multicast address and processes them with its supervisory processor VTP frames are forwarded out trunk links as a special case.

The following global configuration information is distributed in VTP advertisements:

- VLAN IDs (ISL and 802.1Q)
- Emulated LAN names (for ATM LANE)
- 802.10 SAID values (FDDI)
- *VTP domain name*
- *VTP configuration revision number*
- VLAN configuration, including maximum transmission unit (MTU) size for each VLAN
- Frame format

Reference: CCNP Switching Exam Certification Guide: page 115, David Hucaby & Tim Boyles, Cisco Press 2001, ISBN 1-58720 000-7

Incorrect Answers:

A: The total number of trunk links is not advertised.

B: A CFI is not advertised.

D: The TPID is not advertised. The value of 0x8100 is used to identify an 802.1Q trunking tag.

QUESTION NO: 6

VTP switches use advertisements to exchange information with each other. Which of the following advertisement types are associated with VTP? (Select all that apply)

- A. Domain advertisements
- B. Advertisement requests from clients
- C. Subset advertisements
- D. Summary advertisements

Answer: B, C, D

Explanation:

VTP advertisements include:

- Summary Advertisements - These go out every 5 minutes or every time the VLAN topology changes, and lists of information about the management domain (VTP version, domain name, configuration revision number, timestamp, MD5 encryption hash code, & number of subset advertisements incoming). When there is a configuration change, summary advertisements are complemented by or more subset advertisements.

- Subset advertisements – These are sent out by VTP domain servers after a configuration change. They list the specifics of the change (VLAN creation / deletion / suspension / activation / name change / MTU change) and the VLAN parameters (VLAN status, VLAN type, MTU, VLAN name, VLAN number, SAID value).
- Advertisement Requests from Clients- VTP clients request specific VLAN information that they're lacking (ie. Client switch is reset and loses its database, or VTP domain membership changes) so they can be responded by summary and subset advertisements.

Reference: CCNP Switching Exam Certification Guide: pages 116-117 David Hucaby & Tim Boyles, Cisco Press 2001, ISBN 1-58720 000-7

QUESTION NO: 7

Switch TK1 is part of the TestKing VTP domain. What's true of VTP Pruning within this domain? (Select all that apply)

- A. it does not prune traffic from VLANs that are pruning-ineligible
- B. VLAN 1 is always pruning-eligible
- C. it will prune traffic from VLANs that are pruning-ineligible
- D. VLAN 2 is always pruning-ineligible
- E. None of the above.

Answer: A, B

Explanation:

By definition, pruning-ineligible VLANs can not be pruned. You can make specific VLANs pruning ineligible with the **clear vtp pruneeligible vlan_range** command. By default, VLANs 2-1000 are pruning-eligible. Since the default VLAN for any switch port in a Catalyst switch is VLAN 1, it is not eligible for pruning.

Incorrect Answers:

- C: The opposite is true.
- D: By default, VLANs 2-1000 are eligible to be pruned.

QUESTION NO: 8

What action should you execute if you wanted to enable VTP pruning on your entire management domain?

- A. Enable VTP pruning on any client switch in the management domain.
- B. Enable VTP pruning on any switch in the management domain.
- C. Enable VTP pruning on every switch in the management domain.
- D. Enable VTP pruning on a VTP server in the management domain.

E. Disable VTP pruning on a VTP server in the management domain.

Answer: D

Explanation:

Enabling VTP pruning on a VTP server allows pruning for the entire management domain. Enabling this on the VTP server will mean that the VTP pruning configuration will be propagated to all VTP client switches within the domain. VTP pruning takes effect several seconds after you enable it. By default, VLANs 2 through 1000 are pruning-eligible.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 117

QUESTION NO: 9

Switch TK1 is configured with VTP. Which two VTP modes will make TK1 capable of creating and deleting VLANs on itself? (Select two)

- A. Client
- B. Server
- C. Transparent
- D. Pass-through
- E. No-negotiate

Answer: B, C

Explanation:

VTP Modes

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) 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. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive out their trunk interfaces.

If you configure the switch as VTP transparent, you can create and modify VLANs but the changes affect only the individual switch.

Incorrect Answers:

A: Clients can not modify, add, or delete any VLAN information.

D, E: These options are not valid VTP modes.

QUESTION NO: 10

When the Catalyst switch TK1 is enabled to use VTP, which information does it advertise on its trunk ports? (Select two)

- A. VTP mode
- B. STP root status
- C. Negotiation status
- D. Management domain
- E. Configuration revision number

Answer: D, E

Explanation:

The VTP protocol maintains VLAN configuration consistency throughout the network by distributing VLAN information to the network. VLAN information is sent to network devices in advertisements that contain the VTP management domain name, the current configuration revision number, the VLANs that the server knows about, and certain VLAN parameters. Any time you change a VLAN, VTP automatically sends an advertisement to update all other network devices.

The following global configuration information is distributed in VTP advertisements:

- VLAN IDs (ISL and 802.1Q)
- Emulated LAN names (for ATM LANE)
- 802.10 SAID values (FDDI)
- *VTP domain name*
- *VTP configuration revision number*
- VLAN configuration, including maximum transmission unit (MTU) size for each VLAN
- Frame format

Reference: Cisco, Configuring VTP

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_6_1/config/vtp.htm

QUESTION NO: 11

Is the following statement True or False?

MLS requires that MLS components be in the same VTP domain.

- A. False
- B. There is not enough information to determine
- C. True

D. It could be true or false, depending on the type of switch.

Answer: C

Explanation:

MLS requires that MLS components, including the end stations, must be in the same Virtual Trunking Protocol (VTP) domain. VTP is a Layer 2 protocol used for managing VLANs on several Catalyst switches from a central switch. It allows an administrator to create or delete a VLAN on all switches in a domain without having to do so on every switch in that domain. The MultiLayer Switching Protocol (MLSP), which the MLS-SE and the MLS-RP use to communicate with one another, does not cross a VTP domain boundary.

QUESTION NO: 12

Which of the following VTP modes receives and forwards VTP updates, but does NOT participate in VTP synchronization?

- A. Client
- B. Server
- C. Transparent
- D. Pass-through

Answer: C

Explanation:

Transparent—VTP transparent network devices do not participate in VTP. A VTP transparent network device does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2, transparent network devices do forward VTP advertisements that they receive out their trunking LAN ports.

QUESTION NO: 13

How can VTP pruning enhance network bandwidth?

- A. By limiting the spreading of VLAN information.
- B. By reducing unnecessary flooding of traffic to inactive VLANs.
- C. By disabling periodic VTP updates.
- D. By restricting unicast traffic to across VTP domains.
- E. By updating unicast traffic periodically.

Answer: B

Explanation:

VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices.

Reference:

http://www.cisco.com/en/US/products/hw/routers/ps368/products_configuration_guide_chapter09186a0080091636.html#wp1020444

QUESTION NO: 14

Which of the following statements is NOT true regarding VTP?

- A. Switches in VTP transparent mode will simply forward advertisements without processing them.
- B. VTP reduces administrative overhead.
- C. VTP pruning reduces overall network traffic.
- D. VTP pruning is on by default.
- E. All of the above are true statements.

Answer: D

Explanation:

By default, VTP pruning is disabled.

For VTP pruning to be effective, all devices in the management domain must either support VTP pruning or, on devices that do not support VTP pruning, you must manually configure the VLANs allowed on trunks.

Incorrect Answers:

A: This statement is true. Transparent VTP switches do not participate in the VTP process, but they do forward VTP information to other switches.

B, C: VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices.

QUESTION NO: 15

By what condition can a VTP version 2 switch operate in the same domain as a switch running VTP version 1?

- A. VTP version 2 is disabled on the VTP version 2-capable switch
- B. VTP version 2 is enabled on the VTP version 2-capable switch
- C. VTP version 1 is disabled on the VTP version 2-capable switch
- D. None of the above. VTP version 1 and version 2 are not compatible.

Answer: A

Explanation:

A VTP version 2-capable switch can operate in the same VTP domain as a switch running VTP version 1 provided VTP version 2 is disabled on the VTP version 2-capable switch (VTP version 2 is disabled by default). With VTP version 2 disabled, the switch will revert to version 1 to become backward compatible.

QUESTION NO: 16

Is the following statement True or False?

If you modified a VTP transparent switch, the changes you implement will affect all the switches in the network?

- A. True
- B. There is not enough information to determine
- C. False

Answer: C

Explanation:

According to Cisco:

If you configure the switch as VTP transparent, you can create and modify VLANs but the changes affect only the individual switch. Transparent switches do not participate in VTP. Only changes made to a VTP switch in server mode will be propagated to all the other client switches within the network.

QUESTION NO: 17

Which of the following message types are associated with VTP header fields on a Cisco switched network? (Select all that apply)

- A. Summary advertisements
- B. Advertisement requests
- C. VTP Join messages
- D. Subset advertisement
- E. None of the above.

Answer: A, B, C, D

Explanation:

The format of the VTP header can vary depending on the type of VTP message. However, they all contain the following fields in the header:

VTP protocol version: 1 or 2

VTP Message Types: Summary advertisements, Subset advertisement, Advertisement requests, VTP join messages, Management domain length, and Management domain name.

QUESTION NO: 18

In order for the TestKing network to use VTP, which of the following conditions have to be met? (Select all that apply)

- A. Trunking must be enabled between all Catalyst switches.
- B. The Catalyst switches must be non-adjacent for trunking to be possible between them
- C. The Catalyst switches must be adjacent.
- D. Each Catalyst switch in a domain should be assigned the same VTP domain name.

Answer: A, C, D

Explanation:

According to the online documentation provided by Cisco:

In order to use VTP, you must assign a VTP domain name to each switch. VTP information will remain only within the same VLAN domain. The following are conditions for a VTP domain:

- Each Catalyst switch in a domain should be assigned the same VTP domain name.
- The Catalyst switches must be adjacent.
- Trunking must be enabled between all Catalyst switches.

If any one of the previous conditions is not met, the VTP domain is broken and information will not travel between the two separate parts.

QUESTION NO: 19

The TestKing switches are all VTP enabled. What is true about VTP? (Select all that apply)

- A. VTP version 2 is supported in supervisor engine software release 3.1(1) and later.
- B. you must decide whether to use VTP version 1 or version 2.
- C. VTP version 1 is supported in supervisor engine software release 2.1 or later
- D. VTP version 1 is supported in ATM software release 3.1 or later.

Answer: A, B, C, D

Explanation:

According to Cisco Documentation:

If you use VTP in your network, you must decide whether to use VTP version 1 or version 2. VTP version 1 is supported in supervisor engine software release 2.1 or later and ATM software release 3.1 or later. VTP version 2 is supported in supervisor engine software release 3.1(1) and later.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps679/products_configuration_guide_chapter09186a008007d73d.html

QUESTION NO: 20

In the TestKing switched network VTP pruning has been enabled. What is the purpose of VTP pruning?

- A. Enhancing network integrity
- B. Enhancing network bandwidth use
- C. Deploying AAA
- D. Enhancing network security
- E. Enhancing network load balancing

Answer: B

Explanation:

According to Cisco:

VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices.

QUESTION NO: 21

Switches TK1 and TK2 are both configured for transparent mode in the VTP domain. Which statement accurately describes these transparent VTP switches? (Select all that apply):

- A. They do not synchronize VLAN configuration based on received advertisements
- B. They do not participate in VTP
- C. They do not advertise VLAN configuration
- D. None of the above

Answer: A, B, C

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

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. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive out their trunk ports.

QUESTION NO: 22

Is the following statement True or False?

When VTP v.2 is enabled on a switch, all of the version 2-capable switches in the domain power cycle (restart) automatically?

- A. False
- B. There is not enough information to determine
- C. True

Answer: A

Explanation:

According to Cisco:

Although caution should be taken when enabling VTP version 2 on a switch, doing so will not cause all switches to power cycle. When you enable VTP version 2 on a switch, all of the version 2-capable switches in the domain enable VTP version 2. However, it will not cause all of the switches to reboot themselves.

QUESTION NO: 23

You are configuring switch TK1 for VTP. Which of the following are valid VTP operating modes that can be configured on TK1? (Select all that apply)

- A. Server
- B. Frontend
- C. Client
- D. Transparent
- E. Backbone

Answer: A, C, D

Explanation:

There are only three VTP operating modes:

1) **Server:** These switches have full control in the creation and modification of VLANs. Servers advertise out all the VLAN information they receive, and they configure themselves in accord with whatever information they hear. Switches are in server mode by default.

2) **Client:** These switches listen to VTP advertisements, they modify their configuration as a result of what they hear, and they forward out VTP information to neighbouring switches; but they don't have the ability to: create, change, or delete VLANs.

3) **Transparent:** These switches don't participate in the VTP process. They don't advertise their VLAN configurations and they don't synchronize their database when they receive advertisements. In VTP version 1 a switch doesn't relay information it gets to the other switches but in VTP version 2 they do.

QUESTION NO: 24

The TestKing switches are converting their VTP versions from 1 to version 2. Which of the following describe a benefit of using VTP v.2 over VTP v.1?

- A. To save VLAN configuration memory
- B. To reduce broadcast traffic carried on trunk lines.
- C. To reduce the amount of configuration necessary.
- D. To support token ring VLANs
- E. None of the choices.

Answer: D

Explanation:

VTP version 2 supports the following features not supported in version 1:

- Token Ring support—VTP version 2 supports Token Ring LAN switching and VLANs (Token Ring Bridge Relay Function [TrBRF] and Token Ring Concentrator Relay Function [TrCRF]).
- Unrecognized Type-Length-Value (TLV) Support—A VTP server or client propagates configuration changes to its other trunks, even for TLVs it is not able to parse. The unrecognized TLV is saved in NVRAM.
- Version-Dependent Transparent Mode—In VTP version 1, a VTP transparent network device inspects VTP messages for the domain name and version, and forwards a message only if the version and domain name match. Since only one domain is supported in the supervisor engine software, VTP version 2 forwards VTP messages in transparent mode, without checking the version.
- Consistency Checks—In VTP version 2, VLAN consistency checks (such as VLAN names and values) are performed only when you enter new information through the CLI or SNMP.

QUESTION NO: 25

Which of the following are true regarding the default values of a switch that is configured for VTP pruning? (Select two).

- A. VLAN 1-1000 are pruning-eligible
- B. VLAN 2-1000 are pruning-eligible
- C. VLAN 1 is pruning-eligible
- D. VLAN 1 is pruning-ineligible
- E. VLAN 1-1023 is pruning-eligible
- F. VLAN 1-1023 is pruning-ineligible

Answer: B, D

Explanation:

By default, VLANs 2-1000 are pruning-eligible. Since the default VLAN for any switch port in a Catalyst switch is VLAN 1, it is not eligible for pruning.

QUESTION NO: 26

Which of the following tasks are NOT functions performed by VTP switches? (Select all that apply)

- A. To reduce parallel load sharing
- B. To propagate global VLAN information
- C. To provide routing randomness
- D. To set the trunk priority levels of adjacent switches.
- E. To ensure that there is a trunk operating in the network.

Answer: A, C, D, E

Explanation:

VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs within a VTP domain. A VTP domain (also called a VLAN management domain) is made up of one or more network devices that share the same VTP domain name and that are interconnected with trunks. VTP minimizes misconfigurations and configuration inconsistencies that can result in a number of problems, such as duplicate VLAN names, incorrect VLAN-type specifications, and security violations. The fundamental function of VTP is to manage, maintain, and propagate VLAN information throughout the enterprise network. All of the choices, besides B, describe functions that are not performed by VTP.

Section 10: Explain the operation and purpose of the Spanning-Tree Protocol (STP) on a switched network (35 questions)

QUESTION NO: 1

The TestKing network administrator is fine tuning the STP parameters on the Catalyst switches. In which states does the Spanning Tree protocol (STP) get affected by the forward delay parameter? (Select two)

- A. Forwarding
- B. Listening
- C. Blocking
- D. Disabled
- E. Learning

Answer: B, E

Explanation:

The following states utilize information from the forward delay timer:

Listen – The switch listens for a period of time called the fwd delay (forward delay)

Learn – The switch learns for a period of time called the fwd delay (forward delay)

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 139

QUESTION NO: 2

Is the following statement True or False?

STP prevents redundant links.

- A. False
- B. True
- C. There is not enough information to determine

Answer: A

Explanation:

According to Cisco:

STP runs on bridges and switches that are 802.1d-compliant. There are different flavors of STP, with IEEE 802.1d being the most popular and widely implemented. STP is implemented on

bridges and switches in order to prevent loops in the network. STP should be used in situations where you want redundant links, but not loops.

Incorrect Answers:

B: STP prevents the use of redundant active links (by logically disabling the redundant ports) but redundant links are still supported, since the STP topology changes dynamically with the network. Redundant links can be used, but not more than one at a time since STP provides for no load balancing type of mechanism.

QUESTION NO: 3

If two paths to a root switch share the exact same path cost, what information will spanning tree use to determine the root port?

- A. The lowest time to receive BPDUs.
- B. The lowest Port ID.
- C. The lowest sender bridge ID.
- D. The highest MAC address on the receiving port.
- E. None of the above.

Answer: C

Explanation:

A Root Bridge is chosen based on the results of the BPDU process between the switches. Initially, every switch considers itself the Root Bridge! When a switch first powers up on the network, it sends out a BPDU with its own BID as the Root BID. When the other switches receive the BPDU, they compare the BID to the one they already have stored as the Root BID. If the new Root BID has a lower value, they replace the saved one. But if the saved Root BID is lower, a BPDU is sent to the new switch with this BID as the Root BID. When the new switch receives the BPDU, it realizes that it is not the Root Bridge and replaces the Root BID in its table with the one it just received. The result is that the switch that has the lowest BID is elected by the other switches as the Root Bridge.

Based on the location of the Root Bridge, the other switches determine which of their ports has the lowest path cost to the Root Bridge. These ports are called Root Ports and each switch (other than the current Root Bridge) must have one.

The switches determine who will have Designated Ports. A Designated Port is the connection used to send and receive packets on a specific segment. By having only one Designated Port per segment, all looping issues are resolved!


Designated Ports are selected based on the lowest path cost to the Root Bridge for a segment. Since the Root Bridge will have a path cost of "0", any ports on it that are connected to segments will become Designated Ports. For the other switches, the path cost is compared for a given segment. If one port is determined to have a lower path cost, then it becomes the Designated Port for that segment. If two or more ports have the same path cost, then the switch with the lowest BID is chosen.

QUESTION NO: 4**What is true of a topology change in an STP environment?**

- A. The default aging time for MAC address entries will be reduced for a period of the `max_age` timer plus the `forward_delay` interval.
- B. All ports will transition temporarily to the learning state for a period equal to the `forward_delay` interval.
- C. All ports will temporarily transition to the learning state for a period equal to the `max_age` timer plus the `forward_delay` interval.
- D. The default `hello_timer` for configuration BPDUs will be reduced for the period of the `max_age` timer.

Answer: C**Explanation:**

In normal STP operation, a bridge keeps receiving configuration BPDUs from the root bridge on its root port. However, it never sends out a BPDU toward the root bridge. In order to achieve that, a special BPDU called the topology change notification (TCN) BPDU has been introduced. Therefore, when a bridge needs to signal a topology change, it starts to send TCNs on its root port. The designated bridge receives the TCN, acknowledges it, and generates another one for its own root port. The process continues until the TCN hits the root bridge.



Bridge B notifies a topology change by sending a TCN on its root port. The TCN is acknowledged and forwarded up to the root bridge R.

The TCN is a very simple BPDU that contains absolutely no information that a bridge sends out every `hello_time` seconds (this is locally configured `hello_time`, not the `hello_time` specified in configuration BPDUs). The designated bridge acknowledges the TCN by immediately sending back a normal configuration BPDU with the topology change acknowledgement (TCA) bit set. The bridge that notifies the topology change does not stop sending its TCN until the designated

bridge has acknowledged it. Therefore, the designated bridge answers the TCN even though it does not receive configuration BPDU from its root.

Broadcast the Event to the Network

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

The TC bit is set by the root for a period of max_age + forward_delay seconds, which is 20+15=35 seconds by default.

Reference:

Understanding Spanning-Tree Protocol Topology Changes
<http://www.cisco.com/warp/public/473/17.html>

QUESTION NO: 5

Multiple TestKing switches are connected together, forming a loop in the network to provide redundancy. Which of the following technologies provides loop avoidance?

- A. VTP
- B. MLS-RP
- C. MLS-SE
- D. VTP Pruning
- E. STP
- F. STP Trunking
- G. None of the above

Answer: E

Explanation:

Spanning-Tree Protocol (STP) is a Layer 2 protocol designed to run on bridges and switches. The specification for STP is defined in IEEE 802.1d. The main purpose of STP is to ensure that you do not run into a loop situation when you have redundant paths in your network. STP detects/disables network loops and provides backup links between switches or bridges. It allows the device to interact with other STP compliant devices in your network to ensure that only one path exists between any two stations on the network.

Reference: <http://www.zyxel.com/support/supportnote/ves1012/app/stp.htm>

QUESTION NO: 6

Before a port can participate in the STP process the ports have to change. In which sequence do the STP port states change through?

- A. Initial, Learning, Updating, and Active
- B. Blocking, Listening, Updating, and Active
- C. Initial, Learning, Updating, and Forwarding
- D. Blocking, Listening, Learning, and Forwarding

Answer: D

Explanation: The correct order is: blocking state (not participating), listening, learning (prepares to participate), and Forwarding.

Note: STP states:

- Blocking—The Layer 2 LAN port does not participate in frame forwarding
- Listening—First transitional state after the blocking state when STP determines that the Layer 2 LAN port should participate in frame forwarding
- Learning—The Layer 2 LAN port prepares to participate in frame forwarding
- Forwarding—The Layer 2 LAN port forwards frames
- Disabled—The Layer 2 LAN port does not participate in STP and is not forwarding frames.

QUESTION NO: 7

Is the following statement True or False?

The “show spanning-tree port-priority” command only displays information for ports with an active link?

- A. False
- B. There is not enough information to determine
- C. True

Answer: C

Explanation:

According to Cisco:

The show spanning-tree port-priority command only displays information for ports with an active link. If these conditions are not met, enter a show running-config interface command to verify the configuration.

QUESTION NO: 8

Is the following statement True or False?

STP uses BPDU’s (Bridge Data Units) to communicate and compute the spanning tree topology from each switch and in both directions from the route switch.

- A. False
- B. There is not enough information to determine
- C. True

Answer: A

Explanation:

According to Cisco:

To communicate and compute the spanning tree topology, Bridge Protocol Data Units (BPDUs) are transmitted from each switch (configuration BPDUs) and in one direction from the root switch. Only one direction is used from the switch, not both.

QUESTION NO: 9

Which command would you enter to display a summary of the spanning-tree information? (Type in answer below)

Answer: show spantree summary

Explanation:

According to Cisco: Use the show spantree summary command to display a summary of spanning-tree information.

QUESTION NO: 10

In order for STP to run successfully on the TestKing network, what standard the bridges and switches have to comply with?

- A. 802.1c
- B. 802.1e
- C. 802.1x
- D. 802.1f
- E. 802.1d
- F. 802.11

Answer: E

Explanation:

According to the online documentation provided by Cisco:

STP runs on bridges and switches that are 802.1d-compliant. There are different flavors of STP, with IEEE 802.1d being the most popular and widely implemented. STP is implemented on bridges and switches in order to prevent loops in the network. Use it in situations where you

want redundant links, but not loops. Redundant links are important as backups in case of failover in a network. If your primary fails, the backup links are activated so that users can continue using the network. Without STP on the bridges and switches, such a situation could result in a loop.

QUESTION NO: 11

Switches TK1 and TK2 are exchanging Bridge Protocol Data Unit (BPDU) information. Which of the following can result from a BPDU exchange? (Select all that apply)

- A. One switch is elected as the root switch.
- B. Ports included in the spanning tree are selected.
- C. The shortest distance to the root switch is calculated
- D. A designated bridge for each LAN segment is selected.
- E. A root port is selected.

Answer: A, B, C, D, E

Explanation:

A BPDU exchange between devices results in the following:

One switch is elected as the root switch.

The shortest distance to the root switch is calculated for each switch based on the path cost.

A designated bridge for each LAN segment is selected. This is the switch closest to the root bridge through which frames are forwarded to the root.

A root port is selected. This is the port providing the best path from the bridge to the root bridge. Ports included in the spanning tree are selected.

QUESTION NO: 12

What determines the default spanning tree port path cost of STP devices within the TestKing network?

- A. The server speed settings
- B. The available bandwidth.
- C. The media speed of an interface.
- D. The stored IOS settings
- E. The interface number

Answer: C

Explanation:

The spanning tree port path cost default value is derived from the media speed of an interface. In the event of a loop, spanning tree considers port cost when selecting an interface to put into the forwarding state. You can assign lower cost values to interfaces that you want spanning tree to

select first and higher cost values to interfaces that you want spanning tree to select last. If all interfaces have the same cost value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks other interfaces. The possible cost range is 1 through 200000000 (the default is media specific).

QUESTION NO: 13

Switch TK1 is a non-root switch in the TestKing network. By what method does a non-Root switch choose its Root Port?

- A. It chooses the port with the lowest cumulative Root Path Cost to the Root Bridge.
- B. The port receives an inferior BPDU from a neighboring switch on a shared LAN segment.
- C. It chooses the port with the highest cumulative Root Path Cost to the Root Bridge.
- D. The port receives a BPDU announcing a higher Root Path Cost from a neighboring switch on a shared LAN segment.
- E. None of the above.

Answer: A

Explanation:

The spanning tree Protocol uses the information found in the BPDUs to determine which ports should be forwarding and which should be blocking. If costs are equal, the STP reads through BPDU until it finds a parameter that is not equal. The lower port ID becomes the forwarding port, and the higher port ID is placed in a blocked state. As the BPDU prepares to leave a port, it applies a port cost. The sum of all the port costs is the path cost. Spanning Tree looks first at the path cost to decide which ports should forward and which should block. The port that reports the lowest path cost is chosen to forward.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 155

QUESTION NO: 14

What is the default transition time for a switch in the TestKing switched LAN to move from blocking to forwarding state in the Spanning-Tree protocol?

- A. 5 seconds
- B. 50 seconds
- C. 60 seconds
- D. 90 seconds
- E. 120 seconds

Answer: B

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

The default STP timers are shown below:

From blocking to listening 20 seconds

From listening to learning 15 seconds

From learning to forwarding 15 seconds

From blocking to forwarding state 50 seconds

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 141


QUESTION NO: 15

Drag the Spanning Tree Protocol state in the options column on the left to the matching definition column on the right.

Select from these


Place here

Listening



administratively down

Disabled




Receives BDPUs only

Blocking




Forwarding sends or receives user data

Learning



Builds bridging table

Forwarding



Processes BPDUs but does not forward data

Answer:

Select from these

Place here

Disabled

administratively down

Blocking

Receives BPDUs only

Learning

Builds bridging table

Forwarding

Forwarding sends or receives user data

Listening

Processes BPDUs but does not forward data

Explanation:

Learning State:

A port in the learning state is preparing to participate in frame forwarding. This is the second transitional state through which a port moves in anticipation of frame forwarding. The port enters the learning state from the listening state through the operation of Spanning-Tree Protocol.

A port in the learning state performs the following functions:

Discards frames received from the attached segment.

Discards frames switched from another port for forwarding.

Incorporates station locations into its address database.

Receives BPDUs and directs them to the system module.

Receives, processes, and transmits BPDUs received from the system module.

Receives and responds to network management messages.

QUESTION NO: 16

Is the following statement True or False?

STP uses the port cost value when the interface is configured as an access port and it uses VLAN port cost values when the interface is configured as a trunk port.

- A. There is not enough information to determine
- B. False
- C. True

Answer: C

Explanation:

According to Cisco:

Spanning tree uses the port cost value when the interface is configured as an access port and uses VLAN port cost values when the interface is configured as a trunk port.

QUESTION NO: 17

Switch TK1 is participating in the Spanning Tree Protocol (STP). What is true about STP Path Cost on a particular port of TK1?

- A. It is known only to the local switch where the port resides.
- B. It can be modified to help determine Root Bridge selection.
- C. Modifying it can cause TCN BPDU to be sent to the Root Bridge.
- D. When increased, it can provide higher bandwidth to a connecting port.
- E. None of the above

Answer: A

Explanation:

With STP, first a root bridge is elected. Then, the shortest distance to the root bridge is calculated for each switch based on the path cost. This calculation is done locally on each switch and the path cost for that switch is only used on the local switch.

Incorrect Answers:

B: Adjust the STP port priority, not the port path cost, can be done to influence the election of the root bridge.

C: A bridge considers it a topology change only when one of the following occurs:

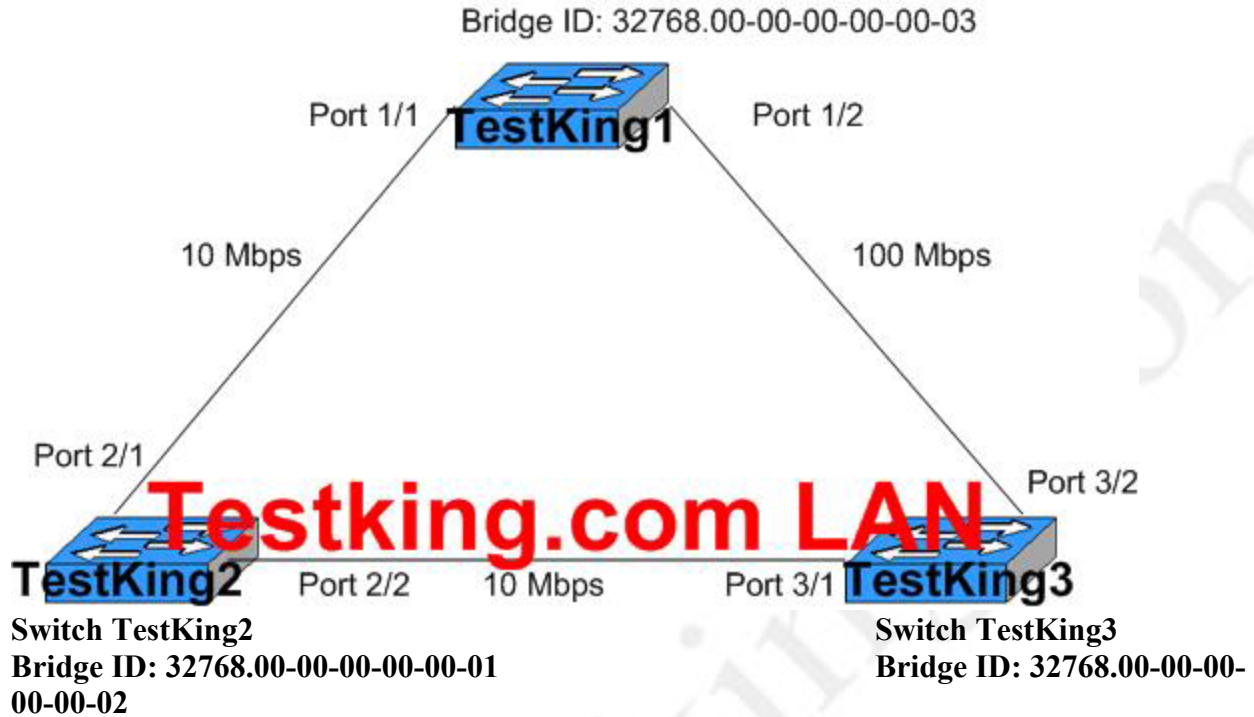
- When a port that was forwarding is going down (blocking for instance).
- When a port transitions to forwarding and the bridge has a designated port. (This means that the bridge is not standalone.)

TCN BPDUs are only sent to other switches within the network if one of the above happens.

D: Simply adjusting the cost value of the port will not make the port faster or provide for additional bandwidth throughput.

QUESTION NO: 18

Study the exhibit below:



Given the network configuration above and assuming that STP is enabled, which port will be elected the non-designated port?

- A. Port 1/1
- B. Port 1/2
- C. Port 2/1
- D. Port 2/2
- E. Port 3/1
- F. Port 3/2

Answer: B:

Explanation:

For each VLAN, the switch with the highest bridge priority (the lowest numerical priority value) is elected as the root bridge. If all switches are configured with the default priority value (32,768), the switch with the lowest MAC address in the VLAN becomes the root bridge.

The spanning tree root bridge is the logical center of the spanning tree topology in a switched network. All paths that are not required to reach the root bridge from anywhere in the switched network are placed in spanning tree blocking mode.

A spanning tree uses the information provided by BPDUs to elect the root bridge and root port for the switched network, as well as the root port and designated port for each switched segment. In this example, since the priorities are set to the default, the switch with the lowest MAC address is used as the tie breaker. In this case, TestKing2 will become the root switch, which means that port 3/1 and 1/1 will become the root ports and must be in the forwarding state. That

leaves the other port on switch TestKing2, port 1/2 as the non-designated port since this switch has the highest MAC address.

QUESTION NO: 19

Which of the following specifications is a companion to the IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) algorithm, and warrants the use multiple spanning-trees?

- A. IEEE 802.1s (MST)
- B. IEEE 802.1Q (CST)
- C. Cisco PVST+
- D. IEEE 802.1d (STP)

Answer: A

Explanation:

MST uses the modified RSTP version called the Multiple Spanning Tree Protocol (MSTP). MST extends the IEEE 802.1w rapid spanning tree (RST) algorithm to multiple spanning trees. This extension provides both rapid convergence and load balancing in a VLAN environment. MST converges faster than PVST+. MST is backward compatible with 802.1D STP, 802.1w (rapid spanning tree protocol [RSTP]), and the Cisco PVST+ architecture. MST allows you to build multiple spanning trees over trunks. You can group and associate VLANs to spanning tree instances. Each instance can have a topology independent of other spanning tree instances. This new architecture provides multiple forwarding paths for data traffic and enables load balancing. Network fault tolerance is improved because a failure in one instance (forwarding path) does not affect other instances (forwarding paths). In large networks, you can more easily administer the network and use redundant paths by locating different VLAN and spanning tree instance assignments in different parts of the network. A spanning tree instance can exist only on bridges that have compatible VLAN instance assignments. You must configure a set of bridges with the same MST configuration information, which allows them to participate in a specific set of spanning tree instances. Interconnected bridges that have the same MST configuration are referred to as an *MST region*.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a008007e71a.html#wp1082480

QUESTION NO: 20

Switch TK1 is powered on for the first time in the TestKing network. Upon initial bootup, which destination address does a TK1 use to send BPDUs?

- A. A well-known multicast address.
- B. The IP address of its default gateway.
- C. The MAC addresses stored in the CAM table.
- D. The MAC address of neighbors discovered via CDP
- E. None of the above

Answer: A

Explanation:

Bridge protocol data units (BPDUs) are used by the spanning tree algorithm to determine information about the topology of the network BPDUs are used to send configuration messages using multicast frames. When STP devices are first powered on, a well known multicast destination MAC address is used to send the BPDU information.

Incorrect Answers:

- B: This would only send the BPDU information to the router. The other switches in the network that are participating in STP need the BPDU information, not the router.
- C: When a switch is first powered up, the CAM table will be empty.
- D: Since STP is standards based, it does not use any Cisco proprietary protocols such as CDP to perform any of its functions. This will ensure inter-operability with switches from other vendors.

QUESTION NO: 21

Switch TK3 is calculating the root path cost to the Root Bridge, TK1. What is true regarding the Root Path cost?

- A. It is the Path Cost of a particular Root Port.
- B. It is the cost sent from the Root Bridge to all non-root bridges.
- C. This value is the cumulative cost of all the links leading to the Root Bridge.
- D. This value is the cumulative cost of all links sent from the Designated Port of the Root Bridge.

Answer: C

Explanation:

The first stage in the STP process is the calculation stage. During this stage, each bridge on the network transmits BPDUs that allow the system to work out:

- The identity of the bridge that is to be the Root Bridge - the central reference point from which the network is configured.
- **The Root Path Costs for each bridge - that is, the cost of the paths from each bridge to the Root Bridge. This value is found by adding up the cost of all of the links to the root bridge.**
- The identity of the port on each bridge that is to be the Root Port - the one that is connected to the Root Bridge using the most efficient path, that is, the one that has the lowest Root Path Cost. Note that the Root Bridge does not have a Root Port.

- The identity of the bridge that is to be the Designated Bridge of each LAN segment - the one that has the lowest Root Path Cost from that segment. Note that if several bridges have the same Root Path Cost, the one with the lowest Bridge Identifier becomes the Designated Bridge.

QUESTION NO: 22

At which layer of the OSI model does the Spanning Tree Protocol (STP) operate at?

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

Answer: D

Explanation:

Spanning-Tree Protocol (STP) is a Layer 2 (L2) protocol designed to run on bridges and switches. The specification for STP is called 802.1d. The main purpose of STP is to ensure that you do not run into a loop situation when you have redundant paths in your network. Loops are deadly to a network.

QUESTION NO: 23

Which of the following specification will allow you to: associate VLAN groups to STP instances so you can provide multiple forwarding paths for data traffic and enable load balancing?

- A. IEEE 802.1d (STP)
- B. IEEE 802.1s (MST)
- C. IEEE 802.1Q (CST)
- D. IEEE 802.1w (RSTP)

Answer: B

Explanation:**IEEE 802.1s MST Overview**

MST extends the IEEE 802.1w rapid spanning tree (RST) algorithm to multiple spanning trees. This extension provides both rapid convergence and load balancing in a VLAN environment. MST converges faster than PVST+. MST is backward compatible with 802.1D STP, 802.1w (rapid spanning tree protocol [RSTP]), and the Cisco PVST+ architecture.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter_09186a008007e71a.html#1050594

QUESTION NO: 24

Is the following statement True or False?

STP prevents loop by forcing certain redundant data paths into a standby (blocked) state, while leaving others in a forwarding state.

- A. False
- B. True
- C. There is not enough information to determine

Answer: B

Explanation:

According to Cisco:

STP forces certain redundant data paths into a standby (blocked) state, while leaving others in a forwarding state. If a link in forwarding state becomes unavailable, STP reconfigures the network and reroutes data paths by activating the appropriate standby path.

QUESTION NO: 25

You want to influence the root switch election process within the TestKing network. When setting up STP in this network, which switch should you configure as the root switch?

- A. The most centralized switch
- B. The most secure switch
- C. The most updated switch
- D. The most powerful switch
- E. The switch that has the longest uptime.

Answer: A

Explanation:

Cisco recommends using the most centralized switch in the network as the root switch.

According to Cisco:

Before configuring STP, you need to select a switch to be the root of the spanning tree. It does not necessarily have to be the most powerful switch; it should be the most centralized switch on the network. All dataflow across the network will be from the perspective of this switch. It is also important that this switch be the least disturbed switch in the network. The backbone

switches are often selected for this function, because they typically do not have end stations connected to them. They are also less likely to be disturbed during moves and changes within the network.

QUESTION NO: 26

You want to influence the Root Bridge election process in the TestKing network. To do this, you are adjusting the Bridge ID values. What is Cisco's philosophy on the STP root selection process regarding the Root ID value?

- A. Smaller is better (more preferred).
- B. Larger is better (more preferred).
- C. The bridge ID should always be zero

Answer: A

Explanation:

In the STP root selection process, a smaller value is preferred over a larger value. If the Root ID on Switch A is advertising an ID that is smaller than the Root ID that its neighbor (Switch B) is advertising, Switch A's information is better. Switch B stops advertising its Root ID, and instead accepts that of Switch A.

QUESTION NO: 27

After you decide which switch should be the root switch, which command would you enter to give it priority in the selection process? (Type in answer below)

Answer: bridge priority

Explanation:

According to Cisco:

After you decide which switch should be the root switch, set the appropriate variables to designate it as the root switch. The only variable you have to set is the bridge priority. If this switch has a bridge priority that is lower than all other switches, it will be automatically selected by the other switches as the root switch.

QUESTION NO: 28

Which of the following factors are NOT used to determine the stable active spanning tree topology of the switched TestKing network?

- A. The port identifier associated with each Layer 2 interface

- B. The port identifier associated with each Layer 3 interface
- C. The spanning tree path cost to the root bridge
- D. The unique bridge ID
- E. All of the above are used.

Answer: B

Explanation:

The Spanning Tree Protocol does not use layer 3 information to determine the overall topology. Layer 3 interfaces do not participate in STP, since spanning tree is a layer 2 technology.

Incorrect Answers:

A, C, D: The stable active spanning tree topology of a switched network is determined by the following:

- The unique bridge ID (bridge priority and MAC address) associated with each VLAN on each switch

- The spanning tree path cost to the root bridge

- The port identifier (port priority and MAC address) associated with each Layer 2 interface

QUESTION NO: 29

The Spanning Tree Protocol is running on numerous TestKing devices. Which of the following equipment types does STP run on in the TestKing network? (Select all that apply):

- A. switches
- B. servers
- C. routers
- D. bridges
- E. None of the above.

Answer: A, D

Explanation:

The Spanning-Tree Protocol (STP) is a Layer 2 (L2) protocol designed to run on bridges and switches. The specification for STP is called 802.1d. The main purpose of STP is to ensure that you do not run into a loop situation when you have redundant paths in your network. Loops are deadly to a network.

Incorrect Answers:

B: Servers generally operate at layers 4-5, since they are used to run applications or to store data. They do not participate in STP.

C: Routers operate at layer 3 and do not participate in STP, since the STP function is used to detect and prevent bridging loops, not routing loops.

QUESTION NO: 30

A failure has occurred in the TestKing switched network, causing a loop. What causes bridging loops to occur in a LAN?

- A. A failure in the route-switch module
- B. A failure in the VLAN tunnel
- C. A failure in the VTP trunk
- D. A failure in the STA

Answer: D

Explanation:

The primary function of the spanning-tree algorithm (STA) is to cut loops created by redundant links in bridged networks. The Spanning-Tree Protocol (STP) operates at Layer 2 of the OSI model and, by the means of bridge protocol data units (BPDUs) exchanged between bridges, elects the ports that will eventually forward or block traffic. This protocol can fail in some specific cases and troubleshooting the resulting situation can be very difficult, depending on the design of the network. We can even say that in this particular area, the most important part of the troubleshooting is done before the problem occurs. A failure in the STA generally leads to a bridging loop (not a spanning tree loop as you don't need STP to have a loop). Most customers calling the TAC for spanning tree problems are suspecting a bug, but experience proves that it is seldom the case. Even if the software is at stake, a bridging loop in a STP environment necessarily comes from a port that should block, but that is forwarding traffic.

QUESTION NO: 31

The bridge priority of switch TK1 is being manually configured. In the STP root selection process, what happens to the switch with the lowest priority in the network?

- A. It is withdrawn from the election process.
- B. It loses the root bridge election process.
- C. It wins the root bridge election process.
- D. None of the above. The bridge priority is not used to determine the root bridge.

Answer: C

Explanation:

As the BPDUs go out through the network, each switch compares the BPDUs it sent out to the one it received from its neighbors. From this comparison, the switches come to an agreement as to who the root switch is. The switch with the lowest priority in the network wins this election process.

QUESTION NO: 32

If a layer 2 interface on switch TK1 uses the Spanning Tree Protocol (STP) which of the following states could it NOT possibly be in at any time?

- A. Forwarding
- B. Learning
- C. Disabled
- D. Blocking
- E. Listening
- F. None of the above

Answer: F

Explanation:

According to Cisco:

Each Layer 2 interface on a switch using spanning tree exists in one of the following five states:

Blocking-The Layer 2 interface does not participate in frame forwarding

Listening-First transitional state after the blocking state when spanning tree determines that the Layer 2 interface should participate in frame forwarding

Learning-The Layer 2 interface prepares to participate in frame forwarding

Forwarding-The Layer 2 interface forwards frames

Disabled-The Layer 2 interface does not participate in spanning tree and is not forwarding frames.

