PassQueen
a.k.a dumpsqueen.com

Brocade Certified Network Engineer 2012

Brocade 150-130

Version Demo

Total Demo Questions: 15

Total Premium Questions: 153

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

What are two valid BPDU types in Spanning Tree 802.1D? (Choose two.)

A. Configuration BPDU

B. Root Override BPDU

C. Failure Notification BPDU

D. Topology Change Notification BPDU

ANSWER: A D

Explanation:

IEEE 802.1D Spanning Tree Protocol uses two BPDU formats: Configuration BPDUs and Topology Change Notification BPDUs. A Configuration BPDU carries the information switches require to elect and identify the root bridge, calculate the path cost toward it, select root and designated ports, and maintain the spanning-tree topology. It includes fields such as the root identifier, root path cost, bridge identifier, port identifier, and timer values. Bridges transmit Configuration BPDUs periodically so that the tree remains current and loops are prevented as the Layer 2 topology changes.

A Topology Change Notification BPDU is generated when a non-edge port changes state in a way that can affect reachability. It is propagated toward the root bridge through root ports. After acknowledgment and notification, the topology-change process causes bridges to age learned forwarding entries more quickly, helping traffic converge to the new physical topology rather than continue using stale MAC-address entries. These are the BPDU types defined for classic 802.1D STP operation. See the IEEE 802.1 working-group overview at IEEE 802.1 and Cisco's STP protocol description at Understanding Spanning-Tree Protocol.

QUESTION NO: 2

Which two statements about ingress extended ACLs are correct? (Choose two.)

A. They use extended ACL IDs 200 - 299.

B. They save network resources by reducing traffic.

C. They eliminate denial of service (DoS) attacks.

D. They block unwanted traffic or users.

ANSWER: B D

Explanation:

Ingress extended access control lists are applied to packets as they enter an interface, allowing the switch to evaluate and filter traffic before it is forwarded further into the network. "They save network resources by reducing traffic." is correct because an ingress rule can discard traffic that does not meet the configured policy at the network edge. Preventing unwanted packets from traversing switching paths, reaching downstream interfaces, or being processed by destination devices conserves bandwidth and related device resources.

"They block unwanted traffic or users." is also correct. Extended access control lists support granular matching criteria, including protocol and transport-layer information as well as source and destination addressing. This makes them suitable for enforcing an access policy that permits authorized communications while denying traffic from prohibited sources or traffic directed at protected services. Applied inbound, that policy takes effect at the earliest practical point in the forwarding path. Brocade FastIron documentation describes access control lists as a mechanism for controlling packet traffic and applying packet-filtering policies; see the FastIron Configuration Guide and the Broadcom switching product documentation portal.

QUESTION NO: 3

An engineer configures a static route for 10.20.30.0/24 with a next hop of 192.0.2.1. Before the route can become active, what must be true?

A. The next hop 192.0.2.1 must be reachable.

B. OSPF must advertise 10.20.30.0/24.

C. The destination network must respond to ping.

D. The route must have an administrative distance of 0.

ANSWER: A

Explanation:

The router must have a **reachable route to 192.0.2.1**. A static route that specifies a next-hop address requires recursive next-hop resolution: the device must be able to determine an outgoing interface and Layer 2 adjacency for the next hop. This can be through a directly connected network or another valid route in the routing table.

A static route does not become usable merely because it is configured. The router must be able to forward packets to its specified next hop. General IPv4 router forwarding requirements are described in RFC 1812.

QUESTION NO: 4

On Brocade FastIron, TurboIron, and NetIron platforms, which command saves the running configuration to the flash's startup configuration?

A. write memory

B. copy running-config startup-config

C. config commit

D. config save

ANSWER: A

Explanation:

`write memory` is the FastIron-family CLI command used to make the active running configuration persistent. Configuration changes entered during a management session take effect immediately in the running configuration, which resides in volatile memory. To ensure those changes survive a reload or power cycle, the administrator must save that active configuration into the startup configuration stored in flash. Issuing `write memory` performs that save operation, allowing the switch to load the saved configuration automatically during its next boot sequence.

This command is part of the traditional Brocade/Foundry FastIron command-line workflow and is appropriate for the FastIron, TurboIron, and NetIron platform families named in the question. Administrators commonly issue it after completing and verifying a configuration change, particularly before planned maintenance or a reboot. The saved startup configuration is the durable boot-time configuration, while the running configuration remains the currently applied state. See the FastIron Management Configuration Guide and the FastIron Command Reference for the configuration-save command and startup configuration behavior.

