PassQueen
a.k.a dumpsqueen.com

CCIE SP Cable Qualification Exam

Cisco 350-021

Version Demo

Total Demo Questions: 20

Total Premium Questions: 208

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

In order to avoid loops when sending routing updates, what is the correct technique to prevent a network from being forwarded on the same interface it is learned?

A. Poison Reverse

B. The use of access-lists used with distribute-list

C. Split Horizon

D. This is not a problem, since this cannot happen.

ANSWER: C

Explanation:

Split Horizon is the loop-prevention technique that suppresses advertisement of a route back out through the same interface on which that route was learned. In a distance-vector routing environment, this rule prevents a neighboring router from receiving an update that could cause it to believe the advertising router has an alternate path to a network that the neighbor originally advertised. By stopping that immediate return advertisement, split horizon reduces the opportunity for routing loops and count-to-infinity behavior after a topology change or failure. The mechanism is applied per outgoing interface: routes learned on an interface are omitted from routing updates transmitted through that interface. This behavior is especially relevant on shared-media and hub-and-spoke designs, where routing updates can otherwise be heard by devices that supplied the original route information. Cisco documents split horizon as a RIP loop-avoidance feature and describes its operation as preventing routes learned through an interface from being advertised back out that same interface. See Cisco's RIP routing and loop-prevention documentation and the Cisco IOS RIP route-control guide.

QUESTION NO: 2

The Hold-Down Timer used in Distance Vector protocols:

A. Sets the time limit of how long a router may keep a packet in its buffer, if the routing protocol is in the process of converging

B. Determines the number of seconds a router will wait before sending another update to neighbor routers

C. Sets a specific period for routers to neither accept or advertise a route from a destination where an original route was recently lost

D. Sets a duration where routes are not accepted from the neighbor router that caused a routing loop

ANSWER: C

Explanation:

"Sets a specific period for routers to neither accept or advertise a route from a destination where an original route was recently lost" is correct because a hold-down timer is a distance-vector loop-prevention mechanism used after a router determines that a route has become unavailable. During the hold-down interval, the router treats the route as unreachable and suppresses potentially inconsistent routing information about that destination. This gives failure information time to propagate through the network before the router considers a newly received route advertisement, helping avoid acceptance and redistribution of stale information that could contribute to a transient routing loop.

In RIP, when a route times out or is otherwise declared unreachable, the route is advertised with an infinite metric while it remains in the routing database for a period. Hold-down behavior is specifically intended to improve stability after topology changes by preventing immediate reaction to possibly outdated advertisements. The exact operational behavior can vary with implementation and configuration, but the essential function is to temporarily suppress route changes for a recently failed destination. Cisco documents RIP timer behavior in its RIP configuration guide; the RIP protocol's handling of invalid routes and route timers is also specified in RFC 2453.

QUESTION NO: 3

What is NOT an ATM class of service?

A. CBR

B. VBR-t

C. ABR

D. UBR

E. CAR

ANSWER: E

Explanation:

CAR is correct because Committed Access Rate is a Cisco traffic-policing mechanism, not an Asynchronous Transfer Mode service category. CAR classifies traffic and applies a configured rate limit, typically using committed and excess burst parameters. Packets that conform to the configured rate can be transmitted or remarked, while traffic exceeding the policy can be dropped or have its precedence modified. This function is applied as part of QoS policy enforcement on an interface; it does not define the cell-transfer behavior or bandwidth/service guarantees of an ATM virtual circuit.

ATM service categories are negotiated characteristics for an ATM connection and describe the traffic behavior the network supports, such as fixed-rate, variable-rate, available-bit-rate, or unspecified-bit-rate delivery. These categories enable the ATM network to allocate resources and handle traffic according to its delay, loss, timing, and bandwidth requirements. CAR can be used with many types of traffic and transport technologies, including traffic carried over ATM, but it remains an edge policing feature rather than an ATM class of service. Cisco describes CAR as a rate-limiting and traffic-policing feature in its QoS documentation: Cisco IOS QoS: Committed Access Rate. ATM traffic-management service categories are specified in ITU-T Recommendation I.371.

