DUMPSARENA

Extreme Networks Specialist

Extreme Networks EW0-300

Version Demo

Total Demo Questions: 17

Total Premium Questions: 178

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

The following modes of operation are supported with the Network login feature. (choose all that apply)

A. Secure mode

B. Campus mode

C. Quick Mode

D. ISP mode

ANSWER: A B C D

Explanation:

ExtremeXOS Network Login (NetLogin) provides several operational modes so authentication behavior can be adapted to the type of access environment being deployed. Secure mode is intended for tightly controlled access and restricts forwarding until the applicable authentication requirements have been met. Campus mode supports the common enterprise-access design, where a user device connects to an access port and is authenticated for network access. ISP mode is designed for environments in which more than one authenticated client may be present behind a physical port, such as provider or shared-access deployments. Quick Mode is also a supported NetLogin mode and is used where a simpler, faster login behavior is appropriate. These modes are selected on NetLogin-enabled ports to align the switch's authentication and access-control behavior with the deployment model. Extreme's documentation portal provides the applicable ExtremeXOS software guides and command references, including the NetLogin configuration material: Extreme Networks Documentation. The product documentation catalog is also available through Extreme Networks Support Documentation, where the relevant ExtremeXOS version can be selected for the precise syntax and mode behavior.

QUESTION NO: 2

All Extreme switch products support the standard 802.1p priority bits that are part of a tagged Ethernet frame.

A. True

B. False

ANSWER: A

Explanation:

True is correct. IEEE 802.1p priority is carried in the three-bit Priority Code Point (PCP) field within an IEEE 802.1Q VLAN tag. The field provides eight traffic priorities, allowing a switch to classify tagged traffic and apply quality-of-service handling such as mapping frames to forwarding queues, preserving priority markings, and using priority information in QoS policy decisions. Support for these standard priority bits is a fundamental Ethernet switching and QoS capability across Extreme Networks switch platforms. Although the precise QoS commands, queue mappings, and supported policy features can differ between operating systems and hardware families, the standard tagged-frame priority field itself is recognized and supported. This statement is specifically about support for the 802.1p bits present in a tagged Ethernet frame; it does not require identical QoS configuration syntax or identical queue behavior on every platform. IEEE identifies 802.1Q as the standard covering bridged networks and VLAN tagging, including traffic-class handling associated with the tag. See the IEEE 802.1Q standard overview and Extreme Networks' product documentation portal for platform-specific QoS and VLAN-tagging implementation details.

QUESTION NO: 3

Which of the following statements about VRRP advertisements are correct? (choose all that apply)

A. They are sent by the Master router

B. They are used by Backup routers to detect Master failure

C. They are sent only after an OSPF adjacency becomes Full

D. They use the virtual MAC address as the source IP address

ANSWER: A B

Explanation:

VRRP advertisements are sent by the current Master router to announce that it owns the virtual router and remains operational. Backup routers receive these advertisements and use them to maintain their master-down timer. If the expected advertisements cease, an eligible Backup router can transition to Master. The virtual MAC address is shared by the virtual router participants, but it is not used as the source IP address of VRRP advertisements. VRRP is a gateway-redundancy protocol and does not depend on OSPF adjacency formation. See RFC 5798, Virtual Router Redundancy Protocol.

QUESTION NO: 4

Which of the following statements describe OSPF virtual links? (choose all that apply)

A. It provides logical connectivity to area 0.

B. It is configured between two Area Border Routers.

C. It uses a non-backbone transit area.

D. It can use a stub area as its transit area.

ANSWER: A B C

Explanation:

OSPF virtual links provide logical backbone connectivity through a non-backbone transit area. They are configured between two Area Border Routers and are commonly used when an area cannot be physically attached directly to area 0, or as a temporary solution during a topology migration. The virtual-link endpoints must be ABRs, and the transit area cannot be a stub area because the virtual link requires the normal inter-area routing information carried through that area. A virtual link does not make a non-ABR router into an ABR, and it is not a replacement for an OSPF neighbor adjacency on every interface. OSPF virtual-link behavior is defined in RFC 2328. Refer to the Extreme Networks Documentation portal for Extreme switch configuration syntax.

QUESTION NO: 5

The normal behavior of a Router is to forward Unicast packets and block broadcast traffic.

A. True

B. False

ANSWER: A

Explanation:

True is correct. A router operates at Layer 3 and makes forwarding decisions for unicast IP packets by examining the destination IP address and consulting its routing table. When it has a matching route and the packet is otherwise permitted, it forwards that traffic from one network interface to another, allowing communication between distinct IP subnets.