QUESTION NO: 5

Which two statements are true about QoS when implemented on a Brocade FastIron platform?

(Choose two.)

A. Packets with lower priority classifications are given a higher precedence for forwarding.

B. Packets with higher priority classifications are given a higher precedence for forwarding.

C. Stackable devices use CoS 0 for stacking communication.

D. Stackable devices use CoS 7 for stacking communication.

ANSWER: B D

Explanation:

FastIron QoS uses packet classification to assign traffic to service classes, then applies queue scheduling so that traffic assigned to a higher priority class receives preferential forwarding treatment during periods of contention. Therefore, “Packets with higher priority classifications are given a higher precedence for forwarding.” correctly describes the fundamental purpose of QoS on the platform. Priority is meaningful when egress resources are oversubscribed: the switch services the more important traffic ahead of traffic placed in lower-priority queues, helping latency-sensitive applications receive the service level intended by the QoS policy.

“Stackable devices use CoS 7 for stacking communication.” is also correct. In FastIron stacking environments, Class of Service 7 is reserved for stack-control traffic. Reserving the highest CoS value for this internal communication ensures that stack-management and control messages retain priority and that stack operation remains reliable even when ordinary user traffic is congested. QoS configuration and queue behavior must account for this reserved use rather than treating every CoS value as generally available for application traffic. RUCKUS provides the FastIron QoS documentation and platform manuals at FastIron Documentation and product support resources at RUCKUS Support.

QUESTION NO: 6

You are configuring a keep-alive VLAN for a Multi-Chassis Trunking deployment on two new

Brocade MLXe chassis. You notice that STP and VLANs are not valid configuration options.

What have you overlooked during the configuration?

A. You need to include no route-only.

B. You need to include trunk deploy.

C. You need to include deploy.

D. You need to include route-only.

ANSWER: A

Explanation:

You need to include `no route-only`. In an MLXe Multi-Chassis Trunking configuration, the keep-alive connection must operate as a Layer 2 VLAN path so that the peer chassis can exchange MCT keep-alive traffic independently of routed forwarding. The `route-only` setting places the applicable interface or context into routed-only operation, which prevents Layer 2 configuration features—including VLAN and Spanning Tree Protocol commands—from being available. Entering `no route-only` restores Layer 2 capability, allowing the keep-alive VLAN to be created and its required STP behavior to be configured. This is important because the keep-alive VLAN is part of the MCT peer-health design: it provides a dedicated mechanism for the two chassis to detect connectivity and peer status. After Layer 2 operation is enabled, configure the designated keep-alive VLAN consistently on both MCT peers and ensure the physical links carrying it are available as required by the deployment design. Brocade/CommScope product documentation and software guides are available through the CommScope documentation portal; MLXe platform resources can also be accessed from the CommScope networking systems page.

QUESTION NO: 7

Which two Brocade products are recommended for use in a campus access environment?

(Choose two.)

A. BrocadeMLXe

B. Brocade TurboIron 24X

C. Brocade FCX

D. Brocade ICX 6430

ANSWER: C D

Explanation:

Brocade FCX and Brocade ICX 6430 are appropriate campus-access products because they were designed as fixed-configuration Ethernet switches for connecting endpoint devices at the network edge. The Brocade FCX family provides campus access switching capabilities such as Power over Ethernet for phones and wireless access points, Layer 2 and Layer 3 services, resilient stacking, and centralized management. These characteristics support the high port density, endpoint connectivity, and operational simplicity expected in wiring-closet deployments.

The Brocade ICX 6430 is likewise a compact, stackable access-switch family intended for small to midsize campus access deployments. It provides Gigabit Ethernet access ports, PoE/PoE+ models, uplink connectivity, and stacking capabilities, allowing multiple switches to operate as a logical unit. This makes it well suited to extending access connectivity while maintaining consistent policy and management. The product name in the supplied choice has been corrected to “Brocade ICX 6430.” Brocade’s FCX product information and the ICX 6430 documentation describe these access-layer-oriented capabilities: Brocade FCX Switches Data Sheet and ICX 6430 documentation.

QUESTION NO: 8

You have been informed a group of remote employees are unable to authenticate with the RADIUS server. From the Brocade FCX router, you enable the debug dot1x events command.

Which three events will be viewed? (Choose three.)