QUESTION NO: 4

What is the principal purpose of the cable modem Maximum CPE setting?

A. To limit the number of customer devices that can be learned behind a cable modem

B. To set the maximum upstream transmit power of a cable modem

C. To limit the number of downstream channels that a cable modem can bond

D. To reserve a fixed number of upstream minislots for each cable modem

E. To set the maximum number of service flows provisioned for a cable modem

ANSWER: A

Explanation:

To limit the number of customer devices that can be learned behind a cable modem is correct. The CMTS learns subscriber host MAC addresses associated with a cable modem and enforces the configured maximum CPE value. This lets an operator control how many customer-premises devices may obtain service through one modem, helping enforce the subscriber service policy and limiting unauthorized sharing of a residential connection.

The setting does not control the number of bonded channels, cable modem transmit power, or the number of service flows. Those functions are determined by PHY operation, ranging, and the DOCSIS configuration and service-flow policy. Cisco cable documentation describes cable modem host-learning and subscriber management functions in its cBR Series configuration documentation.

QUESTION NO: 5

Which addresses below would be valid IP addresses of hosts on the Internet? (multiple answer)

A. 235.1.1.1

B. 223.20.1.1

C. 10.100.1.1

D. 127.0.0.1

E. 24.15.1.1

ANSWER: B E

Explanation:

223.20.1.1 and **24.15.1.1** are valid publicly routable IPv4 unicast address values. Both fall within IPv4 address space that is not designated by IANA as a special-purpose range such as private-use, loopback, multicast, or reserved-for-future-use space. Therefore, either address can be allocated through the normal Internet address-registration hierarchy and assigned to an Internet-connected host or interface. In this context, “valid IP addresses of hosts on the Internet” means addresses that are eligible for global unicast routing; it does not guarantee that a particular address is currently assigned, reachable, or responding to traffic. Actual reachability also depends on allocation status, routing advertisements, network policy, and host configuration. IANA maintains the authoritative registry for IPv4 special-purpose address blocks, which identifies address ranges with non-global purposes and supports determining whether an address is eligible for ordinary Internet use. The private IPv4 ranges are additionally defined in RFC 1918. See the IANA IPv4 Special-Purpose Address Registry and RFC 1918.

QUESTION NO: 6

According to the MGCP call flow shown, from what line can it be deduced that Call agent requests notification about an on hook condition?

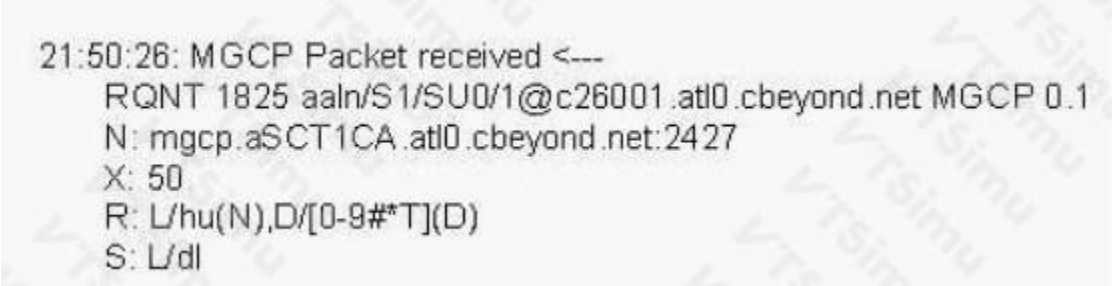

A. N

B. X

C. R

D. S

E. RQNT

ANSWER: C

Explanation:

R is correct because it identifies the Request for Notification (RQNT) transaction in the depicted MGCP exchange where the call agent tells the endpoint which events it must monitor and report. An on-hook state is an endpoint event, conventionally represented by the MGCP line-package event name `L/hu` (hook up). By placing that event in the `RequestedEvents` parameter of an RQNT command, the call agent explicitly requests an asynchronous notification when the subscriber

handset returns to the on-hook state. The gateway subsequently reports occurrence of a requested event to the call agent with a Notify (NTFY) command; however, the instruction to monitor for the event originates in the RQNT. Thus, the line labeled R, which carries the RQNT request containing the on-hook event request, is the evidence that the call agent has requested notification. MGCP defines RQNT as the command used by a call agent to direct an endpoint to detect specified events and notify the call agent when those events occur. See RFC 3435, Media Gateway Control Protocol, particularly the Request for Notification command and the line package event definitions. Cisco's MGCP configuration guidance also describes MGCP call-agent control of gateway endpoints: Cisco IOS MGCP Configuration Guide.

QUESTION NO: 7

Cisco recommends the use of a headend-grade coaxial cable or a quad-shield coaxial cable with a minimum of 60% + 40% braid and double foil insulation to connect the cable modem cards to the HFC network. The center conductor must be straight and extend 1/8 inch (3.2 mm) beyond the end of the connector, and the connector should be securely crimped to the cable. Select three recommended cable types that can be used to connect the Cisco uBR7200 series universal broadband routers: (multiple answer)

A. RG 58

B. RG 59

C. RG 59/U

D. RG 6

E. RG 7

ANSWER: B C D

Explanation:

RG 59, RG 59/U, and RG 6 are the recommended coaxial cable types for connecting Cisco uBR7200 cable modem cards to the hybrid fiber-coaxial network. These cable families are used in broadband cable environments because suitable headend-grade or quad-shield versions provide the shielding and impedance characteristics needed for reliable RF connections. In particular, the installation requirement calls for cable with double foil shielding and at least 60% plus 40% braid coverage, which helps limit ingress and egress and preserves signal quality in the cable plant. Correct connector preparation is equally important: the center conductor must remain straight, project 1/8 inch (3.2 mm) beyond the connector, and the connector must be securely crimped. Those mechanical requirements ensure a consistent RF contact and reduce the risk of attenuation, reflections, or intermittent connectivity at the router's cable interfaces. Cisco's uBR7200 documentation identifies RG 6, RG 59, and RG 59/U for these connections; cable selection must also satisfy the stated shielding and connector-installation requirements. See the Cisco uBR7200 Series support documentation and the Cisco uBR7200 product documentation.

QUESTION NO: 8

To restrict SNMP access to a router, what configuration command could be used?

A. snmp-server community

B. snmp-server enable

C. snmp-server log

D. snmp-server host

ANSWER: A

Explanation:

The correct command is `snmp-server community`. On Cisco IOS, this command creates an SNMPv1 or SNMPv2c community string and can associate that community with an access control list, allowing the router to restrict which

management-station source addresses may use that community. For example, `snmp-server community NMS-RO RO 10` grants read-only SNMP access only to sources permitted by ACL 10. The command can also apply a named standard ACL and specify a read-only or read-write community permission. This makes it a direct mechanism for limiting SNMP management access at the device. In production, Cisco recommends using a restrictive ACL, read-only access whenever write capability is unnecessary, and preferably SNMPv3 because it provides authenticated and encrypted management sessions rather than community-string-based access. Cisco documents the community configuration syntax and optional ACL parameter in its IOS SNMP command reference: Cisco IOS SNMP Command Reference. Additional SNMP security guidance is available from Cisco SNMP Security.

QUESTION NO: 9

What is true about the DLCI field in the Frame Relay header?