QUESTION NO: 33

Is the following statement True or False?

Cisco recommends manually configuring the hello time, forward delay time, and maximum age time after configuring the switch as the root bridge for optimal performance.

- A. True
- B. There is not enough information to determine
- C. False

Answer: C

Explanation:

According to Cisco:

We recommend that you avoid manually configuring the hello time, forward delay time, and maximum age time after configuring the switch as the root bridge.

QUESTION NO: 34

When a network engineer designs a switch topology, they assign higher priority values to interfaces that they want spanning tree to select first and lower priority values to interfaces that they want spanning tree to select last. However, if multiple interfaces have equal priority values, spanning tree puts the interface with the _____ interface number in the forwarding state.

- A. Neutral
- B. Highest
- C. Lowest
- D. Random
- E. First

Answer: C

Explanation:

In the event of a loop, spanning tree considers port priority when selecting an interface to put into the forwarding state. You can assign higher priority values to interfaces that you want spanning tree to select first and lower priority values to interfaces that you want spanning tree to select last. If all interfaces have the same priority value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks other interfaces.

QUESTION NO: 35

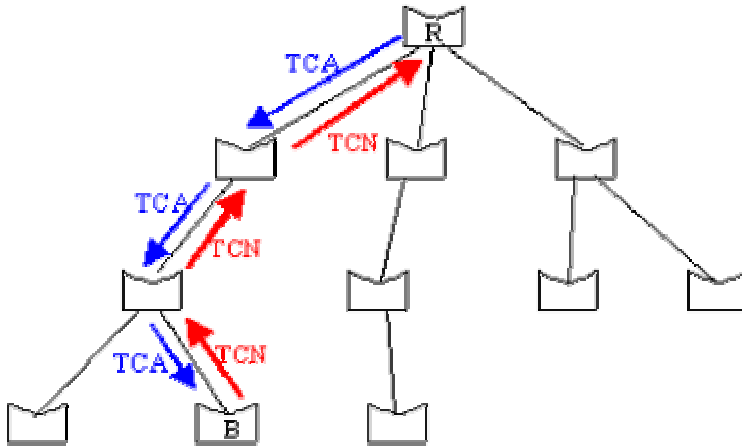
When STP operation is functioning normally and after the topology has converged, which statement is true about BPDUs?

- A. A bridge sends configuration BPDUs towards the root bridge every two seconds.
- B. A bridge sends configuration BPDUs towards the root bridge every 15 seconds.
- C. A bridge sends a TCN BPDU to the root bridge once upon initial configuration, followed by configuration BPDUs every two seconds.
- D. A bridge sends only TCN BPDUs to the root bridge and no configuration BPDUs.

Answer: D

Explanation:

In normal STP operation, a bridge keeps receiving configuration BPDUs from the root bridge on its root port. However, it never sends out a BPDU toward the root bridge. In order to achieve that, a special BPDU called the topology change notification (TCN) BPDU has been introduced. Therefore, when a bridge needs to signal a topology change, it starts to send TCNs on its root port. The designated bridge receives the TCN, acknowledges it, and generates another one for its own root port. The process continues until the TCN hits the root bridge.



Bridge B notifies a topology change by sending a TCN on its root port. The TCN is acknowledged and forwarded up to the root bridge R.

The TCN is a very simple BPDU that contains absolutely no information that a bridge sends out every `hello_time` seconds (this is locally configured `hello_time`, not the `hello_time` specified in configuration BPDUs). The designated bridge acknowledges the TCN by immediately sending back a normal configuration BPDU with the topology change acknowledgement (TCA) bit set. The bridge that notifies the topology change does not stop sending its TCN until the designated bridge has acknowledged it. Therefore, the designated bridge answers the TCN even though it does not receive configuration BPDU from its root.

Reference:

http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a0080094797.shtml

Section 11: Identify the specific types of Cisco route switch processors, and provide implementation details (18 questions)

QUESTION NO: 1

You are in the midst of configuring a Cisco 5000 Catalyst switch named TK1 and you want to control the amount of broadcasts on the network. To control these broadcasts, you create some VLANs on TK1. Which hardware would you use to configure inter-VLAN communication in this scenario?

- A. MLS
- B. RSM
- C. MSFC
- D. VLAN bandwidth

Answer: B

Explanation:

You can view a Route Switch Module (RSM) as an external router that has several interfaces directly connected into the different VLANs of a Catalyst 5000 switch. The RSM is the internal routing processor that lies within the Catalyst 5000 switch. To provide for inter-vlan communication, traffic must pass through the RSM or an external router.

Reference: Troubleshooting InterVLAN Routing on a Catalyst 5000 Switch with RSM
<http://www.cisco.com/warp/public/473/56.html>

Incorrect Answers:

- A: MLS is Multilayer Switching (MLS) and is not used in the Catalyst 5000.
- C: The Multilayer Switch Feature Card (MSFC) is a Route Processor (RP).
- D: This choice does not apply

QUESTION NO: 2

The TestKing switch TK1 must recognize the router as _____ for MLS to function in the network.

- A. A netflow card
- B. An MLS-RP
- C. An MLS-SE
- D. An MLS-RE

Answer: B

Explanation:

For MLS to function, the switch must recognize the router as an MLS-RP. Internal MLS-RPs (the RSM or RSFC in a Catalyst 5000 family member and the MSFC in a Catalyst 6000 family member) are automatically recognized by the MLS-SE in which they are installed. For external MLS-RPs, one must explicitly inform the switch of the router's address. This address is not actually an IP address, although on external MLS-RPs it is chosen from the list of IP addresses configured on the router's interfaces. It is simply a router ID. For internal MLS-RPs, the MLS-ID is normally not even an IP address configured on the router. Since internal MLS-RPs are included automatically, it is commonly a loopback address (127.0.0.x). For MLS to function, include on the MLS-SE the MLS-ID found on the MLS-RP.

QUESTION NO: 3

You need to quickly confirm the status of an interface in the TestKing switch. What show command could you use to confirm whether or not an MLS-RP interface is an 'up/up' state on a router?

- A. show ip interface brief
- B. show ip brief
- C. show interface brief ip
- D. show interface brief

Answer: A

Explanation:

The correct command is (show ip interface brief). All the other choices are incorrect commands. The command (show ip interface) is a valid command, but it gives you much more detail. Since you are only interested on confirming the up/up states, the (show ip interface brief) command is ideal.

QUESTION NO: 4

Is the following statement True or False regarding a Catalyst 6000 switch named TK1?

The Catalyst 6000 family of switches supports the use of an external MLS-RP?

- A. There is not enough information to determine
- B. True
- C. False

Answer: C

Explanation:

According to Cisco:

The Catalyst 6000 family of switches does not support an external MLS-RP at this time. The MLS-RP must be an MSFC.

QUESTION NO: 5

The TestKing network is adding devices that are capable of MLS. A Cisco Catalyst 5000/5500 switch with a NetFlow Feature card can perform MLS with which three devices? (Select three.)

- A. RSM
- B. RSFC
- C. Catalyst 8500
- D. Catalyst 2948G-L3
- E. Catalyst 2900

Answer: A, B, C

Explanation:

A: Cisco IOS running on the RSM has the ability to instruct the NFFC hardware.

B: The RSFC (Route Switch Feature Card) or Route Switch Module (RSM) performs the route processing on the Catalyst switch with a NFFC-II

C: Routing Processing services can also be provided by an externally attached Catalyst 6000 with an MSM (currently supports unicast MLS only), Catalyst 8500, Cisco 7500, 7200, 4700, 4500, 3640 or 3620.

Note: The NetFlow Feature Card increases IP Multilayer Switching (MLS) performance.

Reference: Product Bulletin, No. 909, Catalyst 4000 and 5000 Family Supervisor Engine Software

QUESTION NO: 6

Is the following statement True or False?

The MLS-RP can be internal but can not be external?

- A. True
- B. False
- C. There is not enough information to determine

Answer: B

Explanation:

The MLS-RP can be internal (installed in a switch chassis) or external (connected via a cable to a trunk port on the switch). Examples of internal MLS-RPs are the Route Switch Module (RSM) and the Route Switch Feature Card (RSFC), which are installed in a slot or supervisor of a Catalyst 5000 family member, respectively.

QUESTION NO: 7

You are the administrator of the TestKing network and you are configuring one of the switches. While doing so, you add the following configuration command to switch TK1:

“mls rp ip”

What is the purpose of this command?

- A. For enabling MLSP
- B. For placing an external route processor in the interface of the VTP domain switch
- C. For assigning VLAN ID to route processor interface
- D. For enabling the RSM interface
- E. For entering into the router interface

Answer: A

Explanation:

Multi-Layer Switching (MLS) has become a highly desired method of accelerating routing performance through the use of dedicated Application Specific Integrated Circuits (ASICs). Traditional routing is done through a central CPU and software. MLS offloads a significant portion of routing (packet rewrite) to hardware, and thus has also been termed switching. MLS and Layer 3 switching are equivalent terms.

To enable the Multilayer Switching Protocol (MLSP), use the **mls rp ip** global configuration command. MLSP is the protocol that runs between the switches and routers. Use the **no** form of this command to disable MLS.

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/switch_r/xrmls.htm#1017390

QUESTION NO: 8

While logged into a TestKing multilayer switch, you type in the following command:

interface

What is this command for?

- A. for enabling the RSM interface
- B. for entering into the router interface
- C. for placing an external route processor in the interface of the VTP domain switch
- D. for assigning VLAN ID to route processor interface
- E. for enabling MLSP

Answer: B

Explanation:

This is a trick question with a very simple answer. If you want to configure the actually MLS Route Processor you have to get out of global mode and into interface mode Switch(config-if)#

QUESTION NO: 9

You are the administrator of the TestKing network and you are configuring one of the switches. While doing so, you add the following configuration command to switch TK1:

“Mls rp vlan-id”

What is this command used for?

- A. for enabling MLSP
- B. for assigning VLAN ID to route processor interface
- C. for enabling the RSM interface
- D. for placing an external route processor in the interface of the VTP domain switch
- E. for entering into the router interface

Answer: B

Explanation:

Use this command to assign a VLAN ID to an interface. RSM VLAN interfaces or ISL-encapsulated interfaces do not require the VLAN ID to be assigned.

To assign a VLAN ID, use the **mls ip vlan-id** interface configuration command.

The following example assigns a VLAN ID of 23 to the current interface:

“mls rp vlan-id 23”

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/switch_r/xrmls.htm#1017566

QUESTION NO: 10

Which hardware component do standard and extended access lists use to deny traffic at wire speed?

- A. NetFlow Feature Card
- B. Catalyst Switch Supervisor Engine III
- C. Multilayer Switch Feature Card
- D. MultiLayer Switching Route Processor
- E. None of the above.

Answer: D

Explanation:

MLS allows you to enforce access lists on every packet of the flow without compromising MLS performance. When you enable MLS, the MLS-SE handles standard and extended access list permit traffic at wire speed.

Note Access list deny traffic is always handled by the MLS-RP, not the MLS-SE.

Route topology changes and the addition or modification of access lists are reflected in the MLS switching path automatically on the MLS-SE. The techniques for handling route and access list changes apply to both the RSM and directly attached external routers.

For example, when Station A wants to communicate with Station B, it sends the first packet to the MLS-RP. If an access list is configured on the MLS-RP to deny access from Station A to Station B, the MLS-RP receives the packet, checks the access list to see if the packet flow is permitted, and discards the packet based on the access list. Because the first packet for this flow does not return from the MLS-RP, an MLS cache entry is not established by the MLS-SE

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps679/products_configuration_guide_chapter_09186a008007eadd.html

QUESTION NO: 11

MLS is running on the TestKing network. Which MLS component do the MLS-RP and the MLS-SE use to communicate with one another? (Type in answer below)

Answer: MLSP

Explanation:

MLSP is the protocol that runs between the MLS-SE and the MLS-RP. MLSP is utilized by the MLS-RP and the MLS-SE to communicate with one another; tasks include enabling MLS; installing, updating or deleting flows (cache information); and managing and exporting flow statistics (Netflow Data Export is covered in other documentation). MLSP also allows the MLS-SE to learn the Media Access Control (MAC, Layer 2) addresses of the MLS-enabled router interfaces, check the flowmask of the MLS-RP (explained later in this document), and confirm

that the MLS-RP is operational. The MLS-RP sends out multicast "hello" packets every 15 seconds using MLSP; if three of these intervals are missed, then the MLS-SE recognizes that the MLS-RP has failed or that connectivity to it has been lost.

QUESTION NO: 12

Match the correct definition on the right to the switching term on the left.

Switching term	Definition	Use these
MLS Flow	place here	Protocol used to communicate MLS information
MLS-RP	place here	Cisco device with a route processor that support MLS.
MLS-SE	place here	Sequence of packets that share Layer 3 and Layer 4 information
MLSP	place here	Maintains a Layer 3 switching table (Layer 3 MLS cache)

Answer:

Switching term	Definition	Use these
MLS Flow	Sequence of packets that share Layer 3 and Layer 4 information	
MLS-RP	Cisco device with a route processor that support MLS.	
MLS-SE	Maintains a Layer 3 switching table (Layer 3 MLS cache)	
MLSP	Protocol used to communicate MLS information	

Explanation:

MLS components:

- **Multilayer Switching Engine (MLS-SE)** – The switching entity that handles the function of moving and rewriting packets.
- **Multilayer Switching Route Processor (MLS-RP)** – A route switch module or an externally connected Cisco series router with software that supports multilayer switching.
- **Multilayer Switching Protocol (MLSP)** – This protocol operates between the MLS-SE and MLS-RP to enable multilayer switching.
- **MLS Flow** - The PFC maintains a Layer 3 switching table (the Layer 3 MLS cache) for Layer 3-switched flows. The cache also includes entries for traffic statistics that are updated in tandem with the switching of packets. After the MLS cache is created, packets

identified as belonging to an existing flow can be Layer 3 switched based on the cached information. The MLS cache maintains flow information for all active flows. An MLS cache entry is created for the initial packet of each flow. Upon receipt of a packet that does not match any flow currently in the MLS cache; a new IP MLS entry is created.

Note:

IP MLS Flows

Layer 3 protocols, such as IP and Internetwork Packet Exchange (IPX), are connectionless—they deliver every packet independently of every other packet. However, actual network traffic consists of many end-to-end conversations, or flows, between users or applications. A flow is a unidirectional sequence of packets between a particular source and destination that share the same protocol and transport-layer information. Communication from a client to a server and from the server to the client are separate flows. For example, Telnet traffic transferred from a particular source to a particular destination comprises a separate flow from File Transfer Protocol (FTP) packets between the same source and destination. Flows are based only on Layer 3 addresses, which allow IP traffic from multiple users or applications to a particular destination to be carried on a single flow if only the destination IP address is used to identify a flow.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 219 + 220
http://www.cisco.com/en/US/products/hw/routers/ps368/products_configuration_guide_chapter09186a008007c87c.html#1020417

QUESTION NO: 13

A new MLS-SE switch is being used on the TestKing network. What's the definition of a MLS-SE?

- A. It is a switch with special MLS IOS.
- B. It is a switch with special hardware.
- C. It is a switch with special ports.
- D. It is a switch with special software.
- E. None of the above.

Answer: B

Explanation:

The MLS-SE is a switch with special hardware. For a member of the Catalyst 5000 family, MLS requires that the supervisor have a Netflow Feature Card (NFFC) installed. The Supervisor IIG and IIIG have one by default. In addition, a bare minimum of Catalyst OS 4.1.1 software is also required. Note that the 4.x train is now in General Deployment (GD), or passed rigorous end-user criteria and field-experience targets for stability, so check Cisco's website for the latest releases. IP MLS is supported and automatically enabled for Catalyst 6000 hardware and software with

the MSFC/PFC (other routers have MLS disabled by default). Note that IPX MLS and MLS for multicasting may have different hardware and software (Cisco IOS and Catalyst OS) requirements. More Cisco platforms do/will support the MLS feature. Also, MLS must be enabled in order for a switch to be an MLS-SE.

QUESTION NO: 14

Switch TK1 contains 2 supervisor cards for redundancy. What's true of a Catalyst switch with dual supervisor cards? (Select two)

- A. The supervisor engines must be the same model.
- B. The active supervisor is selected by priority.
- C. The active supervisor controls the system bus.
- D. The relevant protocols are active in the standby supervisor.
- E. The supervisor engines perform load sharing.

Answer: A, C

Redundant supervisor engines must be of the same type with the same model feature card. The active supervisor is responsible for controlling the system bus and all line cards.

QUESTION NO: 15

Switch TK1 is using MLS for switching packets. Which of the following are valid flow masks for MLS-SE? (Select all that apply)

- A. source-destination-ip
- B. ip-sum
- C. ip-bypass
- D. destination-ip
- E. ip-flow
- F. None of the above.

Answer: A, D, E

Explanation:

The three flow masks are as follows:

Destination-IP: The least-specific flow mask. The MLS-SE maintains one MLS entry for each destination IP address. All flows to a given destination IP address use this MLS entry. This mode is used if there are no access lists configured on any of the MLS-RP interfaces.

Source-Destination-IP: The MLS-SE maintains one MLS entry for each source and destination IP address pair. All flows between a given source and destination use this MLS entry regardless of the IP protocol ports. This mode is used if there is a standard access list on any of the MLS-RP interfaces.

IP-flow: The most-specific flow mask. The MLS-SE creates and maintains a separate MLS cache entry for every IP flow. An ip-flow entry includes the source IP address, destination IP address, protocol, and protocol ports. This mode is used if there is an extended access list on any of the MLS-RP interfaces.

QUESTION NO: 16

What is true about the Multi-layer Switching (MLS) cache?

- A. MLS cache entries support unidirectional flows.
- B. The MLS-RP stores routing information in the MLS cache.
- C. The MLS-SE deletes a cache entry when it detects a TCP FIN ACK
- D. The MLS-RP creates MLS cache entries based on known data flows.

Answer: A

Explanation:

An MLS cache entry is created for the initial packet of each flow. A flow is a unidirectional sequence of packets between a particular source and destination that share the same protocol and transport-layer information. Communication from a client to a server and from the server to the client are separate flows. For example, Hypertext Transfer Protocol (HTTP) Web packets from a particular source to a particular destination are a separate flow from File Transfer Protocol (FTP) file transfer packets between the same pair of hosts.

Incorrect Answers:

- B: Routing information is not stored in the MLS cache.
- C: The state and identity of the flow are maintained while packet traffic is active; when traffic for a flow ceases, the entry ages out.
- D: The MLS-SE, not the MLS-RP creates MLS cache entries.

QUESTION NO: 17

Which one of the following answer choices describes a hardware-based PDU header rewriting and forwarding based on specific information obtained from multiple OSI layers?

- A. Multiplayer switching
- B. Cisco express routing
- C. Multilayer switching

- D. Multilayer routing
- E. Router express forwarding

Answer: C

Explanation:

To determine the best path is the primary function of routing protocols, and this can be a CPU-intensive process. Thus, there is a significant performance increase with the offload of a portion of this function to switching hardware. This performance increase is the goal of the MLS feature. Two of the three major components of MLS are the MLS route processor (MLS-RP) and the MLS switching engine (MLS-SE). The MLS-RP is the MLS-enabled router, which performs the traditional function of routing between subnets/VLANs. The MLS-SE is a MLS-enabled switch, which normally requires a router to route between subnets/VLANs. However, with special hardware and software, MLS-SE can handle the rewrite of the packet. When a packet transverses a routed interface, the change (rewrite) of non-data portions of the packet occurs as the packet heads to the destination, hop by hop. Confusion can arise here because a Layer 2 device appears to take on a Layer 3 task. Actually, the switch only rewrites Layer 3 information and "switches" between subnets/VLANs. The router is still responsible for standards-based route calculations and best-path determination. You can avoid much of this confusion if you mentally keep the routing and switching functions separate, especially when they are within the same chassis (as with an internal MLS-RP). Think of MLS as a much more advanced form of route cache, with a separation of the cache from the router on a switch. MLS requires both the MLS-RP and the MLS-SE, along with respective hardware and software minimums

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 219

QUESTION NO: 18

Some Cisco Supervisor Engines have modular uplink ports so that a module can be selected to provide an appropriate level of bandwidth. Which of the following two features are valid on 2-port 1000Base SX or 1000Base LX modular uplink modules? (Select two)

- A. 10/100 Autosensing
- B. Fast EtherChannel
- C. ISL
- D. 802.1Q
- E. L2 rewrite
- F. L3 rewrite

Answer: C, D

Explanation:

Two ways that Ethernet trunking can be implemented are:

- InterSwitch Link (ISL) (Cisco proprietary protocol)

- 802.1q (IEEE standard)

Trunking can only be performed on fast ethernet and Gigabit ethernet links. In this case, either trunking method can be performed over the 2 port Gig E fiber links.

Incorrect Answers:

A: 1000 Base SX and 1000 base LX links are gigabit ethernet, and do not perform the auto-sensing functions that a 10/100 interface performs.

B: Although Gigabit links can be bonded using channeling, the correct term would be gigabit ethernetchannel in this case, not fast EtherChannel.

Section 12: List and describe the operation of the key components required to implement interVLAN routing (5 questions)

QUESTION NO: 1

The TestKing network needs to pass traffic between VLANs. Which device should be used to accomplish this?

- A. Hub
- B. Switch
- C. Router
- D. Bridge

Answer: C

Explanation:

A VLAN is a virtual LAN contained within a switch, so for it to pass information into a different VLAN within the same switch it has to leave that switch and re-enter via a router. VLANs contain local traffic only, so in order to reach users in another VLAN the traffic must go through a router or a layer 3 routing processor.

QUESTION NO: 2

What command would you enter into a Cisco device if you wanted to allow IOS to handle IP datagrams in the TestKing network with the source routing header option? (Type in the answer below):

Answer: ip source-route

Explanation:

To allow the Cisco IOS software to handle IP datagrams with source routing header options, use the ip source-route global configuration command. To have the software discard any IP datagram containing a source-route option, use the no form of this command.

ip source-route
no ip source-route

QUESTION NO: 3

You are configuring a Cisco multilayer switch for the TestKing network. Which command would you use to configure a port to act as a routed interface?

- A. ip routing
- B. switchport mode trunk
- C. no switchport
- D. switchport trunk native vlan 1

Answer: C

To turn a switch-port into a router interface, it is simply a matter of turning off the switch-port functionality.

```
Switch(config)#interface fa 0/1
Switch(config-if)#no switchport
```

QUESTION NO: 4

At what state does an HSRP-configured router have to be in to perform packet transfer?

- A. Listening
- B. Active
- C. Standby
- D. Queuing
- E. Waiting
- F. Speaking and listening

Answer: B

Explanation:

At any time, HSRP-configured routers are in one of the following states:

Active-The router is performing packet-transfer functions.

Standby-The router is prepared to assume packet-transfer functions if the active router fails.

Speaking and listening-The router is sending and receiving hello messages.

Listening-The router is receiving hello messages.

Only the active HSRP router actually forwards packets.

QUESTION NO: 5

Inter-VLAN routing has been implemented in the TestKing network. In VLAN routing, what are some of the disadvantages of designing a router-on-stick configuration? (Select three)

- A. InterVLAN routing cannot be filtered by the router.
- B. The router becomes a single point of failure for the network.
- C. Routers will not route STP BPDUs.
- D. There is a possibility of inadequate bandwidth for each VLAN.
- E. Additional overhead on the router can occur.
- F. NetFlow Switching is required for InterVLAN accounting.

Answer: B, D, E

Explanation:

A router connected to a switch via a single trunk link is better known as router-on-stick or even a one armed router. Since there's only one router, if that router were to go down there'd be no backup. Since there's only one router, that router would have to handle all the bandwidth of every VLAN so there's a chance it could be overloaded, as with the overhead problems of being responsible for too much.

Because traffic routed between the VLANs traverse a single physical port, there is the potential to not provide for enough bandwidth for a VLAN at any given time.

Inter-VLAN routing also does indeed require additional configuration, management, and overhead.

Incorrect Answers:

A: This is not true since routers can indeed filter traffic that is routed between the VLAN subinterfaces.

C: This is not an advantage. Since BPDU's are local to the VLAN, there is generally no need to route this traffic between the VLANs.

F: This does not apply as a disadvantage to inter-VLAN routing.

Section 13: Explain the types of redundancy in a multilayer switched network including hardware and software redundancy (6 questions)

QUESTION NO: 1

Across the TestKing LAN, Root Link Query messages are being transmitted. Which technology uses a Root Link Query Bridge Protocol Data Unit (BPDU)?

- A. BackboneFast
- B. PortFast
- C. UplinkFast
- D. STP standard
- E. None of the above

Answer: A

Explanation:

If the local switch has blocked ports, BackboneFast begins to use the Root Link Query (RLQ) protocol to if upstream switches have stable connections to the Root Bridge.

Reference: Cisco Press CCNP BCMSN, ISBN 1-58720-077-5, by David Hucaby - Chapter 10 Page 254

QUESTION NO: 2

You are the network administrator at TestKing and are overlooking a Cisco switch with a redundant power supply of the same wattage. What is the total power available to the switch when both of the power supplies are operating normally? (Select all that apply.)

- A. Total power of one supply.
- B. Total combined power of both supplies.
- C. Total power is the sum of one-half of total power of both supplies.
- D. Total power required is shared nearly equally by both supplies.

Answer: C, D

Explanation:

Specifying the **redundant** keyword enables redundancy. In a redundant configuration, the total power drawn from both supplies is at no time greater than the capability of one supply. If one supply malfunctions, the other supply can take over the entire system load. When you install and turn on two power supplies, each concurrently provides approximately half of the required power

to the system. Load sharing and redundancy are enabled automatically; no software configuration is required.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a008007e6f6.html

QUESTION NO: 3

What is the purpose of MST, according to the IEEE 802.1s standard?

- A. It is the spanning-tree implementation used by non-Cisco 802.1Q switches.
- B. It runs a separate instance of STP for each VLAN.
- C. It allows a VLAN bridge to use multiple spanning trees to prevent Layer 2 loops.
- D. It creates a single loop-tree structure that spans the entire Layer 2 network.

Answer: C

Explanation:

IEEE 802.1s MST Overview

MST allows you to build multiple spanning trees over trunks. You can group and associate VLANs to spanning tree instances. Each instance can have a topology independent of other spanning tree instances. This new architecture provides multiple forwarding paths for data traffic and enables load balancing. Network fault tolerance is improved because a failure in one instance (forwarding path) does not affect other instances (forwarding paths).

In large networks, you can more easily administer the network and use redundant paths by locating different VLAN and spanning tree instance assignments in different parts of the network. A spanning tree instance can exist only on bridges that have compatible VLAN instance assignments. You must configure a set of bridges with the same MST configuration information, which allows them to participate in a specific set of spanning tree instances. Interconnected bridges that have the same MST configuration are referred to as an *MST region*.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a008007e71a.html#1050594

QUESTION NO: 4

Which of the following port states are classified in IEEE 802.1w RSTP (Rapid Spanning Tree Protocol)?

- A. Listening, Learning, Forwarding, Blocking, Disabled
- B. Learning, Forwarding, Discarding
- C. Listening, Forwarding, Active, Blocking
- D. Learning, Active, Block

Answer: B

Explanation:

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

RSTP Port States

The port state controls the forwarding and learning processes and provides the values of discarding, learning, and forwarding. Table 15-4 provides a comparison between STP port states and RSTP port states.

Operational Status	STP Port State	RSTP Port State	Port Included in Active Topology
Enabled	Blocking ¹	Discarding ²	No
Enabled	Listening	Discarding	No
Enabled	Learning	Learning	Yes
Enabled	Forwarding	Forwarding	Yes
Disabled	Disabled	Discarding	No

¹ IEEE 802.1D port state designation.

² IEEE 802.1w port state designation. Discarding is the same as blocking in RSTP and MST.

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/12_1e/swconfig/spantree.htm#wp1063598

QUESTION NO: 5

Is the following statement True or False?

Unidirectional links frequently causes bridging loops.

- A. True
- B. There is not enough information to determine
- C. False

Answer: A

Explanation:

According to Cisco:

Unidirectional link is a very frequent cause for a bridging loop. Unidirectional links are often caused by a failure not detected on a fiber link for instance, or a problem with a transceiver. Anything that can lead a link to stay up while providing a one-way communication is very

dangerous as far as STP is concerned. In order to prevent the problems associated with this problem, Cisco developed a method to detect and block unidirectional switching links.

QUESTION NO: 6

The TestKing network has just procured a new Catalyst 6500 to use on the LAN backbone. On a Catalyst 6500 with dual supervisor modules running Native IOS, which two features provide supervisor module redundancy? (Select two)

- A. Route Processor Redundancy (RPR)
- B. Route Processor Redundancy Plus (RPR+)
- C. Single Route Mode (SRM)
- D. Dual Router Mode (DRM)

Answer: A, B

Explanation:

Catalyst 6500 series switches support fault resistance by allowing a redundant supervisor engine to take over if the primary supervisor engine fails. RPR supports a switchover time of 2 to 4 minutes and RPR+ supports a switchover time of 30 to 60 seconds.

When RPR+ mode is used, the redundant supervisor engine is fully initialized and configured, which shortens the switchover time. The active supervisor engine checks the image version of the redundant supervisor engine when the redundant supervisor engine comes online. If the image on the redundant supervisor engine does not match the image on the active supervisor engine, RPR redundancy mode is used.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a00800da705.html

Section 14: Explain how IP multicast operates on a multilayer switched network, including PIM, CGMP and IGMP (32 questions)

QUESTION NO: 1

The TestKing network is using IGMP version 2 within their IP multicast network. In IGMPv2, how is the designated querier elected?

- A. The first router to appear on a subnet is designated.
- B. The host that responds first to the election query is designated.
- C. The router with the lowest IP address on subnet is designated.
- D. The host with the lowest MAC address on a segment is designated.

Answer: C

Explanation:

Unlike IGMP Version 1, in which the DR and the IGMP querier are typically the same router, in IGMP Version 2, the two functions are separated. The DR and the IGMP querier are selected based on different criteria and may be different routers on the same subnet. The DR is the router with the highest IP address on the subnet, whereas the IGMP querier is the router with the lowest IP address.

Note: When an IGMP querier is configured for a VLAN, the switch sends out IGMP general query messages every 125 seconds and listens for general query messages from other switches. If the switch receives a general query, a querier election starts. A querier election across switches is based either on IP address or MAC address. For an inbound query, if the source IP address is nonzero, the election is based on the IP address, and the switch with the lower source IP address becomes the querier. If the source IP address is zero for an inbound query, then the election is based on the source MAC address, and the switch with the lower MAC address wins the election and becomes the querier. The switch that becomes the nonquerier maintains an "other querier interval" timer. When this timer expires, the switch elects itself as the querier.

QUESTION NO: 2

What do you have to do in order to configure PIM on router TK1 for IP multicast routing? (Select two)

- A. Have TK1 Join a multicast group.
- B. Enable CGMP on TK1.

- C. Enable IP multicast routing on TK1.
- D. Configure the TTL threshold on TK1.
- E. Enable PIM on an interface of TK1.

Answer: C, E

Explanation

PIM stands for (protocol independent multicast) and it can come in sparse mode (PIM SM), dense mode (PIMDM), or sparse dense mode.

To configure TK1 for PIM, first you enable IP multicast routing:

TK1#ip multicast-routing

Then you Enable PIM on the interface

TK1#interface *type number*

TK1#ip pim sparse-dense-mode

QUESTION NO: 3

What is true about an Ethernet MAC address that is mapped to a Layer 3 multicast IP address? (Select two)

- A. The first 3 bytes of the Ethernet multicast MAC address are 01:00:5E.
- B. The last 3 bytes of the Ethernet multicast MAC address are 01:00:5E.
- C. When assigning a Layer 3 multicast address, an Ethernet Layer 2 address is automatically generated from the hardcoded MAC address.
- D. The multicast address copies the last 23 bits of the IP address into the last 23 bits of the Ethernet multicast MAC address.
- E. The Ethernet multicast address assigned the last 24 bites of the MAC address to all Fs.
- F. The Ethernet multicast address assigns the first 24 bits of the MAC address to all Fs.

Answer: A, D

Explanation:

The prefix 01-00-5e identifies the frame as multicast: the next bit is always 0, leaving only 23 bits for the multicast address. Because IP multicast groups are 28 bits long, the mapping cannot be one-to-one. Only the 23 least-significant bits for the IP multicast group are placed in the frame. The remaining five high-order bits are ignored resulting in 32 different multicast groups, being mapped to the same Ethernet address.

Reference: CCNP Switching Exam Certification Guide: David Hucaby & Tim Boyles, Cisco Press 2001, ISBN 1-58720 000-7; page 345

QUESTION NO: 4

What was added to IGMP version 3 to enhance it that was not previously available in earlier versions?

- A. Membership query message
- B. Membership report message
- C. Leave group message
- D. Source filtering
- E. Destination filtering

Answer: D

Explanation:

IGMP Version 3 (IGMPv3) adds support for "source filtering," which enables a multicast receiver host to signal to a router which groups it wants to receive multicast traffic from, and from which source(s) this traffic is expected. This membership information enables Cisco IOS software to forward traffic only from those sources from which receivers requested the traffic.”

Reference:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t5/dtigmpv3.htm>

QUESTION NO: 5

Switch TK1 has just received a multicast frame. By default, how does a Layer 2 switch handle IP multicast traffic?

- A. It blocks multicast traffic on all ports.
- B. It delivers multicast traffic to all ports.
- C. It delivers multicast traffic only to ports that subscribe to it.
- D. It delivers multicast traffic only to clients that are member of a multicast group.

Answer: B

Explanation:

When you take a layer two switch out of the box (default configuration), it will by send out multicast traffic out all ports, because a Layer 2 switch doesn't know where the multi cast is destined to, so it just sends it out all ports and hopes it is received. The default switch configuration treats an IP multicast packet like any other packet. When a switch receives a frame destined for a device that it is unaware of or that is not yet learned, it will forward the data out all ports.

QUESTION NO: 6

What are the valid defaults for the Fast Switching of IP Multicasts? (Select all that apply)

- A. enabled and supported over X.25 encapsulated interfaces
- B. enabled by default on all interfaces
- C. disabled by default on all interfaces
- D. disabled and not supported over X.25 encapsulated interfaces
- E. None of the above.

Answer: B, D

Explanation:

Fast switching allows higher throughput by switching a packet using a cache created by the initial packet sent to a particular destination. Destination addresses are stored in the high-speed cache to expedite forwarding. Routers offer better packet-transfer performance when fast switching is enabled. Fast switching is enabled by default on all interfaces that support fast switching. Fast switching of IP multicast packets is enabled by default on all interfaces (including GRE and DVMRP tunnels), with one exception: It is disabled and not supported over X.25 encapsulated interfaces.

QUESTION NO: 7

Router TK1 is configured for multicast. How would you display PIM information cached in the routing table of TK1?

- A. show ip pim [group-name] [mapping]
- B. show ip rp [group-name | group-address] [mapping]
- C. show ip pim rp [group-name | group-address] [mapping]
- D. show ip rp pim [group-name | group-address] [mapping]
- E. show ip mroute [Hostname | group-address]
- F. None of the above

Answer: E

Explanation:

The "show ip mroute [Hostname | Group]" command is used to display the multicast routing table. The show ip mroute command displays all groups and sources.

The following is sample output from the show ip mroute command for a router operating in sparse mode:

```
Router# show ip mroute
```

IP Multicast Routing Table

Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
 L - Local, P - Pruned, R - RP-bit set, F - Register flag,
 T - SPT-bit set, J - Join SPT, M - MSDP created entry,
 X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
 U - URD, I - Received Source Specific Host Report, Z - Multicast Tunnel,
 Y - Joined MDT-data group, y - Sending to MDT-data group

Timers: Uptime/Expires

Interface state: Interface, Next-Hop, State/Mode

(* , 224.0.255.3), uptime 5:29:15, RP is 192.168.37.2, flags: SC
 Incoming interface: Tunnel0, RPF neighbor 10.3.35.1, Dvmrp
 Outgoing interface list:
 Ethernet0, Forward/Sparse, 5:29:15/0:02:57

(192.168.46.0/24, 224.0.255.3), uptime 5:29:15, expires 0:02:59, flags: C
 Incoming interface: Tunnel0, RPF neighbor 10.3.35.1
 Outgoing interface list:
 Ethernet0, Forward/Sparse, 5:29:15/0:02:57

Incorrect Answers:

A: This is an invalid command

B: This is an invalid command.

C: To display active rendezvous points (RPs) that are cached with associated multicast routing entries, use the show ip pim rp command. The following is sample output from the show ip pim rp command:

Router# show ip pim rp

Group:227.7.7.7, RP:10.10.0.2, v2, v1, next RP-reachable in 00:00:48

D: This command is invalid

QUESTION NO: 8

In the TestKing IP multicast network, which of the following MAC address corresponds to the multicast IP address of 224.0.1.55?

- A. 00-00-00-00-01-55
- B. 00-01-E0-00-01-37
- C. E0-00-01-37-FF-FF
- D. 01-00-5E-00-01-37
- E. None of the above.

Answer: D

Explanation: Once an application determines the class D IP multicast address it will utilize, that address must be mapped into a MAC layer multicast for delivery across any LAN based system. This process is outlined as follows:

Step 1: Using the Class D address, identify the low order 23 bits of the class D address.

Step2: Map those 23 bits into the low order 23 bits of a MAC address with the fixed high order 25 bits of the IANA's IEEE multicast addressing space prefixed by 01:00:5E.

In this example, the 0.1.55 translates to 00-01-37 in hex so when added to the IEEE multicast MAC address of 01-00-5E the end result is 01-00-5E-00-01-37.

QUESTION NO: 9

In the TestKing IP multicast network, the multicast IP address 224.0.16.111 translates to which MAC address?

- A. 00-01-E0-00-10-6F
- B. 01-00-5E-00-10-6F
- C. 00-01-5E-00-10-6F
- D. 00-00-00-00-10-6F
- E. E0-00-10-6F-FF-FF
- F. None of the above.

Answer: B

Explanation:

Once an application determines the class D IP multicast address it will utilize, that address must be mapped into a MAC layer multicast for delivery across any LAN based system. This process is outlined as follows:

Step 1: Using the Class D address, identify the low order 23 bits of the class D address.

Step2: Map those 23 bits into the low order 23 bits of a MAC address with the fixed high order bits of the IANA's IEEE multicast addressing space prefixed by 01:00:5E.

In this scenario this translates to 01-00-5E-00-10-6F

Incorrect Answers:

The Multicast MAC prefix is always 01-00-5e.

Reference: Whitepaper, Enterasys LAN Switching, Deploying IP Multicast Switching Method of assuring globally unique MAC address mappings in an IP multicast environment.

QUESTION NO: 10

A network administrator at TestKing assigns a multicast address of 239.255.8.5 to an application running on a device with an Ethernet MAC address of 01.b2.7d.05.f1.80. Which Layer 2 multicast address will this particular device use?

- A. 01.00.5e.7f.08.05
- B. 01.b2.7d.05.f1.80
- C. 01.b2.7d.0a.08.05
- D. 01.00.5e.05.f1.80
- E. ff.ff.ff.ff.ff.ff

Answer: A

Explanation:

Once an application determines the class D IP multicast address it will utilize, that address must be mapped into a MAC layer multicast for delivery across any LAN based system. This process is outlined as follows:

Step 1: Using the Class D address, identify the low order 23 bits of the class D address.

Step2: Map those 23 bits into the low order 23 bits of a MAC address with the fixed high order bits of the IANA's IEEE multicast addressing space prefixed by 01:00:5E.

In this scenario the lower 23 bits translates to 127.8.5 (leaving out the first bit in 255 leaves us with 127) which translates to 7F-08-05 in hex so the final multicast MAC address is 01-00-5E-7F-08-05.

QUESTION NO: 11

Which of the following layer-2 hardware addresses is a valid multicast MAC address?

- A. 00-00-00-FA-11-67
- B. 01-00-E0-56-AE-3C
- C. 00-01-E0-AB-B2-C1
- D. 01-00-5E-0A-08-CF
- E. FF-FF-FF-FF-FF-FF

Answer: D

Explanation:

Media Access Control (MAC) layer addresses within Ethernet are 48 bit addresses. These 48 bits comprise 24 bits for the Organizational Unit Identifier (OUI) and 24 bits for serial number of the card, which becomes the remainder of the unique address.

The address of a multicast group does not relate to a physical device, but rather to a transient group of devices; therefore, the MAC address format uses a special OUI. The OUI for IPv4 Multicast is 01:00:5E with the Least Significant Bit Most Significant Byte set. Only half of this address space was allocated for IP Multicast. Therefore, all MAC addresses that start with 01-00-5E are multicast MAC addresses.

QUESTION NO: 12

Which of the following technologies manages layer 2 multicast traffic by configuring Layer 2 LAN interfaces dynamically to forward multicasts only to the interfaces that want to receive it?

- A. IGMP
- B. IGMP snooping
- C. PIM-DIM
- D. DVMRP
- E. MOSPF

Answer: B

Explanation:

Understanding IGMP Snooping:

In subnets where you have configured either IGMP (see "[Configuring IP Multicast Layer 3 Switching](#)") or the IGMP querier (see the "[Enabling the IGMP Querier](#)" section), IGMP snooping manages multicast traffic at Layer 2 by configuring Layer 2 LAN interfaces dynamically to forward multicast traffic only to those interfaces that want to receive it.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a00800f4ff2.html

QUESTION NO: 13

Which of the following multicast address is reserved for sending out to all the hosts on a subnet?

- A. 224.0.0.1
- B. 224.0.0.2
- C. 224.0.0.255
- D. 239.0.0.255

Answer: A

Explanation:

224.0.0.1 is the *all-hosts* group. If you ping that group, all multicast capable hosts on the network should answer, as every multicast capable host *must* join that group at start-up on all its multicast capable interfaces. Some of the well known IP multicast addresses are displayed in the following table:

Class D Address	Purpose
224.0.0.1	All hosts on a subnet
224.0.0.2	All routers on a subnet
224.0.0.4	All DVMRP routers
224.0.0.5	All MOSPF routers
224.0.0.9	Routing Information Protocol (RIP)-Version 2
224.0.1.1	Network Time Protocol (NTP)
224.0.1.2	SGI Dogfight
224.0.1.7	Audio news
224.0.1.11	IETF audio
224.0.1.12	IETF video
224.0.0.13	Protocol Independent Multicasting

Incorrect Answers:

B: The multicast IP address 224.0.0.2 is reserved for All Routers on this Subnet

C: 224.0.1.27-224.0.1.255 Unassigned

D: 239.0.0.255 does not apply.

Reference:

RFC 1458, Host Extensions for IP Multicasting

RFC 1700, ASSIGNED NUMBERS

<http://www.tldp.org/HOWTO/Multicast-HOWTO-2.html>

QUESTION NO: 14

You work as a network administrator at TestKing and you built an IP multicast domain using PIM.

Your CTO asks you if you know which mode assumes that the number of actual end-users of the multicast traffic is relatively small. How would you respond?

- A. PIM-DM
- B. PIM-SM
- C. PIM-RP
- D. CGMP
- E. IGMP snooping

Answer: B

Explanation:

PIM-SM stands for protocol independent multicast – sparse mode; and it is designed for ‘receiver initiated multicast group membership’, so it is indeed sparse in its multicasts.

PIM Dense Mode (PIM-DM) uses a fairly simple approach to handle IP multicast routing. The basic assumption behind PIM-DM is that the multicast packet stream has receivers at most locations. An example of this might be a company presentation by the CEO or President of a company. By way of contrast, PIM Sparse Mode (PIM-SM) assumes relatively fewer receivers. An example would be the initial orientation video for new employees.