A. authentications failed

B. authentications succeeded

C. ACLs requested by the RADIUS server.

D. standard ACL requests

E. extended ACL requests

ANSWER: A B C

Explanation:

The `debug dot1x events` command provides event-level diagnostic output for the IEEE 802.1X authentication process on a FastIron/FCX device. This output includes authentication state results, so **authentications failed** is visible when an exchange does not complete successfully and **authentications succeeded** is visible when the supplicant is successfully authenticated through the RADIUS process. Seeing both outcomes is especially useful during troubleshooting because it lets the administrator correlate a user’s result with the relevant 802.1X and RADIUS interaction.

The event debug also reports authorization information delivered with RADIUS processing, including **ACLs requested by the RADIUS server.** A RADIUS server can return authorization attributes that instruct the switch to apply a policy after the user has authenticated. Reporting the ACL request in the event trace helps confirm that the FCX received the authorization policy and supports troubleshooting cases in which authentication succeeds but the user receives unexpected access permissions. Brocade FastIron documentation groups these diagnostics under its 802.1X and RADIUS configuration and monitoring features. See the Ruckus documentation portal and the Ruckus support document library for FastIron command and security configuration references.

QUESTION NO: 9

On Brocade FastIron, TurboIron, and NetIron platforms, which command saves the running configuration to the flash's startup configuration?

A. write memory

B. copy running-config startup-config

C. config commit

D. config save

ANSWER: A

Explanation:

`write memory` is the FastIron, TurboIron, and NetIron CLI command used to save the active running configuration as the startup configuration in flash memory. Device configuration changes made during a CLI session affect the running configuration immediately, but they are not automatically retained through a reboot. Issuing `write memory` copies those current settings into the startup configuration file, allowing the switch or router to load the saved configuration during its next initialization. This is therefore the appropriate command after completing configuration work that must persist. Brocade documentation describes the `write memory` command as saving the current configuration to the startup configuration in flash. See the Brocade FastIron configuration documentation at FastIron Administration Guide and the Brocade NetIron command reference at NetIron documentation.

QUESTION NO: 10

When is an OSPF virtual link used?

A. When an area has no direct connection to the backbone area.

B. When an area needs access to an ASBR.

C. When an area is configured as a stub area.

D. When an area needs to connect to a non-backbone area.

ANSWER: A

Explanation:

An OSPF virtual link is used when an OSPF area does not have a direct physical connection to the backbone, Area 0. OSPF requires all non-backbone areas to maintain logical connectivity to Area 0 so that inter-area routing information can be exchanged correctly. A virtual link provides that logical backbone connection by joining two area border routers across a shared non-backbone transit area. Operationally, the virtual link behaves as though the participating routers have an Area 0 adjacency, even though the packets traverse the transit area. This design is commonly used as a temporary solution during migration or where topology constraints prevent a direct backbone attachment. The transit area must be a normal area that can carry the required OSPF information, and the virtual-link endpoints are configured using the router IDs of the two ABRs. The OSPF specification describes virtual links as a means of restoring backbone continuity and attaching an area that is otherwise disconnected from the backbone. See RFC 2328, section 15: Virtual Links and the related backbone-routing requirements in RFC 2328, section 3.7.

QUESTION NO: 11

You have configured your Brocade device as shown in the exhibit.

Brocade#sh ip

Switch IP address: 10.200.128.116

Subnet mask: 255.255.255.0

Default router address: 10.200.128.1

TFTP server address: 10.200.46.57

Configuration filename: Brocade

Image filename: None

IP MTU: 1500

Brocade#traceroute 10.200.128.1

Type Control-c to abort

Tracing the route to IP node 10.200.128.1 from 1 to 30 hops

1 < 1 ms < 1 ms < 1 ms 10.200.128.1

Which statement is correct regarding interface Ethernet 10 on this device having inbound traffic from 172.17.1.1/24 to 172.17.2.1/24 on port 21?

A. It would be permitted.

B. It would be denied by the implicit deny statement.

C. It would be denied by FTP port 21.

D. This ACL only applies to outbound traffic.

ANSWER: A

Explanation:

It would be permitted. The access-control-list entry represented by the exhibit matches the stated inbound flow: traffic sourced from the 172.17.1.0/24 network and destined for the 172.17.2.0/24 network. On Brocade FastIron devices, an IP ACL is evaluated in sequence for packets entering an interface when it is applied in the inbound direction. Once a packet matches a permit entry, processing stops and the switch forwards the packet normally, subject to the usual Layer-3 forwarding decision. The source host address 172.17.1.1 is within the specified source prefix, and the destination address 172.17.2.1 is within the specified destination prefix. Port 21 identifies FTP control traffic, but it does not prevent this packet from being forwarded when the applicable ACL match permits the IP flow. The management IP information and successful one-hop traceroute demonstrate reachability to the configured default router; they do not alter ACL matching for the Ethernet 10 packet. Brocade FastIron documentation describes ACL binding and packet filtering behavior in its security configuration guidance. See the Brocade FastIron Security Configuration Guide and the Brocade networking documentation portal.