A. It consists of two portions, source and destination, which map data to a logical channel.

B. It generally has significance only between the local switch and the DTE device.

C. It is an optional field in the ITU-T specification.

D. It is present only in data frames sent through the network.

ANSWER: B

Explanation:

It generally has significance only between the local switch and the DTE device. A Data-Link Connection Identifier (DLCI) identifies the logical Frame Relay virtual circuit at the customer access interface, specifically on the link between a data terminal equipment (DTE) device, such as a router, and the provider's Frame Relay switch. The service provider assigns the DLCI value for that local interface and uses it to associate frames received from the DTE with the appropriate permanent virtual circuit (PVC) or switched virtual circuit (SVC). DLCIs are therefore locally significant: the same numerical value may be used on an entirely different customer interface without representing the same end-to-end virtual circuit. This local scope allows the provider network to map a customer-facing DLCI to its internal forwarding information and to the DLCI used at the remote access interface. In the Frame Relay header, the DLCI is carried in the address field and enables the access device and local switch to identify the intended logical connection. Cisco's Frame Relay documentation describes DLCIs as identifiers of virtual circuits that are locally significant, while RFC 1490 defines the Frame Relay encapsulation and its use of the Frame Relay address field. See Cisco: Frame Relay DLCIs and RFC 1490.

QUESTION NO: 10

Unity gain is being used for alignment. The first amplifier from the node has 25 dB of gain and is 2000 feet away with cable loss of 1 dB per 100 feet. The node has an output of 35 dBmV at 55.25 MHz. What is the input at the amplifier?

A. 13dBmV

B. 3dBmV

C. 15dBmV

D. 27dBmV

ANSWER: C

Explanation:

15dBmV is the correct input level. The cable span is 2,000 feet, and its specified attenuation is 1 dB per 100 feet at 55.25 MHz. Dividing 2,000 feet by 100 feet gives 20 cable-loss sections. Therefore, total loss between the node output and the amplifier input is 20 dB. Signal levels expressed in dBmV can be calculated through a passive cable section by subtracting attenuation from the upstream level: 35 dBmV − 20 dB = 15 dBmV. Thus, the signal arriving at the first amplifier input is 15 dBmV.

The stated 25 dB amplifier gain is relevant to the following active stage and to confirming the aligned output level, but it does not change the signal level that arrives at the amplifier input. Under the unity-gain alignment approach, the cascade is designed so that the active-device gain compensates for the intended preceding or following plant losses while maintaining the planned operating levels. Cisco cable-network materials describe the RF access network as a combination of optical nodes, coaxial cable, and active amplification stages, where level planning accounts for cable attenuation and amplifier gain. See Cisco's Cable solutions overview and service-provider architecture resources.

QUESTION NO: 11

Which are the correct ways to release IBGP from the condition that all IBGP neighbors need to be fully meshed? (multiple answer)

A. Configure local preference

B. Configure route reflectors

C. Configure IBGP neighbors several hops away

D. Configure confederations

ANSWER: B D

Explanation:

Configuring route reflectors and configuring confederations are the two scalable mechanisms that remove the requirement to build a full internal BGP peer mesh. In a conventional internal BGP deployment, routes learned from one internal peer are not advertised to another internal peer. This split-horizon behavior prevents routing loops, but it means every internal speaker would otherwise need a direct session with every other speaker. A route reflector deliberately relaxes that rule for defined clients: it can advertise routes received from a client to other clients and to nonclient peers, while using originator and cluster attributes to prevent loops. This permits a hierarchical design with far fewer sessions. Confederations divide one public autonomous system into multiple internal member autonomous systems. Internal BGP is used within each member, while sessions between member autonomous systems use confederation external BGP behavior. To outside networks, the confederation still presents a single autonomous system. Both approaches substantially improve scalability while preserving loop-prevention controls. Cisco documents these designs in its BGP route reflector overview and BGP confederation overview.

QUESTION NO: 12

Which characteristics describe a DOCSIS service flow? (multiple answer)

A. It is identified by a service flow ID.

B. It can define QoS parameters for classified traffic.

C. It identifies the downstream RF frequency used by the cable modem.

D. It is equivalent to an IEEE 802.1Q VLAN.

E. It can be associated with an upstream scheduling service.

ANSWER: A B E

Explanation:

A DOCSIS service flow is identified by a service flow ID (SFID), so this statement is correct. The SFID is used by the CMTS and cable modem to identify the service flow and associate it with the applicable QoS parameters.