Reference: <http://www.netcraftsmen.net/welcher/papers/multicast02.html>

QUESTION NO: 15

Which of the following statements is true about the default IP multicast settings on a Cisco 5000 switch?

- A. IGMP snooping enabled and PIM enabled
- B. IGMP snooping enabled and CGMP enabled
- C. IGMP enabled and IGMP snooping disabled
- D. IGMP snooping enabled and CGMP disabled
- E. IGMP snooping disabled and CGMP disabled.

Answer: E

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

IGMP snooping must be enabled. It is disabled by default. CGMP is also disabled by default as well.

Incorrect Answers

B: CGMP is not enabled by default. It must be manually enabled. Furthermore, you cannot enable CGMP on a switch if IGMP snooping is already enabled on that switch.

C: CGMP is disabled by default.

A, D: IGMP snooping is disabled by default.

QUESTION NO: 16

You are an administrator at TestKing. Apprentice Tess tells you that there's a Cisco switching technology that can control multicast traffic so it's only delivered to the switch ports that are enabled as multicast clients. Which technology is she talking about?

- A. PIM
- B. VTP
- C. IGMP
- D. CGMP

Answer: D

Explanation:

The key to this question lies with the fact that we are looking for a switching technology. Cisco Group Management Protocol (CGMP) is a Cisco Proprietary protocol which enables IP multicasting at layer 2 on Cisco's Catalyst switches that do not distinguish between IP multicast data packets and IGMP Report messages, which are both MAC level addressed to the same group address. CGMP performs same tasks as IGMP, and was designed to work with other layer 3 multicast protocols.

Incorrect Answers:

A: PIM (Protocol-Independent Multicast) handles the transmission of multicast packets to all hosts in the multicast group while preventing loops and wasted bandwidth. The two main PIM modes are sparse and dense. PIM is fundamentally used with routers.

B: VTP is the VLAN Trunking Protocol. It shares VLAN information between switches. It is not used for multicast.

C: Routers use the Internet Group Management Protocol (IGMP) to learn whether members of a group are present on directly attached subnets, and whether or not to forward multicast packets onto those networks. Hosts join multicast groups by sending IGMP report messages.

QUESTION NO: 17

What is true about CGMP?

- A. IGMP snooping must be disabled before you can enable CGMP.

- B. PIM must be disabled on an interface where CGMP is enabled.
- C. CGMP ensures that all switch ports receive all multicast packets.
- D. A CGMP-enabled switch summarizes IGMP information for connected routers.

Answer: A

Explanation:

Before you enable CGMP on a switch, you must disable IGMP snooping if it is enabled. If you try to enable CGMP without first disabling IGMP snooping, an error message is generated.

Note: CGMP (Cisco Group Management Protocol) was first implemented by Cisco to restrain multicast traffic in a layer 2 network.

You can configure the switch to either snoop on Protocol Independent Multicast/Distance Vector Multicast Routing Protocol (PIM/DVMRP) packets or to listen to CGMP self-join packets.

Incorrect Answers:

B: There is no requirement to disable PIM (Protocol Independent Multicast) on the CGMP interface.

C: The idea of CGMP is to restrain multicast traffic.

D: This is not the way it works.

In CGMP the multicast router is considered to be the server since it has done all the work and the layer 2 switch is the CGMP client that uses the router's information to construct its forwarding tables (CAM).

Reference: Configuring Multicast Services

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_4_2/config/multi.htm

QUESTION NO: 18

CGMP has been configured on switch TK1. What CGMP information do CGMP-enabled switches and routers exchange?

- A. CAM table changes
- B. summarized IGMP information
- C. multicast join and leave events
- D. multicast group to port assignment

Answer: C

Explanation:

When the CGMP-capable router receives an IGMP control packet, it creates a CGMP packet that contains the request type (either join or leave), the multicast group address, and the Media Access Control (MAC) address of the host. The router sends the CGMP packet to a well-known address to which the CGMP-enabled switches listen.

Reference: Configuring Multicast Services

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_4_2/config/multi.htm

QUESTION NO: 19

You want to configure CGMP on switch TK1. Which of the following statements is correct for configuring Cisco Group Management Protocol on this switch?

- A. PIM must be configured on the CGMP router.
- B. Directed broadcasts must be disabled on all CGMP switches.
- C. CGMP must be enabled separately for each VLAN where it is desired.
- D. The switch must be configured with the ip addresses of all neighboring routers.

Answer: A

Explanation:

Protocol Independent Multicast (PIM) is one of the required elements for multicast configuration. It enables IGMP on the router and allows it to receive and forward traffic on the specified interface. PIM must be enabled on every interface that is to participate in the multicast network.

Incorrect Answers:

B: There is no requirement to disable directed broadcasts.

C: CGMP is enabled for a device, not for each VLAN.

The “set cgmp enable” command is used to enable CMGP support for IP multicast on a switch.

D: This is not required.

Reference: Sybex CCNP 640-504, Enabling PIM on an interface, page 391

QUESTION NO: 20

The TestKing network is migrating from IGMP version 1 to IGMP version 2. What message type was added in IGMPv2?

- A. Heartbeat
- B. Join request
- C. Leave report
- D. Status report

Answer: C

Explanation:

A new IGMP Type was created for the IGMPv2 Leave Group message. Explicit leave messages were not available in IGMP version 2.

Note: The Leave and Group-Specific messages work together to allow a host to remove itself from the multicast group immediately without interrupting the state of the interface on the multicast router.

Reference: RFC 2236, Internet Group Management Protocol version 2 (IGMPv2), Appendix I - Changes from IGMPv1.

QUESTION NO: 21

The TestKing network is migrating from IGMP version 1 to IGMP version 2. Internet Group Management Protocol (IGMP) Version 2 (RFC 2236) has improvements over IGMP (RFC 1112). What was improved in version 2?

- A. Leave and join latencies
- B. The potential for infinite loops
- C. Mapping multicast IP addresses to MAC addresses
- D. Lack of coordination between Layer 2 and Layer 3 devices.

Answer: A

Explanation:

The Leave process in version 2 was included to avoid long time-outs that are experienced in version 1. With leave messages, hosts can immediately inform network devices that they no longer wish to receive multicast traffic for a particular session. In version 1, hosts would simply leave and the network devices would need to time out before traffic stopped being forwarded to these hosts.

Reference: RFC2236, Internet Group Management Protocol Version 2

QUESTION NO: 22

IP multicast is being utilized in the TestKing network. Which of the following statements best describes the way a multicast session operates?

- A. A server sends one copy of each packet to a Class D address.
- B. A web server transmits separate content to each client.
- C. The application server services each client connection individually.
- D. A router sends the protocol information to all clients on the network.
- E. None of the above.

Answer: A

Explanation:

IP Multicast uses Class D address as destination. The sending device sends traffic destined to these class D IP addresses. This is used to improve efficiencies, since now a single sender can send traffic to multiple receivers.

Incorrect Answers:

- B: Multicast does not include a web server. Furthermore, all selected clients receive the same information.
- C: All selected clients receive the same information.
- D: Only selected clients receive the information, not every client.

QUESTION NO: 23

Which of the protocols below has these features?

- It multicasts information.
- It is a Cisco switching technology.
- It forwards multicast packets by using multicast information obtained from routers to improve efficiency.

- A. PIM
- B. CDP
- C. IGMP
- D. CGMP

Answer: D

Explanation:

CGMP (Cisco Group Management Protocol) was first implemented by Cisco to restrain multicast traffic in a layer 2 network. CGMP operates between the switch and the router and is able to use information obtained from routers.

Incorrect Answers:

- A: You can configure the switch to either snoop on Protocol Independent Multicast/Distance Vector Multicast Routing Protocol (PIM/DVMRP) packets or to listen to CGMP self-join packets. PIM is not Cisco proprietary.
- B: CDP (Cisco Discovery Protocol) is used by Cisco devices to learn about the neighboring Cisco devices. It is a layer two technology used to discover neighbour information, such as neighbour IP address, device name, and device type. It is not used for multicasts.
- C: IGMP is not a Cisco technology.

References: Configuring Multicast Services

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_4_2/config/multi.htm

RFC 1112, Host Extensions for IP Multicasting

QUESTION NO: 24

CGMP has been enabled on switch TKS1 as well as router TKR1. How does Cisco Group Management Protocol (CGMP) operate?

- A. The router broadcasts CGMP frames to all CGMP-enabled interfaces.
- B. The router forwards all IGMP control packets to CGMP-enabled switches.
- C. The router adds each multicast host MAC address to its CAM table for the destined port.
- D. The router forwards CGMP packets to a well-known address to which all CGMP switches listen.

Answer: D

Explanation:

Cisco Group Management Protocol (CGMP) limits the forwarding of IP multicast packets to only those ports associated with IP multicast clients. These clients automatically join and leave groups that receive IP multicast traffic, and the switch dynamically changes its forwarding behavior according to these requests. CGMP was first implemented by Cisco to restrain multicast traffic in a layer 2 network. CGMP frames are Ethernet frames with the destination MAC address: 01-00-0c-dd-dd-dd.

Reference: Multicast in a Campus Network: CGMP and IGMP Snooping
<http://www.cisco.com/warp/public/473/22.html#CGMP>

QUESTION NO: 25

How can IGMP snooping affect performance on a Catalyst 2950 switch?

- A. Low performance occurs when inbound bandwidth is exceeded.
- B. Low performance occurs when outbound bandwidth is exceeded.
- C. Low performance occurs when heavy traffic is present.
- D. IGMP snooping is performed at wire speed and does not affect switch performance.

Answer: D

Explanation:

With a switching fabric of 13.6 Gbps and a maximum forwarding bandwidth of 13.6 Gbps, Cisco Catalyst 2950 Series switches deliver wire-speed performance on all ports in connecting end stations and users to the company LAN. Cisco Catalyst 2950 Series switches with basic services support performance-boosting features such as Cisco Fast EtherChannel[®] to provide high-performance bandwidth between Cisco Catalyst switches, routers, and servers.

To provide efficient use of resources for bandwidth-hungry applications like multicasts, Cisco Catalyst 2950 Series switches support Internet Group Management Protocol Version 3 (IGMPv3) snooping in hardware. Through the support and configuration of IGMP snooping through the Cisco CMS software, these Cisco Catalyst 2950 Series switches deliver outstanding performance and ease of use in administering and managing multicast applications on the LAN.

The IGMPv3 snooping feature allows the switch to "listen in" on the IGMP conversation between hosts and routers. When a switch hears an IGMP join request from a host for a given multicast group, the switch adds the host's port number to the group destination address list for that group. And when the switch hears an IGMP leave request, it removes the host's port from the content-addressable memory (CAM) table entry.

QUESTION NO: 26

You want to disable CGMP on switch TK1. Which command would you enter if you had to disable CGMP on a non-IOS command-based switch? (Type in answer below)

Answer: set cgmp disable

Explanation:

Remember, set based commands are used in non-IOS based switches. CGMP was first implemented by Cisco to restrain multicast traffic in a layer 2 network. Because a switch is, by essence, not capable of looking at layer 3 packets, it cannot distinguish an IGMP packet. With CGMP, the router provides the interface between the hosts. The routers "talk" IGMP and the switches "talk" CGMP.

To enable CGMP use the "set cgmp enable" command. To disable is, use the "set cgmp disable" command.

QUESTION NO: 27

CGMP is being used on the TestKing network. For CGMP to operate correctly on a switch, what must it be connected to?

- A. A switch running EIGRP
- B. A switch running IGRP
- C. A switch running EGMP
- D. A router running IGRP
- E. A router running CGMP

Answer: E

Explanation:

CGMP was first implemented by Cisco to restrain multicast traffic in a layer 2 network. Because a switch is, by essence, not capable of looking at layer 3 packets, it cannot distinguish an IGMP packet. With CGMP, the router provides the interface between the hosts. The routers "talk" IGMP, and the switches "talk" CGMP with the routers.

QUESTION NO: 28

Before an administrator can set up a multicast, you will need to specify an addressing scheme. Which of the following statements is true concerning IP addressing schemes? (Select all that apply)

- A. Class E addresses are reserved
- B. Class E addresses are allocated dynamically
- C. Class D addresses are allocated dynamically
- D. Class D addresses are allocated manually
- E. Class D addresses are reserved

Answer: C, E

Explanation: Class D addresses are reserved for multicasts using the range 224.0.0.1-239.255.255.255, but they can also be allocated dynamically. (Statically allocated addresses are reserved for specific protocols that require well known address.)

Reference: CCNP Switching Exam Certification Guide David Hucaby & Tim Boyles, pages 343-344

QUESTION NO: 29

CGMP is being utilized on the TestKing network. What kind of information do CGMP-enabled routers and switches exchange?

- A. Summarized IGMP information.
- B. Multicast group to port assignments.
- C. Multicast join and leave events.
- D. CAM table changes.

Answer: C

Explanation:

CGMP is based on a client/server model. The router is considered a CGMP server, with the switch taking on the client role. The basis of CGMP is that the IP multicast router sees all IGMP packets and therefore can inform the switch when specific hosts join or leave multicast groups. The switch then uses this information to construct a forwarding table.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 306

QUESTION NO: 30

Which protocol(s) prevents switches from flooding multicast traffic out of every port, except the source port?

- A. Internet Group Management Protocol Version 1 (IGMPv1)
- B. Protocol Independent Multicast (PIM)
- C. IP Multicast Routing
- D. Cisco Group Management Protocol (CGMP)
- E. Internet Group Management Protocol Version 2 (IGMPv2)

Answer: D


Explanation:

CGMP is a Cisco-developed protocol that allows Catalyst switches to leverage IGMP information on Cisco routers to make Layer 2 forwarding decisions. *CGMP* must be configured both on the multicast routers and on the Layer 2 switches. The net result is that with *CGMP*, IP multicast traffic is delivered only to those Catalyst switch ports that are interested in the traffic. All other ports that have not explicitly requested the traffic will not receive it.

The basic concept of *CGMP* is shown in Figure 43-7. When a host joins a multicast group (part A), it multicasts an unsolicited IGMP membership report message to the target group (224.1.2.3, in this example). The IGMP report is passed through the switch to the router for the normal IGMP processing. The router (which must have *CGMP* enabled on this interface) receives this IGMP report and processes it as it normally would, but in addition it creates a *CGMP* join message and sends it to the switch.

The switch receives this *CGMP* join message and then adds the port to its content addressable memory (CAM) table for that multicast group. Subsequent traffic directed to this multicast group will be forwarded out the port for that host. The router port is also added to the entry for the multicast group. Multicast routers must listen to all multicast traffic for every group because the IGMP control messages are also sent as multicast traffic. With *CGMP*, the switch must listen only to *CGMP* join and *CGMP* leave messages from the router. The rest of the multicast traffic is forwarded using its CAM table exactly the way the switch was designed.

Figure 43-7: Basic CGMP Operation



CGMP allows for switches to selectively determine which ports to send the IP multicast data, rather than flooding the multicast traffic to all ports except the source port which is the normal multicast behavior of a switch.

Note: IGMP snooping always prevents this as described in this question and would also be a valid answer, had it been a choice in this question.

QUESTION NO: 31

Your goal is to enable Cisco Group Management Protocol (CGMP) on an interface on your Catalyst 5000 switch. So you connect to the switch and enter interface configuration mode. Which IOS command should you use next?

- A. cgmp
- B. ip cgmp
- C. cgmp enable
- D. ip igmp enable cgmp
- E. set cgmp enable

Answer: B

Explanation:

In this scenario CGMP is already enabled on the switch. We must enable it on an interface. The **ip cgmp** interface configuration command is used to enable Cisco Group Management Protocol (CGMP) on an interface. This command is used only on IOS based devices. For native catalyst operation, the “set cgmp enable” command would be used on the switch.

QUESTION NO: 32

You are tasked with determining the best multicast routing protocol to use on the TestKing network. Which of the following routing protocols are classified as dense-mode multicast routing protocols? (Select three.)

- A. PIM-SM
- B. PIM-DM
- C. MOSPF
- D. OSPF
- E. DVMRP

Answer: B, C, E

Explanation:

Dense mode routing protocols include the following:

- Distance Vector Multicast Routing Protocol (DVMRP)
- Multicast Open Shortest Path First (MOSPF)
- Protocol-Independent Multicast Dense Mode (PIM DM)

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 313

Section 15: Describe the quality issues with voice traffic on a switched data network, including jitter and delay (3 questions)

QUESTION NO: 1

Jitter is causing problems with the VOIP application in the TestKing network. What causes network jitter?

- A. Variable queue delays
- B. Packet drops
- C. Transmitting too many small packets
- D. Compression

Answer: A

Delay variation or jitter is the difference in the delay times of consecutive packets. A jitter buffer is often used to smooth out arrival times, but there are instantaneous and total limits on buffering ability. Any type of buffering used to reduce jitter directly increases total network delay. In general, traffic requiring low latency also requires a minimum variation in latency.

Note: Jitter in Packet Voice Networks:

Jitter is defined as a variation in the delay of received packets. At the sending side, packets are sent in a continuous stream with the packets being spaced evenly apart. Due to network congestion, improper queuing, or configuration errors, this steady stream can become lumpy, or the delay between each packet can vary instead of remaining constant.

QUESTION NO: 2

VOIP is being implemented on the TestKing network. In a properly designed network, what is the maximum amount of time a voice package should spend crossing a network?

- A. 90 milliseconds
- B. 120 milliseconds
- C. 150 milliseconds
- D. 240 milliseconds

Answer: C

Explanation:

Delay is the time it takes for VoIP packets to travel between two endpoints and you should design networks to minimize this delay. However, because of the speed of network links and the processing power of intermediate devices, some delay is expected. The human ear normally accepts up to about 150 milliseconds (ms) of delay without noticing problems (the ITU standard recommends no more than 150 ms of one-way delay).

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps5014/products_feature_guide09186a00800880e7.html

QUESTION NO: 3

Which of the following network problems would indicate a need to implement QoS features? (Select three)

- A. mis-routed packets
- B. excess jitter
- C. delay of critical traffic
- D. packet loss due to congestion
- E. data link layer broadcast storms
- F. ftp connections unsuccessful

Answer: B, C, D

Explanation:

Loss, jitter, and delay are the three reasons for implementing QoS features on modern networks. Loss is when a packet disappears on a network. Jitter is a timing mismatch between two way traffic, and delay is when a packet takes too long to get somewhere.

Incorrect Answers:

A: This would indicate a routing problem, or packets being “black-holed.” QoS would not help in this situation.

E: Broadcast storms indicate a problem on a LAN segment, such as a babbling host, too many hosts, a segment that is too large, a bad application, etc. QoS would not help in this situation.

F: If only FTP sessions were having issues, then the FTP application or FTP server should be corrected. Normally, FTP sessions are not delay sensitive due to the re-transmission nature of TCP and do not require QoS.

Section 16: Describe the QoS solutions that address voice quality issues (6 questions)

QUESTION NO: 1

Which QoS mechanisms can you use on a converged network to improve VoIP quality? (Select three)

- A. The use of a queuing method that will give VoIP traffic strict priority over other traffic.
- B. The use of RTP header compression for the VoIP traffic.
- C. The proper classification and marking of the traffic as close to the source as possible.
- D. The use of 802.1QinQ trunking for VoIP traffic.
- E. The use of WRED.

Answer: A, C, E

Explanation:

In order to optimize the quality of VOIP calls, QoS should be implemented to ensure that VOIP traffic is prioritized over other traffic types.

By providing a strict queue for VOIP traffic, you will ensure that voice calls take precedence over the other traffic types.

In order to properly provide for QoS across the network, the voice traffic should be marked to give priority as close to the source as possible. This will ensure that the traffic is prioritized end to end.

Finally, WRED (Weighted Random Early Detection) could be configured to prevent congestion. WRED can be used to selectively drop less important traffic types, instead of dropping the voice packets when links become busy.

Incorrect Answers:

B: Compression can be used to lower the bandwidth required to transmit VOIP calls, but it will not help with improving the voice quality. In general, compression of any kind lowers the quality of VOIP.

D. The trunking method used will have no bearing on the VOIP quality.

QUESTION NO: 2

The TestKing is rolling out Cisco's Architecture for Voice, Video and Integrated Data (AVVID). Which of the following choices represent the fundamental intelligent network services in Cisco's AVVID? (Select all that apply.)

- A. Quality of Service (QoS)

- B. Intelligent platforms
- C. Mobility and scalability
- D. Security
- E. High availability

Answer: A, C, D, E

Explanation:

By creating a robust foundation of basic connectivity and protocol implementation, Cisco AVVID Network Infrastructure addresses five primary concerns of network deployment:

- High availability
- Quality of service (QoS)
- Security
- Mobility and
- Scalability

Reference:

http://www.cisco.com/en/US/netsol/netwarch/ns19/ns24/networking_solutions_audience_business_benefit09186a008009d678.html

QUESTION NO: 3

Which of the characteristics below is associated with the (QoS) Integrated Services Model?

- A. QoS classified at layer 3 using IP precedence or DSCP.
- B. Guaranteed rate service.
- C. Implemented using FIFO queues.
- D. All traffic has an equal chance of being dropped.

Answer: B

Explanation:

Cisco IOS QoS includes the following features that provide controlled load service, which is a kind of integrated service:

- ◆ Resource Reservation Protocol (RSVP) can be used by applications to signal their QoS requirements to the router.
- ◆ Intelligent queuing mechanisms can be used with RSVP to provide the following kinds of services:
 - Guaranteed Rate Service, which allows applications to reserve bandwidth to meet their requirements. For example, a Voice over IP (VoIP) application can reserve 32 Mbps end to end using this kind of service. Cisco IOS QoS uses weighted fair queuing (WFQ) with RSVP to provide this kind of service.
 - Controlled Load Service, which allows applications to have low delay and high throughput even during times of congestion. For example, adaptive real-time applications

such as playback of a recorded conference can use this kind of service. Cisco IOS QoS uses RSVP with Weighted Random Early Detection (WRED) to provide this kind of service.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_configuration_guide_chapter09186a008007ff07.html#1000946

QUESTION NO: 4

Which of the following QoS technologies classifies traffic entering a queue and services traffic based on dynamically observed traffic flows?

- A. FIFO
- B. WFQ
- C. Custom Queuing
- D. Priority Queuing

Answer: B

Explanation:

WFQ stands for weighted fair queuing. WFQ is one of Cisco's premier queuing techniques. It is a flow-based queuing algorithm that does two things simultaneously: It schedules interactive traffic to the front of the queue to reduce response time, and it fairly shares the remaining bandwidth between high bandwidth flows. By doing this, it can ensure fair treatment for all traffic types, while ensuring that small transmissions are serviced in a timely manner, rather than waiting for a long data transfer to finish.

Incorrect Answers:

A: This is the First In First Out method, which services traffic just as the name implies.

C, D: These methods service the traffic flows based on the configuration parameters. These configured settings are static, and do not change dynamically.

Reference:

http://www.cisco.com/en/US/tech/tk543/tk544/tk718/tech_protocol_home.html

QUESTION NO: 5

Which Cisco strategy employs Storage Networking?

- A. content delivery
- B. directory service
- C. AVVID
- D. File service

Answer: C

Explanation:

AVVID - As an element of Cisco AVVID (Architecture for Voice, Video and Integrated Data), Cisco Storage Networking allows customers to adopt a strategy for accessing, managing and

protecting their growing information resources across a consolidated IP, Gigabit Ethernet, Fiber Channel, and optical network infrastructure.

Reference: http://www.cisco.com/en/US/partners/pr46/pr13/partners_pgm_brochure.html

QUESTION NO: 6

The TestKing network is using WRED for some QoS functions. Which characteristic best describes weighted random early detection (WRED)?

- A. It is well suited for UDP packets because of their connectionless state.
- B. It takes advantage of the TCP congestion control mechanism and drops packets selectively based on IP precedence.
- C. It should be used in multi-protocol environments where IPX/SPX packet can be dropped instead of UDP packets.
- D. It ensures that flows that use the least bandwidth are more likely to have packets dropped.
- E. None of the above.

Answer: B

Explanation:

RED and WRED are congestion avoidance mechanisms. Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network and internetwork bottlenecks before it becomes a problem. These techniques are designed to provide preferential treatment for premium (priority) class traffic under congestion situations while concurrently maximizing network throughput and capacity utilization and minimizing packet loss and delay. WRED and DWRED are the Cisco IOS QoS congestion avoidance features.

Random Early Detection (RED) is a congestion avoidance mechanism that takes advantage of TCP's congestion control mechanism. Weighted RED (WRED) generally drops packets selectively based on IP precedence. Packets with a higher IP precedence are less likely to be dropped than packets with a lower precedence. WRED is a derivative of RED that also uses TCP's congestion control mechanism.

Incorrect Answers:

A: WRED and RED take advantage of the inherent congestion control mechanisms of TCP, not UDP.

C: WRED is primarily useful only on IP networks, since it takes advantage of TCP functionality.

D: This choice best describes the functionality of weighted fair queuing (WFQ), not WRED.

Section 17: Describe the features and operation of network analysis modules on Catalyst switches to improve network traffic management (8 questions)

QUESTION NO: 1

What is true about a SPAN (switch port analyzer) session?

- A. Affects switching traffic on source ports.
- B. Associates multiple source interfaces with a single destination interface.
- C. Eliminates multiple copies of packets.
- D. Associates a source interface with multiple destination interfaces.
- E. None of the above.

Answer: B

Explanation:

The Switched Port Analyzer (SPAN) feature, sometimes called port mirroring or port monitoring, selects network traffic for analysis by a network analyzer such as a SwitchProbe device or other Remote Monitoring (RMON) probe. It operates by mirroring information from a source port or multiple source ports, and sending the copied information to the port that is defined as the analyzer port. A single SPAN port can be used to monitor the traffic from multiple source ports.

Incorrect Answers:

- A: Traffic on the source port is merely copied, so original traffic is not affected.
- C: SPAN can be used to create multiple copies, not eliminate them.
- D: Only one port can be configured as the SPAN destination port.

QUESTION NO: 2

Which of the following statements are true with SPAN? (Select three)

- A. A destination port can be a source port.
- B. A destination port can participate in only one SPAN session at a time.
- C. A destination port can be an EtherChannel group.
- D. A source port can be monitored in multiple SPAN sessions.
- E. A source port can be an EtherChannel group.
- F. A source port can be a destination port.

Answer: B, D, E

When configuring SPAN:

There can only be one destination port, and this destination port can only participate in one SPAN session at a single time. However, multiple SPAN sessions can be configured, and a source port can participate in multiple SPAN sessions.

An EtherChannel does not form if one of the ports in the bundle is a SPAN destination port. If you try to configure this, the switch tells you:

*“Channel port cannot be a Monitor Destination Port
Failed to configure span feature”*

A port in an EtherChannel bundle can be used as a SPAN source port.

Incorrect Answers:

A, F: A port cannot double as a destination port and a source port. It must be one or the other.

C: The following restrictions apply for ports that have port-monitoring capability:

- A monitor port cannot be in a Fast EtherChannel or Gigabit EtherChannel port group.
- A monitor port cannot be enabled for port security.
- A monitor port cannot be a multi-VLAN port.
- A monitor port must be a member of the same VLAN as the port monitored. VLAN membership changes are disallowed on monitor ports and ports being monitored.
- A monitor port cannot be a dynamic-access port or a trunk port. However, a static-access port can monitor a VLAN on a trunk, a multi-VLAN, or a dynamic-access port. The VLAN monitored is the one associated with the static-access port.
- Port monitoring does not work if both the monitor and monitored ports are protected ports

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps700/products_tech_note09186a008015c612.shtml

QUESTION NO: 3

You are a professor at the TestKing Academy, and a student asks you to describe ‘NetFlow traffic flow’ to her. How would you respond?

- A. It is a sequence of packets between a particular source and destination.
- B. It is a uni-directional sequence of packets between a particular source and destination.
- C. It is a bi-directional sequence of packets between a particular source and destination.
- D. It is a multi-directional sequence of packets between a particular source and destination.

Answer: A**Explanation:**

A NetFlow export-enabled device is one that has been configured to operate with Cisco IOS NetFlow Services software in a way that enables the device to export information about traffic flows between communicating end nodes in a network.

For NetFlow data export, traffic flows in a network have the following attributes in common:

- Source and destination autonomous system (AS) numbers
- Source and destination IP addresses
- Source and destination application port numbers
- Input interfaces
- IP type of services (ToS)
- IP protocol

Reference:

http://www.cisco.com/en/US/products/sw/netmgmtsw/ps1974/products_installation_guide_chapter09186a0080080774.html

QUESTION NO: 4

Which of the following SPAN implementations is designed to support source ports, source VLANs, and destination ports across different switches?

- A. RVSPAN
- B. MSPAN
- C. VSPAN
- D. RSPAN

Answer: D

RSPAN is an implementation of SPAN designed to support source ports, source VLANs, and destination ports across different switches.

Remote SPAN (RSPAN): Some source ports are not located on the same switch as the destination port. This is an advanced feature that requires a special VLAN to carry the traffic being monitored by SPAN between switches. RSPAN is not supported on all switches, so check the respective release notes or configuration guide to see if it can be used on the switch you are deploying.

Incorrect Answers:

A, B: This is an invalid SPAN type.

C: This is the VLAN-Based SPAN (VSPAN). On a given switch, the user can choose to monitor all the ports belonging to a particular VLAN in a single command.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps700/products_tech_note09186a008015c612.shtml#topic8-4

QUESTION NO: 5

Leading the way in IT testing and certification tools, www.testking.com

You are a CCNP in the midst of configuring a switching solution on a Token Ring network with some Catalyst 5000 series switches and you need to configure SPAN. Which of the following is true regarding the configuration of the token ring port?

- A. The source port must not be a single Token Ring port if the SPAN destination port is a Token Ring port
- B. There is nothing special to consider
- C. The source port must be a single Token Ring port if the SPAN destination port is a non-Token Ring port
- D. The source port must be a single Token Ring port if the SPAN destination port is a Token Ring port
- E. None of the above.

Answer: D

Explanation:

According to Cisco, follow these guidelines when configuring SPAN:

If the SPAN destination port is a Token Ring port, then the source port must be a single Token Ring port.

In software releases prior to 4.2, if the SPAN destination port is connected to another device, the port always receives incoming packets for the VLAN it is assigned to but does not participate in spanning tree for that VLAN. To avoid creating spanning tree loops, assign the SPAN destination port to an unused VLAN.

In software release 4.2 and later, incoming traffic on the SPAN destination port is disabled by default. You can enable it using the `inpkts enable` keywords. However, while the port receives traffic for its assigned VLAN, it does not participate in spanning tree for that VLAN. To avoid creating spanning tree loops with incoming traffic enabled, assign the SPAN destination port to an unused VLAN.

You cannot disable the reception of incoming packets on the destination SPAN port (using the `inpkts disable` keywords) on Token Ring SPAN destination ports.

QUESTION NO: 6

What is the purpose of a SPAN session on a switch?

- A. To identify the destination portion of a MAC source address.
- B. To identify the destination of ISL packets on the outbound switch.
- C. To identify the port that mirrors traffic to a protocol analyzer.
- D. To identify the destination of ISL packets on all other switches.
- E. To select network traffic for analysis
- F. To identify the destination for the spanning-tree BPDU.

Answer: E

Explanation:

SPAN (Switch Port Analyzer) selects network traffic for analysis by a Catalyst switch Network Analysis Module, a SwitchProbe device, or other RMON probe. SPAN mirrors traffic from one or more source ports (Ethernet, Fast Ethernet, Token Ring, or FDDI) on any VLAN to a destination port for analysis

QUESTION NO: 7

You are a CCNP and are configuring a switching solution with a set based (non IOS) switch. Which command would you enter to disable SPAN? (Type in answer below):

Answer: set span disable

Explanation:

SPAN selects network traffic for analysis by a Catalyst 5000 family Network Analysis Module, a SwitchProbe device, or other RMON probe. SPAN mirrors traffic from one or more source ports (Ethernet, Fast Ethernet, Token Ring, or FDDI) on any VLAN to a destination port for analysis. To disable SPAN, perform this task in privileged mode:

```
set span disable [dest_mod/dest_port | all]
```

QUESTION NO: 8

Switch TK1 has IP accounting enabled. Which command would you enter if you wanted to define filters to control the hosts for which IP accounting information is kept? (Type in answer below)

Answer: ip accounting-list

Explanation:

To define filters to control the hosts for which IP accounting information is kept, use the ip accounting-list global configuration command. To remove a filter definition, use the no form of this command.

```
“ip accounting-list ip-address wildcard”
```

```
“no ip accounting-list ip-address wildcard”
```

Section 18: Describe Transparent LAN Services and how they are implemented in a service provider network (1 question)

QUESTION NO: 1

An ISP provides transparent LAN services to interconnect five different TestKing locations.

What is true regarding this solution? (Select two)

- A. Broadcasts are sent to all sites.
- B. It is difficult to implement.
- C. Routers do not have to peer with each other.
- D. There are scalability issues with this solution.

Answer: A, C

Explanation:

Transparent LAN Service (TLS) provides Ethernet connectivity among geographically separated customer locations, creating a VLAN that spans those locations. Like traditional ethernet networks, broadcasts are sent across the network to all locations. Typically, enterprises deploy TLS within a metro area to interconnect multiple enterprise locations. However, TLS also can be extended to locations worldwide. Used this way, TLS converts wide-area connectivity into a VLAN so that the enterprise customer does not need to own and maintain customer premises equipment (CPE) with wide-area interfaces. Customers are freed from the burden of managing—or even knowing anything about—the WAN connection that links their separate LANs. The need for routers can be eliminated, since the network is one big LAN.

TLS is much less expensive and simpler to implement on Ethernet than on a Frame Relay or ATM network. The lower cost results primarily from lower equipment costs. Cost savings are a primary reason that TLS accounted for more than 63 percent of Metro Ethernet revenue in 2002 and will account for 60.3 percent in 2007, according to IDC.

Another advantage of implementing TLS on Ethernet is that service providers gain the flexibility to provision more bandwidth, and with varying quality-of-service (QoS) capabilities and service-level agreements (SLAs.)

Because TLS is a low-cost service, the service provider can use it as a draw for bundled services, which increase margins and strengthen the customer relationship. Typical value-added services include Ethernet interface to the Internet, storage transport, and data-center connectivity.

Topic 2: IMPLEMENTATION (135 questions)

Section 1: Convert CatOS to native IOS on Catalyst switches and manage native IOS images according to best practices (2 questions)

QUESTION NO: 1

On a Catalyst 2924XL, which command would you enter to set the port duplex?

- A. duplex
- B. set duplex
- C. port duplex
- D. set port duplex
- E. set duplex port

Answer: A

Explanation:

To configure the duplex operation on an interface, use the duplex command. Use the no form of this command to return the system to half-duplex mode.

Command:

duplex {full | half}

no duplex

Syntax Description

full	Specifies full-duplex operation.
half	Specifies half-duplex operation.

Default is half

Command Modes: Interface configuration

Incorrect Answers:

B, D, E: Since the 2924XL switch uses Cisco IOS, the “set” commands are not used.

C: This is an invalid command.

QUESTION NO: 2

You need to download a software image into switch TK1. Which of the following elements are needed to be able to download this image? (Select all that apply)

- A. network connection to a TFTP server
- B. the File Transfer Protocol
- C. the Trivial File Transfer Protocol
- D. network connection to a FTP server
- E. None of the above

Answer: A, C

Explanation:

You can download system software images to the switch using the Trivial File Transfer Protocol (TFTP). TFTP allows you to download system image files over the network from a TFTP server. Some modules, such as Catalyst 5000 family FDDI and ATM modules, have their own onboard Flash. When you download a software image file, the switch checks the header of the image file to determine the type of software image.

Section 2: Configure access ports for static and multi-VLAN membership (20 questions)

QUESTION NO: 1

VLAN maps have been configured on switch TK1. Which of the following actions are taken in a VLAN map that does not contain a match clause?

- A. Implicit deny feature at end of list.
- B. Implicit deny feature at start of list.
- C. Implicit forward feature at end of list
- D. Implicit forward feature at start of list.

Answer: A

Explanation:

Each VLAN access map can consist of one or more map sequences, each sequence with a match clause and an action clause. The match clause specifies IP, IPX, or MAC ACLs for traffic filtering and the action clause specifies the action to be taken when a match occurs. When a flow matches a permit ACL entry the associated action is taken and the flow is not checked against the remaining sequences. When a flow matches a deny ACL entry, it will be checked against the next ACL in the same sequence or the next sequence. If a flow does not match any ACL entry and at least one ACL is configured for that packet type, the packet is denied.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps700/products_configuration_guide_chapter09186a008007f4d4.html

QUESTION NO: 2

You have been tasked with configuring trunks throughout the TestKing switched LAN. Which mode must you choose if you want it to be in permanent trunking mode? (Select all that apply)

- A. No negotiate
- B. On
- C. Auto
- D. Desirable
- E. Off

Answer: A, B

Explanation:

The following table describes the various Trunking modes and their functions:

Ethernet Trunking Modes Mode	Function
on	Puts the port into permanent trunking mode and negotiates to convert the link into a trunk link. The port becomes a trunk port even if the neighboring port does not agree to the change.
off	Puts the port into permanent nontrunking mode and negotiates to convert the link into a nontrunk link. The port becomes a nontrunk port even if the neighboring port does not agree to the change.
desirable	Makes the port actively attempt to convert the link to a trunk link. The port becomes a trunk port if the neighboring port is set to on, desirable, or auto mode.
auto	Makes the port willing to convert the link to a trunk link. The port becomes a trunk port if the neighboring port is set to on or desirable mode. This is the default mode for all Ethernet ports.
nonegotiate	Puts the port into permanent trunking mode but prevents the port from generating DTP frames. You must configure the neighboring port manually as a trunk port to establish a trunk link.

QUESTION NO: 3

What is a characteristic of assigning a static VLAN membership?

- A. VMPS server lookup is required
- B. Easy to configure
- C. Easy of adds, moves, and changes
- D. Based on MAC address of the connected device

Answer: B

Explanation:

Static port VLAN membership on the switch is assigned manually by the administrator on a port-by-port basis.

Characteristics of static VLAN configurations include the following:

1. Secure
2. Easy to configure
3. Straight forward to monitor
4. Works well in networks where moves, adds, and changes are rare.

Incorrect Answers:

A: VMPS server lookups are a function of dynamic VLANs and are not used with statically assigned VLANs.

C: Moves, adds, and changes, would require a network administrator to change the configuration of the switch every time a change is required.

D: This would describe a function of dynamic VLAN configurations, where the MAC address of the end user determines the VLAN that it belongs in, instead of the physical port.

QUESTION NO: 4

Static VLANs are being used on the TestKing network. What is true about static VLAN's?

- A. Devices use DHCP to request their VLAN.
- B. Attached devices are unaware of any VLANs.
- C. Devices are assigned to VLANs based on their MAC addresses.
- D. Devices are in the same VLAN regardless of which port they attach to.

Answer: B

Explanation:

LAN port VLAN membership can be assigned manually on a port-by-port basis. When you assign LAN ports to VLANs using this method, it is known as port-based, or static, VLAN membership.

Attached devices will be unaware of any VLANs.

Incorrect Answers:

A: The DHCP service is not involved in VLAN assignment.

C: Devices are not assigned to VLAN based on their MAC addresses. This is a function of dynamic VLANs.

D: Static VLANs are configured on a port by port basis.

Reference: Configuring VLANs

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/121_8aex/swconfig/vlans.htm

QUESTION NO: 5

You are the network administrator at TestKing and have just applied this VLAN access map on one of your switches:

```
Router(config)#vlan access-map thor 10
```

```
Router(config-access-map)#match ip address net_10
Router(config-access-map)#action forward
Router(config-access-map)#exit
Router(config)#vlan filter thor vlan-list 12-15
```

What will this configuration result in?

- A. All VLAN 12 through 15 IP traffic matching net_10 is forwarded and all other IP packets are dropped.
- B. IP traffic matching net_10 is dropped and all other IP packets are forwarded to VLANs 12 through 15.
- C. IP traffic matching vlan-list 12-15 is forwarded and all other IP packets are dropped.
- D. All VLAN 12 through 15 IP traffic is forwarded, other VLAN IP traffic matching net_10 is dropped.

Answer: A

Explanation:

- **vlan access-map thor 10** Defines the VLAN access map. Optionally, you can specify the VLAN access map sequence number.
- **match ip address net_10** Configures the match clause in a VLAN access map sequence.
- **action forward** Configures the action clause in a VLAN access map sequence.
- **vlan filter thor vlan-list 12-15** Applies the VLAN access map to the specified VLANs. VLAN access maps can be applied to VLANs.

Each VLAN access map can consist of one or more map sequences, each sequence with a match clause and an action clause. The match clause specifies IP, IPX, or MAC ACLs for traffic filtering and the action clause specifies the action to be taken when a match occurs. When a flow matches a permit ACL entry, the associated action is taken and the flow is not checked against the remaining sequences. When a flow matches a deny ACL entry, it will be checked against the next ACL in the same sequence or the next sequence. If a flow does not match any ACL entry and at least one ACL is configured for that packet type, the packet is denied.

To use access-control for both bridged and routed traffic, you can use VACLs alone or a combination of VACLs and ACLs. You can define ACLs on the VLAN interfaces to use access-control for both the input and output routed traffic. You can define a VACL to use access-control for the bridged traffic.

Reference:

http://www.cisco.com/en/US/products/hw/routers/ps368/products_configuration_guide_chapter09186a008016113d.html

QUESTION NO: 6

When configuring a VLAN for dynamic membership; which of the following guidelines NOT required?

- A. Configure a VMPS server
- B. Turn off trunking on the port
- C. Turn off 802.1x port security
- D. Configure the spanning-tree PortFast feature
- E. All of the above are required for dynamic VLAN configuration

Answer: C

Reference:

Explanation:

Turning port security on or off is not necessary for enabling dynamic VLANs.

These guidelines and restrictions apply to dynamic port VLAN membership:

- **You must configure VMPS before you configure ports as dynamic.**
- **When you configure a port as dynamic, spanning-tree PortFast is enabled automatically for that port.** Automatic enabling of spanning-tree PortFast prevents applications on the host from timing out and entering loops caused by incorrect configurations. You can disable spanning-tree PortFast mode on a dynamic port.
- If you reconfigure a port from a static port to a dynamic port on the same VLAN, the port connects immediately to that VLAN. However, VMPS checks the legality of the specific host on the dynamic port after a certain period.
- Static secure ports cannot become dynamic ports. You must turn off security on the static secure port before it can become dynamic.
- Static ports that are trunking cannot become dynamic ports. **You must turn off trunking on the trunk port before changing it from static to dynamic.**

QUESTION NO: 7

What is true about access control on bridged and routed VLAN traffic? (Select three)

- A. Router ACLs can be applied to the input and output directions of a VLAN interface.
- B. Bridged ACLs can be applied to the input and output directions of a VLAN interface.
- C. Only router ACLs can be applied to a VLAN interface.
- D. VLAN maps and router ACLs can be used in combination.
- E. VLAN maps can be applied to a VLAN interface

Answer: A, B, D

Router ACLs are applied on interfaces as either inbound or outbound.

To filter both bridged and routed traffic, VLAN maps can be used by themselves or in conjunction with router ACLs.

VLAN ACLs, also called VLAN maps, which filter both bridged and routed packets. VLAN maps can be used to filter packets exchanged between devices in the same VLAN.

QUESTION NO: 8

Switch TestKing1 needs to have a port assigned to an existing VLAN. Which IOS command could you use to assign a switch port to a VLAN?