Broadcast traffic has a fundamentally different scope. A Layer 2 broadcast is intended for all devices in its local broadcast domain, so routers do not normally bridge it from one routed interface to another. This separation is an important function of routing: each router interface defines the boundary of an IP broadcast domain. Consequently, broadcasts such as those used for local address resolution remain local rather than being propagated throughout an internetwork. The IETF router requirements specify that routers generally must not forward IP limited broadcasts, reinforcing this normal routed-network behavior. Directed broadcasts are a specialized case and are commonly disabled for security; they do not alter the stated normal behavior.

See RFC 1812: Requirements for IP Version 4 Routers and the Extreme Networks Documentation Portal for platform-specific routing and broadcast configuration guidance.

QUESTION NO: 6

How are multicast packets forwarded in a Layer 2 environment with IGMP snooping enabled?

A. The packets flood to all switch ports.

B. The switch forwards the packets to member ports only.

C. IGMP snooping does not affect the forwarding of multicast packets.

D. None of the above.

ANSWER: B

Explanation:

With IGMP snooping enabled, the switch examines IGMP membership reports and leave messages exchanged between multicast receivers and the multicast router. It uses that information to build multicast forwarding state for the VLAN, associating a multicast group with the switch interfaces on which interested receivers are present. Consequently, multicast data for that group is forwarded only through the relevant member-facing ports rather than being treated as unknown multicast and sent broadly throughout the VLAN. This reduces unnecessary traffic on devices that have not joined the group, preserves available bandwidth, and prevents endpoints from receiving multicast streams they did not request. The forwarding state is refreshed through normal IGMP query and report activity, so active group membership remains current. In operational deployments, the switch also maintains the appropriate path toward the multicast router so membership control and multicast delivery can continue across the Layer 2 domain. This behavior is the purpose of IGMP snooping and aligns with Layer 2 multicast-snooping guidance in RFC 4541; Extreme platform-specific configuration and operational documentation is available through the Extreme Networks Documentation Portal.

QUESTION NO: 7

Which of the following actions can a QoS policy apply to traffic that matches its classification criteria? (choose all that apply)

A. Assign traffic to a forwarding queue

B. Remark a packet QoS value

C. Apply a rate limit or meter

D. Create an OSPF adjacency

E. Form a load-sharing group

ANSWER: A B C

Explanation:

A QoS policy associates a traffic classification with a defined forwarding treatment. Depending on the switch platform and configured QoS profile, matching traffic can be assigned to a forwarding queue, marked or remarked with a QoS value, subject to a meter or rate limit, and given an appropriate bandwidth or priority treatment. These actions are implemented in the switch forwarding hardware on supported platforms so that traffic can be handled at high speed. A QoS policy does not dynamically create IP routing adjacencies or establish a Layer 2 load-sharing group; those are separate networking functions. The available policy actions vary by hardware family and software release. Refer to the Extreme Networks Documentation portal for the QoS feature guide and command reference for the deployed platform.

QUESTION NO: 8

Manual replication of saved images and configurations from the master MSM64i to the slave MSM64i on the BlackDiamond 6804 is done with the synchronize command. However execution of this command does not replicate the run-time configuration.

You must use the save configuration command to store the run-time configuration first.

A. True

B. False

ANSWER: A

Explanation:

True is correct. On a redundant BlackDiamond management configuration, the `synchronize` command is intended to copy the stored software image and saved configuration information from the active/master MSM to its standby/slave peer. The running configuration is the active, in-memory state used by ExtremeXOS; it is not itself a saved configuration file that can be transferred as part of the normal synchronization process. Therefore, any changes made since the last configuration save must first be committed to the switch's saved configuration by using `save configuration`. Once saved, that configuration becomes available for replication with `synchronize`, allowing the standby MSM to have a usable, current startup configuration if it later assumes the active role. This distinction is operationally important: saving before synchronization ensures that intentional runtime changes are preserved through both a reboot and a management-module failover. Extreme's documentation resources for ExtremeXOS commands, configuration management, and chassis redundancy are available through the Extreme Networks Documentation portal and the Extreme Networks Support Portal.

QUESTION NO: 9

How does an IGMP router learns about the existence of hosts in multicast groups

A. By receiving a response on an IGMP membership query

B. By looking into the ARP cache

C. There is no dynamic learning process, instead static FDB entries are used

D. None of the above

ANSWER: A

Explanation:

An IGMP router learns that hosts are interested in receiving traffic for a multicast group by sending IGMP Membership Queries and receiving IGMP Membership Reports in response. Hosts that have joined a multicast group respond to a query, allowing the router to maintain group-membership state for the network segment. This state is used by the multicast-routing process, or by an IGMP-snooping switch working with the router's query process, to determine where multicast traffic should be forwarded. Membership is refreshed periodically through the query/report exchange. Hosts can also send unsolicited membership reports when they initially join a group, which lets the router learn of a new membership without waiting for the next general query. In IGMPv2, hosts may additionally send a Leave Group message when departing, after which the router uses group-specific queries to verify whether any members remain. Thus, receiving a response on an IGMP membership query is the fundamental mechanism described by the question. The IETF's IGMPv2 specification defines the router query and host report behavior in RFC 2236. For the current IGMPv3 protocol behavior and membership-reporting model, see RFC 3376.