A service flow can also contain QoS parameters such as maximum sustained traffic rate, maximum traffic burst, minimum reserved traffic rate, and scheduling type. These parameters define the forwarding and scheduling treatment that the DOCSIS system applies to packets classified into the flow. For example, an unsolicited grant service flow can be used for traffic requiring periodic, low-latency upstream grants.

A service flow is not simply a VLAN, and it does not identify an RF frequency. DOCSIS service-flow concepts are specified in the CableLabs DOCSIS MAC and Upper Layer Protocols Interface specification.

QUESTION NO: 13

When using IS-IS for IP routing, Dual IS-IS defined by RFC 1195, what is true? (multiple answer)

A. It is necessary to configure a NSAP address.

B. It is not possible to perform both IP and CLNS routing wih the same process.

C. IP address and subnet information is carried in the TLV field on the L-1/L-1 LSPs.

D. Dual IS-IS does not support VLSM information.

ANSWER: A C

Explanation:

Dual (Integrated) IS-IS, specified in RFC 1195, uses the OSI IS-IS control-plane framework while carrying IP reachability information in IS-IS link-state PDUs. Therefore, an IS-IS router requires an OSI Network Service Access Point address, typically configured as a Network Entity Title. The NET identifies the router's IS-IS instance and supplies the system ID used in IS-IS operations; it is required even when the deployment is routing IP only.

IP reachability is advertised through IP-specific TLVs in Level 1 and Level 2 LSPs. RFC 1195 defines IP Internal Reachability Information and IP External Reachability Information TLVs, which include an IP address and subnet mask along with the associated metric data. Carrying the subnet mask explicitly permits the protocol to advertise variable-length prefixes, rather than being restricted to classful network boundaries. Thus, the statements concerning the required NSAP address and carriage of IP address/subnet information in Level 1 and Level 2 LSP TLVs are accurate. See RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments and Cisco's IS-IS overview.

QUESTION NO: 14

Using the above diagram, which are valid BGP AS_Path Attributes received at Net1 for a route originating from AS65100? (multiple answer)

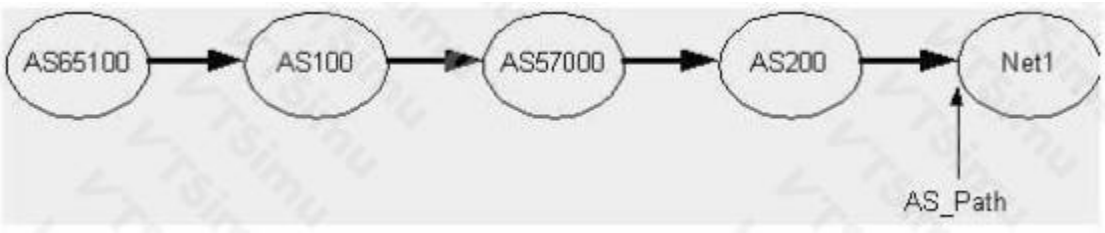

A. 200 57000 100 65100

B. 200 57000 100 100

C. 100 57000 200

D. 200 57000 100

E. 65100 100 57000 200

ANSWER: A D

Explanation:

For an eBGP-learned route, each autonomous system that advertises the route prepends its own ASN to the left side of the AS_PATH. Therefore, when the route originates in AS65100 and is propagated through AS100, AS57000, and AS200

before reaching Net1, the path received at Net1 is **200 57000 100 65100**. The leftmost ASN, 200, identifies the immediately adjacent advertising AS, while the rightmost ASN identifies the origin AS.

