PassQueen
a.k.a dumpsqueen.com

Cisco Service Provider Routing Field Engineer Exam

Cisco 500-230

Version Demo

Total Demo Questions: 5

Total Premium Questions: 52

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cisco Service Provider Routing Architecture	19
Topic 2, Cisco Service Provider Routing Protocols	11
Topic 3, Cisco Service Provider MPLS	22
Total	52

QUESTION NO: 1

Why is an OSPF sham link configured in an MPLS Layer 3 VPN deployment?

A. To make the VPN path appear as an intra-area OSPF path between PE routers

B. To advertise MPLS labels through the OSPF backbone

C. To establish an LDP session between CE routers

D. To redistribute VPNv4 routes into the global OSPF table

ANSWER: A

Explanation:

To make the VPN path appear as an intra-area OSPF path between PE routers is correct. When a customer site has both an OSPF backdoor link and an MPLS VPN path between two sites in the same OSPF area, OSPF can prefer the backdoor path because it is learned as an intra-area route. A sham link is a logical OSPF point-to-point link between PE routers in the same VRF. It causes routes through the VPN backbone to be treated as intra-area routes, allowing normal OSPF cost selection to determine whether the MPLS VPN path or the customer backdoor path is preferred. Cisco documents sham links as the method used to address OSPF route preference in MPLS Layer 3 VPN environments. See the Cisco MPLS Layer 3 VPN OSPF configuration guide.

QUESTION NO: 2

Which route-reflector rule applies to a route learned from a nonclient peer?

A. It can be reflected only to route-reflector clients.

B. It can be reflected only to nonclient peers.

C. It must be reflected to every iBGP peer.

D. It cannot be reflected by a route reflector.

ANSWER: A

Explanation:

It can be reflected only to route-reflector clients is correct. A route reflector breaks the normal iBGP rule that prevents an iBGP speaker from advertising routes learned through iBGP to another iBGP peer. However, route reflection has defined propagation rules to reduce the possibility of loops. A route received from a nonclient peer can be reflected to clients, but it must not be reflected to other nonclients. Routes received from a client can be reflected to both clients and nonclients, subject to normal BGP policy and loop-prevention attributes. RFC 4456 defines the route-reflection rules and use of the ORIGINATOR_ID and CLUSTER_LIST attributes. See RFC 4456, BGP Route Reflection.

QUESTION NO: 3

What priority values do the WDSYSMON Ticker and Watcher threads utilize?

A. Ticker 11, Watcher 254

B. Ticker 10, Watcher 63

C. Ticker 11, Watcher 63

D. Ticker 10, Watcher 254

ANSWER: B

Explanation:

WDSYSMON uses a Ticker thread priority of 10 and a Watcher thread priority of 63. WDSYSMON is part of the Cisco IOS system-monitoring and watchdog infrastructure, where distinct threads perform periodic timing work and monitor the system's ability to continue scheduling and servicing critical activity. The Ticker thread runs at priority 10 for its recurring timing function, while the Watcher thread runs at priority 63 so that its watchdog-monitoring work receives the intended scheduling precedence. These numeric values are implementation-specific thread priorities shown in Cisco IOS process and thread diagnostic information; they are not user-configurable QoS, routing, or interface-priority values. When validating operational behavior, engineers commonly use IOS process and platform diagnostic commands to inspect process activity and scheduling state, correlating the thread names with their expected priorities. Cisco's IOS process-management command documentation provides the relevant operational context for examining processes, while Cisco's service-provider routing training catalog identifies the platform and troubleshooting focus associated with this exam domain. See Cisco IOS Process Commands and Cisco 500-230 exam information.

QUESTION NO: 4

Which is required on the CE router in an L3VPN solution?

A. Configure the PE-CE Routing protocol

B. Define the VPN by creating a VRF

C. Assign the PE facing interface to a VRF

D. Configure MPLS forwarding on the interface to the PE router

ANSWER: A

Explanation:

Configure the PE-CE Routing protocol is correct because the customer-edge router must exchange customer-route reachability information with its directly connected provider-edge router. This CE-to-PE routing relationship allows the provider-edge router to learn prefixes for the customer site and advertise the appropriate reachable prefixes back toward the CE. In an MPLS Layer 3 VPN, the CE operates as a conventional IP router at the service-provider boundary: it forwards customer traffic toward the PE using IP routing and does not need to participate in the provider MPLS core. Cisco supports several PE-CE routing choices, including static routing and dynamic protocols such as BGP, OSPF, EIGRP, and RIP; exam terminology commonly describes this required function as configuring the PE-CE routing protocol. The selected routing method must provide the required customer-prefix exchange and correctly support the intended VPN connectivity. Cisco's MPLS L3VPN documentation describes PE-CE routing as the mechanism by which customer routes are exchanged between the customer and provider networks. See the Cisco IOS MPLS Layer 3 VPN Configuration Guide and Cisco's MPLS VPN overview.

QUESTION NO: 5

What is the default MP-BGP address-family on IOS-XR?

A. IPv6 unicast

B. There is no default address-family

C. IPv4 unicast

D. VPNv4

ANSWER: C

Explanation:

IPv4 unicast is the default BGP address family on Cisco IOS XR. When a BGP process is created, IOS XR uses the IPv4 unicast AFI/SAFI as its default context for exchanging conventional IPv4 routing information. This behavior provides the base

BGP routing capability without requiring an additional multiprotocol address-family selection for ordinary IPv4 prefixes. In IOS XR configuration, address-family submodes are used to enable and apply policy, redistribution, network advertisement, and neighbor-specific settings for particular route families. IPv4 unicast is therefore the expected default address family, while specialized multiprotocol families are configured when the deployment requires their specific route types and semantics. This distinction matters operationally because BGP neighbor activation and policy must be applied in the relevant address-family context; a session may be established, but routes are exchanged only for families configured and activated for that peer. Cisco's IOS XR BGP documentation describes BGP address-family configuration and the IPv4 unicast context, and Cisco's IOS XR software support portal provides the applicable release-specific configuration guides. See Cisco IOS XR BGP Configuration Guide and Cisco IOS XR Software documentation.