- A. switchport mode access
- B. switchport trunk access
- C. switchport access vlan
- D. switchport vlan

Answer: C

Explanation:

To assign a switchport to the VLAN, you would use the **switchport access vlan** interface configuration command.

Reference: CCNP Switching Exam Certification Guide: page 104, David Hucaby & Tim Boyles, Cisco Press 2001, ISBN 1-58720 000-7

QUESTION NO: 9

What's true with VLAN port associations? (Select all that apply)

- A. ASIC enhances the performance of the association
- B. VLAN membership is based on Port through port-to-VLAN association.
- C. Routing table enhances the performance of the association
- D. VLAN membership is based on Port through port-to-WAN ID association.

Answer: A, B

Explanation:

ASIC (Application Specific Integrated Circuits), layer 2 switches have ASIC chips to help them with wire speed hardware switching. With ASIC, the performance of this association is very high, and is more desirable than the complex routing table lookup type of operation.

VLAN membership is based on Port through port-to-VLAN association.

Incorrect Answers:

C: Routing tables are not consulted when transferring VLAN traffic since VLANs are handled at layer 2 and routing occurs at layer 3.

D: VLAN associations deal with layer two VLANs, not layer 3 WAN IDs.

QUESTION NO: 10


Which Cisco switch command would you use to map VLANs 10 to 20 to MST instance 1?

- A. Switch(config)#vlan 10-20 instance 1
- B. Switch(config)#instance 1 vlan 10-20
- C. Switch(config-mst)#vlan 10-20 instance 1
- D. Switch(config-mst)#instance 1 vlan 10-20
- E. None of the above

Answer: D

Explanation:

Beginning in privileged EXEC mode, follow these steps to specify the MST region configuration and enable MSTP. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst configuration	Enter MST configuration mode.
Step 3	instance <i>instance-id</i> vlan <i>vlan-range</i>	Map VLANs to an MST instance. <ul style="list-style-type: none"> • For <i>instance-id</i>, the range is 1 to 15. • For <i>vlan vlan-range</i>, the range is 1 to 4094. When you map VLANs to an MST instance, the mapping is incremental, and the VLANs specified in the command are added to or removed from the VLANs that were previously mapped. To specify a VLAN range, use a hyphen; for example, instance 1 vlan 1-63 maps VLANs 1 through 63 to MST instance 1. To specify a VLAN series, use a comma; for example, instance 1 vlan 10, 20, 30 maps VLANs 10, 20, and 30 to MST instance 1.
Step 4	name <i>name</i>	Specify the configuration name. The <i>name</i> string has a maximum length of 32 characters and is case sensitive.
Step 5	revision <i>version</i>	Specify the configuration revision number. The range is 0 to 65535.
Step 6	show pending	Verify your configuration by displaying the pending configuration.
Step 7	exit	Apply all changes, and return to global configuration mode.
Step 8	spanning-tree mode mst	Enable MSTP. RSTP is also enabled. 

		<p>Caution Changing spanning-tree modes can disrupt traffic because all spanning-tree instances are stopped for the previous mode and restarted in the new mode.</p> <hr/> <p>You cannot run both MSTP and PVST+ or both MSTP and rapid PVST+ at the same time.</p>
Step 9	end	Return to privileged EXEC mode.
Step 10	show running-config	Verify your entries.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps5528/products_configuration_guide_chapter09186a0080211377.html

QUESTION NO: 11

Token Ring VLANs are being used in some locations within the TestKing network. What's the default VLAN value on a token ring with default port assignments?

- A. VLAN 0
- B. VLAN 1
- C. VLAN 1003
- D. VLAN ON
- E. VLAN A
- F. None of the above

Answer: C**Explanation:**

As a rule on the Catalyst 3900, TrCRFs cannot span separate switches or stacks of switches. One exception to this rule is the *default* TrCRF. The default TrCRF can contain ports located on separate switches. By default, the Token Ring VLAN configuration on the Catalyst 3900 and the Catalyst 5000 series Token Ring modules has all ports assigned to the default TrCRF (1003). In turn, this default TrCRF is associated with the default TrBRF (1005), which can span switches via ISL. If a user does not configure the ports of a Token Ring module to be associated with a new TrCRF, traffic is passed between the default TrCRFs located on separate switches that are connected via ISL.

Reference:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/trsr2/vlan.htm>

QUESTION NO: 12

Which Cisco IOS command assigns a Catalyst switch port to VLAN 10?

- A. switchport mode vlan 10
- B. switchport trunk vlan 10
- C. switchport access vlan 10
- D. switchport mode access vlan 10

Answer: C

Explanation:**Switchport access:**

Use the switchport access interface configuration command to configure a port as a static-access port. The port operates as a member of the configured VLAN.

Use the no form of this command to reset the access mode to the default VLAN for the switch.

Syntax

```
switchport access vlan vlan-id  
no switchport access vlan vlan-id
```

Syntax Description

vlan *vlan-id*

ID of the VLAN. Valid IDs are from 1 to 1005. Do not enter leading zeroes.

Defaults

All ports are in static-access mode in VLAN 1.

Command Modes

Interface configuration.

Usage Guidelines

An access port can be assigned to only one VLAN.

When the no switchport access vlan form is used, the access mode is reset to static access on VLAN 1.

Example

The following example shows how to assign a port to VLAN 2 (instead of the default VLAN 1):

```
Switch(config-if)# switchport access vlan 2
```

You can verify the previous command by entering the **show interface** *interface-id* switchport command in privileged EXEC mode and examining information in the Administrative Mode and Operational Mode rows.

QUESTION NO: 13

You are connected to a Catalyst switch via a console cable as shown below:



You work as a systems administrator at the TestKing.com main office in the greater Toronto area. The number of employees on your floor has exceeded the infrastructure of your current network equipment. Your CTO has ordered a new switch chassis, but it's going to be another 6-8 working days until it arrives. In the meantime you can to connect 24 new workstations to an old Cisco Catalyst 2950, which your junior administrator has just finished erasing, and rebooting (to purge old VLAN information).

Your tasks are to:

- disable VTP
- Ensure that all non-trunking interfaces do not participate in Spanning Tree by default by globally configuring PortFast.

For the following two tasks, you are required to use global commands to configure the ports:

- Ensure all FastEthernet interfaces are in permanent non-trunking mode.
- Place FastEthernet interfaces 0/12 through 0/24 in VLAN 20.

Start by clicking on host TestKiA.

Answer:

```
enable
configure terminal
Switch(config)#vtp mode transparent           (disable vtp)
Switch(config)#spanning-tree portfast default (Globally, enable portfast on all ports)
Switch(config)#interface range fa0/1 - 24    (select interfaces)
Switch(config-if)#switchport mode access     (set ports for access mode, NOT Trunking)
exit
Switch(config-if)#interface range fa0/12 - 24 (select interfaces)
Switch(config-if)#switchport access vlan 20  (assign ports to vlan 20)
end
copy running-config startup-config (save configuration)
```

QUESTION NO: 14

Leading the way in IT testing and certification tools, www.testking.com

You are the network administrator of a network with the active VLANs: 1, 2, 3, 4, 10, 20, and 50. However, you only need to carry VLANs 1,2,10 and 20 on a trunk. Which of the following commands should you use to fulfil this requirement? (Select all that apply.)

- A. switchport trunk allowed vlan remove 3,4,50
- B. switchport trunk allowed vlan except 3,4,50
- C. switchport trunk allowed vlan except 1,2,10,20
- D. switchport trunk allowed vlan add 1,2,10,20
- E. switchport trunk disallowed vlan remove 3,4,50
- F. switchport trunk disallowed vlan add 3,4,50

Answer: A, D

Explanation:

switchport trunk allowed vlan *vlan_list*

The *vlan_list* format is **all** | **none** | [**add** | **remove** | **except**] *vlan_atom*[,*vlan_atom*...], where:

- **all** specifies all VLANs from 1 to 4094. This keyword is not supported on commands that do not permit all VLANs in the list to be set at the same time.
- **none** indicates an empty list. This keyword is not supported on commands that require certain VLANs to be set or at least one VLAN to be set.
- **add** adds the defined list of VLANs to those currently set, instead of replacing the list.
- **remove** removes the defined list of VLANs from those currently set, instead of replacing the list.
- **except** lists the VLANs that should be calculated by inverting the defined list of VLANs.
- *vlan_atom* is either a single VLAN number from 1 to 4094 or a continuous range of VLANs described by two VLAN numbers, the lesser one first, separated by a hyphen.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps663/products_command_reference_chapter09186a0080144e79.html#1022676

QUESTION NO: 15

Is the following statement True or False?

For each VLAN, if all switches are configured with the default priority, the switch with the highest MAC address in the VLAN will become the root switch.

- A. There is not enough information to determine
- B. False
- C. True

Answer: B

Explanation:

For each VLAN, the switch with the highest bridge priority (the lowest numerical priority value) is elected as the root switch. If all switches are configured with the default priority (32768), the switch with the lowest MAC address in the VLAN becomes the root switch.

QUESTION NO: 16

Is the following statement True or False?

The selection of the root switch is not an important issue.

- A. False
- B. True
- C. There is not enough information to determine

Answer: A**Explanation:**

The selection of the root switch for a particular VLAN is very important. You can choose it, or you can let the switches decide on their own using default values. The second option is risky because there may be sub-optimal paths in your network if the root selection process is not controlled by you.

Before configuring STP, you need to select a switch to be the root of the spanning-tree. It does not necessarily have to be the most powerful switch; it should be the most centralized switch on the network. All dataflow across the network will be from the perspective of this switch. It is also important that this switch be the least disturbed switch in the network. The backbone switches are often selected for this function, because they typically do not have end stations connected to them. They are also less likely to be disturbed during moves and changes within the network.

Reference: <http://www.cisco.com/warp/public/473/5.html>

QUESTION NO: 17

When you're building up a VLAN, which of the following are part of the sequence? (Select all that apply)

- A. Assign ports.
- B. Create VLAN.
- C. Create VLAN naming scheme
- D. Configure ports for trunking.
- E. Remove the trunk when the trunk is no longer needed.
- F. Set the baud rate for ports
- G. Verify configuration.

Answer: A, B, D, E, G

Explanation:

D, E: To create VLANs on a Catalyst switch, you must first enable the VLAN Trunking Protocol (VTP). The switch must be in VTP server or transparent mode to do this. VTP clients can not create VLANs.

B: The next step is to create the VLAN.

A: Once the VLAN is created, the final step is to assign individual ports to the VLAN.

G: After everything is configured, it should be verified.

Incorrect Answers:

C: A VLAN naming scheme isn't necessary because VLANs are numbered by default when they're created. Note that a VTP domain name must be created, but the VLANs themselves are not required to be named.

F: The baud rate doesn't have to be set for ports. Setting baud rate is for out-of-band console connections.

QUESTION NO: 18

What is true regarding the deployment of a VLAN? (Select all that apply)

- A. All VLAN hosts are members of the same subnet domain
- B. All VLAN hosts are members of the same IP domain
- C. All VLAN hosts are members of the same broadcast domain
- D. We use VLANs to establish separate broadcast domains to enjoy efficient bandwidth utilization

Answer: A, B, C, D

Explanation:

VLAN's are virtual LAN's and share the same characteristics. Devices on them belong to the same broadcast domains, and they are used for the sake of providing more efficient bandwidth utilization.

VLANs are often associated with IP subnetworks. For example, all the end stations in a particular IP subnet belong to the same VLAN. Traffic between VLANs must be routed. All members of a VLAN belong to the same IP subnet.

QUESTION NO: 19

Switch TK1 is running the Catalyst OS software. What do have to do to create an Ethernet VLAN on TK1? (Select all that apply)

- A. set vlan vlan_num [name name] [said said] [mtu mtu] [translation vlan_num]

- B. set vlanID vlan_num [name name] [said said] [mtu mtu] [translation vlan_num]
- C. Go into interface mode
- D. set vlanint vlan_num [name name] [said said] [mtu mtu] [translation vlan_num]
- E. Go into privileged mode

Answer: A, E

Explanation:

The “set vlan vlan-name” command is used to configure VLANs on CAT OS switches. In order to make any configuration changes to these switches, you must first be in privileged enable mode.

Incorrect Answers:

B, D: These are invalid commands that are using the wrong syntax.

C: In switches running CAT OS, there is no interface configuration mode. All configuration commands are done from the global configuration mode.

QUESTION NO: 20

What does Cisco recommend you use VLAN 1 for on a switch? (Select two)

- A. security
- B. load balancing
- C. troubleshooting
- D. management

Answer: C, D

Explanation:

The default VLAN in a switch for all ports is VLAN 1. It is recommended to create other VLANs within the switch and assign user ports to these new VLANs. However, VLAN 1 should be kept for troubleshooting and management purposes.

Incorrect Answers:

A: Since VLAN 1 is the default VLAN assigned to all ports, it is recommended that users be placed in different VLANs for security purposes. If a user tries to gain unauthorized access into the network, VLAN 1 will be the first VLAN that this user will try to use.

B: The use of VLANs alone will not provide for any load balancing functionality.

Section 3: Configure and verify 802.1Q trunks (5 questions)**QUESTION NO: 1**

Assuming you have an IOS based switch; which command would you execute if you wanted to specify IEEE 802.1Q encapsulation on a trunked port?

- A. Switch(config-if)#switchport trunk encapsulation dot1q
- B. Switch(config-if)#switchport encapsulation dot1q
- C. Switch(config-if)#switchport trunk encapsulation isl
- D. Switch(config)#switchport 0/1 trunk encapsulation isl

Answer: A

Explanation:

Ethernet Trunk Encapsulation Types:

- **switchport trunk encapsulation isl** - Specifies ISL encapsulation on the trunk link.
- **switchport trunk encapsulation dot1q** - Specifies 802.1Q encapsulation on the trunk link.
- **switchport trunk encapsulation negotiate** - Specifies that the interface negotiate with the neighboring interface to become an ISL (preferred) or 802.1Q trunk, depending on the configuration and capabilities of the neighboring interface.

The trunking mode, the trunk encapsulation type, and the hardware capabilities of the two connected interfaces determine whether a link becomes an ISL or 802.1Q trunk.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps700/products_configuration_guide_chapter_09186a008007f659.html

QUESTION NO: 2

Two TestKing switches are connected as shown below:



```
SwitchTK1(config)#inter fa 0/1
SwitchTK1#switchport trunk encapsulation dot1q
SwitchTK1(config-if)#switchport mode trunk
```

```
SwitchTK2(config)#inter fa 0/1
SwitchTK2(config)#switchport trunk encapsulation dot1q
SwitchTK2(config-if)#switchport mode trunk
```

Which statements are true regarding the configuration of the above pair of switches?
(Select two)

- A. The trunk is currently using the ISL trunking protocol.
- B. The trunk is currently using the 802.1q trunking protocol.
- C. By default, all VLANs will be transmitted across this trunk.
- D. By default, SwitchTK1 and SwitchTK2's Fast Ethernet 0/1 port will not generate DTP messages.
- E. By default, the trunk can only support one VLAN, and only that single VLAN is transmitted across the trunk.

Answer: B, C

Explanation:

The second line in each configuration (`#switchport trunk encapsulation dot1q`) proves that B is correct, as dot1q is Cisco IOS for 802.1q trunking.

Since the interface fa/01 is configured (`#interface fa 0/1`) and the mode is set to trunk (`#switchport mode trunk`) in both switches, there is no need for dynamic trunking protocol since the trunk is already set. By default, all VLANs will be able to cross the trunk, unless explicitly configured not to do so.

QUESTION NO: 3

Switches TestKing1 and TestKing2 are connected as shown in the diagram below:



Use the following output taken from each port

TestKing1:

show config:

```
interface GigabitEthernet0/1
```



```
switchport trunk encapsulation dot1q  
switchport mode dynamic auto  
no ip address
```

show interface gig0/1 switchport:

```
Name: Gi0/1  
Switchport: Enabled  
Administrative Mode: dynamic auto  
Operational Mode: down  
Administrative Trunking Encapsulation: dot1q  
Negotiation of Trunking: On  
Access Mode VLAN: 1 (default)
```

TestKing2:

show interface gig0/1 switchport:

```
Name: Gi0/1  
Switchport Enabled  
Administrative Mode: dynamic auto  
Operational Mode: down  
Administrative Trunking Encapsulation: dot1q  
Negotiation of Trunking: On  
Access Mode VLAN: 1 (default)  
Trunking Native Mode VLAN: 1 (default)
```

In accordance with the above exhibit: what's preventing the two switches from trunking on the link between them?

- A. There is no IP address denied.
- B. **no shutdown** needs to be entered on the interfaces.
- C. Both sides are in auto negotiation mode.
- D. ISL should be used instead of 802.1q.
- E. Access mode VLAN must be different from native mode VLAN.

Answer: C

According to Cisco table Auto & Auto results in NO trunk formation. At least one end of the trunk should be set to on or desirable in order for the trunk to operate correctly.

QUESTION NO: 4

You have a Cisco Catalyst 3500XL switch within the TestKing LAN and you want to configure a trunk port on it. Which IOS command should you issue?

- A. Switch(config)#vtp mode
- B. Switch(config-if)#set trunk
- C. Switch(config-if)#encapsulation
- D. Switch(config-if)#switchport trunk encapsulation

Answer: D

Explanation:

The **switchport trunk encapsulation** command is used to specify the trunk encapsulation mode for a port.

Note: Ethernet Trunk Encapsulation Types

Encapsulation	Function
switchport trunk encapsulation isl	Specifies ISL encapsulation on the trunk link.
switchport trunk encapsulation dot1q	Specifies 802.1Q encapsulation on the trunk link.
switchport trunk encapsulation negotiate	Specifies that the interface negotiate with the neighboring interface to become an ISL (preferred) or 802.1Q trunk, depending on the configuration and capabilities of the neighboring interface.

Note: Syntax: **set trunk mod/port {on | off | desirable | auto | nonegotiate} [vlans] [isl | dot1q | negotiate]**

Incorrect Answers:

- A: VTP can run in three modes: server, client, and transparent. The vtp mode command is used to set the VTP mode.
- B: The set trunk command to configure trunk ports and to add VLANs to the allowed VLAN list for existing trunks.
- C: The **encapsulation** command would only list is used to specify the encapsulation of the VLAN.

QUESTION NO: 5

The TestKing network is using 802.1Q for all of the trunks. By default, which statement is correct when an IEEE 802.1Q trunk port receives an untagged frame?

- A. The frame is considered in the native VLAN and forwarded to the ports associated with that VLAN.
- B. The frame is encapsulated and tagged as in the native VLAN.

- C. The frame is broadcast on all ports regardless of VLAN association
- D. The frame is dropped

Answer: A

Explanation:

The IEEE 802.1Q standard is extremely restrictive to untagged frames. The standard provides only a per-port VLANs solution for untagged frames. For example, assigning untagged frames to VLANs takes into consideration only the port from which they have been received. Each port has a parameter called a permanent virtual identification (Native VLAN) that specifies the VLAN assigned to receive untagged frames.

The main characteristics of IEEE 802.1Q are as follows:

- Assigns frames to VLANs by filtering.
- The standard assumes the presence of a single spanning tree and of an explicit tagging scheme with one-level tagging.

Each physical port has a parameter called PVID. Every 802.1Q port is assigned a PVID value that is of its native VLAN ID (default is VLAN 1). All untagged frames are assigned to the LAN specified in the PVID parameter. When a tagged frame is received by a port, the tag is respected. If the frame is untagged, the value contained in the PVID (native VLAN) is considered as a tag. Untagged frames are then forwarded to the ports associated with this native VLAN.

Section 4: Configure and verify ISL trunks (6 questions)**QUESTION NO: 1**

You have just configured an ISL trunk line over Ethernet media between two Cisco Switches, each switch having identical modules, software revisions, and VLAN configurations. Which of the following variables are NOT necessary for the ISL trunk to operate properly? (Select all that apply)

- A. Identical trunk negotiation parameters at each end of the link
- B. Identical duplex at each end of the link
- C. Identical speed at each end of the link
- D. Identical native VLAN parameters at each end of the link

Answer: A, D

Explanation:

In order for a trunk connection to function properly, it is not necessary for the trunking negotiation parameters to be identical. For example, one end could be configured as “on” and the other could be configured for “auto-negotiate” and the trunk would be operational. Similarly, it is not necessary for the native VLAN parameters to be the same at each end.

Incorrect Answers:

B, C: One of the requirements for trunking to work is to have both sides of the trunk agree on the speed and duplex settings. Both sides must be configured with identical speed and duplex settings.

QUESTION NO: 2

You must configure an ISL trunk between a Catalyst 5000 Switch and a Catalyst 6000 switch.

What parameters have to be identical for an ISL trunk to work properly? (Select two)

- A. an identical VTP mode
- B. an identical speed/duplex
- C. an identical trunk negotiation parameter
- D. an identical trunk encapsulation parameter

Answer: B, D

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

Speed, duplex, and trunk encapsulation have to be identical at each end. If one end of the trunk is configured for ISL encapsulation and the other is set for 802.1Q encapsulation the trunk will not come up.

Incorrect Answers:

A: The trunk modes have to be compatible. They don't have to be identical.

C: The Trunk negotiation parameter does not have to be identical. For example, one end could be configured as "on" and the other could be configured for "auto-negotiate" and the trunk would be operational.

QUESTION NO: 3

An ISL trunk connects switches TK1 and TK2. What is the numerical range of user-configurable ISL VLANs on these switches?

- A. 1-1001
- B. 0-4095
- C. there is no range
- D. 0 - 1000
- E. None of the above

Answer: A

Explanation:

The valid range of user-configurable ISL VLANs is 1-1001. The valid range of VLANs specified in the IEEE 802.1Q standard is 0-4095. In a network environment with non-Cisco devices connected to Cisco switches through 802.1Q trunks, you must map 802.1Q VLAN numbers greater than 1000 to ISL VLAN numbers. 802.1Q VLANs in the range 1-1000 are automatically mapped to the corresponding ISL VLAN. 802.1Q VLAN numbers greater than 1000 must be mapped to an ISL VLAN in order to be recognized and forwarded by Cisco switches.

QUESTION NO: 4

An ISL trunk connects switches TK1 and TK2. What is true about the Inter-Switch Link (ISL) protocol? (Select two)

- A. ISL can be used between Cisco and non-Cisco switch devices.
- B. ISL calculates a new CRC field on top of the existing CRC field.
- C. ISL adds 4 bytes of protocol-specific information to the original Ethernet frame.
- D. ISL adds 30 bytes of protocol-specific information to the original Ethernet frame.

Answer: B, D

Explanation:

- B: A second frame check sequence (FCS) field lies at the end of the frame.
 D: ISL is an external tagging process: new 26-byte ISL header is added to the original Ethernet frame. A second 4-byte frame check sequence (FCS) field is added at the end of the frame so 30 bytes of total overhead is added.

Incorrect Answers:

- A: Cisco's propriety version of frame tagging is ISL. ISL can only be used between Cisco routers.
 C: 30 bytes are added to the Ethernet frame, not 4 bytes. 4 bytes are added using 802.1Q encapsulation.

QUESTION NO: 5

Which of the commands below enables a trunking protocol that appends a four byte CRC to the packet when applied to the TestKing switch?

- A. Switch(config-if)#switchport trunk encapsulation dot1q
 B. Switch(config-if)#switchport trunk encapsulation ietf
 C. Switch(config-if)#switchport trunk encapsulation fddi
 D. Switch(config-if)#switchport trunk encapsulation isl

Answer: D

Explanation:

ISL is made up of three major components: a header, the original Ethernet frame, and a frame check sequence (FCS) at the end. With ISL, an Ethernet frame is encapsulated with a header that transports VLAN IDs between switches and routers. The 26-byte header containing a 10-bit VLAN ID is added to each frame. In addition, a 4-byte tail is added to the frame to perform a cyclic redundancy check (CRC). This CRC is in addition to any frame checking that the Ethernet frame performs.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 99

QUESTION NO: 19

Which statement is true regarding the configuration of ISL trunks?

- A. All catalyst switches support ISL trunking.
 B. A Catalyst switch will report giants if one side is configured for ISL while the other side is not.
 C. ISL trunking requires that native VLANs match.
 D. A Catalyst switch cannot have ISL and IEEE 802.1q trunks enabled.

Answer: B

Explanation:

The 802.1q tag is 4 bytes; hence the resulting ethernet frame can be as large as 1522 bytes. The minimum size of the Ethernet frame with 802.1q tagging is 68 bytes.

ISL frames are the standard MTU used in Ethernet frames, which is 1518 bytes. If one end of the trunk is configured for ISL frames of up to 1518 bytes will be expected on it, while the other end will be sending frames up to 1522 bytes in length. On the ISL configured end, these incoming frames will be considered as giants. This is just one of the reasons why ISL and 802.1Q are incompatible.

Incorrect Answers:

A: Not every Cisco switch model supports ISL.

C: In ISL, it is not necessary for the native VLANs to match.

D: Although it is true that each end of a trunk should be configured using the same encapsulation types, it is possible for a switch to have an ISL trunk configured on one port and an 802.1Q trunk on another port.

Section 5: Configure VTP domains in server, client and transparent modes (23 questions)

QUESTION NO: 1

What command would you enter onto a Cisco router if you wanted to add an IP MLS interface to a VTP domain named sales? (Type in the command below)

Answer: mls rp vtp-domain sales

Explanation:

According to the online documentation provided by Cisco:

To add an IP MLS interface to a VTP domain, perform this task in interface configuration mode:

To add an IP MLS interface to a VTP domain: mls rp vtp-domain [domain_name]

Example: This example shows how to add an IP MLS interface to a VTP domain name engineering:

```
Router(config-if)#mls rp vtp-domain engineering
```

QUESTION NO: 2

Which interface configuration mode command would you enter if you wanted to assign a route processor interface to a VTP domain on switch TK1?

- A. mls rp vlan-id *domain-name*
- B. set mls domain *domain-name*
- C. mls rp vtp-domain *domain-name*
- D. set mls vtp-domain *domain-name*
- E. None of the above

Answer: C

Explanation:

The **mls rp vtp-domain** command, applied in interface configuration mode, is used to assign a Multilayer Switching (MLS) interface to a specific Virtual Trunk Protocol (VTP) domain on the Multilayer Switching-Route Processor.

Incorrect Answers:

A: The **mls rp vlan-id** command is used to assign a virtual LAN (VLAN) identification number to an MLS interface.

B, D: There are no such commands.

QUESTION NO: 3

You have a non IOS switch named TK1. On this switch you enter the following command:

set vtp domain

What is the purpose of this command?

- A. For determining management domain name
- B. For enabling VTP pruning.
- C. For selecting VTP version.
- D. For verifying configuration set.
- E. For verifying configuration.

Answer: A

Explanation: Use the **set vtp** command to set the options for VTP.

set vtp [domain *domain_name*] [mode {client | server | transparent}] [passwd *passwd*]
[pruning {enable | disable}] [v2 {enable | disable}]

domain <i>domain_name</i>	(Optional) Keywords to define the name that identifies the VLAN management domain. The <i>domain_name</i> can be from 1 to 32 characters in length.
-------------------------------------	---

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cmd_refr/set_v.htm

QUESTION NO: 4

Switch TK1 is a Cisco Catalyst 3500XL switch. Which VTP command would you use if you wanted to set the management domain name on a Catalyst 3500XL switch?

- A. Switch(vlan)#vtp domain *domain-name*
- B. Switch(config)#vtp domain *domain-name*
- C. Switch(vlan)#set vtp domain *domain-name*
- D. Switch(enable) set vtp domain *domain-name*

Answer: A

Explanation:

The **vtp domain name** command is used to assign a name to the VTP management domain on a Catalyst 3500XL switch. Furthermore the prompt would look like: **Switch(vlan)#** on a switch of this type.

Catalyst 2900 and 3500 switches utilize a VLAN database configuration mode for making changes to the VLAN database parameters.

VLAN Database Mode

The VLAN database commands allow you to modify VLAN parameters. Enter the **vlan database** command to access VLAN database mode:

```
Switch> vlan database
```

```
Switch(vlan)#
```

From this mode, enter the “vtp domain” command to set the VTP domain name:

vtp domain

Use the **vtp domain** VLAN database command to configure the VLAN Trunking Protocol (VTP) administrative domain.

Incorrect Answers:

B: This command would be correct on a 1900 series switch, not on a Catalyst 3500XL switch.

C: The **set vtp domain name** command must be issued in enable mode.

D: This command would be correct on a Catalyst 5000 switch, not on a Catalyst 3500XL switch.

QUESTION NO: 5

You need to configure switch TK1 for pruning. Which VTP command would you use if you wanted to allow pruning?

- A. show vtp
- B. set vtp
- C. set vtp domain
- D. set vtp pruneeligible
- E. None of the above.

Answer: D**Explanation:**

Use the **set vtp** command to set the options for VTP.

```
set vtp [domain domain_name] [mode {client | server | transparent}] [passwd passwd]
[pruning {enable | disable}] [v2 {enable | disable}]
```

The **pruning** keyword is used to enable or disable VTP pruning for the VTP domain. VTP pruning causes information about each pruning-eligible VLAN to be removed from VTP updates if there are no stations belonging to that VLAN out a particular switch port. Use the **set vtp pruneeligible** and **clear vtp pruneeligible** commands to specify which VLANs should or should not be pruned when pruning is enabled for the domain.

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cmd_refr/set_v.htm

QUESTION NO: 6

If you have just configured a Catalyst switch to operate in VTP mode, and that switch is configured to not advertise VLAN configuration information. Which VTP mode has been configured on this switch?

- A. Client
- B. Server
- C. Host
- D. Transparent
- E. Native

Answer: D

Explanation:

You can configure a switch to operate in any one of these VTP modes:

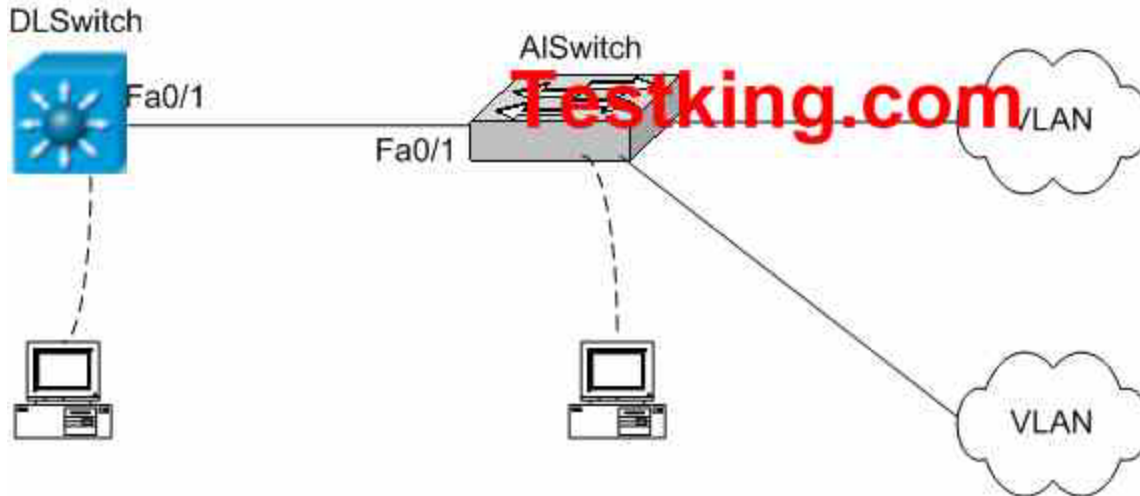
Server-In VTP server mode, you can create, modify, and delete VLANs and specify other configuration parameters (such as VTP version and VTP pruning) for the entire VTP domain. VTP servers advertise their VLAN configuration to other switches in the same VTP domain and synchronize their VLAN configuration 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. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive out their trunk ports.

QUESTION NO: 7

The TestKing network is displayed in the diagram below:



You have just been hired by TestKing.com to help their main office expand. The main offices have enhanced their wiring closets with some Layer 3 switches. The new distribution layer switch has been installed and a new access layer switch cabled next to it. Your task is to configure the distribution layer and access layer switch with VTP to share VLAN information, then to configure inter-VLAN routing on the distribution layer switch to route traffic between the different VLANs that are configured on the access layer switches.

VTP Domain	Distribution	
VLAN Ids	20	31
IP Addresses	172.16.71.1/24	172.16.132.1/24

These are your specific tasks:

- Configure the VTP information with the distribution layer switch as the VTP server
- Configure the VTP information with the access layer switch as a VTP client
- Configure VLANs on the distribution layer switch
- Configure inter-VLAN routing on the distribution layer switch
- Specific VLAN port assignments will be made as users are added to the access layer switches in the future.
- All VLANs and VTP configurations are to be completed in the global configuration

To configure the switch click on the host icon that is connected to the switch by way of a serial console cable.

Answer:

LAB configuration:

```
switch#conf t
```

```
switch(config)#vtp mode server
```

```
switch(config)#vtp domain CISCO
```

```

switch(config)#vlan 20
switch(config)#vlan 31
switch(config)#int vlan 20
switch(if-config)#ip add 172.64.20.1 255.255.255.0
switch(if-config)#no shut
switch(if-config)#int vlan 31
switch(if-config)#ip add 192.162.31.1 255.255.255.0
switch(if-config)#no shut
switch(if-config)#exit
switch#ip routing
switch#sh run
switch#copy run start
switch#conf t
switch(config)#vtp mode client
vtp domain CISCO
switch(config)#exit
switch#show run
switch#copy run start

```

Alternative #1

VTP Domain	Distribution	
VLAN Ids	20	31
IP Addresses	172.16.16.1/24	172.16.193.1/24

Alternative #12

VTP Domain	Distribution	
VLAN Ids	30	21
IP Addresses	172.16.203.1/24	172.16.93.1/24

QUESTION NO: 8

The following commands were entered on a TestKing switch:

```

Switch(config)# vtp mode transparent
Switch(config)# vtp version 2

```

What is the result of these commands?

- A. VLAN configuration information is saved in RAM only.
- B. VLANs cannot be created, modified or deleted via command line interface.
- C. VLAN configuration information received via VTP advertisements are forwarded to other switches within the management domain.

- D. VLAN configuration information is synchronized with information within VTP advertisements received from other switches in the management domain.

Answer: C

Explanation:

VTPv2 will allow the switch to be in transparent mode which will forward VTP info. The command series above put the switch in VTP transparent mode. This TestKing switch does not actively participate in VTP, it doesn't advertise its VLAN configuration to other switches, and when other switches advertise their VLAN configuration it doesn't consider that information. It will, however, pass incoming VLAN information that was received to other switches within the VTP domain.

QUESTION NO: 9

You are a CCNP in the midst of configuring a switching solution on a switch that participates in multilayer switching. What show command would you use to view the MLS interfaces for a specific VTP domain? (Type in the answer below):

Answer: show mls rp vtp-domain


Explanation:

To show MLS interfaces for a specific VTP domain, use the show mls rp vtp-domain EXEC command.

Command: show mls rp vtp-domain [domain-name]

QUESTION NO: 10

The TestKing network is displayed in the diagram below:



The network in the above exhibit is configured with VLANs 1,2,3,4, & 5 and 802.1 Q. trunking is enabled between all switches. However, access ports for TestKing3 and TestKing6 are the only access ports for VLAN 3. What could an administrator do to make sure that other switches don't receive unnecessary broadcast packets destined for VLAN 3, while still allowing all the other VLAN packets to cross?

- A. Configure VTP pruning.
 - B. Configure TestKing3 and TestKing6 as transparent switches.
 - C. Configure TestKing1, TestKing2, TestKing4 and TestKing5 as transparent switches.
 - D. Nothing is required.
- Only TestKing3 and TestKing6 will receive VLAN3 packets by default.

Answer: A

Explanation:

VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices. By configuring VTP pruning, traffic will not flow to switches destined for VLANs that they are not attached to.

QUESTION NO: 11

Leading the way in IT testing and certification tools, www.testking.com

You are configuring VTP on a non IOS switch named TK1, and you enter the following command:

set vtp pruneeligible

What is this command useful for?

- A. For determining management domain name
- B. For verifying configuration.
- C. For enabling VTP pruning.
- D. For selecting VTP version.
- E. For verifying configuration set

Answer: C

Explanation:

Use the **set vtp** command to set the options for VTP.

```
set vtp [domain domain_name] [mode {client | server | transparent}] [passwd passwd]
[pruning {enable | disable}] [v2 {enable | disable}]
```

The **pruning** keyword is used to enable or disable VTP pruning for the VTP domain. VTP pruning causes information about each pruning-eligible VLAN to be removed from VTP updates if there are no stations belonging to that VLAN out a particular switch port. Use the [set vtp pruneeligible](#) and [clear vtp pruneeligible](#) commands to specify which VLANs should or should not be pruned when pruning is enabled for the domain.

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cmd_refr/set_v.htm

QUESTION NO: 12

While verifying your configuration on the non-IOS based switch named TK1, you issue the following command:

show trunk

What is this command useful for?

- A. For verifying configuration.
- B. For enabling VTP pruning.
- C. For verifying configuration set
- D. For selecting VTP version.
- E. For determining management domain name

Answer: A

Explanation:

Cisco documentation on the use of this command is as follows:

Use the show trunk command to display trunking information for the switch.

show trunk [mod_num[/port_num]] [detail]mod_num (Optional) Number of the module.

/port_num (Optional) Number of the port.

detail (Optional) Keyword to show detailed information about the specified trunk port.

QUESTION NO: 13

Switch TK1 is configured as a VTP server. What is true when you enable VTP pruning on a VTP server?

- A. It is not possible without a root re-election
- B. It enables pruning for the entire management domain.
- C. It cannot be done on a VTP server
- D. It enables pruning for the individual switch.

Answer: B

Explanation:

Enabling VTP pruning on a VTP server enables pruning for the entire management domain. VTP pruning takes effect several seconds after you enable it. By default, VLANs 2 through 1000 are pruning-eligible. VTP pruning does not prune traffic from VLANs that are pruning-ineligible. VLAN 1 is always pruning-ineligible; traffic from VLAN 1 cannot be pruned.

QUESTION NO: 14

One of the configurable VTP commands is displayed below:

```
Clear vtp pruneeligible vlan_range
```

What is the purpose of this above command?

- A. Verify the VTP pruning configuration.
- B. Make specific VLANs pruning-eligible on the device.
- C. Make specific VLANs pruning-ineligible on the device.
- D. Enable VTP pruning in the management domain.
- E. Verify that the appropriate VLANs are being pruned on trunk ports.

Answer: C

Explanation:

This command makes specific VLANs pruning-ineligible on the device. (By default, VLANs 2-1000 are pruning-eligible.)

Note: VLAN 1 is not pruning eligible.

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_5_2/config/vtp.htm

QUESTION NO: 15

When setting up multiple VTP domains, what should be considered in order to maintain VLAN database consistency? (Select two)

- A. Do not configure any switches as a VTP server
- B. Ensure that all switches not authorized to make changes are in client mode
- C. Always configure switches using VTP server mode when adding them to the existing network
- D. Allow only one VTP server in each domain so that adding and deleting VLANs can be centralized to one location.

Answer: B, D

Explanation:

B: Switches not authorized to make changes should be run as VTP clients. VTP clients receive information from VTP servers and send and receive updates, but they cannot make any changes.

D: You need at least one server in your VTP domain to propagate VLAN information throughout the domain. You are able to use several VTP servers in a domain. However, only allowing one VTP server would help keep the VLAN database consistent.

Incorrect Answers:

A: Switches can very well be used as VTP servers. VTP server mode is the default for all Catalyst switches.

C: It is more prudent to configure switches using VTP client mode. They will not be able to update information in the VLAN domain database.

QUESTION NO: 16

VTP is running on the TestKing network. In which VTP modes can a full list of all VLANs be maintained? (Select two)

- A. VTP Bypass
- B. VTP Client
- C. VTP Transparent
- D. VTP Restore
- E. VTP Server

Answer: B, E

Explanation:

VTP-capable devices can be configured to operate in the following three modes:

The VTP Server maintains a full list of all VLANs within the VTP domain. Information is stored in nonvolatile random-access memory (NVRAM). The server can add, delete, and rename VLANs.

The VTP Client also maintains a full list of all VLANs. However, it will not store in NVRAM. The client can not add, delete, or rename VLANs. Any changes made must be received from a VTP server advertisement.

The VTP Transparent mode does not participate in VTP. However, it will pass on a VTP advertisement. VLAN, as defined, is only local to the switch and is stored in NVRAM.

QUESTION NO: 17

You wish to configure VTP on switch TK1. What do you have to do before you can create a VLAN on a VTP server?

- A. The VTP server ID must be cleared
- B. The VTP membership list must be refreshed
- C. The priority must be cleared
- D. The management domain name must be specified

Answer: D

Explanation:

By default, the switch is in VTP server mode and is in the no-management domain state until the switch receives an advertisement for a domain over a trunk link or you configure a management domain. You cannot create or modify VLANs on a VTP server until the management domain name is specified or learned.

QUESTION NO: 18

What must you do if you wish to configure VTP in secure mode within the TestKing LAN?

- A. Assign a management domain password to the VTP Server in the domain.
- B. Assign a management domain password to each switch in the domain.
- C. Assign a management domain password to the root switch in the domain.
- D. None of the above.

Answer: B

Explanation:

If you configure VTP in secure mode, the management domain will not function properly if you do not assign a management domain password to each switch in the domain. All switches must be configured with the password in order for VTP to function properly in a network.

QUESTION NO: 19

If you configure a switch as a VTP server offline, then connect it to a network, what could happen to the network?

- A. Cause a loss of VLAN information
- B. Destabilize the spanning tree
- C. Revert to simplex mode
- D. Revert to duplex mode
- E. Ignore the configuration revision numbers created on the other VTP servers
- F. Revert to client mode

Answer: A

Explanation:

When connecting a new switch to your network you can accidentally change your current VLAN database if the new switch has a higher VLAN Trunking Protocol (VTP) revision number. If the newly inserted switch has no VLANs configured and the revision number is higher and is configured as a VTP server, it will override the configuration of the other switches within the network, deleting all of the configured VLANs. To avoid this, you must clear the VTP revision number on the new switch. The easiest way is to change the VTP domain name to "something_else" and back to "your_VTP_domain" on the new switch. This sets the VTP revision number to 0 and you can connect the switch to the network without any problem.

QUESTION NO: 20

Which of the following are true if you configure a password for VTP? (Select all that apply)

- A. It is carried in all summary-advertisement VTP packets
- B. It needs to be the same on all switches in the VTP domain
- C. It needs to be configured on all switches in the VTP domain
- D. It is translated using an algorithm in a 24 bytes word
- E. None of the above

Answer: A, B, C

Explanation:

According to the online documentation provided by Cisco:

If you configure a password for VTP, it needs to be configured on all switches in the VTP domain and it needs to be the same password. The VTP password you configure is translated using an algorithm in a 16 bytes word (MD5 value) carried in all summary-advertisement VTP packets.

Incorrect Answers:

D: The algorithm uses a 16 byte word, not 24 bytes.

QUESTION NO: 21

Is the following statement True or False?

With VTP, if an administrator makes configuration changes centrally on one or more switches, those changes will be automatically communicated to all the other switches on the network.

- A. There is not enough information to determine
- B. True
- C. False

Answer: B**Explanation:**

This statement is true. Before you create virtual LANs (VLANs), you must decide whether to use VTP in your network. With VTP, you can make configuration changes centrally on one or more switches and those changes are automatically communicated to all the other switches in the network. Changes made to VTP servers are propagated to all other switches within the VTP domain.

QUESTION NO: 22

You have to enter a new switch into the existing TestKing VTP domain without altering the configurations of the systems currently on this domain. Which of the following answer choices describes one of the conditions required to ensure that the new switch will not change the existing VTP domain configuration?