200 57000 100 is also valid when the private ASN 65100 has been removed before the route is advertised onward. AS65100 falls within the private 16-bit ASN range of 64512 through 65534. BGP implementations can be configured to remove private ASNs from outbound advertisements, allowing the public-AS transit sequence to remain while omitting AS65100. The ordering of the remaining public ASNs is unchanged because AS_PATH processing prepends ASNs rather than reversing them. This behavior follows the AS_PATH propagation rules in RFC 4271; the private ASN range and its intended handling are specified in RFC 6996.

QUESTION NO: 15

Given the upstream ports above are combined 2:1, which combinations of ports will work in a live network? (multiple answer)

ubr7246#sch
Upstream Port Poll Missed Min Missed Hop Hop Corr Uncor

Port	Status	Rate Poll (ms)	Poll Count	Poll Sample	Poll Pcnt	Thres Pcnt	Period (sec)	FEC Errors	FEC Errors
Cable5/0/U0	16.000 Mhz	416	***	set to fixed frequency	***			5	5
Cable5/0/U1	20.000 Mhz	416	***	set to fixed frequency	***			0	5
Cable5/0/U2	24.000 Mhz	416	***	set to fixed frequency	***			0	2
Cable5/0/U3	30.000 Mhz	416	***	set to fixed frequency	***			0	4
Cable5/0/U4	30.000 Mhz	416	***	set to fixed frequency	***			0	6
Cable5/0/U5	20.000 Mhz	416	***	set to fixed frequency	***			0	0

A. U0 & U2, U1 & U5, U3 & U4

B. U0 & U1, U2 & U3, U4 & U5

C. U0 & U5, U1 & U2, U3 & U4

D. U0 & U4, U1 & U3, U1 & U5

E. U0 & U5, U1 & U3, U2 & U4

ANSWER: B E

Explanation:

Both *U0 & U1, U2 & U3, U4 & U5* and *U0 & U5, U1 & U3, U2 & U4* are valid because each selection creates three independent, non-overlapping upstream port pairs that conform to the 2:1 combining paths shown in the upstream combining diagram. In a live cable network, every physical upstream port must be assigned once within the selected pairing arrangement; this preserves the intended combining relationship and avoids attempting to feed a port into more than one upstream combination simultaneously. The first valid arrangement follows the direct paired layout of the ports, while the second follows the alternate valid paths provided by the combining matrix. Both arrangements therefore deliver a complete and supported mapping for all six upstream ports under 2:1 operation. Cisco cable-platform design and hardware guidance emphasizes using supported RF/port configurations and maintaining the defined upstream signal paths when configuring channel groups and RF combining. See Cisco's Cisco cBR-8 Hardware Installation Guide and CableLabs' DOCSIS specifications for upstream RF and channel-operation context.

QUESTION NO: 16

A network administrator is using debug commands to check the performance of a network.

What steps can the administrator take to ensure that the "debug" will not require too much CPU, or at least that she will not have to reboot the router to disable debug? (multiple answer)

A. Make the debug command as specific as possible

B. Use the max-time parameter of the debug command

C. In configuration mode, enter scheduler interval 15

D. Configure a loopback to channel debug traffic

ANSWER: A C

Explanation:

Making the debug command as specific as possible is essential because debugging can generate substantial control-plane processing and console output. Cisco recommends limiting debugging to the protocol, interface, neighbor, address, or condition being investigated, rather than enabling broad, high-volume debugging. This reduces the number of packets or events examined and logged, lowering CPU consumption while still providing useful diagnostic information. Cisco's guidance on using debug commands emphasizes using them carefully and selectively, especially on production routers: Cisco: Using Debug Commands.

Entering `scheduler interval 15` in configuration mode reduces the maximum interval before IOS gives scheduled processes an opportunity to run. During a CPU-intensive debug, this helps preserve processor time for essential IOS processes and for management-plane activity, such as accepting a command to disable debugging remotely. A smaller scheduler interval therefore improves the administrator's ability to regain control of the router without requiring a reboot. Cisco documents the IOS process scheduler and its configuration commands in its IOS command-reference material: Cisco IOS Process and Scheduler Commands.