QUESTION NO: 10

Which of the following are valid reasons to configure an EAPS protected VLAN in an EAPS domain? (choose all that apply)

A. To provide rapid restoration for user traffic carried by the VLAN

B. To keep the VLAN loop-free during normal ring operation

C. To carry EAPS health-check packets as the control VLAN

D. To make the VLAN available on non-ring ports automatically

ANSWER: A B

Explanation:

Protected VLANs are the customer or data VLANs for which EAPS provides loop prevention and rapid restoration within an EAPS domain. They are associated with the domain so that the master and transit nodes can control forwarding correctly following a ring failure. The control VLAN is not a protected VLAN; it carries EAPS control traffic and must remain available around the ring. Similarly, an unrelated VLAN that has not been associated with the domain does not receive EAPS protection from that domain. Extreme Networks EAPS configuration guides are available through the Extreme Networks Documentation portal.

QUESTION NO: 11

When configuring a Summit 5i for OSPF routing, it is recommended to manually configure the OSPF router-id rather than setting the router-id set to automatic.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. An OSPF router ID is a 32-bit value that uniquely identifies the router in the OSPF domain and is used in OSPF neighbor relationships, link-state advertisements, and the router-LSA. On a Summit 5i, explicitly configuring this value is recommended because it gives the router a stable, predictable identity independent of interface-address changes, interface availability, or changes to the address that might otherwise be selected automatically. A stable router ID is especially important in operational networks: a router-ID change causes the OSPF process to use a different identity, which can require adjacency and LSDB recalculation and makes troubleshooting, topology interpretation, and log correlation more difficult. Selecting a deliberate, unique address-like value for each router also makes the OSPF design easier to document and maintain. OSPF defines the router ID as a fundamental identifier carried in protocol operation, as described in RFC 2328, OSPF Version 2. Extreme Networks documentation and software resources are available through the Extreme Networks Documentation Portal; consult the Summit/ExtremeWare OSPF configuration material for the platform-specific router-ID command syntax.

QUESTION NO: 12

Which of the following are benefits of using RADIUS with Network Login? (choose all that apply)

A. Centralized user authentication

B. Centralized authorization of access sessions

C. Session accounting information

D. Encryption of all endpoint data traffic after login

ANSWER: A B C

Explanation:

RADIUS provides centralized authentication, authorization, and accounting services for Network Login deployments. A switch can send endpoint credentials to a RADIUS server instead of relying only on locally configured user information, allowing a common identity store and policy source to be used across many access switches. RADIUS authorization can

return attributes that influence the access session, including VLAN or policy-related information where supported by the configured Network Login method and software release. RADIUS accounting can record session events for operational and audit purposes. RADIUS does not encrypt all user traffic after successful login; endpoint data protection requires separate controls such as MACsec, IPsec, or application-layer encryption. For RADIUS protocol details, see RFC 2865. Extreme Network Login documentation is available from the Extreme Networks Documentation portal.

QUESTION NO: 13

In SLB the forwarding mode is the method a switch uses to forward traffic to the virtual servers. On the Summit 7i switch four forwarding modes are supported, select the two forwarding modes that are hardware based.

A. Transparent mode

B. Translation mode

C. Port Translation mode

D. GoGo mode

ANSWER: A D

Explanation:

Transparent mode and GoGo mode are the hardware-based Server Load Balancing forwarding modes on the Summit 7i. In transparent mode, the switch forwards traffic toward the selected real server while preserving the client-facing IP addressing model. This design allows the SLB forwarding decision and the associated Layer 2 forwarding changes to be performed by the switch hardware, providing high-throughput handling for eligible flows. GoGo mode is also designed for hardware forwarding. It uses a forwarding model in which traffic is sent to the selected server through hardware-based address handling, while the server and network are configured to support the virtual service addressing and return-path behavior required by this mode. Because both modes use forwarding operations supported in the Summit 7i hardware path, they are the two modes identified as hardware based in the platform SLB documentation. Correctly identifying the hardware modes matters when designing for performance, because they avoid relying on the software forwarding path for each load-balanced packet. Extreme Networks maintains product documentation through its Documentation Portal; legacy Summit software and documentation are also available through the Extreme Networks Support portal.