- A. The switch must be in client mode.
- B. The switch must be in a mode other than the client mode.
- C. The VTP domain must not have a password assigned to it.
- D. The trunk links must not be configured for ISL
- E. None of the above

Answer: A

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. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive out their trunk ports.

Only by using the client or transparent modes can you ensure that the other switches within the domain are left unaffected.

QUESTION NO: 23

A brand new stand alone Catalyst 3550 switch is being installed. Multiple VLANs will be configured on the switch. What needs to be configured before adding any VLAN to the VLAN database if it is in VTP server mode?

- A. VTP pruning
- B. VTP domain name
- C. VTP version number
- D. ISL or IEEE 802.1Q trunking

Answer: B

Explanation:

In order to configure any VLANs and assign ports to them, a VTP server must first have a VTP domain name configured.

Configure the Switch as a VTP Server:

When a switch is configured as a VTP server, you must define a VTP domain before you can create VLANs.

To configure a switch as a VTP server, perform these tasks in privileged mode:

	Task	Command
Step 1	Assign a name to the VTP management domain.	set vtp domain <i>name</i>
Step 2	Set the VTP mode.	set vtp mode server
Step 3	Verify the VTP configuration.	show vtp domain

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps4324/products_configuration_guide_chapter09186a008007d8ff.html#1015043

www.testking.com

Section 6: Enable Spanning Tree on ports and VLANs (7 questions)**QUESTION NO: 1**

Refer to the output shown on switch TK1 below:

*VLAN 1 bridge priority set to 8192.
VLAN 1 bridge max aging time set to 20.
VLAN 1 bridge hello time set to 2.
VLAN 1 bridge forward delay set to 15.
Switch is now the root switch for active VLAN 1.*

What command would you enter to reproduce this output? (Type in answer below)

Answer: set spantree root 1

Explanation:

According to Cisco:

The default priority for switches is 32768. This command setting means that the switch will be selected as the root switch because it has the lowest priority. This command will set the bridge priority to 8192, unless another switch on the network is already configured with a priority value less than 8192. If this is the case, the priority will be set to one less than this value, ensuring that it will become the root switch.

Note: In STP, a lower bridge priority is preferred over a higher value.

QUESTION NO: 2

Refer to the output shown on switch TK1 below:

*Warning: Span tree port fast start should only be enabled on ports connected to a single host.
Connecting hubs, concentrators, switches, bridges, etc. to a fast start port can cause temporary spanning tree loops. Use with caution.
Span tree ports 4/1-24 fast start enabled.*

What command could you enter to reproduce this output? (Type in answer below)

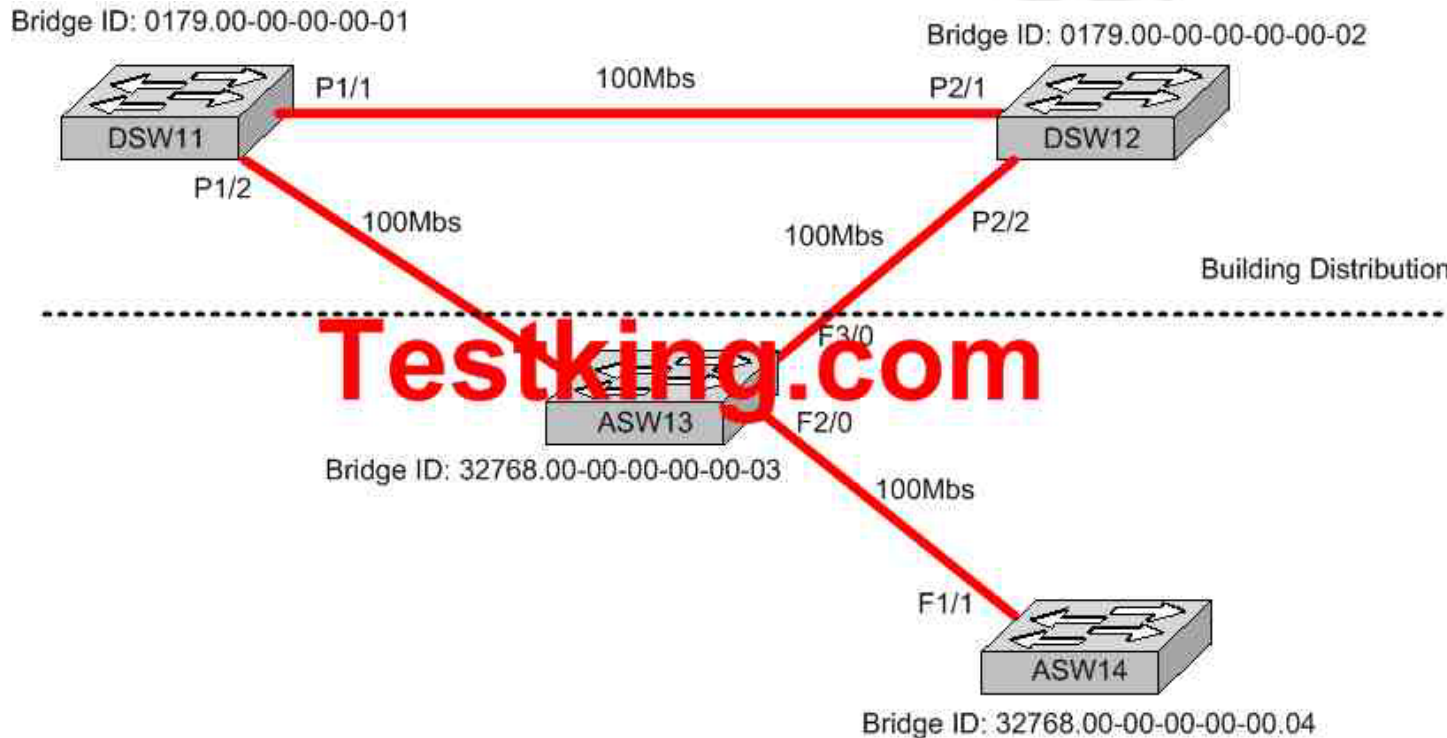
Answer: set spantree portfast 4/1-24 enable

Explanation:

The output shown in this question is the result of the “set spantree portfast” command. This setting should be configured only on ports that are connected to workstations or PCs. Do not enable portfast on any port connected to another switch.

QUESTION NO: 3

The TestKing switched LAN is displayed in the diagram below:



Based on the assumption that STP is enabled on all the switch devices, which of the following statements are true? (Choose two)

- A. DSW11 will be elected the root bridge.
- B. DSW12 will be elected the root bridge.
- C. ASW13 will be elected the root bridge.
- D. P1/1 will be elected the nondesignated port.
- E. P2/1 will be elected the nondesignated port.
- F. F3/0 will be elected the nondesignated port.

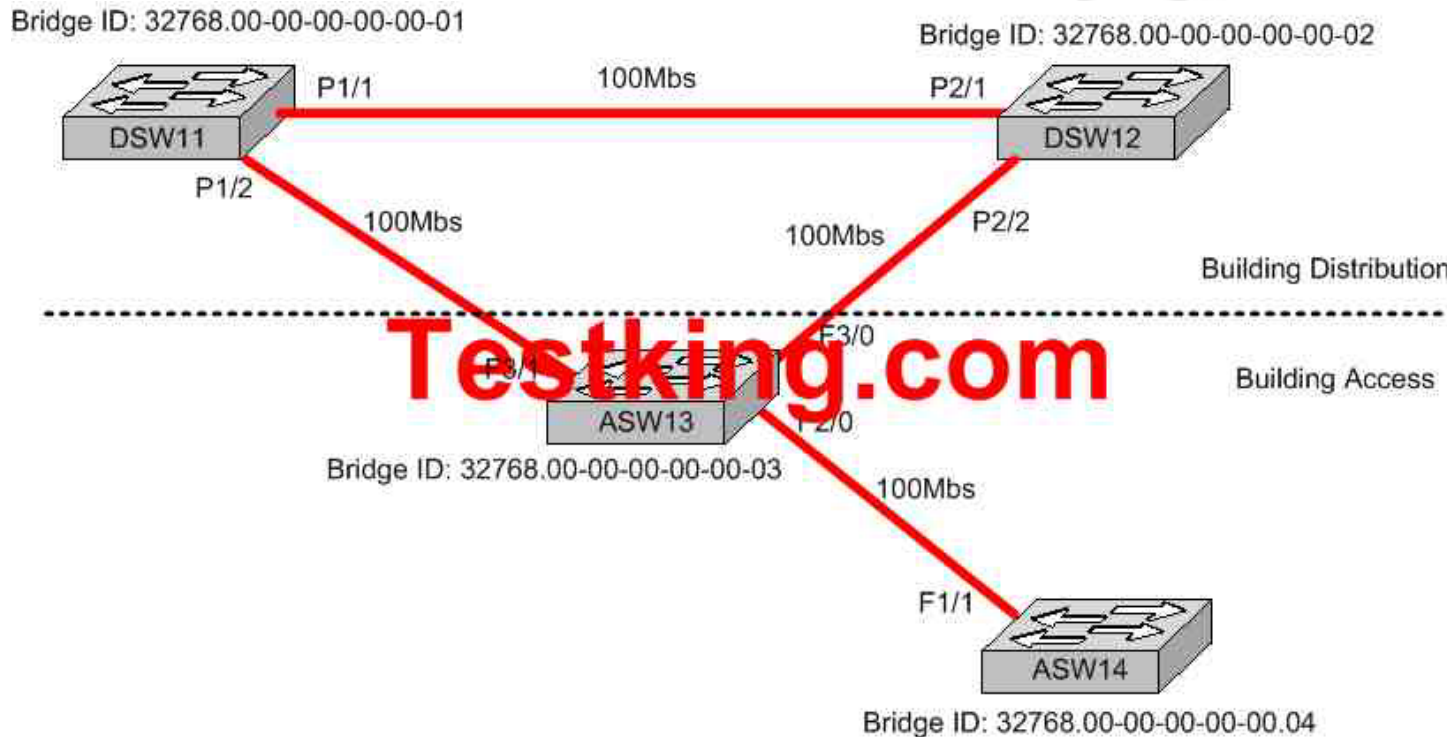
Answer: A, F

The root bridge should be placed as close to the core as possible and should be the most centrally located. By default, the switch with the lowest bridge ID will become the root bridge, assuming

all other parameters are left as default. This makes DSW11 the root bridge. Also, all ports directly connected to the root bridge will become designated ports, since they are closest to the root bridge. In this case, port F3/0 will become the non-designated port.

QUESTION NO: 4

The TestKing switched LAN is displayed below:



Your junior network administrator has just finished installing the above switched network using Cisco 3550s and would like to manipulate the root bridge election. Which switch should he configure as the root bridge and with which command?

- A. DSW11(config)# spanning-tree vlan 1 priority 4096
- B. DSW12(config)# set spanning-tree priority 4096
- C. ASW13(config)# spanning-tree vlan 1 priority 4096
- D. DSW11(config)# set spanning-tree priority 4096
- E. DSW12(config)# spanning-tree vlan 1 priority 4096
- F. ASW13(config)# set spanning-tree priority 4096

Answer: F

Explanation:

Before configuring STP, you need to select a switch to be the root of the spanning-tree. It does not necessarily have to be the most powerful switch; it should be the most centralized switch on the network. All dataflow across the network will be from the perspective of this switch. It is also important that this switch be the least disturbed switch in the network. The backbone switches are often selected for this function, because they typically do not have end stations connected to them. They are also less likely to be disturbed during moves and changes within the network. In this case, switch ASW13 is the most centrally located switch so it should have its bridge priority lowered to become the root.

Note: In the network shown above, if no configuration changes are made, switch DSW11 will become the root by default, since it has the lowest Bridge ID.

QUESTION NO: 5

Which three items are configured in MST configuration submode? (Select three)

- A. Region name
- B. Configuration revision number
- C. VLAN instance map
- D. IST STP BPDU hello timer
- E. CST instance map
- F. PVST+ instance map

Answer: A, B, C

Explanation:

spanning-tree mst configuration:

Use the spanning-tree mst configuration command to enter the MST configuration submode. Use the no form of this command to return to the default MST configuration.

Defaults:

The default value for the MST configuration is the default value for all its parameters:

- No VLANs are mapped to any MST instance (all VLANs are mapped to the CIST instance).
- The region name is an empty string.
- The revision number is 0.

Usage Guidelines:

The MST configuration consists of three main parameters:

- Instance VLAN mapping (see the [instance](#) command)
- Region name (see the [name](#) command)
- Configuration revision number (see the [revision](#) command)

QUESTION NO: 6

You are connected to switch TestKing1 as displayed in the diagram below:

Leading the way in IT testing and certification tools, www.testking.com



You work as a network engineer at TestKing.com. The TestKing.com Toronto office is installing a temporary Catalyst 3550 in an IDF to connect 24 additional users. To prevent network corruption, it is important to have the correct configuration prior to connecting to the production network. It will be necessary to ensure the switch does not participate in VTP but forwards VTP advertisements received on trunk ports.

All interfaces should transition immediately to the forwarding state of Spanning-Tree due to errors that have been experienced on office computers. Also, configure the user ports (All FastEthernet ports) so that the ports are permanently non-trunking.

You will configure FastEthernet ports 0/12 through 0/24 for users who belong to VLAN 20. Also, all VLAN and VTP configurations are to be completed in global configuration mode as VLAN database mode is being deprecated by Cisco.

You are required to accomplish the following tasks:

- Ensure the switch does not participate in VTP but forwards VTP advertisements received on trunk ports.
- Ensure all non-trunking interfaces (Fa0/1 to Fa0/24) transition immediately to the forwarding state of Spanning-Tree.
- Ensure all FastEthernet interfaces are in a permanent non-trunking mode.
- Place FastEthernet interfaces 0/12 through 0/24 in VLAN 20

Answer:

```
enable
configure t
vtp mode transparent
```

```
spanning-tree portfast default
```

```
interface range fa0/1 - 24
switchport mode access
copy run start
```

```
end
```

- or -

```
enable
configure t
vtp mode transparent
```

```
interface range fa0/1 - 24
switchport mode access
spanning-tree portfast
copy run start
```

```
end
```

QUESTION NO: 7

If the root bridge fails, configuration BPDUs will no longer be sent. Which STP timer will have to expire before the other switches can actively restore connectivity with topology change procedure of STP?

- A. hello timer
- B. BPDU timer
- C. Forward_delay timer
- D. Max_age timer
- E. Dead timer
- F. Wait timer

Answer: D

Explanation:

Once a stable network topology has been established, all bridges listen for Hello BPDUs (Bridge Protocol Data Units) transmitted from the root bridge. If a bridge does not get a Hello BPDU after a predefined interval (Max Age), the bridge assumes that the link to the root bridge is down. This bridge then initiates negotiations with other bridges to reconfigure the network to re-establish a valid network topology.

Max age takes into account that the switch at the periphery of the network should not time out the root information under stable condition (that is, if the root is still alive). This is the value that max age needs to take into account the total BPDU propagation delay and the message age overestimate. As such, the formula for max age is as follows:

```
Max_age
= End-to-end_BPDU_propa_delay + Message_age_overestimate
= 14 + 6
= 20 sec
```

This explains how IEEE reaches the default recommended value for max age.

Reference: <http://www.zyxel.com/support/supportnote/ves1012/app/stp.htm>

www.testking.com

Section 7: Configure Spanning Tree parameters including: port priority, VLAN priority, root bridge, BPDU guard, PortFast and UplinkFast (17 questions)

QUESTION NO: 1

You have been promoted to CTO at TestKing, Inc. and you are looking for ways to increase the efficiency of your network administration. What can your administrators do to improve the Spanning Tree Protocol's operation?

- A. Properly place the Root Bridge to ensure an optimal STP topology.
- B. Configure access switches as Root Bridges to ensure optimal workstation access to the network.
- C. Load balance on redundant links through the use of technologies such as BackboneFast.
- D. Provide for efficient workstation access through the use of BackboneFast.

Answer: A

Explanation:

One of the most important decisions that must be made in the Spanning tree network is the location(s) of the root bridge. Proper placement of the root bridge optimizes the path that is chosen by the Spanning-Tree Protocol.

The root bridge should be placed as close to the core of the network as possible, or in a centrally located position within the LAN.

Incorrect Answers:

B: Core or distribution layer switches should be used as the root of the STP, not access layer switches.

C: Backbone fast does not provide for load balancing. The STP does not provide for any load balancing mechanisms, since its function is to detect and prevent loops.

D: Backbone fast is a Cisco proprietary feature that, once enabled on all switches of a bridge network, can save a switch up to 20 seconds (max_age) when recovering from an indirect link failure. It does not provide for more efficient workstation access.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 151

QUESTION NO: 2

You are an independent network consultant, and you've just been contracted to troubleshoot a multilayer switched network which is going through some intermittent end-station accessibility issues. The network has a redundant topology and core layer and

access layer switches (the root bridge is on one of the core switches). After going through your systematic troubleshooting methodology, you realize that the problem has to do with STP convergence and the accessibility issues depend on the individual ports' STP state. In this scenario you need to decrease STP convergence time. What's the best way of doing it?

- A. Enable PortFast on the core switched to accelerate the choice of a new root port.
- B. Enable UplinkFast on the WLAN that is having the most accessibility problems.
- C. Enable UplinkFast on the wiring closet switches at the edge of the network.
- D. Enable PortFast on all IOS based switches that have been configured for bridge priority.

Answer: C

Explanation:

If a switch loses connectivity, it begins using the alternate paths as soon as STP selects a new root port. When STP reconfigures the new root port, other interfaces flood the network with multicast packets, one for each address that was learned on the interface.

By using STP UplinkFast, you can accelerate the choice of a new root port when a link or switch fails or when STP reconfigures itself. The root port transitions to the forwarding state immediately without going through the listening and learning states, as it would with normal STP procedures. UplinkFast also limits the burst of multicast traffic by reducing the max-update-rate parameter (the default for this parameter is 150 packets per second). However, if you enter zero, station-learning frames are not generated, so the STP topology converges more slowly after a loss of connectivity.

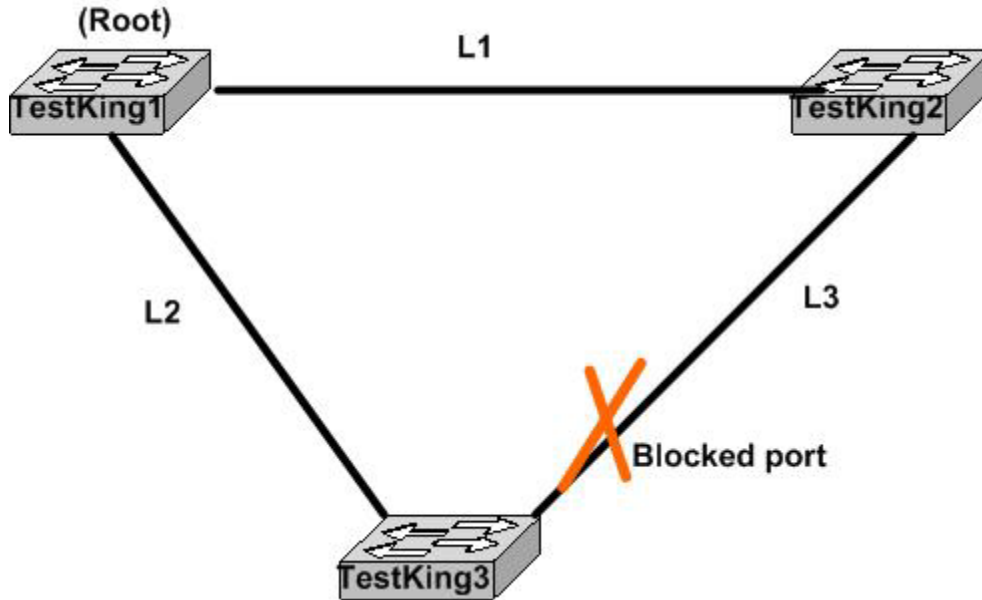
Note: UplinkFast is most useful in wiring-closet switches at the access or edge of the network. It is not appropriate for backbone devices.

Reference:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat2950/1216ea2b/scg/swgstp.htm>

QUESTION NO: 3

Three TestKing switches are connected as shown below:



Switch TestKing3 is configured with UplinkFast.
 If L2 were to fail, how much time will pass before Switch TestKing3 activates the port connection to L3?

- A. 1-5 seconds
- B. 15-20 seconds
- C. 30-35 seconds
- D. 45 seconds
- E. 60 seconds
- F. None of the above.

Answer: A

Explanation:

If Switch C detects a link failure on the currently active link L2 (a *direct* link failure), UplinkFast unblocks the blocked port on Switch C and transitions it to the forwarding state immediately, without transitioning the port through the listening and learning states. This switchover takes approximately one to five seconds.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps679/products_configuration_guide_chapter_09186a008007ef24.html

QUESTION NO: 4

Which of the following commands would you use if you wanted a Layer 2 access port to bypass the listening and learning states and move directly to the forwarding state?

- A. **spanning-tree uplinkfast**
- B. **spanning-tree port-priority**
- C. **spanning-tree portfast**
- D. **spanning-tree vlan vlan-id reset primary**

Answer: C

Explanation:

Spanning Tree PortFast causes an interface configured as a Layer 2 access port to enter the forwarding state immediately, bypassing the listening and learning states. You can use PortFast on Layer 2 access 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. If the interface receives a bridge protocol data unit (BPDU), which should not happen if the interface is connected to a single workstation or server, spanning tree puts the port into the blocking state.

To enable PortFast on a Layer 2 access port to force it to enter the forwarding state immediately, perform this procedure:

Task Command

Step 1 Specify an interface to configure. Switch(config)# **interface** {{fastethernet | gigabitethernet} slot/ port} | {port-channel port_channel_number}

Step 2 Enable PortFast on a Layer 2 access port connected to a single workstation or server.

You can use the **no** keyword to disable PortFast.

Switch(config-if)# [**no**] **spanning-tree portfast**

QUESTION NO: 5

You are a network troubleshooter and you've been called into the TestKing to manually put a switch port back into service after it was put into the error disabled state upon receipt of Spanning Tree messages. Which of the following STP features puts a switch port into an error-disabled state when it receives Spanning Tree data messages?

- A. BPDU Filtering
- B. Root Guard
- C. BPDU Guard
- D. Port Fast
- E. Loop Guard
- F. None of the above

Answer: C

Explanation:**Understanding BPDU Guard**

The BPDU guard feature can be globally enabled on the switch or can be enabled per interface, but the feature operates with some differences.

At the global level, you can enable BPDU guard on Port Fast-enabled ports by using the **spanning-tree portfast bpduguard default** global configuration command. Spanning tree shuts down ports that are in a Port Fast-operational state. In a valid configuration, Port Fast-enabled ports do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled port signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the port in the error-disabled state.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps5206/products_configuration_guide_chapter09186a008017ff97.html

QUESTION NO: 6

What's true about the UplinkFast feature? (Select two)

- A. It must be used with the PortFast feature enabled.
- B. When enabled, it is enabled for the entire switch and all VLANs.
- C. It should be configured on all switches, including the Root Bridge.
- D. When the primary Root Port uplink fails, another blocked uplink can be immediately brought up for use.

Answer: B, D


Explanation:

UplinkFast provides fast convergence in the network access layer after a spanning-tree topology change 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 (not including self-looped ports). The uplink group provides an alternate path in case the currently forwarding link fails.

Note UplinkFast is most useful in wiring-closet switches with a limited number of active VLANs. This enhancement might not be useful for other types of applications and should not be enabled on backbone or distribution layer switches.


[Figure 9-1](#) shows an example UplinkFast network topology. Switch A, the root switch, is connected directly to Switch B over link L1 and to Switch C over link L2. The port on Switch C that is connected to Switch B over link L3 is in blocking state.

Figure 9-1: UplinkFast Example Before Direct Link Failure



If Switch C detects a link failure on the currently active link L2 (a *direct* link failure), UplinkFast unblocks the blocked port on Switch C and transitions it to the forwarding state immediately, without transitioning the port through the listening and learning states (as shown in [Figure 9-2](#)). This switchover takes approximately one to five seconds.

Figure 9-2: UplinkFast Example After Direct Link Failure



As soon as the switch transitions the alternate port to the forwarding state, the switch begins transmitting dummy multicast frames on that port, one for each entry in the local EARL table (except those entries associated with the failed root port). By default, approximately 15 dummy multicast frames are transmitted per 100 milliseconds.

Each dummy multicast frame uses the station address in the EARL table entry as its source MAC address and a dummy multicast address (01-00-0C-CD-CD-CD) as the destination MAC address. Switches receiving these dummy multicast frames immediately update their EARL table entries for each source MAC address to use the new port, allowing the switches to begin using the new path almost immediately.

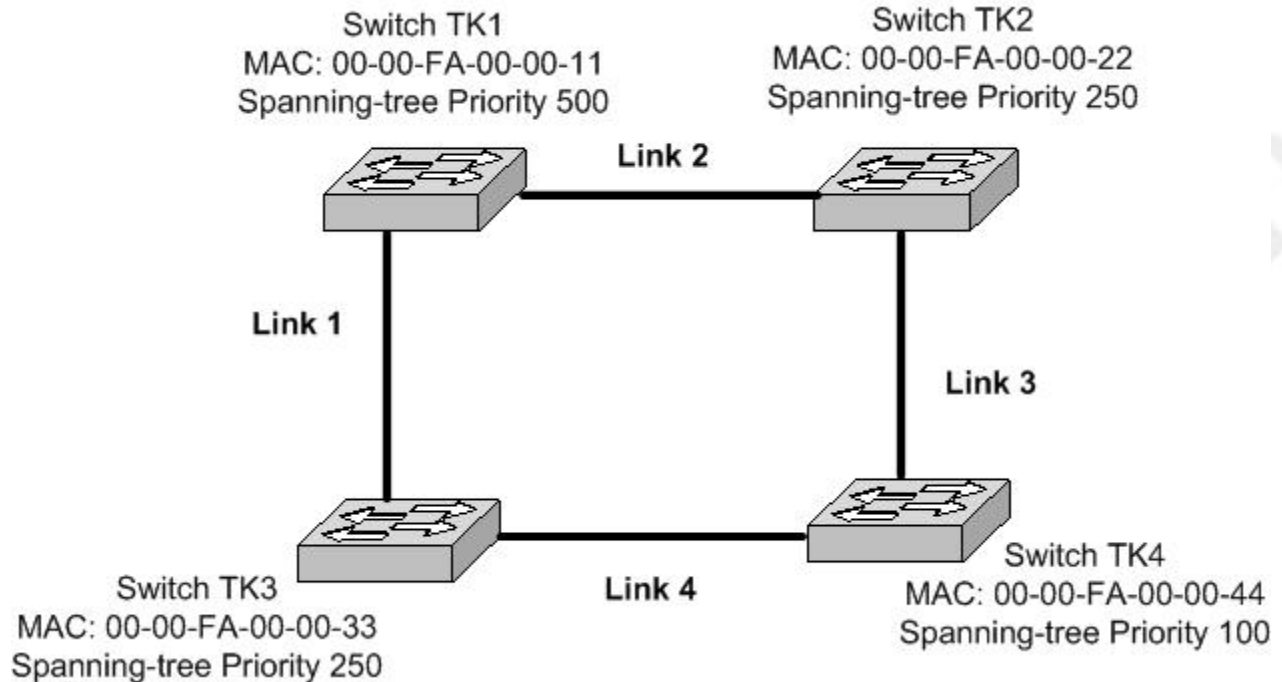
If connectivity on the original root port is restored, the switch waits for a period equal to twice the forward delay time plus 5 seconds before transitioning the port to the forwarding state in order to allow the neighbor port time to transition through the listening and learning states to the forwarding state.

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_5_4/config/stp_enha.htm#xtocid25869

QUESTION NO: 7

Four TestKing switches are connected together as shown below:



All the ports in the diagram all have the same spanning-tree cost and they're all configured as access ports. If you wanted to use UplinkFast to improve convergence time after link failure; where should you configure UplinkFast?

- A: Switch TK1
- B: Switch TK2
- C: Switch TK3
- D: Switch TK4

Answer: A

Explanation:

Note: The UplinkFast feature provides fast convergence in the network access layer after a spanning tree topology change by using uplink groups. UplinkFast accelerates the choice of a new root port when a link or switch fails or when STP reconfigures itself.

The UplinkFast feature is designed to run in a switched environment when the switch has at least one alternate/backup root port (port in blocking state), that is why Cisco recommends that UplinkFast be enabled only for switches with blocked ports, typically at the access-layer. Do not use on switches without the implied topology knowledge of an alternative/backup root link typically to distribution and core switches in Cisco multilayer design.

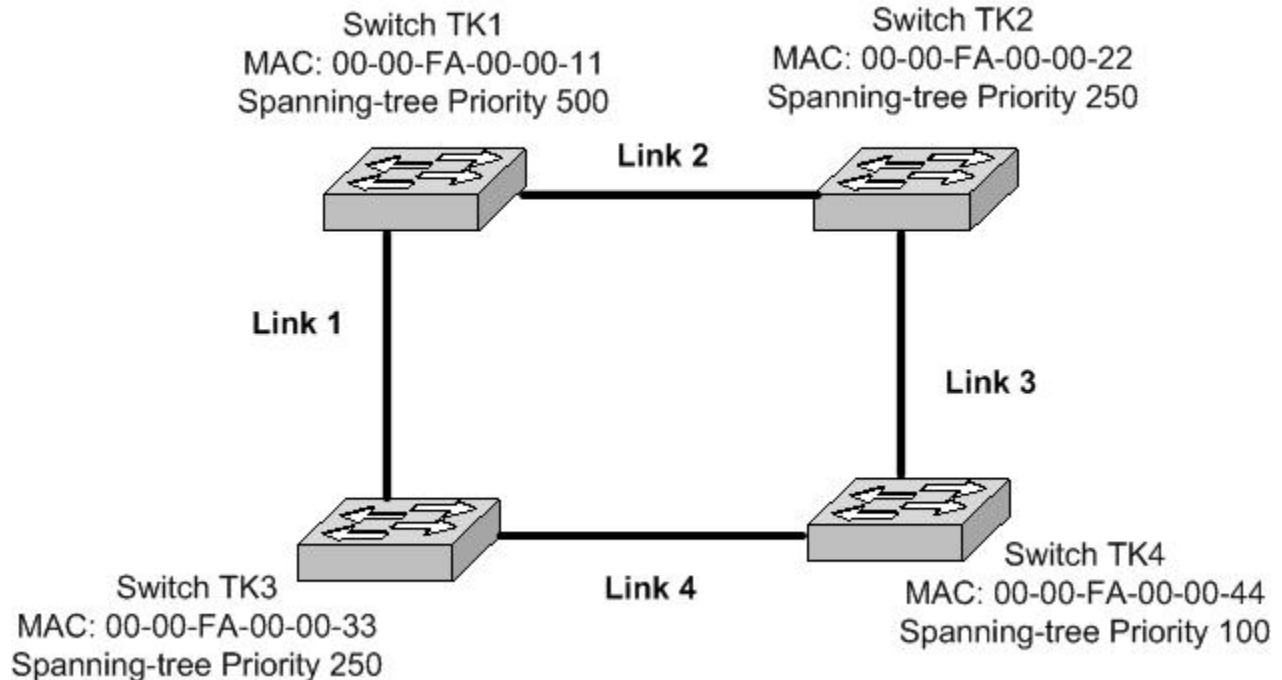
In this example, switch TK4 will become the root switch, since it has the lowest bridge priority. Therefore, uplink fast should be configured on switch TK1.

Reference:

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

QUESTION NO: 8

The TestKing network is depicted in the diagram below:



All the ports are configured as access ports and they all have the same spanning-tree cost. So you want to configure BackboneFast to improve convergence time if a link were to fail. Where should you configure the BackboneFast?

- A: Switch TK4 only.
- B: Switch TK2 and switch TK3 only.
- C: Switch TK2, switch TK3 and switch TK4 only.
- D: Switch TK1, switch TK2, switch TK3, and switch TK4.

Answer: D**Explanation:**

Backbone fast is a Cisco proprietary feature that, once enabled on all switches of a bridge network, can save a switch up to 20 seconds (max_age) when recovering from an indirect link failure.

You must enable BackboneFast on all switches in the network.

Note: The BackboneFast feature provides fast convergence in the network backbone after a spanning tree topology change occurs.

Reference:

http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a00800c2548.shtml

QUESTION NO: 9

What command would you enter, if you wanted to enable the spanning tree feature that causes a port to immediately switch from blocking to forwarding mode? (Type in answer below)

Answer: portfast

Explanation:

If you are connecting a workstation or a server with a single NIC card 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. Cisco added the PortFast or fast-start feature, which means the STP for this port will assume that the port is not part of a loop and will immediately move to the forwarding state, without going through the blocking, listening, or learning states. This command does not turn STP off. This command makes STP skip a few (unnecessary in this circumstance) steps in the beginning on the selected port.

The portfast variable, when enabled on a port, causes the port to immediately switch from blocking mode to forwarding mode. This helps prevent time-outs on clients that use Novell Netware or that use Dynamic Host Configuration Protocol (DHCP) to obtain an IP address. However, it is important that you do not use this command when you have switch-to-switch connection. It could potentially result in a loop. The 30-60 second delay that occurs when transitioning from blocking to forwarding mode transition prevents a temporal loop condition in the network when connecting two switches.

QUESTION NO: 10

Which of the following are FALSE, when configuring load sharing by using STP path costs? (Select all that apply)

- A. All priorities must be set to 100
- B. Load-sharing links can connect to different switches
- C. All priorities must be set to 0
- D. The switch must be restarted for the second time
- E. Both load-sharing links must connect to the same switch

Answer: A, C, D, E

Explanation:

The only true statement in the choices above is choice B. Load sharing divides the bandwidth supplied by parallel trunks connecting switches. To avoid loops, Spanning-Tree Protocol (STP) normally blocks all but one parallel link between switches. With load sharing, you divide the traffic between the links according to which VLAN the traffic belongs to. There are two ways to configure load sharing by using trunk ports: using STP port priorities or using STP path costs. If you configure load sharing using STP port priorities, both load-sharing links must be connected to the same switch. If you configure load sharing using STP path costs, each load-sharing link can be connected to the same switch or to two different switches.

QUESTION NO: 11

If a switch is configured as a secondary root, what is the new default spanning tree bridge priority value? (Type in the answer below)

Answer: 16384

Explanation:

When you configure a switch as the secondary root, the spanning tree bridge priority is modified from the default value (32768) to 16384. This means that the switch is likely to become the root bridge for the specified VLANs if the primary root bridge fails (assuming the other switches in the network use the default bridge priority of 32768).

QUESTION NO: 12

PortFast is being configured on switch TK1. What should you take into consideration when configuring a switch with PortFast? (Select two choices below that are true statements)

- A. It increases the forward delay time interval to 30 seconds.
- B. It should be enabled on ports connecting to hubs and routers.
- C. It should not be enabled on ports with redundant links to another switch.
- D. It enables fast connectivity to be established on the access layer port to a booting workstation.

Answer: C, D

Explanation:

C: Portfast on redundant links could cause network loops when improperly used.

D: PortFast is used to make a point-to-point port almost immediately enter into forwarding state by decreasing the time of the listening and learning states.

Incorrect Answers:

A: PortFast decreases the forward delay time.

B: Ideally PortFast should only be used on point-to-point links connected only to workstations or servers.

QUESTION NO: 13

What should you take into consideration when configuring a switch with UplinkFast?

(Select two.)

- A. It must be used with the PortFast feature enabled.
- B. When enabled, it is enabled for the entire switch and all VLANs.
- C. It should be configured on all switches, including the root bridge.
- D. When the primary Root Port uplink fails, another blocked uplink can be immediately brought up for use.

Answer: B, D

Explanation:

B: All VLANs on the switch are affected and you cannot configure UplinkFast on individual VLANs.

D: When a link fault occurs on the primary root link, UplinkFast transitions the blocked port to a forwarding state. UplinkFast changes the port without passing through the listening and learning phases.

Incorrect Answers:

A: These two features can be used independently from each other.

C: This is true of the Backbone fast, not of the uplink fast feature.

QUESTION NO: 14

Which three conditions need to be present for UplinkFast to trigger a fast reconfiguration?

Choose three.

- A. The switch must have at least one unblocked port.
- B. The switch must have UplinkFast enabled.
- C. The switch must be configured for one VLAN.
- D. The switch must have at least one blocked port.
- E. The failure must be on the root port.
- F. The switch must be enabled on a VLAN with switch priority configured.


Answer: B, D, E

Explanation:

Uplink Fast Theory of Operation:

The UplinkFast feature is based on the definition of an uplink group. On a given switch, the uplink group consists in the root port and all the ports that provide an alternate connection to the root bridge. If the root port is failing (that is, if the primary uplink fails), a port with next lowest cost from the uplink group is selected to immediately replace it.

The following diagram helps to explain what the UplinkFast feature is based on:



In this diagram, root ports are represented with a blue R and designated ports are represented with a green d. The green arrows represent the BPDUs generated by the root bridge and retransmitted by the bridges on their designated ports. Without entering a formal demonstration, we can determine the following about BPDUs and ports in a stable network:

- When a port is receiving a BPDU, it has a path to the root bridge. This is because BPDUs are originated from the root bridge. In this diagram, check switch A: three of its ports are receiving BPDUs, and three of its ports lead to the root bridge. The port on A that is sending BPDU is designated and not leading to the root bridge.
- On any given bridge, all ports receiving BPDUs are blocking, except the root port. A port receiving a BPDU is leading to the root bridge. If we had a bridge with two ports leading to the root bridge, we would have a bridging loop.
- A self-looped port does not provide an alternate path to the root bridge. See switch B in the diagram. Switch B's blocked port is self-looped, which means that it cannot receive its own BPDUs. In this case, the blocked port is not providing an alternate path to the root.

On a given bridge, the root port and all blocked ports that are not self-looped form the uplink group. The following section describes step-by-step how UplinkFast achieves fast convergence using an alternate port from this uplink group.

Note: UplinkFast is only working when the switch has blocked ports. The feature is typically designed for an access switch having redundant blocked uplinks. When you enable UplinkFast, it is enabled for the entire switch and cannot be enabled for individual VLANs.

QUESTION NO: 15

Three TestKing switches are connected together as shown in the diagram below:



Switch TestKing1, the root bridge, is connected directly to Switch TestKing2 over Link L1 and to Switch TestKing3 over link L2. The Layer 2 LAN interface on TestKing3 that is connected directly to TestKing2 is in the blocking state. If TestKing3 detects a link failure on the currently active link L2 on the root port, which spanning tree enhancement will allow Switch TestKing3 to unblock the blocked port and transition it to forwarding without going through the listening and learning states?

- A. PortFast
- B. UplinkFast
- C. BackboneFast
- D. Rapid spanning tree
- E. None of the above

Answer: B


Explanation:

The following example explains the Uplink fast procedure step by step:

Uplink Failure With Uplink Fast Enabled

This section details the steps for UplinkFast recovery. We will use the network diagram that was introduced at the beginning of the document.


Immediate Switch Over to the Alternate Uplink



Follow these steps for an immediate switch over to the alternate uplink:

1. A's uplink group consists of P1 and its nonself-looped blocked port, P2.
2. When the link between D1 and A fails, A detects a link down on port P1.


It knows immediately that its unique path to the root bridge is lost (other paths are via the uplink group, for example, port P2, which is blocked).



Time=0

3. A places port P2 in forwarding mode immediately, thus violating the standard STP procedures.

We know that there will be no loop in the network, as the only path to the root bridge is currently down. Therefore, recovery is almost immediate.



Time=0+

Reference: <http://www.cisco.com/warp/public/473/51.html>

QUESTION NO: 16

Leading the way in IT testing and certification tools, www.testking.com

Which statement is true about the STP Path Cost on a particular port?

- A. It is known only to the local switch where the port resides.
- B. It can be modified to help determine Root Bridge selection.
- C. Modifying it can cause TCN BPDU to be sent to the Root Bridge.
- D. When increased, it can provide higher bandwidth to a connecting port.
- E. None of the above
- F. All of the above.

Answer: A

Explanation:

When two ports on a switch are part of a loop, the spanning tree port priority and port path cost setting determine which port is put in the forwarding state and which port is put in the blocking state. The spanning tree port priority value represents the location of an interface in the network topology and how well located it is to pass traffic. The spanning tree port path cost value represents media speed.

The path cost is known locally to the switch only, as the path cost information is not advertised to other switches within a network.

Incorrect Answers:

- B: Modifying the cost may help determine which port is a root port, but it will not aid in the selection of the root switch since the path cost value is only known locally.
- C: The cost information is only stored locally on the switch, so changes made are not propagated to other switches.
- D: Although the cost is a direct reflection of the speed on an interface, it does not affect the actual bandwidth or throughput of an interface.

QUESTION NO: 17

Switch TK1 has been configured with the root guard feature. What statement is true if the spanning tree enhancement Root Guard is enabled?

- A. If BPDUs are not received on a non-designated port, the port is moved into the STP loop-inconsistent blocked state
- B. If BPDUs are received on a PortFast enabled port, the port is disabled.
- C. If superior BPDUs are received on a designated port, the interface is placed into the root-inconsistent blocked state.
- D. If inferior BPDUs are received on a root port, all blocked ports become alternate paths to the root bridge.

Answer: C

Explanation:

Root guard is configured on a per-port basis, and does not allow the port to become a STP root port. This means that the port is always STP-designated. If there is a better BPDU received on this port, root guard will put the port into root-inconsistent STP state, rather than taking the BPDU into account and electing a new STP root. Root guard needs to be enabled on all ports where the root bridge should not appear. In a way one can configure a perimeter around part of network where STP root is allowed to be located.

Reference:

http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a00800ae96b.shtml

Section 8: Implement IP technology on a switched network with auxiliary VLANs (7 questions)

QUESTION NO: 1

What command would you enter if you wanted to enable 'IP accounting' on one of your interfaces? (Type in answer below):

Answer: ip accounting

Explanation:

To enable IP accounting on an interface, use the ip accounting interface configuration command. To disable IP accounting, use the no form of this command.

ip accounting [access-violations]
no ip accounting [access-violations]

QUESTION NO: 2

VLANs are being implemented on the TestKing network. What does Cisco recommend as the ideal ratio of VLANs to IP subnets?

- A. One-to-one
- B. Many-to-one
- C. One-to-many
- D. VLANs are mapped to MAC addresses
- E. None of the above

Answer: A

Explanation:

Cisco Systems recommend a one-to-one correspondence between VLANs and IP subnets. There should be a separate IP subnet for each VLAN.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 90.

QUESTION NO: 3

Which command could you use on a Catalyst 3550 switch if you wanted to set a port to operate as a nontrunking, single VLAN, Layer 2 interface so it will send and receives non-encapsulated (non-tagged) frames?

- A. switchport
- B. switchport mode access
- C. switchport nonegotiate
- D. switchport access vlan dynamic
- E. None of the above

Answer: B

Explanation:

To set up a single vlan non trunking port you'll need to create a static VLAN membership. "The **switchport mode access command** configures the port for static VLAN membership."

Reference: CCNP Switching Exam Certification Guide: page 104, David Hucaby & Tim Boyles, Cisco Press 2001, ISBN 1-58720 000-7

QUESTION NO: 4

While logged into a TestKing switch you issue the following command:

```
TestKingSwitch(config-mst)# instance 10 vlan 11-12
```

What does this command accomplish?