QUESTION NO: 17

Which of the following CGMP (Cisco Group Management Protocol) statements is correct?

(multiple answer)

A. CGMP manages multicast traffic in Catalyst 5000 series switches by allowing directed switching of IP multicast traffic.

B. CGMP will switch IP multicast packets to all ports in one specific VLAN.

C. CGMP filtering requires a network connection from the Catalyst 5000 series switch to a router running CGMP.

D. CGMP handles ARP, SAP, UDP, SSAP and DSAP.

ANSWER: A C

Explanation:

"CGMP manages multicast traffic in Catalyst 5000 series switches by allowing directed switching of IP multicast traffic." is correct because Cisco Group Management Protocol is a Cisco-proprietary protocol designed for the Catalyst 5000 family and related legacy Layer 2 switching environments. Its purpose is to prevent a switch from treating IP multicast as traffic that must be flooded throughout a VLAN. Instead, the switch builds multicast forwarding information and forwards multicast frames only toward ports where interested receivers are connected. This is commonly described as directed multicast switching.

"CGMP filtering requires a network connection from the Catalyst 5000 series switch to a router running CGMP." is also correct. CGMP relies on a multicast router to learn receiver membership through IGMP and to communicate CGMP join and leave information to the connected Catalyst switch. The router therefore supplies the membership knowledge that enables the switch to create, maintain, and remove the relevant Layer 2 multicast forwarding entries. The router-to-switch relationship is fundamental to CGMP operation; it allows the switch to constrain multicast distribution while the router continues to perform multicast control-plane functions. Cisco documents CGMP as a legacy multicast-management mechanism and describes its router-driven membership notification model in its CGMP configuration documentation. Additional multicast switching context is available in Cisco's multicast technology support documentation.

QUESTION NO: 18

Look at the router configuration above. If this router has a configuration-register setting of

0x102, select the proper boot sequence:

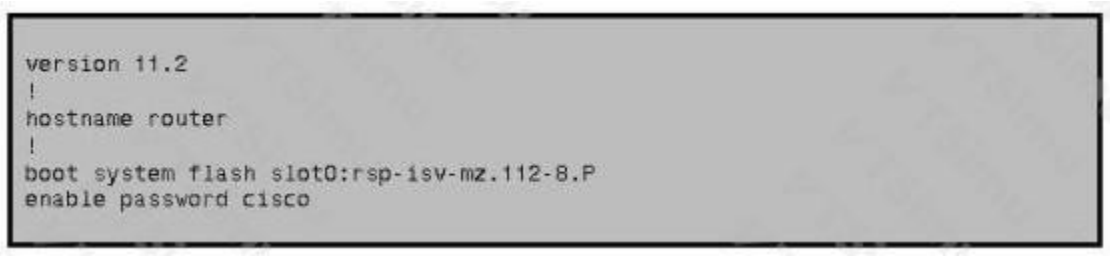

A. The router will try to use the image "rsp-isv-mz.112-8.P" on slot 0, then attempt to boot from a network server, and finally boot from ROM.

B. The router will try use the image "rsp-isv-mz.112-8.P" on slot 0, then attempt to boot from any other valid image in flash, and finally boot from ROM.

C. The router will try to use the image "rsp-isv-mz.112-8.P" on slot 0, and then it will boot from ROM.

D. The router will try to use the image "rsp-isv-mz.112-8.P" on slot 0, and then attempt to boot from a network server.

ANSWER: A

Explanation:

The router will try to use the image "rsp-isv-mz.112-8.P" on slot 0, then attempt to boot from a network server, and finally boot from ROM. This follows the boot-field value of `0x2` in the configuration register `0x102`. A boot field of 2 directs Cisco IOS to use the configured `boot system` commands rather than entering ROM monitor mode or unconditionally loading the first Flash image. Therefore, the explicit image identified in the configuration is the first boot candidate.