QUESTION NO: 14

Which of the following statements fit the description of PIM sparse mode? (choose three that apply)

A. PIM sparse mode is suitable for network with minimal multicast deployment.

B. PIM sparse mode builds shared trees based on a common Rendez-Vous Point

C. In PIM sparse mode the RPF check depends on the tree type

D. An explicit join is not required in PIM sparse mode.

ANSWER: A B C

Explanation:

PIM sparse mode is designed for multicast environments in which receivers are relatively sparse and traffic should be forwarded only where it has been explicitly requested. This makes it appropriate for networks with limited or selective multicast deployment, rather than assuming that every part of the network needs multicast traffic. PIM-SM initially uses a shared distribution tree rooted at a Rendez-Vous Point (RP). Sources register with the RP, and receivers join the group through the RP, allowing multicast distribution to begin without network-wide flooding. PIM-SM can subsequently transition to a shortest-path tree rooted at the source when that is beneficial. Reverse Path Forwarding behavior therefore depends on the active tree: packets received on the RP tree are checked against the route toward the RP, while packets received on a source tree are checked against the route toward the multicast source. These mechanisms provide loop-free forwarding while supporting both the shared-tree and source-tree phases of sparse-mode operation. The PIM sparse-mode protocol

behavior, including RP trees, source trees, joins, and RPF rules, is specified in RFC 4601. Extreme Networks product documentation and configuration resources are available through the Extreme Networks Documentation Portal.

QUESTION NO: 15

IS-IS is a designed to work in conjunction with

A. ISO 8473 (The ISO Connectionless Network Layer Protocol)

B. ISO 9542 (The ISO End System to Intermediate System Protocol)

C. Both A and B

D. None of these

ANSWER: C

Explanation:

Both A and B is correct because IS-IS originated as the routing protocol for the ISO Connectionless Network Protocol environment. ISO 8473 defines CLNP, the connectionless network-layer protocol that provides the packet format and network-layer service model within the original OSI architecture. IS-IS supplies the intermediate-system routing function for that environment: it distributes topology and reachability information so intermediate systems can calculate paths and forward CLNP packets.

ISO 9542, commonly called ES-IS, complements this design by defining communication between end systems and intermediate systems. It supports functions such as end-system and router discovery and helps an end system identify an appropriate intermediate system for traffic outside its local area. Thus, the original ISO routing architecture uses IS-IS alongside both the CLNP specification and the ES-IS protocol, rather than treating IS-IS as an isolated protocol. IS-IS was later extended for IP routing through Integrated IS-IS, while retaining its core link-state routing architecture. The historical relationship and protocol context are documented in RFC 1142; IS-IS operation and its IP extensions are further specified in RFC 1195.

QUESTION NO: 16

What is true of a Rendezvous Point in PIM-DM?

A. It is the meeting point for PIM sources and receivers.

B. It is the meeting point for PIM sources and the bootstrap router.

C. In PIM-DM you don't need a Rendezvous point.

D. It is the distribution point for multicast streams from the Mbone into an AS.

ANSWER: C

Explanation:

In PIM-DM you don't need a Rendezvous point. Protocol Independent Multicast–Dense Mode operates using a dense-mode, source-based multicast distribution model. When a multicast source begins transmitting, PIM-DM initially floods multicast traffic throughout the PIM-enabled network. Routers that have no downstream receivers then prune the unwanted traffic, and they can later rejoin when receivers appear. This behavior builds source-specific forwarding state based on the location of the source and the multicast group, rather than building a shared tree rooted at a Rendezvous Point. A Rendezvous Point is a core function of PIM Sparse Mode, where it supports initial shared-tree forwarding and source registration for multicast groups. Because PIM-DM neither uses PIM-SM shared trees nor source registration to an RP, an RP does not have to be configured or discovered for dense-mode multicast operation. Extreme's multicast configuration guidance identifies PIM-DM and PIM-SM as distinct operating modes and describes RP configuration in the PIM-SM context. For additional background, see the Extreme Networks Documentation portal and the PIM Dense Mode specification (RFC 3973).

QUESTION NO: 17 - (SIMULATION)

SIMULATION

Given router interfaces 10.0.0 1.24, 10.10.10.10/24, 192.168.0.1/24 and 192.168.255.254/24 on a Summit 7i switch. What will be the OSPF router ID upon switch reboot, given the following CLI command config ospf routerid automatic?

ANSWER: See the explanation for the answer

Explanation:

OSPF router ID: `192.168.255.254`

With `configure ospf routerid automatic`, the Summit selects the highest IP address assigned to its interfaces when it initializes OSPF. Of the listed interface addresses—`10.0.0.1`, `10.10.10.10`, `192.168.0.1`, and `192.168.255.254`—the numerically highest address is `192.168.255.254`.