- A. It enables a PVST+ instance of 10 for vlan 11 and vlan 12
- B. It enables vlan 11 and vlan 12 to be part of the MST region 10
- C. It maps vlan 11 and vlan 12 to the MST instance of 10.
- D. It creates an Internal Spanning Tree (IST) instance of 10 for vlan 11 and vlan 12
- E. It create a Common Spanning Tree (CST) instance of 10 for vlan 11 and vlan 12
- F. It starts two instances of MST, one instance for vlan 11 and another instance for vlan 12.

Answer: C

Explanation:

MST extends the IEEE 802.1w rapid spanning tree (RST) algorithm to multiple spanning trees. This extension provides both rapid convergence and load balancing in a VLAN environment. MST converges faster than Per VLAN Spanning Tree Plus (PVST+) and is backward compatible with 802.1D STP, 802.1w (Rapid Spanning Tree Protocol [RSTP]), and the Cisco PVST+ architecture.

MST allows you to build multiple spanning trees over trunks. You can group and associate VLANs to spanning tree instances. Each instance can have a topology independent of other spanning tree instances. This architecture provides multiple forwarding paths for data traffic and

enables load balancing. Network fault tolerance is improved because a failure in one instance (forwarding path) does not affect other instances.

Map the VLANs to an MST instance.

If you do not specify the **vlan** keyword, you can use the **no** keyword to unmap all the VLANs that were mapped to an MST instance.

If you specify the **vlan** keyword, you can use the **no** keyword to unmap a specified VLAN from an MST instance.

```
Switch(config-mst)# instance
instance_number vlan vlan_range
```

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps663/products_configuration_guide_chapter09186a00800dde9e.html#36881

QUESTION NO: 5

What happens when you apply an outgoing access list to an interface of a Catalyst switch?

- A. it will purge any entries for flows on that interface and records no new entries
- B. it will generate excessive MLSP messages
- C. it will record packets only if the administrator sets the MLS RP IP ACL command on the interface
- D. it will result in no action taken

Answer: A

Explanation:

According to Cisco: Traditionally, switches operated at Layer 2 only; switches switched traffic within a VLAN and routers routed traffic between VLANs. Catalyst 6000 family switches with the Multilayer Switch Feature Card (MSFC) can accelerate packet routing between VLANs by using Layer 3 switching (Multilayer Switching [MLS]). The switch first bridges the packet, the packet is then routed internally without going to the router, and then the packet is bridged again to send it to its destination. During this process, the switch can access control all packets it switches, including packets bridged within a VLAN. IOS ACLs access control routed traffic between VLANs, and VLAN ACLs (VACLs) access control all packets. Standard and extended IOS ACLs are used to classify packets. Classified packets can be subject to a number of features such as access control (security), encryption, policy-based routing, and so on. Standard and extended IOS ACLs are only configured on router interfaces and applied on routed packets.

QUESTION NO: 6

Some Cisco switches process Access Control Lists (ACL's) in their hardware. What would happen if the hardware reaches its maximum storage capacity of ACLs? (Select all that apply.)

- A. Packets are dropped.
- B. Packet filtering will be accomplished.
- C. Performance is increased.
- D. Performance is decreased.
- E. None of the above

Answer: B, D

Explanation:

Determining if the ACL Configuration Fits in Hardware

As previously stated, ACL processing in the Catalyst 3550 switch is mostly accomplished in hardware. However, if the hardware reaches its capacity to store ACL configurations, the switch software attempts to fit a simpler configuration into the hardware. This simpler configuration does not do all the filtering that has been configured, but instead sends some or all packets to the CPU to be filtered by software. In this way, **all configured filtering will be accomplished**, but **performance is greatly decreased** when the filtering is done in software.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps646/products_configuration_guide_chapter_09186a008007e701.html

QUESTION NO: 7

Switch TK1 is a Catalyst 3550. What kinds of access-control lists can you use on TK1 to filter traffic? (Select two)

- A. CBAC
- B. VLAN Maps
- C. Router ACLs
- D. Reflexive ACL

Answer: B, C

Explanation:

VLAN Maps and conventional access control lists are how the Catalyst 3550 switch filters traffic.

Incorrect Answers:

A, D: CBAC (Contact Based Access Control) and reflexive access lists are beyond the scope of this test, and are available on routers using specialized IOS versions.

www.testking.com

Section 9: Configure and verify router redundancy using HSRP, VRRP, GLBP, SRM, and SLB (20 questions)

QUESTION NO: 1

HSRP is being set up between two TestKing devices. In what three states is it possible for an HSRP router to be in? (Select three)

- A. Standby
- B. Established
- C. Active
- D. Idle
- E. Backup
- F. Init

Answer: A, C, F

Explanation:

With HSRP, a set of routers work together to present the illusion of a single virtual router to the hosts on the LAN. This set is known as an HSRP group or a standby group. A single router elected from the group is responsible for forwarding the packets that hosts send to the virtual router. This router is known as the Active router. Another router is elected as the Standby router. In the event that the Active router fails, the Standby assumes the packet-forwarding duties of the Active router. Although an arbitrary number of routers may run HSRP, only the Active router forwards the packets sent to the virtual router. Before a router becomes the active or standby router, it will be in the Init (initial) state.

Reference:

http://www.cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a0080094a91.shtml

QUESTION NO: 2

What three tasks must a network administrator perform to properly configure Hot Standby Routing Protocol (HSRP)? (Select three)

- A. Define the encapsulation type.
- B. Define the standby router.
- C. Define the standby IP address.
- D. Enable the standby priority.

Answer: B, C, D

Explanation:

Three of the required configuration commands needed for enabling HSRP is to define the standby routing process, define the HSRP IP address, and configure the HSRP priority.

Configuring HSRP:

- Configuring an interface to participate in an HSRP standby group
- Assigning HSRP standby priority
- Configuring HSRP standby pre-empt
- Configuring HSRP over trunk links
- Configuring hello message timers
- HSRP interface tracking
- Displaying the status of HSRP

Incorrect Answers:

A: There are no encapsulation options for enabling HSRP.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 272

QUESTION NO: 3

To protect against first-hop router failure, four protocols were developed to ensure IP routing redundancy. Which of the following are they? (Select four)

- A. HSRP
- B. IRDP
- C. ICMP
- D. VRRP
- E. MSTP
- F. GLBP

Answer: A, B, D, F

Explanation:

A: HSRP is the Hot Standby Routing Protocol, which is the Cisco proprietary method for automatic failover and provides for redundant default gateways for hosts.

B: Some newer IP hosts use ICMP Router Discovery Protocol (IRDP) (RFC 1256) to find a new router when a route becomes unavailable. A host that runs IRDP listens for hello multicast messages from its configured router and uses an alternate router when it no longer receives those hello messages.

D: VRRP is the Virtual Router Redundancy Protocol, which is similar in many ways to HSRP. One key difference is that VRRP is standards based, where HSRP is Cisco developed.

F: Gateway Load Balancing Protocol (GLBP) protects data traffic from a failed router or circuit, like Hot Standby Router Protocol (HSRP) and Virtual Router Redundancy Protocol (VRRP), while allowing packet load sharing between a group of redundant routers.

QUESTION NO: 4

HSRP is compatible over which of the following networks? (Select all that apply)

- A. Banyan VINES
- B. IP
- C. IBM DLC
- D. Novell IPX
- E. AppleTalk

Answer: A, B, D, E

Explanation:

According to the online documentation provided by Cisco:

You can configure HSRP in networks that, in addition to IP, run AppleTalk, Banyan VINES, and Novell IPX. AppleTalk and Novell IPX continue to function when the standby router becomes the active router, but they take time to adapt to topology changes. In general, AppleTalk hosts discover a new active router in less than 30 seconds. Novell 4.x hosts discover a new active router in 10 seconds, on average. Novell 2.x or Novell 3.x hosts might require more time to adapt.

QUESTION NO: 5

HSRP has been configured between two TestKing devices. What kind of message does an HSRP configured router send out every 3 seconds? (Select all that apply)

- A. Retire
- B. Coup
- C. Resign
- D. Send
- E. Hello

Answer: E

Explanation:

Hello-The hello message conveys to other HSRP routers the router's HSRP priority and state information. By default, an HSRP router sends hello messages every three seconds.

Incorrect Answers:

A, D: These messages are not used by HSRP.

B: Coup-When a standby router assumes the function of the active router, it sends a coup message. This message is used by HSRP, but it is not sent out every 3 seconds.

C: Resign-A router that is the active router sends this message when it is about to shut down or when a router that has a higher priority sends a hello message. This message is only sent before it resigns, not every 3 seconds.

QUESTION NO: 6

HSRP has been configured between two TestKing devices. Which of the following describe reasons for deploying HSRP? (Select all that apply)

- A. HSRP provides redundancy and fault tolerance
- B. HSRP allows one router to automatically assume the function of the second router if the second router fails
- C. HSRP allows one router to automatically assume the function of the second router if the second router starts
- D. HSRP provides redundancy and load balancing

Answer: A, B, D

Explanation:

One way to achieve near-100 percent network uptime is to use HSRP, which provides network redundancy for IP networks, ensuring that user traffic immediately and transparently recovers from first hop failures in network edge devices or access circuits.

By sharing an IP address and a MAC (Layer 2) address, two or more routers can act as a single "virtual" router. The members of the virtual router group continually exchange status messages. This way, one router can assume the routing responsibility of another, should it go out of commission for either planned or unplanned reasons. Hosts continue to forward IP packets to a consistent IP and MAC address, and the changeover of devices doing the routing is transparent. Through the use of multiple HSRP standby groups, traffic can be load balanced between the HSRP routers. For example, users on one VLAN could use one router as the primary HSRP router, and users on another VLAN can use the other HSRP router as the primary.

QUESTION NO: 7

Which one of the statements below correctly describes the Virtual Router Redundancy Protocol (VRRP)?

- A. A VRRP group has one active and one or more standby virtual routers.
- B. A VRRP group has one master and one or more backup virtual routers.
- C. A VRRP group has one active and one or more standby virtual routers.
- D. A VRRP group has one master and one redundant virtual router.

Answer: B

Explanation:

The Virtual Router Redundancy Protocol (VRRP) feature can solve the static configuration problem. VRRP enables a group of routers to form a single *virtual router*. The LAN clients can

then be configured with the virtual router as their default gateway. The virtual router, representing a group of routers, is also known as a VRRP group.

In a topology where multiple virtual routers are configured on a router interface, the interface can act as a master for one virtual router and as a backup for one or more virtual routers.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1612/products_feature_guide09186a0080080a60.html

QUESTION NO: 8

You want to allow Router TK1 to immediately become the active router if its priority is highest than the active router fails. What command would you use if you wanted to configure this?

- A. en standby preempt
- B. standby preempt enable
- C. standby preempt
- D. hot standby preempt

Answer: C

Explanation:

The HSRP preemption feature enables the router with highest priority to immediately become the Active router. Priority is determined first by the priority value that you configure, and then by the IP address. In each case a higher value is of greater priority.

When a higher priority router preempts a lower priority router, it sends a coup message. When a lower priority active router receives a coup message or hello message from a higher priority active router, it changes to the speak state and sends a resign message. To configure preemption, use the “standby *standby-number* preempt” command.

QUESTION NO: 9

Routers TK1 and TK2 are configured for HSRP as shown below:

```
Router TK1:
interface ethernet 0
 ip address 20.6.2.1 255.255.255.0
 standby 35 ip 20.6.2.21
 standby 35 priority 100
interface ethernet 1
 ip address 20.6.1.1.2 255.255.255.0
 standby 34 ip 20.6.1.21
```

```

Router TK2:
interface ethernet 0
  ip address 20.6.2.2 255.255.255.0
  standby 35 ip 20.6.2.21
interface ethernet 1
  ip address 20.6.1.1.1 255.255.255.0
  standby 34 ip 20.6.1.21
  standby 34 priority 100

```

You have configured the routers TK1 & TK2 with HSRP. While debugging router TK2 you notice very frequent HSRP group state transitions. What is the most likely cause of this?

- A. physical layer issues
- B. no spanning tree loops
- C. use of non-default HSRP timers
- D. failure to set the command `standby 35 preempt`

Answer: A

Explanation: TK2 is not able to from the standby state to reach the active state. This could be caused by missing HSRP hello messages. There are several possible causes for HSRP packets to get lost between the peers. The most common problems are Physical Layer Problems or excessive network traffic caused by Spanning-Tree Issues.

Note:

Hot Standby Routing Protocol (HSRP) is a Cisco proprietary protocol used for allowing redundant connections. It can keep core connectivity if the primary routing process fails. HSRP defines six states in which an HSRP router may run: initial, learn, listen, speak, standby, and active.

Incorrect Answers:

- B: Spanning tree loops does not affect this problem.
- C: Not a likely cause. Besides, in the example here the default values were indeed used.
- D: If the Preempt option is set, then an election of the Active router will take place. This process is called a coup. However, an election would take place by default.

Reference:

Understanding and Troubleshooting HSRP Problems in Catalyst Switch Networks
<http://www.cisco.com/warp/public/473/62.shtml>
 RFC 2281, Cisco Hot Standby Router Protocol (HSRP)

QUESTION NO: 10

Which type of scheme describes the default operation of Gateway Load Balancing Protocol (GLBP)?

- A. per host using a round robin scheme
- B. per host using a strict priority scheme
- C. per session using a round robin scheme
- D. per session using a strict priority scheme
- E. per GLBP group using a round robin scheme
- F. per GLBP group using a strict priority scheme

Answer: A

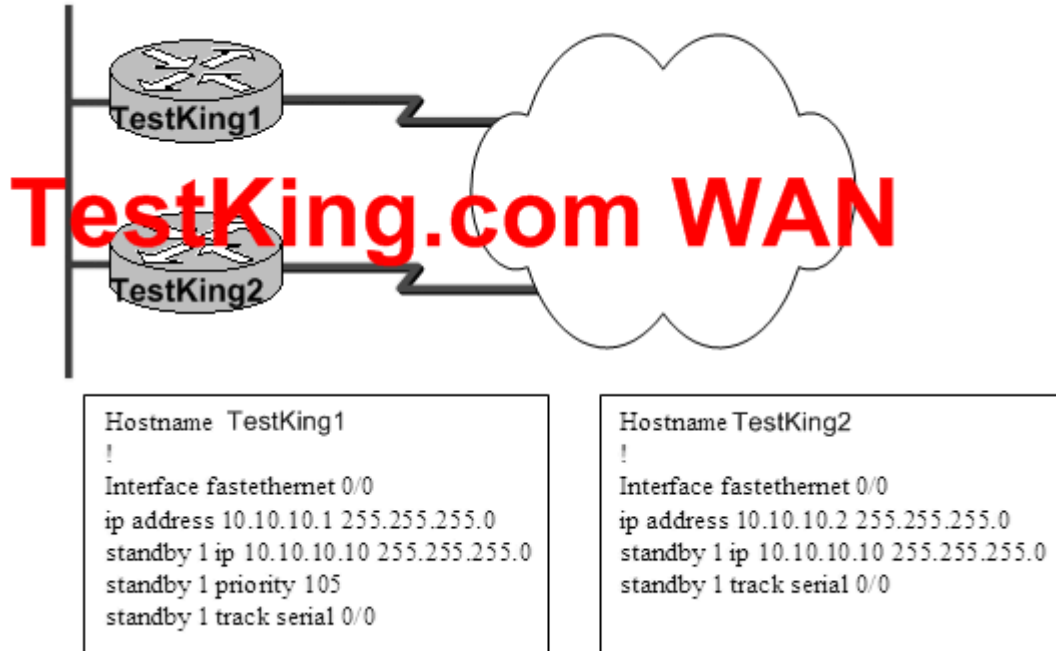
Explanation:

The Gateway Load Balancing Protocol feature provides automatic router backup for IP hosts configured with a single default gateway on an IEEE 802.3 LAN. Multiple first hop routers on the LAN combine to offer a single virtual first hop IP router while sharing the IP packet forwarding load. Other routers on the LAN may act as redundant GLBP routers that will become active if any of the existing forwarding routers fail.

GLBP performs a similar, but not identical, function for the user as the HSRP and the VRRP. HSRP and VRRP protocols allow multiple routers to participate in a virtual router group configured with a virtual IP address. One member is elected to be the active router to forward packets sent to the virtual IP address for the group. The other routers in the group are redundant until the active router fails. These standby routers have unused bandwidth that the protocol is not using. Although multiple virtual router groups can be configured for the same set of routers, the hosts must be configured for different default gateways, which results in an extra administrative burden. GLBP provides load balancing over multiple routers (gateways) using a single virtual IP address and multiple virtual MAC addresses. Each host is configured with the same virtual IP address, and all routers in the virtual router group participate in forwarding packets. In this way, per host load balancing is achieved using a round robin mechanism.

QUESTION NO: 11

The TestKing network is using two routers with HSRP for their Internet access as shown below:



Which command will need to be added to TestKing2 to ensure that it will take over if serial 0/0 on TestKing1 fails?

- A. standby 1 preempt
- B. standby 1 track 10.10.10.1
- C. standby 1 priority 130
- D. standby 1 track fastethernet 0/0
- E. None of the above

Answer: A

Explanation:

When this command is configured, the router is configured to preempt, which means that when the local router has a Hot Standby priority higher than the current active router, the local router should attempt to assume control as the active router. If the hsrp preempt command is not configured, the local router assumes control as the active router only if it receives information indicating that no router currently is in the active state (acting as the designated router).

In this example, TestKing1 was properly configured to lower its own HSRP priority when the serial 0 interface goes down. However, even if this happens, router TestKing2 will not become the active router unless it is configured to pre-empt, or take over, if it suddenly has a higher priority than TestKing1.

QUESTION NO: 12

**On a 3550 EMI switch, which three types of interfaces can be used to configure HSRP?
(Select three)**

- A. Loopback interface
- B. SVI interface
- C. Routed port
- D. Access port
- E. EtherChannel port channel
- F. BVI interface

Answer: B, C, E

Explanation:

This Hot Standby Router Protocol (HSRP) provides routing redundancy for routing IP traffic without being dependent on the availability of any single router. To use this feature, you must have the enhanced multilayer software image installed on your switch. All Catalyst 3550 Gigabit Ethernet switches ship with the enhanced multilayer software image (EMI) installed. Catalyst 3550 Fast Ethernet switches can be shipped with either the standard multilayer software image (SMI) or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the SMI to the EMI.

Only routed interfaces that provide access to hosts can be configured for HSRP. These interfaces include: routed Ethernet, routed fast Ethernet, routed Gigabit Ethernet, SVI, and EtherChannel.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps646/products_configuration_guide_chapter09186a00800c9fec.html

QUESTION NO: 13

Which protocol is an extension of ICMP that allows an IP host on the TestKing network to find a new router when a router becomes unavailable, as defined by RFC 1256?

- A. ICMP (IRDP)
- B. SNMP
- C. HSRP
- D. VRRP
- E. None of the above

Answer: A

Explanation:

RFC 1256 defines the ICMP Router Discovery Protocol. Some newer IP hosts use ICMP Router Discovery Protocol (IRDP) to find a new router when a route becomes unavailable.

ICMP Router Discovery Protocol (IRDP) enables a host to determine the address of a router that it can use as a default gateway. Similar to ES-IS but used with IP.

Router discovery uses Internet Control Message Protocol (ICMP) router advertisements and router solicitation messages to allow a host to discover the addresses of operational routers on the subnet. Hosts must discover routers before they can send IP datagrams outside their subnet. Router discovery allows a host to discover the addresses of operational routers on the subnet. Each router periodically multicasts a router advertisement from each of its multicast interfaces, announcing the IP address of that interface. Hosts listen for advertisements to discover the addresses of their neighboring routers. When a host starts, it can send a multicast router solicitation to ask for immediate advertisements.

The router discovery messages do not constitute a routing protocol. They enable hosts to discover the existence of neighboring routers, but do not determine which router is best to reach a particular destination.

References: CCNP Switching Exam Certification Guide, David Hucaby CCIE #4594 & Tim Boyles. Cisco Press, ISBN 1-58720-000-7 Page 308 *“IRDP is an extension to ICMP that provides a mechanism for routers to advertise useful default routes.”*

<http://www.javvin.com/protocolIRDP.html>

<http://www.faqs.org/rfcs/rfc1256.html>

QUESTION NO: 14

Match the HSRP states on the left with the correct definition on the right.

Select from these

Learn

Listen


Speak

Standby

Active

Initial

Place here




State from which the routers begin the HSRP process




A candidate to become the next active router




The router is still waiting to hear from the active router



The router is currently forwarding packets



Listens for hello messages from the active and standby router



Participates in the election for the active or standby router

Answer:

Select from these

Place here

Initial	State from which the routers begin the HSRP process
Standby	A candidate to become the next active router
Learn	The router is still waiting to hear from the active router
Active	The router is currently forwarding packets
Listen	Listens for hello messages from the active and standby router
Speak	Participates in the election for the active or standby router

QUESTION NO: 15

In the hardware address 0000.0c07.ac0av, what does 07.ac represent?

- A. HSRP well-known physical MAC address
- B. Vendor code
- C. HSRP router number
- D. HSRP group number
- E. HSRP well-known virtual MAC address

Answer: E

Explanation:

HSRP code (HSRP well-known virtual MAC address) – The fact that the MAC address is for an HSRP virtual router is indicated in the next two bytes of the address. The HSRP code is always 07.ac. The HSRP protocol uses a virtual MAC address, which always contains the 07.ac numerical value.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 268

QUESTION NO: 16

The TestKing network needs to enhance the reliability of the network. Which of the following protocols provides network redundancy for IP networks, ensuring that user traffic immediately and transparently recovers from first-hop failures in network edge devices or access circuits, as defined by RFC 2281?

- A. STP
- B. IRDP
- C. ICMP
- D. HSRP

Answer: D

Explanation:

HSRP is defined in RFC 2281. The Hot Standby Router Protocol, HSRP, provides a mechanism which is designed to support non-disruptive failover of IP traffic in certain circumstances. In particular, the protocol protects against the failure of the first hop router when the source host cannot learn the IP address of the first hop router dynamically. The protocol is designed for use over multi-access, multicast or broadcast capable LANs (e.g., Ethernet). HSRP is not intended as a replacement for existing dynamic router discovery mechanisms and those protocols should be used instead whenever possible. A large class of legacy host implementations that do not support dynamic discovery are capable of configuring a default router. HSRP provides failover services to those hosts.

Reference: <http://www.faqs.org/rfcs/rfc2281.html>

QUESTION NO: 17

Which of the following protocols enables a group of routers to form a single virtual router, and then use the real IP address of a router as the gateway address, as defined in RFC 2338?

- A. Proxy ARP
- B. HSRP
- C. IRDP
- D. VRRP
- E. GLBP

Answer: D

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

The Virtual Router Redundancy Protocol (VRRP) feature enables a group of routers to form a single *virtual router*. The LAN clients can then be configured with the virtual router as their default gateway. The virtual router, representing a group of routers, is also known as a VRRP group. VRRP is defined in RFC 2338.

Reference: <http://www.faqs.org/rfcs/rfc2338.html>

QUESTION NO: 18

Routers can send hello messages in three HSRP states. Which ones are they? (Select three)

- A. standby
- B. learn
- C. listen
- D. speak
- E. active

Answer: A, D, E

Explanation:

The various HSRP states are described below:

Listen: The router knows the virtual IP address, but is neither the active router nor the standby router. It listens for hello messages from those routers.

Speak: The router sends periodic hello messages, and is actively participating in the election of the active and/or standby router. A router cannot enter speak state unless it has the virtual IP address.

Standby: The router is a candidate to become the next active router, and sends periodic hello messages. Excluding transient conditions, there would be at most one router in the group in standby state.

Active: The router is currently forwarding packets that are sent to the group's virtual MAC address. The router sends periodic hello messages. Excluding transient conditions, there must be at most one router in active state in the group.

Initial: This is the starting state, and indicates that HSRP is not running. This state is entered via a configuration change, or when an interface first comes up.

Learn: The router has not determined the virtual IP address, and has not yet seen an authenticated hello message from the active router. In this state, the router is still waiting to hear from the active router.

QUESTION NO: 19

Two TestKing routers are configured for HSRP. Cisco's Hot Standby Routing Protocol (HSRP) can provide automatic router backup over which networks?

- A. Ethernet and FDDI
- B. Ethernet, FDDI and Token Ring LANs
- C. Token Ring LANs only
- D. VINES and IPX only
- E. Ethernet and Token Ring LANs

Answer: B

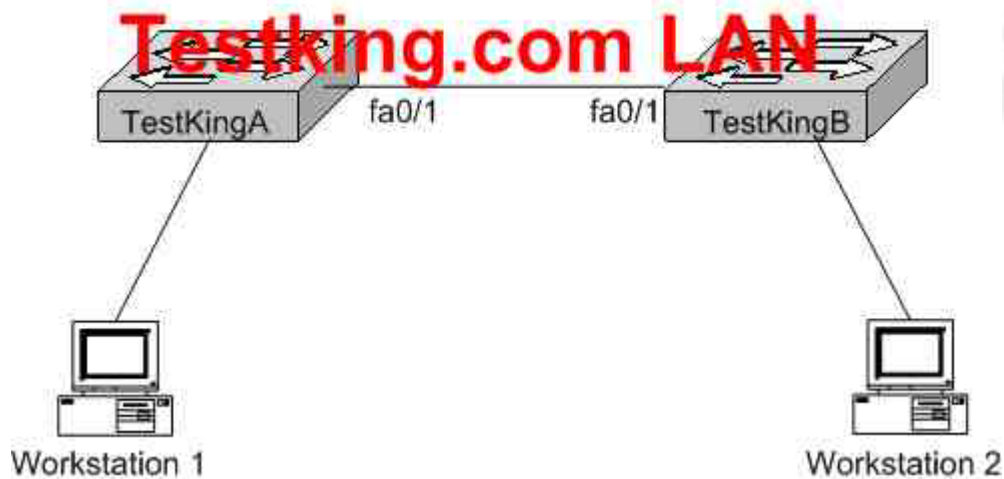
Explanation:

Cisco's Hot Standby Routing Protocol (HSRP) provides automatic router backup when you configure it on Cisco routers that run the Internet Protocol (IP) over Ethernet, Fiber Distributed Date Interface (FDDI), and Token Ring local-area networks (LANs). HSRP is compatible with Novell's Internetwork Packet Exchange (IPX), AppleTalk, and Banyan VINES, and it is compatible with DECnet and Xerox Network Systems (XNS) in certain configurations.

Section 10: Configure QoS features on multilayer switched networks to provide optimal quality and bandwidth utilization for applications and data (4 questions)

QUESTION NO: 1

The TestKing network is displayed in the following topology exhibit:



Workstation 1 traffic is set for cos 5 and the switch TestKingA sends workstation 1 traffic to the switch TestKingB. However, not all of the traffic from Switch TestKingA is from workstation 1.

Switch TestKingA configurations

```
mls qos
interface fa0/1
switchport mode trunk
switchport trunk encapsulation dot 1q
switchport trunk native vlan 1
```

Switch TestKingB Configuration:

```
mls qos
interface fa0/1
switchport trunk mode
switchport trunk encapsulation dot 1q
switchport trunk native vlan 1
```

Frames from Workstation 1 are given the rightful priority through Switch TestKingA, but Switch TestKingB doesn't reciprocate, and treats Workstation 1 frames as if they have no precedence. Which of the following actions will prioritize traffic from Workstation 1?

- Configure **qos all** command under Switch TestKingB fa0/1 interface.
- Configure **mls qos trust cos** command under Switch TestKingB fa0/1 interface.
- Configure **mls qos trust cos 5** command under Switch TestKingB fa0/1 interface.

- D. Configure **qos cos 5** command under Switch TestKingB fa0/1 interface.
- E. Configure **mls qos trust cos** command under Switch TestKingA fa0/1 interface.
- F. Configure **qos cos 5** command under Switch TestKingA fa0/1 interface.

Answer: B

Explanation:

The default action is for a switch with QoS features activated not to trust edge devices and any frames that enter the switch have their CoS re-written to the lowest priority of zero. If the edge device can be trusted, this default behaviour must be overridden and the access switch must be configured to switch the frame, leaving the CoS bits untouched.

The trust is configured on the switch port using the command:
switch(config-if)#mls qos trust cos

QUESTION NO: 2

You are a CCNP at TestKing, and your incompetent junior administrator has incorrectly changed the MTU (maximum transmission unit) settings of the IP packets sent on an interface. What command would you enter if you wanted to restore the MTU size to its default? (Type in answer below)

Answer: no ip mtu

Explanation:

To set the maximum transmission unit (MTU) size of IP packets sent on an interface, use the ip mtu interface configuration command. To restore the default MTU size, use the no form of this command.

ip mtu bytes
no ip mtu

QUESTION NO: 3

Switch TK1 is an IOS-based switch. Which command could an administrator use to establish a traffic policy on this IOS based switch?

- A. traffic-list
- B. route-map
- C. policy-map
- D. policy-list
- E. None of the above

Answer: C

Explanation:

The **policy-map** command is used to create a traffic policy. The purpose of a traffic policy is to configure the QoS features that should be associated with the traffic that has been classified in a user-specified traffic class or classes. A traffic policy contains three elements: a name, a traffic class (specified with the **class** command), and the QoS policies (which are detailed in the "[Configuring the Modular Quality of Service Command-Line Interface](#)" chapter of this book).

The name of a traffic policy is specified in the **policy-map** CLI (for example, issuing the **policy-map class1** command would create a traffic policy named class1).

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800bd908.html

QUESTION NO: 4

Which command would you enter if you had a Cisco 3550 switch, and you wanted to configure priority queuing on your gig0/1 interface?

- A. Under the global configuration, configure "priority-queue out"
- B. Under the global configuration, configure "interface priority-queue gig0/1"
- C. Under the interface gig0/1, configure "priority-queue out"
- D. Priority queuing is on by default

Answer: C

Explanation:

To configure priority queuing on an interface, use the "priority-queue" command in interface mode.

Incorrect Answers:

A, B: Assigning the queuing type is an interface configuration command and needs to be done in interface mode, not in the global configuration mode.

D: On interfaces faster than a T1, the default queuing method is first in, first out (FIFO).

Section 11: Configure Fast EtherChannel and Gigabit EtherChannel to increase bandwidth for interswitch connections (10 questions)

QUESTION NO: 1

Switch TK1 and TK2 are connected via an ethernet channel as shown below:



Switch TK1:

```
interface port channel 1
switchport
switchport access vlan 10
interface fastethernet 0/1
channel-group 1 mode passive
interface fastethernet 0/2
channel-group 1 mode passive
```

Switch TK2:

```
interface port-channel 1
switchport
switchport access vlan 10
interface fastethernet 0/1
channel-group 1 mode passive
interface fastethernet 0/2
channel-group 1 mode passive
```

In accordance with the above configuration, which of the statements below is correct?

- A. PAgP is correctly configured and the EtherChannel will form.
- B. LACP is correctly configured and the EtherChannel will form.
- C. One switch must be in LACP Active mode for the EtherChannel to form.
- D. Only one switch must be in the On mode and the other in the LACP Passive mode for Etherchannel to form.
- E. Each physical port in the EtherChannel must have the command **switchport access vlan 10** for the EtherChannel to form.

Answer: C

Explanation:

Link Aggregation Control Protocol (LACP) is part of an IEEE specification (802.3ad) that allows you to bundle several physical ports together to form a single logical channel. LACP allows a switch to negotiate an automatic bundle by sending LACP packets to the peer. It performs a similar function as Port Aggregation Protocol (PAgP) with Cisco EtherChannel.

To start automatic EtherChannel configuration with LACP, configure at least one end of the link to active mode to initiate channelling, because ports in passive mode passively respond to initiation and never initiate the sending of LACP packets.

QUESTION NO: 2

What does an EtherChannel port do if a VLAN range doesn't match its port list?

- A. The ports will form EtherChannel if they are set to auto mode.
- B. The ports will form EtherChannel if they are all set to the same trunk type.
- C. The ports will not form an EtherChannel.
- D. The ports will form an EtherChannel if the mode is set to on.

Answer: C

Explanation:

An EtherChannel supports the same allowed range of VLANs on all the ports in a trunking EtherChannel. If the allowed range of VLANs is not the same for a port list, the ports do not form an EtherChannel even when set to the **auto** or **desirable** mode with the **set port channel** command.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a00800eb765.html

QUESTION NO: 3

Which command would you enter on your Catalyst 2900XL switch if you wanted to enable an EtherChannel bundle?

- A. Port group
- B. Set port channel on
- C. Port etherchannel enable
- D. Set etherchannel port enable

Answer: A

Explanation

Under the interface command you have to indicate the syntax as port group (group number) for each interface you want to bundle in the etherchannel.

The **port group** command is used to enable an etherchannel bundle on a Catalyst 2900XL switch.

Incorrect Answers:

- B: The **set port channel** command is not used on a Catalyst 2900XL switch. It used on other switches, such as Cisco 5000 series however.
 C, D: These are invalid configuration commands.

QUESTION NO: 4

You have a Catalyst 5000 and you've just configured an Etherchannel bundle. If one of the links were to fail, how long will it take for traffic to be rerouted to a new link?

- A. one minute
- B. a few seconds
- C. a few milliseconds
- D. not until appropriate commands are entered

Answer: C

Explanation:

If a link is lost in an EtherChannel network, traffic is rerouted to one of the other links in just a few milliseconds. This rerouting is automatic, and the time it takes for traffic to get re-routed is normally not noticeable by end users.

QUESTION NO: 5

You are configuring a switching solution and you want to take advantage of the Fast EtherChannel ports. When configuring FastEthernet ports, which precautions can you take to avoid configuration problems which can cause the ports to be automatically disabled? (Select two)

- A. Allow some ports in a channel to be partly disabled.
- B. Configure ALL the ports in a channel as dynamic.
- C. Assign all ports in a channel to the same VLAN
- D. Allow some ports in a channel to be disabled.
- E. Allow all ports in a channel to be disabled.
- F. Configure all ports in a channel to operate at the same speed but in different duplex modes
- G. Assign all ports in a channel to the same VLAN or configure them as trunk ports.

Answer: C, G

Explanation:

Cisco's Fast EtherChannel technology builds upon standards based 802.3 full duplex Fast Ethernet to provide network managers a reliable high speed solution for the campus network

backbone. Fast EtherChannel provides bandwidth scalability within the campus by providing increments from 200 Mbps to 800 Mbps with multi-gigabit capacity in the future. Fast EtherChannel technology not only solves the immediate problem of scaling bandwidth within the network backbone today, but also paves the path for an evolution to standards-based Gigabit Ethernet and beyond, because Fast EtherChannel technology can be applied to support Gigabit EtherChannel.

In order for a channel to function properly, the aggregated links should be in the same VLAN or the links should be assigned as a trunk. In addition, all links should have identical speed and duplex settings.

QUESTION NO: 6

You have just configured an EtherChannel bundle on switch TK1 and it is now operational on a trunk. Which of the following could cause the disabling of the ports in this bundle?

(Select two)

- A. Disabling port security
- B. Excessive errors on one port
- C. Changing the VLAN mode to dynamic
- D. Changing the speed attribute of one port in the bundle.

Answer: C, D

Explanation:

C: Do not configure the ports in an EtherChannel as dynamic VLAN ports. It could adversely affect switch performance.

D: All ports in an EtherChannel should be configured to operate at the same speed and duplex mode (full or half duplex).

Reference: Cisco, Configuring Fast EtherChannel and Gigabit EtherChannel

QUESTION NO: 7

Switch TK1 is a Catalyst 5000 switch. Which of the following set commands would you use to enable Fast EtherChannel on this switch?

- A. "set channel fast"
- B. "set port channel"
- C. "set link channel"
- D. "set etherchannel"
- E. None of the above

Answer: B

Explanation:

In order to configure ports on a switch to belong to an etherchannel bundle, use the “set port channel” configuration command. For example: **set port channel 3/1-2** will configure ports 3/1 and 3/2 to belong to the channel.

Reference: Cisco Application Note, Understanding and Designing Networks Using Fast EtherChannel Technology

QUESTION NO: 8

Cisco switches use a logical operation to determine which links to send EtherChannel traffic. What kind of logical operation is it?

- A. OR
- B. AND
- C. XOR
- D. NAND

Answer: C

Explanation:

EtherChannel performs the XOR operation, which works like this:

A	B	C
0	XOR	0 -> 0
0	XOR	1 -> 1
1	XOR	0 -> 1
1	XOR	1 -> 0

QUESTION NO: 9

Which of the following commands would you enter if you wanted to find out whether or not switch TK1 is capable of supporting EtherChannel?

- A. show trunk
- B. show interface
- C. show port channel
- D. show port capabilities

Answer: D

Explanation:

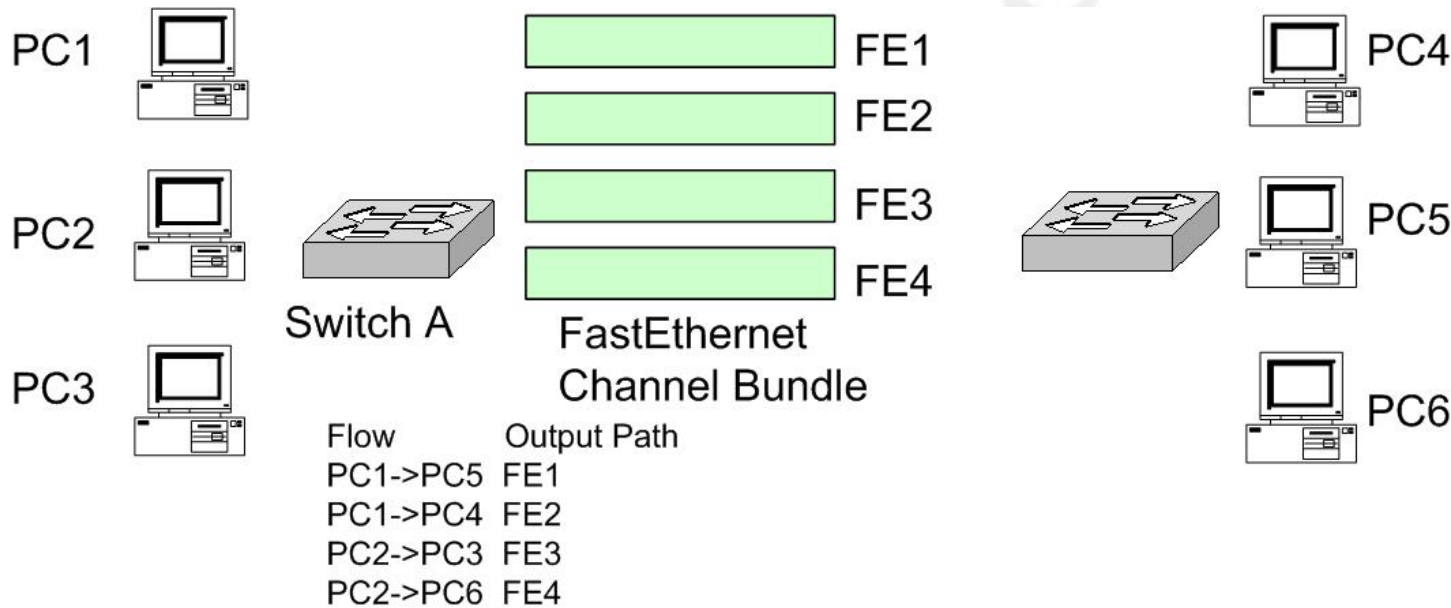
The **show port capabilities** command will show you the capabilities of the modules and ports in a switch. For example, it will display the type, speed, and duplex.

Incorrect Answers:

- A: Trunking information does not include the required information.
 B: Interface configuration information does not include the required information.
 C: Use the **show port channel** command to display EtherChannel information for a specific module or port. It will not show the capabilities of the switch of the switch however.

QUESTION NO: 10

Switch A and Switch B are configured for FastEthernet Channel as shown in the diagram below:



Assuming that Fast EtherChannel was set properly set up; if FE1 were to fail, what would happen with the traffic flow between PC1 & PC5?

- A. Traffic is transferred to FE2.
 B. Traffic is transferred to FE4.
 C. PC1 to PC4 traffic is distributed over the remaining links.
 D. The session is disconnected while spanning tree rebuilds.

Answer: C

Explanation:

If a port within an EtherChannel fails, traffic previously carried over the failed port switches to the remaining ports within the EtherChannel.

Note: Fast/Gigabit EtherChannel allow high-speed redundant links in a spanning tree by allowing dual parallel links to be treated as though they were one link. If a link is lost in a

Fast/Gigabit EtherChannel network, traffic rerouted to one of the other links in just a few milliseconds.

Reference: Configuring Fast EtherChannel and Gigabit EtherChannel

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/rel7_1/config/channel.htm

Topic 3: DESIGN (28 questions)

Section 1: Compare end-to-end and local VLANs, determine when to use each (5 questions)

QUESTION NO: 1

Under what circumstances should an administrator prefer local VLANs over end-to-end VLANs?

- A. Eighty percent of traffic on the network is destined for Internet sites.
- B. There are common sets of traffic filtering requirements for workgroups located in multiple buildings.
- C. Eighty percent of a workgroup's traffic is to the workgroup's own local server.
- D. Users are grouped into VLANs independent of physical location.
- E. None of the above

Answer: A

Explanation:

This geographic location can be as large as an entire building or as small as a single switch inside a wiring closet. In a geographic VLAN structure, it is typical to find 80 percent of the traffic remote to the user (server farms and so on) and 20 percent of the traffic local to the user (local server, printers, and so on).

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 93

QUESTION NO: 2

What are some virtues of implementing end-to-end VLANs? (Choose two)

- A. End-to-end VLANs are easy to manage.
- B. Users are grouped into VLANs independent of a physical location.
- C. Each VLAN has a common set of security and resource requirements for all members.
- D. Resources are restricted to a single location.

Answer: B, C

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

In an end-to-end VLAN, users are grouped into VLANs independent of physical location and dependent on group or job function.

Each VLAN has a common set of security requirements for all members.

Incorrect Answers:

A: End to end VLANs are more difficult to manage than local VLANs, due to the physical distances that they can span.

D: In an end-to-end VLAN, network resources are generally distributed across the entire enterprise wide area network.

QUESTION NO: 3

Which of the following statements is true about the 80/20 rule (Select all that apply)?

- A. 20 percent of the traffic on a network segment should be local
- B. no more than 20 percent of the network traffic should be able to move across a backbone.
- C. no more than 80 percent of the network traffic should be able to move across a backbone.
- D. 80 percent of the traffic on a network segment should be local

Answer: B, D

Explanation:

The 80/20 rule in network design originated from Pareto's Principle. The Italian economist Vilfredo Pareto came up with the discovery that 20% of the people controlled 80% of the wealth and applied the principle of how inputs don't match outputs in real life. Keeping this number in mind, 80% of network traffic should be local to a segment and 20% should move across the backbone.

Note: With the availability of inexpensive bandwidth and centralized data centers, this rule appears to have become obsolete. In fact, most networks have taken on the 20/80 rules, as opposed to the legacy 80/20 rule.

QUESTION NO: 4

Which two factors give merit to the 20/80 LAN traffic model? (Select two)

- A. The Internet
- B. Local servers
- C. Server farms
- D. Localized applications
- E. More powerful desktop PC's

Answer: A, C

Explanation:

Remote services (server farms, Internet, etc.) are factors which contributed to increased backbone traffic.

Also consider:

This geographic location can be as large as an entire building or as small as a single switch inside a wiring closet. In a geographic VLAN structure, it is typical to find 80 percent of the traffic remote to the user (server farms and so on) and 20 percent of the traffic local to the user (local server, printers, and so on).

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 93

QUESTION NO: 5

Which feature assigns switch ports to VLANs dynamically based on the source MAC address of the device connected to the port?