If that specified slot 0 image cannot be loaded—for example, because the file is absent, corrupt, or otherwise unusable—the router proceeds through the normal fallback process for this boot method. It attempts a network boot, allowing an IOS image to be obtained from a configured or discoverable network server. If no usable network image can be loaded, the final recovery fallback is the ROM-resident software. The high-order portion of `0x102` does not alter this boot-field selection; the significant boot decision is the low-order value of 2. Cisco documents configuration-register boot-field behavior in its configuration register reference and describes IOS image boot selection in the Cisco IOS booting documentation.

QUESTION NO: 19

In MPLS what is an LSP?

A. Label Selection Pair

B. Label Switched Path

C. Lightweight Signaling Protocol

D. Large Sampling Path

ANSWER: B

Explanation:

In Multiprotocol Label Switching (MPLS), a **Label Switched Path** is the predetermined forwarding path that labeled packets follow through an MPLS network, from an ingress label edge router to an egress label edge router. At the ingress, the router classifies the packet into a forwarding equivalence class and pushes an MPLS label onto it. Each intermediate label switch router uses that label to look up the appropriate forwarding action, typically swapping it for another label and sending the packet to the next hop. At the egress, the final label is removed and the packet is forwarded using its native Layer 3 header.

LSPs enable efficient forwarding based on short, fixed-length labels rather than performing a full IP lookup at every transit hop. They may follow routes established by the network's routing information or may be explicitly established for traffic engineering, including RSVP-TE signaling. Cisco's MPLS overview describes label switching and the ingress-to-egress forwarding behavior, while RFC 3031 defines MPLS architecture and terminology, including label-switched paths. See Cisco: What Is MPLS? and RFC 3031: Multiprotocol Label Switching Architecture.

QUESTION NO: 20

The address field contains: (multiple answer)

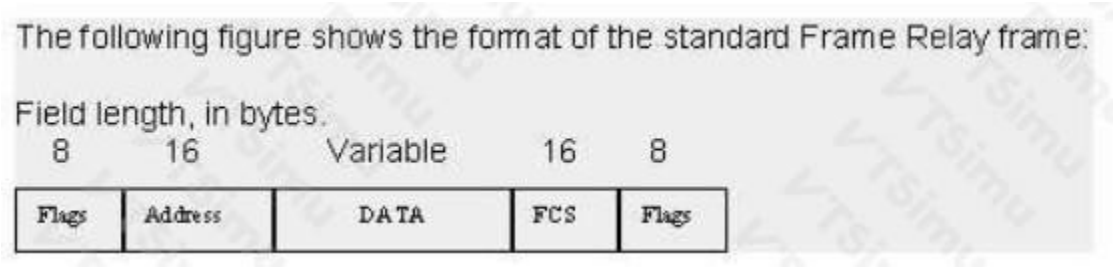

A. The DLCI Value

B. The Extended Address (EA)

C. Congestion Control

D. FCS

ANSWER: A B C

Explanation:

In a Frame Relay frame, the address field identifies the virtual circuit and communicates link-level control information. The DLCI Value is carried in this field and identifies the data-link connection associated with the frame on the local Frame Relay interface. Because DLCIs have local significance, the value selects the appropriate logical PVC or SVC only for that access link.

The Extended Address (EA) bits are also part of the address field. They indicate whether another address octet follows, allowing the Frame Relay address to be represented in its standard two-octet form or extended when required. These bits therefore define the boundary and format of the address information rather than being payload data.

Congestion Control is represented by control bits in the Frame Relay address field, including Forward Explicit Congestion Notification and Backward Explicit Congestion Notification. These bits notify endpoints of congestion in the forward or reverse direction so that traffic sources can respond appropriately. The Discard Eligibility indication is likewise carried in this field and marks lower-priority traffic for potential discard during congestion. Cisco's Frame Relay documentation describes DLCI and congestion-notification operation, while RFC 2427 defines the Frame Relay frame structure and address-field format: Cisco IOS XE Frame Relay Configuration Guide; RFC 2427.