- A. Dynamic VLAN Protocol
- B. Dynamic Trunking Protocol
- C. VLAN Database
- D. VLAN Management Policy Server

Answer: D

Explanation

With VMPS (VLAN Management Policy Server), you can assign switch ports to VLANs dynamically, based on the source Media Access Control (MAC) address of the device connected to the port. When you move a host from a port on one switch in the network to a port on another switch in the network, the switch assigns the new port to the proper VLAN for that host dynamically.

When you enable VMPS, a MAC address-to-VLAN mapping database downloads from a Trivial File Transfer Protocol (TFTP) server and VMPS begins to accept client requests. If you reset or power cycle the switch, the VMPS database downloads from the TFTP server automatically and VMPS is re-enabled.

VMPS opens a User Datagram Protocol (UDP) socket to communicate and listen to client requests. When the VMPS server receives a valid request from a client, it searches its database for a MAC address-to-VLAN mapping.

Section 2: Design a VLAN configuration with VTP to work for a given specific scenario (4 questions)

QUESTION NO: 1

What method could you use to eliminate unicast and broadcast traffic that is flooded to a VLAN unnecessarily?

- A. VTP pruning
- B. MLS-SE
- C. VTP trunking
- D. VTP compression
- E. All of the above

Answer: A

Explanation:

VTP ensures that all switches in the VTP domain are aware of all VLANs. There are occasions, however, when VTP can create unnecessary traffic. All unknown unicasts and broadcasts in a VLAN are flooded all over the VLAN. All switches in the network receive all broadcasts, even in situations where few users are connected in that VLAN. VTP pruning is a feature used to eliminate (prune) this unnecessary traffic.

QUESTION NO: 2

You are a network engineer and you've been assigned the task of configuring an Ethernet media trunk between two Cisco switches. Assuming that these two switches have the same modules, software revisions and VLAN configurations, which of the following are NOT true for the trunk to operate? (Choose all that apply)

- A. The link must be a point to point for ISL encapsulation.
- B. The link must be 100Mbps or slower
- C. The link must use the IEEE 802.1Q trunk protocol.
- D. The link may use the IEEE 802.1Q trunk protocol.

Answer: B, C

Explanation:

Trunks can operate at 10, 100, or 1000 MB interfaces.

The industry standard method of trunking is 802.1Q. As an alternative to this, the Cisco proprietary ISL method is also an option for setting up trunks.

Incorrect Answers:

A: This statement is true. ISL trunks must be configured on point to point links; point-to-multipoint configurations are not supported.

D: This statement is also true. 802.1Q can be used, but it does not have to be.

QUESTION NO: 3

Is the following statement True or False?

MLSP can cross a VTP domain boundary.

- A. False
- B. True
- C. There is not enough information to determine

Answer: A

Explanation:

MLS requires that MLS components, including the end stations, must be in the same Virtual Trunking Protocol (VTP) domain. VTP is a Layer 2 protocol used for managing VLANs on several Catalyst switches from a central switch. It allows an administrator to create or delete a VLAN on all switches in a domain without having to do so on every switch in that domain. The MultiLayer Switching Protocol (MLSP), which the MLS-SE and the MLS-RP use to communicate with one another, does not cross a VTP domain boundary.

QUESTION NO: 4

The TestKing LAN is becoming saturated with broadcasts and multicast traffic. What could you do to help a network with many multicasts and broadcasts?

- A. Creating smaller broadcast domains by implementing VLANs.
- B. Separate nodes into different hubs.
- C. Creating larger broadcast domains by implementing VLANs.
- D. Separate nodes into different switches.
- E. All of the above.

Answer: A

Explanation:

Controlling broadcast propagation throughout the network is important to reduce the amount of overhead associated with these frames. Routers, which operate at Layer 3 of the OSI model, provide broadcast domain segmentation for each interface. Switches can also provide broadcast domain segmentation using virtual LANs (VLANs). A VLAN is a group of switch ports, within a single or multiple switches, that is defined by the switch hardware and/or software as a single broadcast domain. A VLAN's goal is to group devices connected to a switch into logical broadcast domains to control the effect that broadcasts have on other connected devices. A VLAN can be characterized as a logical network.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 8

Section 3: Plan QoS implementation within a multilayer switched network (3 questions)

QUESTION NO: 1

FIFO (First-In-First-Out) queue is associated with which QoS model?

- A. Less than Best Effort Model
- B. Best Effort Model
- C. Differentiated Services Model (DiffServ)
- D. Integrated Services Model (IntServ)
- E. None of the above

Answer: B

Explanation:

There are three QoS Models namely, Integrated Services Model, Best effort Model, and Differentiated Services model. Best effort is a single service model in which an application sends data whenever it must, in any quantity, without requesting permission or first informing the network. For best-effort service, the network delivers data if it can, without any assurance of reliability, delay bounds, or throughput. The Cisco IOS QoS feature that implements best-effort service is FIFO queuing.

QUESTION NO: 2

You are an instructor at TestKing and one of your students has a few questions about QoS. After class she comes into your office, and asks, “If I wanted to configure congestion avoidance, what mechanism randomly drops packets with a certain IP precedence value when the buffers fill to a predefined threshold?”

Which mechanism is she asking about?

- A. WFQ
- B. CQ
- C. LLQ
- D. WRED
- E. tail drop

Answer: D

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

WRED drops packets using IP precedence or DSCP value of the packets; packets with higher precedence are less likely to be dropped. If the default settings are preventing QoS, the precedence value can be used to control how WRED determines when and how often to drop packets.

QUESTION NO: 3

WRED is being configured on different network devices throughout the TestKing network. What is true about WRED?

- A. Cisco Express Forwarding cannot be enabled with WRED.
- B. WRED is useful when the bulk of traffic is not TCP/IP.
- C. The rate of packet-drop decreases as the average queue size increases, until the average queue size reaches the maximum threshold.
- D. Packets not flagged by IP precedence are randomly dropped when the average queue depth is above the minimum threshold.
- E. WRED defines packet-drop probability using an access list.
- F. All of the above.

Answer: D

Explanation:

Weighted Random Early Detection (WRED) generally drops packets selectively based on IP precedence. Packets with a higher IP precedence are less likely to be dropped than packets with a lower precedence. This mechanism is used to avoid network congestion and give priority to certain traffic types.

Section 4: Select multilayer switching architectures, given specific multilayer switching needs (15 questions)

QUESTION NO: 1

You are a CCNP working for a large online studio in Amsterdam. Your network team has redesigned its 'campus style' network in a way that it support three switch blocks, and these switch blocks act as broadcast domains to each individual switch block, while still allowing inter-VLAN routing within and between switch blocks. Which of the following switches is the best distribution layer device for your needs?

- A. Catalyst 2900 series switch
- B. Catalyst 3000 series switch
- C. Catalyst 1900
- D. Catalyst 4000 series switch
- E. Catalyst 5000 series switch with 1 RSM
- F. Catalyst 8500 series switch

Answer: E

Explanation:

Modules for the Catalyst 5000 family chassis --- Catalyst 5500, 5509, 5505, 5000, and 5002--- are designed for complete interoperability and investment protection. New functionalities in the Catalyst 5000 family support multiprotocol NetFlow Switching for scalable convergence of Layer 2 and Layer 3 switching, adding the benefits of multiprotocol, multilayer switching and other Cisco IOS network services. The range of media support in the Catalyst 5000 family enables network managers to deliver high-performance backbone access to accommodate Web browser-based traffic across the intranet. A growing number of interface modules operate in any Catalyst 5000 family switch to deliver dedicated bandwidth to users through high-density group switched 10BaseT or 100BaseT Ethernet; flexible 10/100BASE-T Ethernet, fiber-based Fast Ethernet, and Fast EtherChannel; Token Ring; CDDI/FDDI; ATM LAN Emulation (LANE) and Multiprotocol over ATM (MPOA); the Route/Switch module (based on the Route/Switch Processor for the Cisco 7500 series router); and Gigabit Ethernet. Unique to the Catalyst 5500 Series are the ATM Switch Processor and ATM switch interface modules and port adapters.

QUESTION NO: 2

What does a Catalyst 5000 switch need installed in order for MLS to run on it? (Select all that apply)

- A. IOS V13
- B. IOS V12.11
- C. Netflow Feature Card
- D. The MLS-SE RP Card

Answer: C

Explanation:

The MLS-SE is a switch with special hardware. For a member of the Catalyst 5000 family, MLS requires that the supervisor have a Netflow Feature Card (NFFC) installed. The Supervisor IIG and IIIG have one by default. In addition, a bare minimum of Catalyst OS 4.1.1 software is also required. Note that the 4.x train is now in General Deployment (GD), or passed rigorous end-user criteria and field-experience targets for stability, so check Cisco's website for the latest releases. IP MLS is supported and automatically enabled for Catalyst 6000 hardware and software with the MSFC/PFC (other routers have MLS disabled by default). Note that IPX MLS and MLS for multicasting may have different hardware and software (Cisco IOS and Catalyst OS) requirements. More Cisco platforms do/will support the MLS feature. Also, MLS must be enabled in order for a switch to be an MLS-SE.

QUESTION NO: 3

A technical institute has redesigned its campus network to support three switch blocks of 2,000 users per block. The network administrator wants to control broadcast domains to each individual switch block, while still allowing interVLAN routing within and between switch blocks. What distribution layer device should be used to accomplish this?

- A. Catalyst 4000 series switch
- B. Catalyst 8500 series switch
- C. Catalyst 6000 series switch with an internal MSFC
- D. Catalyst 1900 with a two-port 100Base uplink module.

Answer: C

Explanation:

A Catalyst 6000 series switch with an internal MSFC (Multilayer Switch Feature Card) is the most appropriate solution. Catalyst 6000 series switches use the MSFC and the Policy Feature Card (PFC) to gather and cache header information.

Note:

Catalyst 6000 Family switches equipped with MSFCs provide transparent Web Cache redirection using Cisco's Web Cache Communication Protocol v2 (WCCP). WCCP is the industry's leading web-cache redirection protocol that localizes network traffic and provides network- intelligent load distribution.

QUESTION NO: 4

Your CTO has authorized you to purchase 10 new switches that can support the VLAN Management Policy Server (VMPS) feature. Which three of the following switches can you consider? (Select three)

- A. 2900XL series
- B. 3500XL series
- C. 5000 series
- D. 8500 series

Answer: A, B, C

Explanation:

A, B: Catalyst 2900 and 3500 Series XL Features include VMPS.

C: VMPS functionality is present on all Catalyst 5000 Family switches.

Note: The VLAN Management Policy Server (VMPS) service is used to set up a database of MAC addresses that can be used for dynamic addressing of VLANs. VMPS is a MAC-address-to-VLAN mapping database.

QUESTION NO: 5

You are a salesman at a Cisco network equipment shop, when one of your regular customers, a CTO of a college comes in. He wants to launch a multimedia center to distribute her program information throughout the campus.

- **High availability**
- **Gigabit data transfer speed**
- **Access aggregation**
- **InterVLAN routing between users and Server Farms**

What should you recommend?

- A. Catalyst 2948G-L3
- B. Catalyst 4000 series switch
- C. Catalyst 6000 series switch
- D. Catalyst 7100 series switch

Answer: C

Explanation:

The Catalyst 6000 switches provide high-density Fast Ethernet and Gigabit Ethernet in both campus-backbone and server-aggregation environments. The Catalyst 6006 and the Catalyst

6009 switches have a 32-Gbps switching capacity, while the Catalyst 6506, the Catalyst 6509, the Catalyst 6509-NEB, and the Catalyst 6513 switches support a backplane architecture that scales from 32 Gbps to 256 Gbps.

Incorrect Answers:

- A: The Catalyst 2948G-L3 provides an aggregate throughput of 10 Mbps for Layer 3 switching. However, the requirement is of gigabit speed.
- B: Not adequate.
- D: Catalyst 7100 Series are VPN routers. Packet throughput is 175 Kpps. It is adequate for large branch and central site VPN router, for a dedicated site-to-site VPN solution.

QUESTION NO: 6

Which of the following switches utilize Cisco Express Forwarding (CEF)? (Select two)

- A. Catalyst 8500
- B. Catalyst 2900XL
- C. Catalyst 3500XL
- D. Catalyst 2948G-L3

Answer: A, D

Explanation:

A: Catalyst 8500 switches use a forwarding information base (FIB) for Cisco Express Forwarding.

D: The Catalyst 2948G-L3 uses Cisco Express Forwarding (CEF).

Note: Cisco Express Forwarding (CEF) is advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with large and dynamic traffic patterns, such as the Internet, on networks characterized by intensive Web-based applications, or interactive sessions.

CEF was originally developed for the Cisco 12000 series gigabit switch router (GSR), the Catalyst 8500, and the Cisco 7500.

Incorrect Answers

B, C: These switches do not support CEF.

Reference: Cisco, Cisco Express Forwarding

QUESTION NO: 7

You are a salesman at an authorized Cisco dealer. A customer comes into your store asking for a switch that supports redundant supervisor engines. Which models would you show him? (Select all that apply)

- A. Catalyst 4000
- B. Catalyst 5500

- C. Catalyst 6000
- D. Catalyst 12000

Answer: B, C

Explanation:

- B: The Catalyst 5500 series switches support an optional redundant supervisor engine.
- C: Catalyst 6000 family switches support fault resistance by allowing a redundant supervisor engine to take over if the primary supervisor engine fails.

QUESTION NO: 8

Which of the following switches have an IOS-based user interface?

- A. Catalyst 2924XL
- B. Catalyst 2926
- C. Catalyst 4003
- D. Catalyst 5500

Answer: A

Explanation:

Catalyst 2924XL has an IOS-based user interface.

Note: Switches can either be IOS-based or set-based.

With IOS based user interface you can configure the switch from a CLI (command line interface) that is very similar to the ones used on Cisco routers. Catalyst 1900, 2820, and 2900 switches can be used with an IOS-based CLI.

Set-based user interface uses older, set-based CLI configuration commands. The Cisco switches that use the set-based CLI are the 2926 series, the 1945G, the 4000, the 5000, and the 6000 series.

QUESTION NO: 9

You are a salesperson at an authorized Cisco dealership. A customer comes in wanting a switch, but reveals to you that he's IOS illiterate and only knows how to use the set-based CLI commands? Which models will you show him? (Select three)

- A. 2900XL
- B. 2948G
- C. 3500XL
- D. 4000
- E. 6500

Answer: B, D, E

Explanation:

A 2948G, 4000, 5000, and 6500 series switch uses set based CLI commands.

Incorrect Answers:

2900XL and 3500XL switches do not use set based CLI commands. These switches are IOS based, and are configured similarly to a Cisco router.

QUESTION NO: 10

Your business venture is rapidly expanding and you need to provide access for numerous users and need a high density switch at the lowest possible price. Which of the following switches has the lowest price per port?

- A. Catalyst 1900 series
- B. Catalyst 3500XL series
- C. Catalyst 5000 series
- D. Catalyst 85000 series

Answer: C

Explanation:

A Catalyst 5000 series switch would be the optimal solution in this scenario. Although a 5000 series is not the cheapest switch that Cisco offers, it does have the lowest cost per port due to its high port density.

QUESTION NO: 11

A switched Ethernet costs considerably more to set up than a shared Ethernet using hubs. What justifies the added expense? (Select two.)

- A. It provides greater scalability
- B. It is less complex to manage
- C. It permits full-duplex operation
- D. It simplifies routing between LAN segments

Answer: A, C

Explanation:

A: A switched Ethernet has more collision domains, and each collision domain is smaller compared to Shared Ethernet. A switched Ethernet can therefore support more nodes compared to a shared Ethernet.

C: A switched connection may support full-duplex operation. Shared Ethernet is only half duplex.

Incorrect Answers:

B: A shared Ethernet requires less administration compared to a switched Ethernet.

D: Switches operates at Layer 2, the data link layer. Routing is performed at Layer 3, the network layer.

QUESTION NO: 12

The TestKing LAN is using VLANs, and traffic between them needs to be routed. What kind of network equipment do you need for interVLAN routing?

- A. ISL
- B. switch block
- C. IEEE 802.1Q
- D. access switch
- E. route processor

Answer: E

Explanation:

The traffic has to be routed between the VLANs. A routing mechanism has to be provided for interVLAN communication. This can either be accomplished using a router or a route processor within the switch.

Incorrect Answers:

A, C: Either IEEE 802.1Q or ISL trunking protocol can be used, but these are trunking protocols, not network hardware.

B, D: Neither of these is used for interVLAN routing.

QUESTION NO: 13

You're a salesperson at a network equipment wholesaler. A customer comes in and tells you that he needs a distribution layer switch for a large campus style network, with a high amount of Gigabit Ethernet port density. Which device series would you recommend to him?

- A. 4908G-L3
- B. 5000 series
- C. 6000 series
- D. 8500 series

Answer: C

Explanation:

The Catalyst 6000 can provide up to 384 10/100 Ethernet connections, 192 100FX FastEthernet connections, and 130 Gigabit Ethernet ports.

Incorrect Answers:

A, B: A 4908G-L3 switch or a Catalyst 5000 series switch are not capable of the required performance.

D: The Cisco Catalyst 8500 is a core layer switch that provides high-performance switching. This particular customer is looking for a distribution layer switch, not a core switch.

QUESTION NO: 14

Which of the following mechanisms combines the speed of switching with the scalability of routing?

- A. Layer 3 switching
- B. Fast switching
- C. Layer 2 routing
- D. Process routing
- E. All of the above

Answer: A

Explanation:

Layer 3 switching is hardware-based routing. In particular, packet forwarding is handled by specialized hardware ASICs. A layer 3 switch does everything to a packet that a traditional router does.

Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 18

QUESTION NO: 15

Which three of the following network features are methods used to achieve high availability? (Select all that apply.)

- A. Spanning Tree Protocol (STP)
- B. Delay reduction
- C. Hot Standby Routing Protocol (HSRP)
- D. Dynamic routing protocols
- E. Quality of Service (QoS)
- F. Jitter management

Answer: A, C, D

Explanation:

Because the importance of high availability networks is increasingly being recognized, many organizations are beginning to make reliability/availability features a key selection criteria for network infrastructure products. With this in mind, Cisco Systems engaged ZD Tag to observe and confirm the results of a series of tests demonstrating the high availability features of Cisco Catalyst Layer 2/Layer 3 switches. In order to maximize the relevance of the results, the demonstration was based on a model of a “real world” campus (in one of Cisco’s Enterprise Solution Center labs in San Jose, California).

This switched internetwork consisted of wiring closet, wiring center, and backbone switches and conformed to Cisco’s modular three-tier (Access/Distribution/Core) design philosophy. The testing demonstrated the following high availability and resilience features of Catalyst switches:

- **per-VLAN Spanning Tree (PVST) using Cisco’s InterSwitch Link (ISL) and 802.1Q VLAN Trunking**
- **Cisco Spanning Tree Enhancements, including UplinkFast and PortFast**
- **Cisco Hot Standby Router Protocol (HSRP) and HSRP Track**
- **Cisco IOS per-destination load balancing over equal cost OSPF paths**
- **Cisco IOS fast convergence for OSPF**

Reference: <http://www.cisco.com/warp/public/779/largeent/learn/technologies/campuslan.pdf>

Section 5: Describe the general design models when implementing IP telephony in a switched network environment (2 questions)

QUESTION NO: 1

You are a network administrator of a large investor relations company that uses a switched network to carry both data and IP telephony services. Why should you carry voice traffic on a separate VLAN?

- A. IP phones require inline power and must be in separate VLAN to receive inline power.
- B. IP telephony applications require prioritization over other traffic as they are more delay sensitive.
- C. IP phones can only receive IP addresses through DHCP if they are in separate VLAN.
- D. The CDP frames from the IP phone can only be recognized by the switch if the phone is in an auxiliary vlan.

Answer: B

Explanation:

Voice conversations don't take up a lot of bandwidth, but the bandwidth they do is very delicate. If anything happens with the connection or the integrity of the data transfer in either direction the conversation won't seem natural. To ensure the highest degree of integrity you should put voice traffic on its own separate VLAN and give that VLAN the highest priority.

QUESTION NO: 2

Which of the following switches have inline power support for Cisco IP telephony? (Select three)

- A. 3500 series
- B. 4000 series
- C. 5000 series
- D. 6000 series

Answer: A, B, D

Explanation:

A: With the expansion of inline power needs for IP phones and wireless access points, the Catalyst 3524-PWR XL is the leading choice.

- B: The Cisco Catalyst 4000 Family Inline Power 10/100BaseT Ethernet Switching Module intelligently detects and provides power to IP enabled devices such as Cisco IP Phones.
- D: The Cisco Catalyst 6000 Family Inline Power 10/100BaseT Ethernet Switching Module extends the voice capabilities of the Catalyst backbone to the enterprise wiring closet and branch office.

Note 1: Each Cisco IP Telephone provides Toll-quality audio and doesn't require a companion PC. Because it is an IP-based telephone, it can be installed anywhere on a corporate local or wide area IP network.

Note 2: Inline power is 48-volt DC power provided over standard Category 5 unshielded twisted-pair (UTP) cable up to 100 meters. Instead of requiring wall power, terminal devices such as IP telephones can utilize power provided from the Catalyst Inline Power Patch Panel.

Incorrect Answers:

- C: Catalyst 5000 series switches are used in large campuses to provide access for more than 250 users. They support 10/100/1000Mbps Ethernet switching. They don't support inline power for Cisco IP phones, however.

TOPIC 4: TROUBLESHOOTING (36 questions)

Section 1: Verify and Troubleshoot common VLAN problems on a switched network (7 questions)

QUESTION NO: 1

You're a network administrator and you issue the command (`show port 3/1`) on an Ethernet port. To your surprise you notice a non-zero entry in the 'Giants' column. What could be the cause of this?

- A. IEEE 802.1Q
- B. IEEE 802.10
- C. Misconfigured NIC
- D. User configuration
- E. All of the above

Answer: A

The 802.1Q standard can create an interesting scenario on the network. Recalling that the maximum size for an Ethernet frame as specified by IEEE 802.3 is 1518 bytes, this means that if a maximum-sized Ethernet frame gets tagged, the frame size will be 1522 bytes, a number that violates the IEEE 802.3 standard. To resolve this issue, the 802.3 committee created a subgroup called 802.3ac to extend the maximum Ethernet size to 1522 bytes.

Note: The `show port` command is used to display port status and counters. Giants denote the number of received giant frames (frames that exceed the maximum IEEE 802.3 frame size) on the port.

Reference: Trunking between Catalyst 4000, 5000, and 6000 Family Switches Using 802.1q Encapsulation

<http://www.cisco.com/warp/public/473/27.html>

QUESTION NO: 2

You have a trunk link operating between two switches and you're experiencing problems with frames leaking between the two VLANs. Each switch has identical modules, software revisions and VLAN configuration information. Spanning tree protocol is disabled on all VLANs. What is probably causing this problem? (Select all that apply)?

- A. The link is using IEEE 802.1Q protocol
- B. The link is using IEEE 802.1E protocol
- C. Spanning tree is disabled
- D. Not enough information to determine.
- E. The native VLAN information is identical at each end of the link.
- F. The native VLAN information is different at each end of the link.

Answer: A, F

Explanation:

While internal to a switch, VLAN numbers and identification are carried in a special extended format that allows the forwarding path to maintain VLAN isolation from end to end without any loss of information. Instead, outside of a switch, the tagging rules are dictated by standards such as ISL or 802.1Q.

ISL is a Cisco proprietary technology and is in a sense a compact form of the extended packet header used inside the device: since every packet always gets a tag, there is no risk of identity loss and therefore of security weaknesses.

On the other hand, the IEEE committee that defined 802.1Q decided that because of backward compatibility it was desirable to support the so-called native VLAN, that is to say, a VLAN that is not associated explicitly to any tag on an 802.1Q link. This VLAN is implicitly used for all the untagged traffic received on an 802.1Q capable port.

This capability is desirable because it allows 802.1Q capable ports to talk to old 802.3 ports directly by sending and receiving untagged traffic. However, in all other cases, it may be very detrimental because packets associated with the native VLAN lose their tags, for example, their identity enforcement, as well as their Class of Service (802.1p bits) when transmitted over an 802.1Q link.

For these sole reasons—loss of means of identification and loss of classification—the use of the native VLAN should be avoided.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_white_paper09186a008013159f.shtml

QUESTION NO: 3

What command could you enter to display the trunking status of a module/port in the switch? (Type in the answer below):

Answer: show trunk

Explanation:

Use the show trunk command to display trunking information for the switch.

show trunk [mod_num[/port_num]] [detail]mod_num (Optional) Number of the module.

/port_num (Optional) Number of the port.

detail (Optional) Keyword to show detailed information about the specified trunk port.

QUESTION NO: 4

You are troubleshooting a Catalyst 5000 trunk in the TestKing network. What should you do if there's a disagreement about the VLANs configured to use the trunk?

- A. Reload the active VLAN configuration
- B. Clear the affected port and bring it up again.
- C. Explicitly set the trunk for the VLAN to be on.
- D. Remove all the VLANs set

Answer: B

Explanation:

In this situation you may want to set or clear the VLANs on both ends. A trunk is a point-to-point link between one or more Ethernet switch interfaces and another networking device such as a router or a switch. Trunks carry the traffic of multiple VLANs over a single link and allow you to extend VLANs across an entire network. Two trunking encapsulations are available on all Ethernet interfaces:

Inter-Switch Link (ISL)-ISL is a Cisco-proprietary trunking encapsulation

802.1Q-802.1Q is an industry-standard trunking encapsulation

When a trunk is first brought up using either of these methods, it may be beneficial to clear the port immediately after.

QUESTION NO: 5

You want to check your Catalyst switch port to see if there's an active link state. Which of the following actions would NOT be useful? (Select all that apply)

- A. Switch fan
- B. Port's link LED of the Switching Module.
- C. Switch RP
- D. Switch slot
- E. Switch backplane

Answer: A, C, D, E

Explanation:

To find out if there is an active link state on a Catalyst port, check the port's link LED. As an alternative, using various show commands will also aid in troubleshooting port problems.

However, checking the fan, RP, slot, or backplane will not help when looking for the status of an individual port's link status.

QUESTION NO: 6

Which of the following commands could you use to clear a switch of its current configuration?

- A. the "clear config all" command
- B. the "del config all" command
- C. the "erase config all" command
- D. the "clean config all" command
- E. None of the above

Answer: A

Explanation:

In a Cisco switch, the 'clear config all' command will purge the existing configuration on a switch and start you off with a brand new default configuration. All the other commands: del, erase, and clean; aren't valid IOS commands.

QUESTION NO: 7

Which kind of management can be performed from the console port of a Cisco 6500 switch?

- A. Physical management of the switch.
- B. Logical management of the switch.
- C. In-band management of the switch.
- D. Out-of-band management of the switch.

Answer: D

Explanation:

When you configure a switch or a router from the console, it is considered 'out of band' because you don't get in there from any of the paths that the network device is a part of. Modems are often attached to the console port, providing for remote out of band management of the device.

Section 2: Tune and troubleshoot spanning-tree protocol on a multilayer switched network to enhance network performance, prevent network loops, and minimize downtime (16 questions)

QUESTION NO: 1

What should you do to reduce spanning-tree protocol BPDU traffic during extended periods of instability in your VLANs?

- A. Combine all the VLAN spanning trees into a single spanning tree.
- B. Set forward delay and max-age timers to the maximum possible values.
- C. None of the choices.
- D. Change the router VTP server mode.
- E. Disable the root bridge

Answer: B

Explanation:

There are several STP timers, as listed below:

- **hello:** the hello time is the time between each Bridge Protocol Data Unit (BPDU) that is sent on a port. This is equal to two seconds by default, but can be tuned to be between one and ten seconds.
- **forward delay:** the forward delay is the time spent in the listening and learning state. This is by default equal to 15 seconds, but can be tuned to be between four and 30 seconds.
- **max age:** the max age timer controls the maximum length of time a bridge port saves its configuration BPDU information. This is 20 seconds by default and can be tuned to be between six and 40 seconds.

The STP timers (hello, forward delay, and max age) are included in each BPDU. An IEEE bridge is not concerned about its local configuration of the timers value. It will consider the value of the timers contained in the BPDU that it is receiving. Effectively, that means only a timer configured on the root bridge of the STP is important. Obviously, in case you would lose the root, the new root would start to impose its local timer value to the entire network. So, even if it is not required to configure the same timer value in the entire network, it is at least mandatory to configure any timer changes on the root bridge and on the backup root bridge.

In order to reduce the number of BPDU's in the spanning tree topology, the forward delay and max-age timers should be increased. This will reduce the BPDU traffic, but it will also increase the convergence time during a topology change.

QUESTION NO: 2

You are network consultant troubleshooting a problem at TestKing Inc. The local technician tells you that users can't access the Domain Controllers or DHCP servers from their workstations. To top it off, they aren't seeing their Novel Login Screen and they can't access their AppleTalk network. The customer use Cisco 4000, Cisco 5000, and Cisco 6000 switches. What command could you use to resolve these problems?

- A. spanning-tree portfast
- B. set port connect mod/port
- C. spantree start-forwarding
- D. set spantree portfast mod/port enable

Answer: D

Explanation:

When the switch powers up, or when a device is connected to a port, the port normally enters the spanning tree listening state. When the forward delay timer expires, the port enters the learning state. When the forward delay timer expires a second time, the port is transitioned to the forwarding or blocking state. This delay could cause the problems described in the scenario. We remove the delay with the PortFast feature. We enable PortFast on a switch port connected to a single workstation or server with the **set spantree portfast mod_num/port_num enable** command.

Note: The spanning tree PortFast feature causes a port to enter the spanning tree 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, instead of waiting for the port to transition from the listening and learning states to the forwarding state.

Reference: Cisco, Configuring Spanning Tree PortFast, UplinkFast, and BackboneFast

QUESTION NO: 3

What command should you enter if do you want to find out whether or not the Backbone Fast convergence feature of STP is enabled on switch TK1? (Type in answer below):

Answer: show spantree backbonefast

Explanation:

The following list various commands to use for troubleshooting Catalyst switches:


show spantree vlan_id - Shows the current state of the spanning tree for the "vlan_id" entered from the perspective of the switch on which it is entered.

show spantree summary - Provides a summary of connected spanning tree ports by VLAN.

show spantree statistics - Shows spanning tree statistical information.
 show spantree backbonefast - Displays whether the spanning tree Backbone Fast Convergence feature is enabled.
 show spantree blockedports - Displays only the blocked ports.
 show spantree portstate - Determines the current spanning tree state of a Token Ring port within a spanning tree.
 show spantree portvlancost - Shows the path cost for the VLANs on a port.
 show spantree uplinkfast - Shows the uplinkfast settings.

QUESTION NO: 4

The TestKing network is displayed in the diagram below:



You use the following information for switch TestKingA:

Port	Mode	Encapsulation	Status
fa0/1	Native VLAN desirable 5	n-802.1q	trunking
Port fa0/	VLANs is allowed on trunk 1 1-100, 102-1005		

Leading the way in IT testing and certification tools, www.testking.com

Port VLANs is owned and active in management domain
fa0/1 1-6, 8-100, 102-115, 197-999, 1002-1005

Port VLANs in spanning tree forwarding state and not pruned
fa0/1 1-6, 8-100, 102-105, 108-999, 1002-1005

TestKing users in VLAN 107 complain that they are unable to gain access to the resources through the TestKing1 router.

What is the cause of this problem?

- A. VLAN 107 is not configured on the trunk.
- B. VLAN 107 does not exist on switch TestKingA.
- C. VTP is pruning VLAN 107.
- D. Spanning tree is not enabled on VLAN 107.
- E. None of the above

Answer: C

Explanation:

In this example, VLAN 7, 101, 106, and 107 are being pruned. VLAN 107 is being pruned incorrectly in this case. By disabling VTP pruning, VLAN 107 should be able to once again gain access to the network resources.

Incorrect Answers:

A, B: Based on the output shown above, VLAN 107 is known and active within the management domain. Therefore, it must have been configured and the VLAN is indeed allowed to traverse the trunk. Only VLAN 101 has been configured to not pass along this trunk.

D: By default, STP is enabled on all VLANs.

QUESTION NO: 5

Which of the following commands would you enter if you wanted to display spanning tree statistical information?

- A. show spantree backbonefast
- B. show spantree statistics
- C. show spantree uplinkfast
- D. show spantree blockedports
- E. show spantree portstate
- F. show spantree portvlancost

Answer: B

Explanation:

The command 'show spantree statistics' is the correct IOS command to show spanning tree statistical information and is obviously the correct answer choice.

The following list various commands to use for troubleshooting Catalyst switches:

show spantree vlan_id - Shows the current state of the spanning tree for the "vlan_id" entered from the perspective of the switch on which it is entered.

show spantree summary - Provides a summary of connected spanning tree ports by VLAN.

show spantree statistics - Shows spanning tree statistical information.

show spantree backbonefast - Displays whether the spanning tree Backbone Fast Convergence feature is enabled.

show spantree blockedports - Displays only the blocked ports.

show spantree portstate - Determines the current spanning tree state of a Token Ring port within a spanning tree.

show spantree portvlancost - Shows the path cost for the VLANs on a port.

show spantree uplinkfast - Shows the uplinkfast settings.

QUESTION NO: 6

Is the following statement True or False?

The "show spanning-tree" command only shows information about ports with their red or amber lights on.

- A. True
- B. There is not enough information to determine
- C. False

Answer: C

Explanation:

The show spanning-tree command only displays information for ports with an active link (green light is on). If these conditions are not met, you can issue a show running-configuration command to confirm the configuration.

QUESTION NO: 7

Is the following statement True or False?

For optimal performance you should manually select the root switch.

- A. False
- B. True
- C. There is not enough information to determine

Answer: B

Explanation:

The selection of the root switch for a particular VLAN is very important. You can choose it, or you can let the switches decide on their own. The second option is risky because there may be sub-optimal paths in your network if the root selection process is not controlled by you.

QUESTION NO: 8

On switch TK1 the following output was shown:

```
VLAN1 is executing the ieee compatible Spanning Tree protocol
Bridge Identifier has priority 32768, address 0030.94fc.0a00
Configured hello time 2, max age 20, forward delay 15
Current root has priority 32768, address 0001.6445.4400
Root port is 323 (FastEthernet6/3), cost of root path is 19
Topology change flag not set, detected flag not set
Number of topology changes 2 last change occurred 00:02:19 ago
from FastEthernet6/1
Times: hold 1, topology change 35, notification 2
hello 2, max age 20, forward delay 15
Timers:hello 0, topology change 0, notification 0, aging 300
```

```
Port 323 (FastEthernet6/3) of VLAN1 is forwarding
Port path cost 19, Port priority 128, Port Identifier 129.67.
Designated root has priority 32768, address 0001.6445.4400
Designated bridge has priority 32768, address 0001.6445.4400
Designated port id is 129.67, designated path cost 0
Timers:message age 2, forward delay 0, hold 0
Number of transitions to forwarding state:1
BPDU:sent 3, received 91
```

Which command could you use to reproduce the above output (Type in answer below)

Answer: show spanning-tree vlan 1

Explanation:

This example shows how to display spanning tree information for a specific VLAN:

```
Switch# show spanning-tree vlan 1
VLAN1 is executing the ieee compatible Spanning Tree protocol
Bridge Identifier has priority 32768, address 0030.94fc.0a00
Configured hello time 2, max age 20, forward delay 15
```


We are the root of the spanning tree
Topology change flag not set, detected flag not set
Number of topology changes 5 last change occurred 01:50:47 ago
from FastEthernet6/16
Times: hold 1, topology change 35, notification 2
hello 2, max age 20, forward delay 15

Timers:hello 0, topology change 0, notification 0, aging 300

Port 335 (FastEthernet6/15) of VLAN1 is forwarding
Port path cost 19, Port priority 128, Port Identifier 129.79.
Designated root has priority 32768, address 0030.94fc.0a00
Designated bridge has priority 32768, address 0030.94fc.0a00
Designated port id is 129.79, designated path cost 0
Timers:message age 0, forward delay 0, hold 0
Number of transitions to forwarding state:1
BPDU:sent 6127, received 0

Switch#

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/12_1_12/cmdref/show1.htm#30158

QUESTION NO: 9

What command would you enter if you wanted to display the current state of the spanning tree for the "vlan_id" entered from the perspective of the switch on which it is entered?

- A. show spantree id vlan_id
- B. show spantree vlan_id state
- C. show spantree vlan_id
- D. show spantree state vlan_id
- E. show spantree vlan vlan_id

Answer: C

Explanation:

Commands to Use for Verifying the Configuration is Working:

show spantree vlan_id - Shows the current state of the spanning tree for the "vlan_id" entered from the perspective of the switch on which it is entered.

QUESTION NO: 10

You want to load balance traffic across your LAN. Which of the methods below are NOT the valid ways to configure load sharing with trunk ports? (Select all that apply)

- A. using STP vector metrics
- B. using ISL VLAN
- C. using STP path costs
- D. using STP port priorities
- E. using STP SID

Answer: A, B, E

Explanation:

Load sharing divides the bandwidth supplied by parallel trunks connecting switches. To avoid loops, Spanning-Tree Protocol (STP) normally blocks all but one parallel link between switches. With load sharing, you divide the traffic between the links according to which VLAN the traffic belongs to. There are two ways to configure load sharing by using trunk ports: using STP port priorities or using STP path costs.

Incorrect Answers:

C: If you configure load sharing using STP path costs, each load-sharing link can be connected to the same switch or to two different switches.

D: If you configure load sharing using STP port priorities, both load-sharing links must be connected to the same switch.

QUESTION NO: 11

What switch characteristic(s), aside from the MAC address, determines which switch will become the root bridge (Select all that apply)?

- A. IP address
- B. The port cost
- C. Path cost
- D. Priority number
- E. The port ID

Answer: D

Explanation:

When you configure a switch as the secondary root, the spanning-tree bridge priority is modified from the default value (32768) to 16384 so that the switch is likely to become the root for the specified VLANs if the primary root switch fails (assuming the other switches in the network use the default bridge priority of 32768). The MAC address is also used in the determination as a tie-breaker if two switches have the same priority value.

Note: In STP, lower is better, meaning that the lower bridge priority is preferred over a higher value.

QUESTION NO: 12

In a CLI based switch, what command will display the information comparable to the IOS command “show span”? (Type in answer below)

Answer: show spantree

Explanation:

Use the show spantree command to display spanning-tree information for a VLAN.

show spantree [vlan | mod_num/port_num] [active]vlan (Optional) Number of the VLAN. If the VLAN number is not specified, the default is VLAN 1.

mod_num (Optional) Number of the module.

port_num (Optional) Number of the port on the module.

active (Optional) Keyword that specifies to display only the active ports.

QUESTION NO: 13

You are a network troubleshooter, and you’ve arrived at a jobsite to troubleshoot a Catalyst 5000 switch. After talking with the system administrator you come to suspect that the Root Bridge for VLAN 1 is incorrect. Which command would you enter at the CLI to determine VLAN 1’s root bridge?

- A. show span 1
- B. show spantree
- C. show bridge vlan 1
- D. show spantree root bridge
- E. None of the above

Answer: B

Explanation: By default the **show spantree** command displays the STP information for VLAN 1. The bridge ID, MAC address, and timers are displayed.

Sample output:

```
TestKing> (enable) show spantree
VLAN 1
Spanning tree enabled
Spanning tree type ieee
Designated Root 00-d1-22-24-56-00
<Rest of output deleted>
```

The **Designated Root** value in the output is the MAC address of the root bridge.

QUESTION NO: 14

Which command would you enter to display the blocked ports on a spanning tree environment? (Type in answer below)

Answer: show spantree blockedports

Explanation:

Use the show spantree blockedports command to display only the blocked ports.
 show spantree blockedports [vlan_num]vlan_num (Optional) Number of the VLAN.
 The following list various commands to use for troubleshooting Catalyst switches:

show spantree vlan_id - Shows the current state of the spanning tree for the “vlan_id” entered from the perspective of the switch on which it is entered.

show spantree summary - Provides a summary of connected spanning tree ports by VLAN.

show spantree statistics - Shows spanning tree statistical information.

show spantree backbonefast - Displays whether the spanning tree Backbone Fast Convergence feature is enabled.

show spantree blockedports - Displays only the blocked ports.

show spantree portstate - Determines the current spanning tree state of a Token Ring port within a spanning tree.

show spantree portvllancost - Shows the path cost for the VLANs on a port.

show spantree uplinkfast - Shows the uplinkfast settings.

QUESTION NO: 15

You have a congested Ethernet network with only one Root Bridge on the TestKing network. What can you do to reduce BPDU traffic on this network?

- A. Remove redundant links between switches
- B. Decrease the MaxAger timer on all non-Root Bridges
- C. Increase the BPDU Hello timer only on the Root Bridge
- D. Increase the Path Cost on the Designated Port on all non-Root Bridges

Answer: C

Explanation:

There are several STP timers, as listed below:

- **hello:** the hello time is the time between each Bridge Protocol Data Unit (BPDU) that is sent on a port. This is equal to two seconds by default, but can be tuned to be between one and ten seconds.

- **forward delay:** the forward delay is the time spent in the listening and learning state. This is by default equal to 15 seconds, but can be tuned to be between four and 30 seconds.
- **max age:** the max age timer controls the maximum length of time a bridge port saves its configuration BPDU information. This is 20 seconds by default and can be tuned to be between six and 40 seconds.

The STP timers (hello, forward delay, and max age) are included in each BPDU. An IEEE bridge is not concerned about its local configuration of the timers value.

To reduce BPDU traffic on this network, increase these timers.

Incorrect Answers:

A: Redundant links are not used when STP is in use. STP will block these redundant links automatically.

B: This will be counterproductive, as it will increase the number of BPDU's

D: This will have no effect on the BPDU traffic.

QUESTION NO: 16

While logged in to switch TestKing1 you see the following output:

```
%SPANTREE-2-ROOTGUARDBLOCK: Port 1/1 tried to become non-designated in VLAN 77.
Moved to root-inconsistent state.
```

Which action could have caused the following output to appear on a switch?

- The switch is configured with Loop Guard and stops receiving BPDUs.
- The switch is configured with PortFast and starts receiving BPDUs
- The switch is configured with Loop Guard and stops receiving superior BPDUs
- The switch is configured with Loop Guard and starts receiving inferior BPDUs

Answer: C

Explanation:

The loop guard is intended to provide additional protection against L2 forwarding loops (STP loops). An STP loop is created when an STP blocking port in a redundant topology erroneously transitions to forwarding state. This usually happens because one of the ports of a physically redundant topology (not necessarily the STP blocking port) stopped receiving STP BPDUs. In its operation, STP relies on continuous reception or transmission of BPDUs, depending on the port role (designated port transmits, non-designated port receives BPDUs).

When one of the ports in a physically redundant topology stops receiving BPDUs, the STP conceives the topology as loop free. Eventually, the blocking port from the alternate or backup port becomes designated, and moves to forwarding state, thus creating a loop.

With the loop guard, an additional check is made. If BPDUs are not received any more on a non-designated port and the loop guard is enabled, that port will be moved into the STP loop-

inconsistent blocking state instead of moving to the listening / learning / forwarding state. Without the loop guard, the port would assume the designated port role. The port would move to STP forwarding state, and thus create a loop.

When the loop guard blocks an inconsistent port, the following message is logged.

SPANTREE-2-LOOPGUARDBLOCK: No BPDUs were received on port 3/2 in vlan 3. Moved to loop-inconsistent state.

Once the BPDU is received on a port in a loop-inconsistent STP state, the port will transition into another STP state. According to the received BPDU, this means that the recovery is automatic, and no intervention is necessary. After the recovery, the following message is logged.

SPANTREE-2-LOOPGUARDUNBLOCK: port 3/2 restored in vlan 3.

Reference:

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

Section 3: Identify interVLAN routing performance and scalability issues, and propose solutions (1 question)

QUESTION NO: 1

You are a network troubleshooter, and you've just been contracted by TestKing Inc. to troubleshoot their switched network, which just suddenly started experiencing difficulties after a junior administrator added a switch named TEST1 to the network.

- Their network runs on a VTP domain called 'main1'.
- The network has the active VLANs 1,2,3,4,5,10, & 20.
- No traffic is being passed on VLANs 2,3,4,5,10, &20 (which means the switches are working)

What are the configuration issues on the new switch that could be responsible for the network outage? (Select all that apply)

- A. TEST1 is configured as a VTP server with a different domain name.
- B. TEST1 is not configured to participate in VTP.
- C. TEST1 is configured as a VTP server with the domain name **main1**.
- D. TEST1 has a lower VTP configuration revision than the current VTP revision.
- E. TEST1 has a higher VTP configuration revision than the current VTP revision.
- F. TEST1 is configured with only VLAN1.

Answer: C, E, F

Explanation:

If a VTP server with the same name is added to the VTP domain (C), and the configuration revision number is higher (E), all other switches in the domain will synchronize with the highest number and take on that configuration (F), only VLAN1. In this case, if the new VTP server is added to the network and is only configured with VLAN1, this information will be propagated throughout the network and will delete the other VLANs.

Section 4: Verify and troubleshoot interVLAN routing on a switched network (2 questions)

QUESTION NO: 1

Which IOS command could you use to confirm whether or not routing is enabled on router TK1?

- A. Switch(config)#ip routing
- B. Switch#show ip routing
- C. Switch(config)#routing
- D. Switch#show routing

Answer: B

Explanation:

Use the **show ip routing** command in EXEC mode to display the current state of the IP routing protocols being routed.

QUESTION NO: 2

Is the following statement True or False?

MLS requires that the router have a path to each of the VLANs on the network.

- A. There is not enough information to determine
- B. False
- C. True

Answer: C

Explanation:

It is a basic topology requirement of MLS that the router have a path to each of the VLANs. Remember that the point of MLS is to create a shortcut between two VLANs, so that the "routing" between the two end devices can be performed by the switch, thus freeing the router for other tasks. The switch is not actually routing; it is rewriting the frames so that it appears to the end devices that they are talking through the router. If the two devices are in the same VLAN, then the MLS-SE will switch the frame locally without utilizing MLS, as switches do in such a transparently bridged environment, and no MLS shortcut will be created. One can have

multiple switches and routers in the network, and even multiple switches along the flow path, but the path between the two end devices for which one desires an MLS shortcut must include a single MLS-RP in that VLAN for that path. In other words, the flow from source to destination must cross a VLAN boundary on the same MLS-RP, and a candidate and enabler packet pair must be seen by the same MLS-SE for the MLS shortcut to be created. If these criteria are not met, then the packet will be routed normally without the use of MLS.

Section 5: Identify QoS implementation issues at the network access layer (3 questions)

QUESTION NO: 1

The CEO of your company, Tess King, is complaining about slow network performance on her workstation. While applying your systematic Cisco troubleshooting approach you clear the counters and issue the `show port` command which indicate to you a high number of alignment and FCS errors. What is potentially causing Tess's problem?

- A. There is a speed mismatch.
- B. There is a duplex mismatch.
- C. There is a trunk mode mismatch.
- D. There is a VTP mode mismatch.

Answer: B

Explanation:

The `show port` command is used to display port status and counters. Alignment and FCS errors are frames that do not end with an even number of octets and have a bad CRC. This indicates that a valid connection exists but that there are corrupt frames. This could be caused by a duplex mismatch.

Incorrect Answers

- A: A speed mismatch cannot be the cause of the problem. The speed would automatically be configured to the highest common speed.
- C: A trunk mode mismatch would not allow transfers of any frames at all.
- D: A VTP mode mismatch would not allow transfers of any frames at all.

QUESTION NO: 2

What kinds of losses are observed on healthy networks during brief periods of congestion?

- A. Loss due to jitter.
- B. Loss due to delay.
- C. Loss due to noise.
- D. Loss due to deliberately dropped packets.

Answer: D

Explanation:

Leading the way in IT testing and certification tools, www.testking.com

When a network becomes too congested, some network devices deliberately drop some packets to maintain flow. Most of these packets are TCP/IP file and print services, so if a few packets are dropped there's no major concern. However, if a UDP packet of a video or voice transaction gets dropped, there's a concern. To address this, some congestion avoidance methods were implemented that are supported by Cisco devices, such as RED and WRED.

QUESTION NO: 3

In the three-layer hierarchical network design model; what's associated with the access layer? (Select two)

- A. optimized transport structure
- B. high port density
- C. boundary definition
- D. data encryption
- E. local VLANs
- F. route summaries

Answer: B, E

Explanation:

The access layer is the outermost layer, and it is composed of the least sophisticated network equipment. The most important function of the access layer is high port density, since these devices connect the individual end users. The access layers are also where VLANs are implemented, since VLANs are assigned on a per-port basis.

Section 6: Identify QoS implementation issues at the network distribution and core layers (3 questions)

QUESTION NO: 1

You are a system administrator and you've recently implemented 'tail drops' on your network as a congestion avoidance mechanism. Which QoS technique can you use to avoid the potential problems caused by tail drops?

- A. CAR
- B. WRED
- C. CBWFQ
- D. RSVP

Answer: B

Explanation:

With class-based QoS queuing, you can create up to 64 classes for an interface. (Unlike WFQ, queues are not automatically based on the packet's ToS value.) Class-based QoS queuing also lets you control the drop mechanism used when congestion occurs on the interface. You can use WRED for the drop mechanism, and configure the WRED queues, to ensure that high-priority packets within a class are given the appropriate weight. If you use tail drop, all packets within a class are treated equally, even if the ToS values are not equal.

Reference:

www.cisco.com/en/US/products/sw/cscowork/ps2064/products_user_guide_chapter09186a00800e0a04.html

QUESTION NO: 2

The TestKing network is utilizing QoS techniques to prioritize their mission critical data. What's true about implementing QoS at the access, distribution, and core layers? (Select two)

- A. QoS implementation is the same for the access, distribution, and core layers.
- B. Classification and Marking should be done at the access layer.
- C. No QoS mechanisms are configured at the access layer since access layer switches like the Catalyst 2950 are not QoS capable.
- D. The high speed core layer only requires proper queuing (like LLQ) and dropping (like WRED) configurations.

Answer: B, D

Explanation:

You typically want to classify and mark traffic as close to source as possible (typically access layer unless you have server farms that are grouped at distribution layer). In addition, the network core should be designed to process and pass data as quickly as possible, without the added overhead of QoS configurations.

Incorrect Answers:

A: QoS markings are typically performed on the access layer, with some additional QoS mechanisms used at the distribution layer in many cases. In the core of the network, QoS should be avoided.

C: This is not true, as a number of QoS capabilities exist on these access layer switches.

QUESTION NO: 3

In the three-layer hierarchical network design model, what's associated with the core layer? (Select two)

- A. Access control list
- B. Data encryption
- C. Optimized transport
- D. Address aggregation
- E. Packet switching
- F. Boundary definition

Answer: C, E

Explanation:

A hierarchical network design includes the following three layers:

- The backbone (core) layer that provides optimal transport between sites
- The distribution layer that provides policy-based connectivity
- The local-access layer that provides workgroup/user access to the network

The distribution layer of the network is the demarcation point between the access and core layers and helps to define and differentiate the core. The purpose of this layer is to provide boundary definition and is the place at which packet manipulation can take place. In the campus environment, the distribution layer can include several functions, such as the following:

Address or area aggregation

- Departmental or workgroup access
- Broadcast/multicast domain definition
- Virtual LAN (VLAN) routing
- Any media transitions that need to occur
- Security

The distribution layer can be summarized as the layer that provides policy-based connectivity

Reference: www.alteridem.net/networking/idg4/idgbasic.htm

Topic 5, Mixed Questions (45 Questions)

QUESTION NO: 1

Exhibit



```

Hostname TestKingB
..
Interface fastethernet 0/10
spanningtree vlan 1-5 port priority 10
switchport mode trunk
!
Interface fastethernet 0/12
spanningtree vlan 6-10 port priority 10
switchport mode trunk
  
```

Refer to the exhibit. Which two statements are true about VLAN traffic? Select two.

- A. VLANs 1-5 will be blocked if fa0/10 goes down.
- B. VLANs 1-5 will use fa0/10 as a backup only.
- C. VLANs 6-10 will use fa0/10 as a backup only.
- D. VLANs 6-10 have a port priority of 128 on fa0/0.
- E. VLANs 1-10 are configured to load share between fa0/10 and fa0/12.

Answer: C, E

QUESTION NO: 2

When authentication is required, where must 802.1x be configured in order to connect a PC to a switch?

- A. client PC only
- B. switch port only
- C. switch port and client PC
- D. switch port and RADIUS server

Answer: D

QUESTION NO: 3

Exhibit

switchport mode access

switchport port-security

switchport port-security maximum 2

switchport port-security mac-address 0002.0002.0002

switchport port-security violation shutdown

Give the switch interface configuration in the exhibit, what happens when a host with the MAC address of 0003.0003.0003 is directly connected to the switch port?

- A. The port will shut down.
- B. The host will be allowed to connect.
- C. The host will be refused access.
- D. The host can only connect through a hub/switch where 0002.0002.0002 is already connected.

Answer: B

QUESTION NO: 4

Which two statements are true when the extended system ID feature is enabled? Select two.

- A. THE BID is made up of the bridge priority value (two bytes) and bridge MAC address (six bytes).
- B. THE BID is made up of the bridge priority value (four bits), the system ID (12 bits), and a bridge MAC address (48 bits).
- C. The BID is made up of the system ID (six bytes) and bridge priority value (two bytes).
- D. The system ID value is the VLAN ID (VID).
- E. The system ID value is unique MAC address allocated from a pool of MAC addresses assigned to the switch or module.
- F. The system ID value is a hex number used to measure the preference of a bridge in the spanning-tree algorithm.

Answer: B, D

Expklation:

According to Cisco white paper:Extended System ID

A 12-bit extended system ID field is part of the bridge ID. Chassis that support only 64 MAC addresses always use the 12-bit extended system ID. On chassis that support 1024 MAC

addresses, you can enable use of the extended system ID. STP uses the VLAN ID as the extended system ID. See the "Enabling the Extended System ID" section.

QUESTION NO: 5

A VTP domain named TESTKING has six active VLANs. Without notice, all VLANs except VLAN1 fail. Just prior to the failure, SwitchTestKing2 was added to the network. Which three issues on SwitchTestKing2 could be the cause? Select three.

- A. SwitchTestKing2 is configured for only VLAN1.
- B. SwitchTestKing2 is a VTP server in a different domain.
- C. SwitchTestKing2 is a VTP server in the TESTKING domain.
- D. SwitchTestKing2 is not a VTP domain.
- E. SwitchTestKing2 has a lower VTP configuration revision number than the current VTP revision.
- F. SwitchTestKing2 has a higher VTP configuration revision number than the current VTP revision.

Answer: A, C, F

Explanation: A VTP server in a given domain with the highest revision number will overwrite the VTP configuration of all other switch in the same VTP domain. Cisco best practices advises one to configure the correct VTP domain, VTP password, VTP mode, (server, client, transparent), and VTP revision number before adding any new switch to a network. The default VTP mode is server. A network can have more than one VTP domain. Each VTP domain has its own server(s) that do not influence clients in other VTP domains.

QUESTION NO: 6

After how long does it take for the port to change from blocking to forwarding when spanning-tree PortFast is enabled?

- A. immediately
- B. 15 seconds
- C. 20 seconds
- D. 30 seconds
- E. 50 seconds

Answer: A

QUESTION NO: 7

You work as a network Technician at TestKing.com. A new workstation has consistently been unable to obtain an IP address from the DHCP server when the workstation boots. Older workstations function normally, and the new workstation obtains an address when manually forced to renew its address.

What should be configured on the switch to allow the workstation to obtain an IP address at boot?

- A. UplinkFast on the switch port connected to the server
- B. BackboneFast on the switch port connected to the server
- C. PortFast on the switch port connected to the workstation
- D. trunking on the switch

Answer: C

QUESTION NO: 8

The original frame is encapsulated and an additional header is added before the frame is carried over a trunk link. At the receive end, the header is removed and the frame is forwarded to the assigned VLAN. This describes which technology?

- A. DISL
- B. DTP
- C. IEEE802.1Q
- D. ISL
- E. MPLS

Answer: D

QUESTION NO: 9

Exhibit

```

*Mar 1 00:16:43.095: %LINK-3-UPDOWN: Interface Vlan11, changed state to up
*Mar 1 00:16:43.099: SB: V111 Interface up
*Mar 1 00:16:43.099: SB11: V111 Init: a/HSRP enabled
*Mar 1 00:16:43.099: SB11: V111 Init -> Listen
*Mar 1 00:16:43.295: SB11: V111 Hello in 172.16.11.112 Active pri 50 ip 172.16.11.115
*Mar 1 00:16:43.295: SB11: V111 Active router is 172.16.11.112
*Mar 1 00:16:43.295: SB11: V111 Listen: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar 1 00:16:43.295: SB11: V111 Active router is local, was 172.16.11.112
*Mar 1 00:16:43.299: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state Listen -> Active
*Mar 1 00:16:43.299: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:43.303: SB11: V111 Hello in 172.16.11.112 Speak pri 50 ip 172.16.11.115
*Mar 1 00:16:46.207: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:49.095: SB11: V111 Hello in 172.16.11.112 Speak pri 50 ip 172.16.11.115

```

Based on the debug output shown in the exhibit, which three statements about HSRP are true? Select three.

- A. The final active router is 172.16.11.111
- B. The 172.16.11.111 router has preempt configured.
- C. The 172.16.11.112 router has a more preferred priority than the 172.16.11.111 router does.
- D. 172.16.1.115 is the virtual HSRP IP address.
- E. The 172.16.11.112 router has nonpreempt configured.
- F. The 172.16.11.112 router is using default HSRP priority.

Answer: A, B, D

QUESTION NO: 10

Exhibit



Assuming that VLAN 1 and VLAN 2 traffic is enabled on the above network, what effect will the following command have when entered on port 0/2 on switch TestKingB?

spanning-tree vlan 1 port-priority 16

- A. VLAN 1 traffic will be blocked on Switch TestKingB port 1/1.
- B. VLAN 2 traffic will be blocked on Switch TestKingB port 1/1.
- C. VLAN 2 traffic will be blocked on Switch TestKingA port 0/2.
- D. VLAN 1 and 2 traffic will be blocked on Switch TestKingA port 0/1.

E. VLAN 1 and 2 traffic will be blocked on Switch TestKingA port 0/2.

Answer: A

QUESTION NO: 11

What happens when a trunk port is configured as a SPAN destination port?

- A. The SPAN port will only be able to monitor traffic for the native VLAN of the trunk port.
- B. It allows SPAN to monitor traffic for multiple VLANs and encapsulates it on the trunk.
- C. The SPAN port will only be able to monitor traffic for a single VLAN that is defined in the SPAN command syntax.
- D. A trunk port is not capable of being configured as a SPAN destination port.

Answer: B

QUESTION NO: 12

Which two statements concerning STP state changes are true? Select two

- A. Upon bootup, a port transitions from blocking to forwarding because it assumes itself as root.
- B. Upon bootup, a port transitions from blocking to listening because it assumes itself as root.
- C. Upon bootup, a port transitions from listening to forwarding because it assumes itself as root.
- D. If a forwarding port receives no BPDUs by the max_age time limit, it will transition to listening.
- E. If a forwarding port receives an inferior BPDU, it will transition to listening.
- F. If a blocked port receives no BPDUs by the max_age time limit, it will transition to listening.

Answer: B, F

QUESTION NO: 13

Exhibit



VLAN 1 and VLAN 2 are configured between the switches TestKingA and TestKingB in the exhibit.

What should be done to load balance VLAN traffic between TestKingA and TestKingB?

- A. Lower the port priority for VLAN 1 on port 0/1 for Switch TestKingA.
- B. Lower the port priority for VLAN 1 on port 0/2 for Switch TestKingA.
- C. Make the bridge ID of Switch TestKingB lower than Switch TestKingA.
- D. Enable ports 0/1 and 0/2 on Switch TestKingA and 1/1 and 1/2 on Switch TestKingB to be trunk ports.
- E. Enable HSRP on the access ports.

Answer: A

QUESTION NO: 14

In a static VLAN environment, how does a host join a VLAN?

- A. It must be assigned to a VLAN dynamically by the VLAN server.
- B. It automatically assumes the VLAN of the port.
- C. It will be assigned to a VLAN based on the username.
- D. It will automatically be assigned a VLAN based on its MAC address.

Answer: B

QUESTION NO: 15

How does 802.1q trunking keep track of multiple VLANs?

- A. modifies the port index of a data frame to indicate the VLAN
- B. adds a new header containing the VLAN ID to the data frame
- C. encapsulates the data frame with a new header and frame check sequence
- D. tags the data frame with VLAN information and recalculates the CRC value

Answer: D

QUESTION NO: 16

By default, all VLANs will belong to which MST instance when using Multiple STP?

- A. MST00
- B. MST01
- C. the last MST instance configured
- D. none

Answer: A

QUESTION NO: 17

What are four basic security measures that should be implemented on every device at every layer of the hierarchical model? Select four.

- A. managed remote access
- B. password protection
- C. security surveillance
- D. privilege levels
- E. physical security
- F. inventory audit

Answer: A, B, D, E

QUESTION NO: 18

Which well-defined routing protocol would a network administrator configure on multicast routers when member routers are widely dispersed?

- A. Distance Vector Multicast Routing Protocol (DVMRP)
- B. Protocol Independent Multicast Dense Mode (PIM-DM)
- C. Multicast Open Shortest Path First (MOSPF)
- D. Protocol Independent Multicast Sparse Mode (PIM-SM)
- E. Core-Based Trees (CBT)

Answer: D

QUESTION NO: 19

Which type of IGMP message is sent when a network client wants to join a multicast group?

- A. host member ship query
- B. host membership report
- C. host membership status
- D. host membership notification

Answer: B

QUESTION NO: 20

Which protocol inserts a four byte tag into the Ethernet frame and recalculates CRC value?

- A. VTP
- B. 802.1Q
- C. DTP
- D. ISL

Answer: B

QUESTION NO: 21

Exhibit



Assuming that VLAN 1 and VLAN 2 traffic is enabled on the above network, what effect will the following command have when entered on port 0/2 on switch TestKingA?

spanning-tree vlan 1 port-priority 16

- A. VLAN 1 traffic will be blocked on Switch TestKingB port 1/1.
- B. VLAN 2 traffic will be blocked on Switch TestKingB port 1/1.

642-891

- C. VLAN 2 traffic will be blocked on Switch TestKingA port 0/2.
- D. VLAN 1 and 2 traffic will be blocked on Switch TestKingA port 0/1.
- E. VLAN 1 and 2 traffic will be blocked on Switch TestKingA port 0/2.

Answer: A

QUESTION NO: 22**Exhibit**

```

TestKingA# show standby
Ethernet0/1 Group 1
  State is Active
    2 state changes, last state change 00:30:59
  Virtual IP address is 10.1.0.20
    Secondary virtual IP address 10.1.0.21
  Active virtual MAC address is 0004.4d82.7981
    Local virtual MAC address is 0004.4d82.7981 (bia)
  Hello time 4 sec, hold time 12 sec
    Next hello sent in 1.412 secs
  Preemption enabled, min delay 50 sec, sync delay 40 sec
  Active router is local
  Standby router is 10.1.0.6, priority 75 (expires in 9.184 sec)
  Priority 95 (configured 120)
    Tracking 2 objects, 0 up
      Down Interface Ethernet0/2, pri 15
      Down Interface Ethernet0/3
  IP redundancy name is "HSRP1", advertisement interval is 34 sec

```

Study the router output displayed in the exhibit.

Which two items are correct? Select two.

- A. TestKingA will assume the active stat if its priority is the highest.
- B. If Ethernet 0/2 goes down, the standby router will take over.
- C. When Ethernet 0/3 of TestKingA comes back up, the priority will become 105.
- D. The local IP address of TestKingA is 10.1.0.6.
- E. The local IP address of TestKingA is 10.1.0.20.

Answer: A, C

QUESTION NO: 23

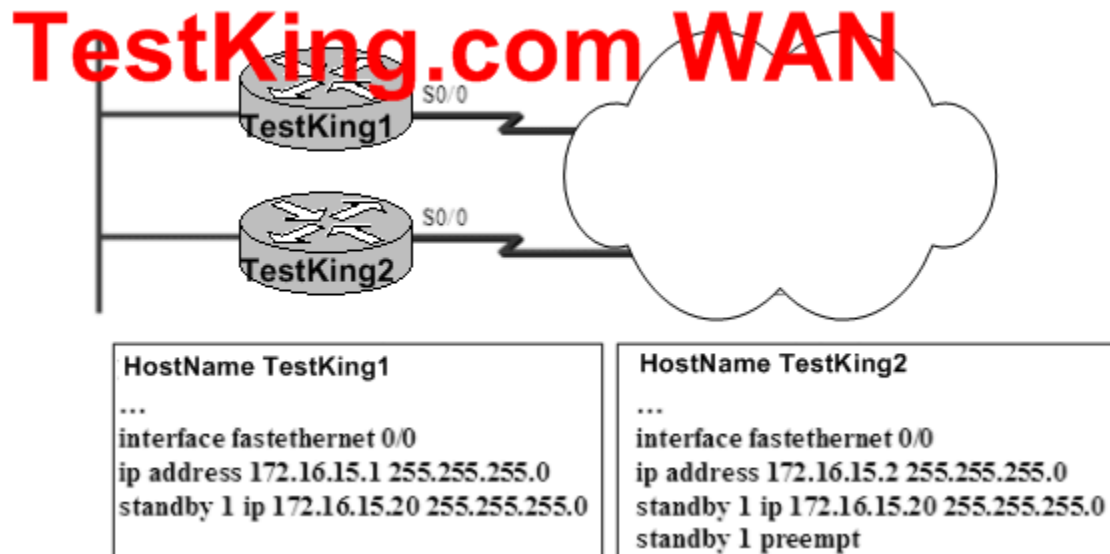
When IP multicast is enabled via PIM, which mode uses the flood and prune method?

- A. PIM sparse-dense
- B. Bidir-PIM
- C. PIM-RP
- D. PIM-DM
- E. PIM-SM

Answer: D

QUESTION NO: 24

Exhibit



Which command will ensure that TestKing2 will be the primary router for traffic using the gateway address of 172.16.15.20?

- A. On TestKing2 add the command **standby 1 priority 80**
- B. On TestKing1 add the command **standby 1 priority 110**
- C. On TestKing1 add the command **standby 1 priority 80**
- D. On TestKing2 remove the command **standby 1 preempt**

Answer: B

Explanation: HSRP default priority is 100 - the highest configured priority will win the standby IP function. The default HSRP priority is 100; any router with a HSRP priority greater than 100 will become the active speaker. If the active router stops speaking, the backup router detects the outage and will assume HSRP standby IP. A HSRP router can only reclaim the standby IP if the preempt statement has been configured. In the case mentioned above the router must be configured with a HSRP priority greater than 100; therefore the only correct answer is 110 which is greater than 100.

QUESTION NO: 25

On a multilayer Catalyst switch, which interface command is used to convert a Layer 3 interface to Layer 2 interface?

- A. switchport
- B. no switchport

- C. switchport mode access
- D. switchport access vlan *vlan-id*

Answer: A

QUESTION NO: 26

Which three statements about STP timers are true? Select three.

- A. STP times values (hello, forward delay, max age) are included in each BPDU
- B. A switch is not concerned about its local configuration of the STP timers values. It will only consider the value of the STP timers contained in the BPDU it is receiving.
- C. To successfully exchange BPDUs between two switches, their STP timers value (hello, forward delay, max age) must be the same.
- D. If any STP timer value (hello, forward delay, max age) needs to be changed, it should at least be changed on the root bridge and backup root bridge.
- E. On a switched network with a small network diameter, the STP hello timer can be tuned to a lower value to decrease the load on the switch CPU.
- F. The root bridge passes the timer information on BPDUs to all routers in the Layer 3 configuration.

Answer: A, B, D

QUESTION NO: 27

Which MST configuration statement is correct?

- A. MST configurations can be propagated to other switches using VTP.
- B. After MST is configured on a Switch, PVST+ operations will also be enabled by default.
- C. MST configurations must be manually configured on each switch within the MST region.
- D. MST configurations only need to be manually configured on the Root Bridge.
- E. MST configurations are entered using the VLAN Database mode on Cisco Catalyst switches.

Answer: C

MST configuration must be manually configured on each switch within the MST region.

QUESTION NO: 28

Exhibit

```
TestKing1#show spanning-tree vlan 200
```

VLAN200

```
Spanning tree enabled protocol ieee
```

```
Root ID      Priority      32968
             Address      000c.ce29.ef00
             Cost        19
             Port        2 (FastEthernet0/2)
             Hello time 10 Sec Max Age 20 sec Forward Delay 30 sec
```

```
Bridge ID    Priority      32968 (priority 32768 sys-id-ext 200)
             Address      000c.ce2a.4180
             Hello Time 2 sec Max Age 20 Sec Forward Delay 15 sec
```

Interface	Role	Sts	Cost	PrioNbr	Type
Fa0/2	Root	FWD	19	128.2	P2p
Fa0/3	Altn	BLK	19	128.3	P2p

Based on the show spanning-tree vlan 200 output shown in the exhibit, which two statements about the STP process for VLAN 200 are true? (Choose two)

- A. BPDUs will be sent out every two seconds.
- B. The time spent in the listening state will be 30 seconds
- C. The time spent in the learning state will be 15 seconds
- D. The maximum length of time that the BPDU information will be saved is 30 seconds.
- E. This switch is the root bridge for VLAN 200.
- F. BPDUs will be sent out every 10 seconds.

Answer: B, F

Changing the Spanning Tree Protocol Timers The STP timers (hello, forward delay, and max age) are included in each BPDU. An IEEE bridge is not concerned about its local configuration of the timers value. It will consider the value of the timers contained in the BPDU that it is receiving. Effectively, that means only a timer configured on the root bridge of the STP is important. Obviously, in case you would lose the root, the new root would start to impose its local timer value to the entire network. So, even if it is not required to configure the same timer value in the entire network, it is at least mandatory to configure any timer changes on the root bridge and on the backup root bridge.

QUESTION NO: 29

You are configuring a switching solution and you want to take advantage of the Fast EtherChannel ports. When configuring FastEthernet ports, which precautions can you take to avoid configuration problems which can cause the ports to be automatically disabled? (Select two)

- A. Allow some ports in a channel to be partly disabled.
- B. Configure ALL the ports in a channel as dynamic.
- C. Configure all ports in a channel to operate at the same speed and duplex mode
- D. Assign all ports in a channel to the same VLAN or configure them as trunk ports.

Answer: C, D

Explanation:

Cisco's Fast EtherChannel technology builds upon standards based 802.3 full duplex Fast Ethernet to provide network managers a reliable high speed solution for the campus network backbone. Fast EtherChannel provides bandwidth scalability within the campus by providing increments from 200 Mbps to 800 Mbps with multi-gigabit capacity in the future. Fast EtherChannel technology not only solves the immediate problem of scaling bandwidth within the network backbone today, but also paves the path for an evolution to standards-based Gigabit Ethernet and beyond, because Fast EtherChannel technology can be applied to support Gigabit EtherChannel.

In order for a channel to function properly, the aggregated links should be in the same VLAN or the links should be assigned as a trunk. In addition, all links should have identical speed and duplex settings.

QUESTION NO: 30

Which process plays a major role in the creation of the CEF adjacency table? (811)

- A. Address Resolution Protocol (ARP)
- B. PDU header rewrite
- C. NetFlow Switching
- D. hello packet exchange

Answer: A

QUESTION NO: 31

Which protocol is an extension to ICMP that provides a mechanism for routers to advertise useful default routes?

- A. IRDP
- B. HSRP
- C. VRRP
- D. Proxy ARP
- E. GLBP

Answer: A

IRDP is an extension to ICMP that provides a mechanism for routers to advertise useful default routes

QUESTION NO: 32

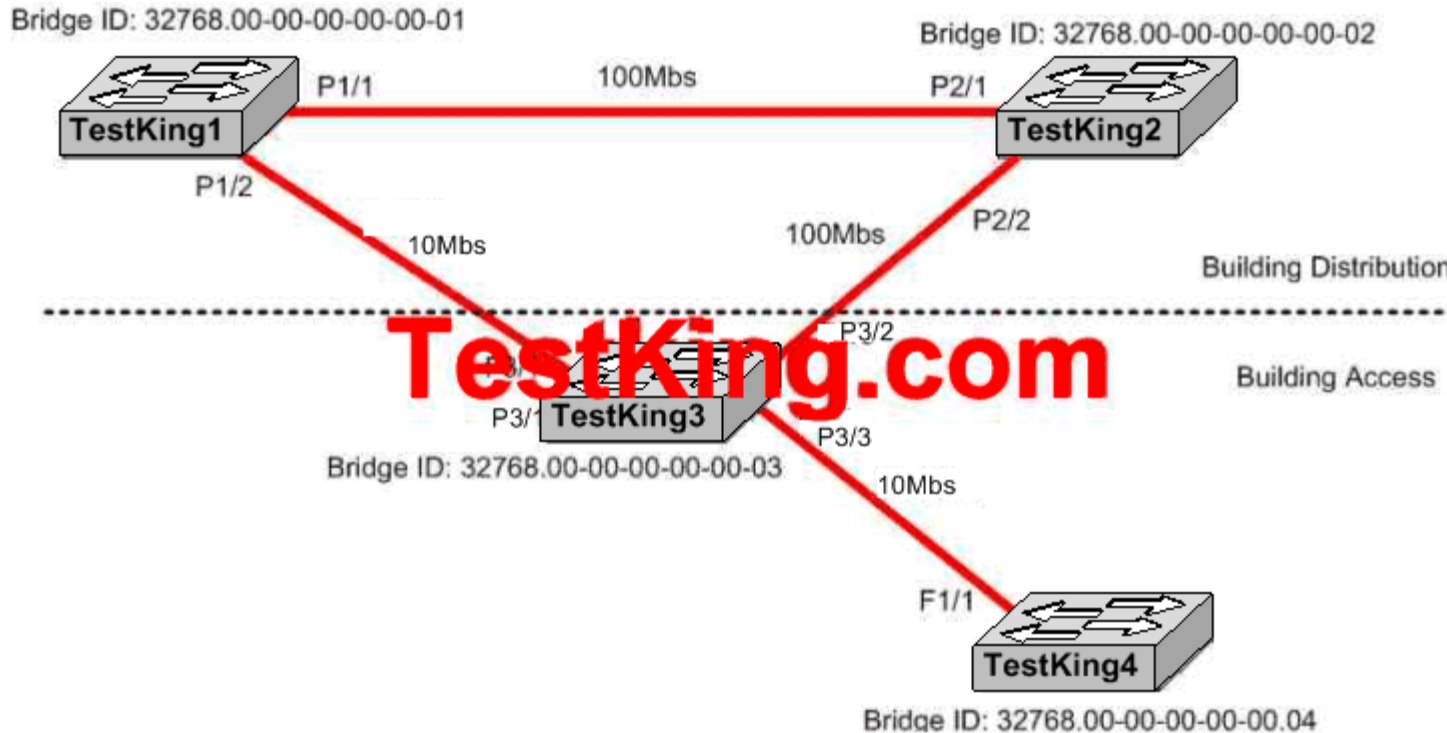
Which router redundancy protocol cannot be configured for interface tracking?

- A. HSRP
- B. GLBP
- C. VRRP
- D. SLB
- E. RPR
- F. RPR+

Answer: C

QUESTION NO: 33

The TestKing switched LAN is displayed below:



Your junior network administrator has just finished installing the above switched network using Cisco 3550s and would like to manipulate the root bridge election. Which switch should he configure as the root bridge and with which command?

- A. TestKing1(config)# spanning-tree vlan 1 priority 4096
- B. TestKing2(config)# set spanning-tree priority 4096
- C. TestKing3(config)# spanning-tree vlan 1 priority 4096
- D. TestKing1(config)# set spanning-tree priority 4096
- E. TestKing2(config)# spanning-tree vlan 1 priority 4096
- F. TestKing3(config)# set spanning-tree priority 4096

Answer: E

QUESTION NO: 34

What is the default priority value assigned to a switch when STP is enabled?

- A. 1
- B. 255
- C. 4096
- D. 32,768
- E. 65,536

Answer: D

QUESTION NO: 35

Which two table types are CEF components? Select two.

- A. forwarding information base
- B. adjacency tables
- C. neighbor tables
- D. caching tables
- E. route tables.

Answer: A, B

QUESTION NO: 36

Which three QoS mechanisms can be configured to improve VoIP quality on a converged network? Select three.

- A. the use of a queuing method that will give VoIP traffic strict priority over other traffic
- B. the use of RTP header compression for the VoIP traffic.
- C. the proper classification and marking of the traffic as close to the source as possible
- D. the use of 802.1QinQ trunking for VoIP traffic
- E. the use of WRED for the VoIP traffic

Answer: A, B, C

QUESTION NO: 37

What protocol specified by RFC 1256 will allow an enabled IP host a new router when a router becomes unavailable?

- A. IRDP
- B. SNMP
- C. HSRP
- D. VRRP

Answer: A

QUESTION NO: 38 Drag and Drop

As a TestKing.com administrator you are required to drag the port states to their correct description.

Description

sends and receives BPDUs to determine root, but does not update the MAC address table	Place here
does not participate in frame forwarding or in STP	Place here
does not participate in frame forwarding	Place here
sends and receives data frames	Place here
populates the MAC address table, but will not forward user data	Place here

Select from these:

- Blocking
- Listening
- Learning
- Forwarding
- Disabled

Answer:

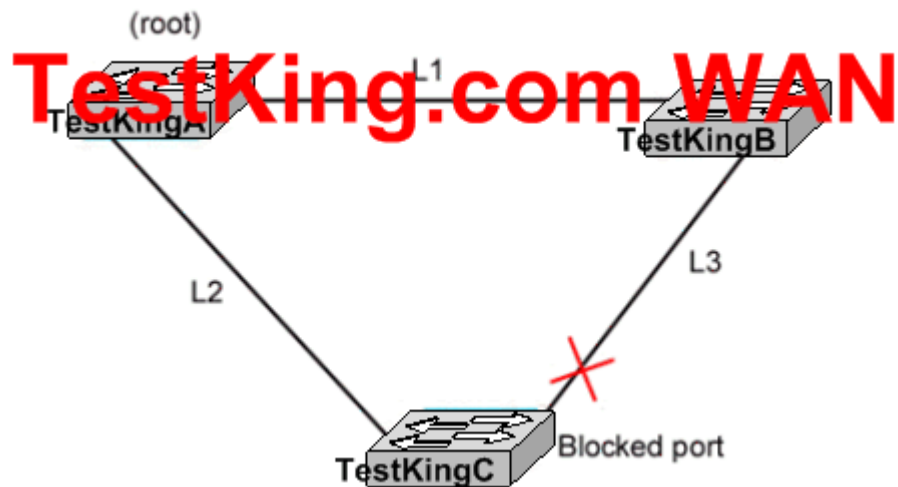
Description

sends and receives BPDUs to determine root, but does not update the MAC address table	Listening
does not participate in frame forwarding or in STP	Disabled
does not participate in frame forwarding	Blocking
sends and receives data frames	Forwarding
populates the MAC address table, but will not forward user data	Learning

Select from these:

QUESTION NO: 39

Exhibit, Network Topology



Switch TestKingC is configured with UplinkFast. How much time will expire after a failure in L2 before Switch TestKingC activates the port connected to L3?

- A. 1-5 seconds
- B. 15 seconds
- C. 30 seconds
- D. 50 seconds

Answer: A

QUESTION NO: 40

Exhibit

```
*Mar 1 00:12:16.871: SB11: V111 Hello in 172.16.11.112 Active pri 50 ip 172.16.11.115
*Mar 1 00:12:16.871: SB11: V111 Active router is 172.16.11.112
*Mar 1 00:12:18.619: %LINK-3-UPDOWN: Interface Vlan11, changed state to up
*Mar 1 00:12:18.623: SB: V111 Interface up
*Mar 1 00:12:18.623: SB11: V111 Init: a/HSRP enabled
*Mar 1 00:12:18.623: SB11: V111 Init -> Listen
*Mar 1 00:12:19.619: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan11, changed state to up
*Mar 1 00:12:19.819: SB11: V111 Hello in 172.16.11.112 Active pri 50 ip 172.16.11.115
*Mar 1 00:12:19.819: SB11: V111 Listen: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar 1 00:12:22.815: SB11: V111 Hello in 172.16.11.112 Active pri 50 ip 172.16.11.115
*Mar 1 00:12:22.815: SB11: V111 Listen: h/Hello rcvd from lower pri Active router
*Mar 1 00:12:25.683: SB11: V111 Hello in 172.16.11.112 Active pri 50 ip 172.16.11.115
*Mar 1 00:12:25.683: SB11: V111 Listen: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar 1 00:12:28.623: SB11: V111 Listen: d/Standby time expires (unknown)
*Mar 1 00:12:28.623: SB11: V111 Listen -> Speak
*Mar 1 00:12:28.623: SB11: V111 Hello out 172.16.11.111 Speak pri 100 ip 172.16.11.115
*Mar 1 00:12:28.659: SB11: V111 Hello in 172.16.11.112 Active pri 50 ip 172.16.11.115
*Mar 1 00:12:28.659: SB11: V111 Speak: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar 1 00:12:31.539: SB11: V111 Hello in 172.16.11.112 Active pri 50 ip 172.16.11.115
*Mar 1 00:12:31.539: SB11: V111 Speak: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar 1 00:12:31.575: SB11: V111 Hello out 172.16.11.111 Speak pri 100 ip 172.16.11.115
*Mar 1 00:12:34.491: SB11: V111 Hello in 172.16.11.112 Active pri 50 ip 172.16.11.115
```

What can be determined about the HSRP relationship from the displayed debug output?

- A. The pre-empt feature is not enable don the 172.16.11.111 router.
- B. The nonpreempt feature is enabled on the 172.16.11.112 router.
- C. Router 172.16.11.111 will be the activate router because its HSRP priority is preferred over router 172.16.11.112.
- D. Router 172.16.11.111 will be the activate router because its HSRP priority is preferred over router 172.16.11.111.
- E. The IP address 172.16.11.111 is the virtual HSRP router IP address.
- F. The IP address 172.16.11.112 is the virtual HSRP router IP address.

Answer: A

QUESTION NO: 41

Leading the way in IT testing and certification tools, www.testking.com

In the use of 802.1X access control, which three products are allowed through the switch port before authentication takes place? Select three.

- A. STP
- B. CDP
- C. EAP MD5
- D. TACACS+
- E. EAP-over-LAN
- F. protocols not filtered by an ACL

Answer: A, B, E

QUESTION NO: 42

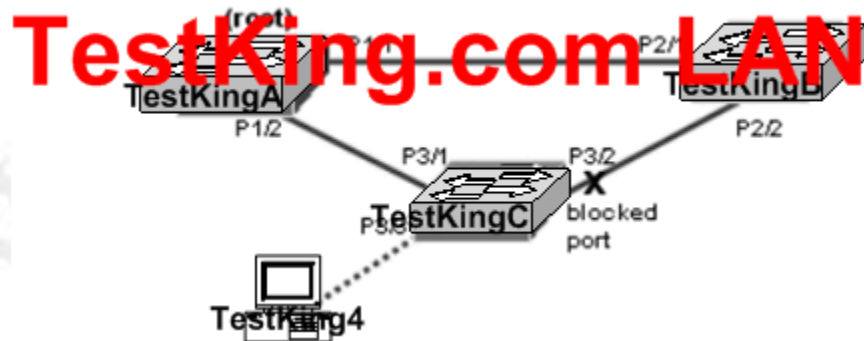
Which protocol does TACACS+ use to communicate between a TACACS+ server and a TACACS+ client?

- A. UDP
- B. TCP
- C. IP
- D. LEAP

Answer: B

QUESTION NO: 43

Exhibit



You work as a technician at TestKing.com. Study the exhibit carefully. Spanning tree is enabled on all devices. Currently either Switch TestKingB or TestKingC can serve as the

root should switch TestKingA fail. A client recently connected to Device TestKing4, a PC running Windows XP SP2 and switching application software, to Switch TestKingC port P3/3. You must configure Root Guard to ensure that the TestKing4 PC does not assume the role of the root. All other parameters must stay the same. On which interface(s) must Root Guard be enabled?

- A. P1/2
- B. P2/2
- C. P3/3
- D. P1/1 and P1/2
- E. P1/2 and P2/2
- F. P1/2, P2/2 and P3/3

Answer: C

QUESTION NO: 44

You work as a technician at TestKing.com. You map VLANs 10 through 20 to MST instance 2.

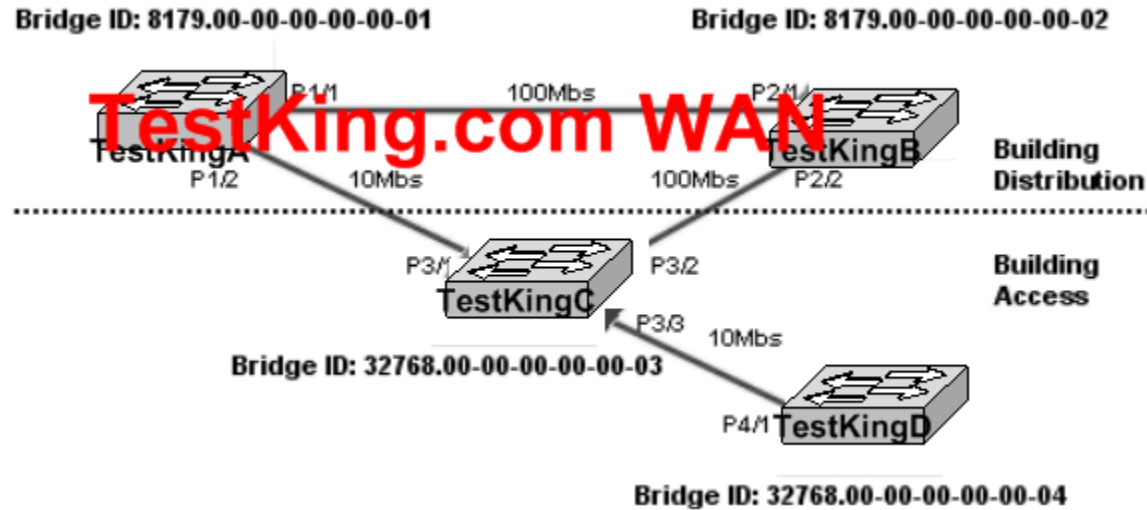
How will this information be propagated to all appropriate switches?

- A. Information will be carried in the RSTP BPDUs.
- B. It will be propagated in VTP updates.
- C. Information is stored in the Forwarding Information Base and the switch will reply upon query.
- D. Multiple Spanning Tree must manually configured on the appropriate switches,.

Answer: D

QUESTION NO: 45

Exhibit



Based on the assumption that STP is enabled on all the switch devices, which of the following statements are true? (Choose two)

- A. TestKingA will be elected the root bridge.
- B. TestKingB will be elected the root bridge.
- C. TestKingC will be elected the root bridge.
- D. P3/1 will be elected the nondesignated port.
- E. P2/2 will be elected the nondesignated port.
- F. P3/2 will be elected the nondesignated port.

Answer: A, D