QUESTION NO: 12

Which two statements about ICL links in Brocade Multi-Chassis Trunk (MCT) implementations are true? (Choose two.)

A. ICL links connect the two MCT chassis and allows MCT clients to see both devices as a single device.

B. ICL links connect the MCT client devices to the MCT logical chassis.

C. ICL links are necessary for the two devices to form an MCT cluster.

D. ICL links allow the MCT logical chassis to communicate with a non-MCT device.

ANSWER: A C

Explanation:

ICL links connect the two MCT chassis and allows MCT clients to see both devices as a single device. In a Brocade MCT design, the two peer switches operate together as an MCT logical chassis from the perspective of a dual-attached downstream client. The client can use one port channel, with member links terminating on separate physical MCT peers, while retaining active forwarding and resiliency. The Inter-Chassis Link carries the control and data-plane coordination required between those peers so that the logical-chassis behavior is maintained.

ICL links are necessary for the two devices to form an MCT cluster. The ICL is the dedicated inter-peer connection used by the MCT peers to synchronize forwarding-related information and to handle traffic that must traverse between the chassis. Consequently, it is a foundational component of an MCT deployment rather than a client-facing attachment. Brocade documentation describes MCT as a technology for connecting a client through a multi-chassis trunk to two VDX switches, with the peer relationship and inter-switch connectivity providing the basis for this logical topology. See the Brocade VDX product documentation resources at Broadcom VDX 6740 and the archived configuration-guide listings at Broadcom Ethernet Switching downloads.

QUESTION NO: 13

You have inherited the routing redundancy configuration shown in the exhibit.

Router_A

interface ethernet 1

ip address 192.53.5.1 255.255.255.0

ip vrrp-extended vrid 1

backup priority 110 track-priority 20

ip-address 192.53.5.3

track-port e 14

enable

!

Router_B

interface ethernet 1

ip address 192.53.5.2 255.255.255.0

ip vrrp-extended vrid 1

ip-address 192.53.5.3

track-port e 11

enable

!

Router A is the current VRRP-E master. Your network is not failing over to router B. What needs to be added to configure router B as a backup and maintain the current master?

A. backup priority 100 track-priority 20

B. backup priority 110 track-priority 20

C. backup priority 120 track-priority 20

D. backup priority 130 track-priority 20

ANSWER: A

Explanation:

The required addition is **backup priority 100 track-priority 20**. Router A already has a configured VRRP-E backup priority of 110. Assigning Router B a priority of 100 establishes it as the lower-priority member of VRID 1, so Router A remains the elected master during normal operation while Router B assumes the backup role. Both routers use the same virtual router identifier and virtual IP address, which is essential because they must participate in the same VRRP-E redundancy group.

The `track-priority 20` value enables the configured tracked Ethernet port to affect the router's effective VRRP-E priority. This permits a router to relinquish or assume the master role appropriately when its tracked uplink state changes, rather than relying only on the Ethernet 1 interface state. Matching the tracking decrement on both peers creates consistent failover behavior, while the 110-versus-100 base-priority relationship preserves Router A as master until a relevant failure occurs. This follows VRRP's priority-based master election model, documented in RFC 5798; Brocade FastIron Layer 3 configuration documentation covers the platform's VRRP-E implementation and tracking commands in the FastIron Layer 3 Configuration Guide.

QUESTION NO: 14

Your Brocade network monitoring solution includes the requirements for random sampling of packets and time-based sampling of counters.

You have been asked to determine an appropriate solution.

What will fulfill the specified requirements?

A. sFlow

B. RMON1

C. RMON2

D. Netflow

ANSWER: A

Explanation:

sFlow fulfills both stated monitoring requirements because it is specifically designed to export two complementary kinds of sampled telemetry. Its flow-sampling component randomly selects packets according to a configured sampling rate and sends packet-header and interface-context information to an sFlow collector. This provides a statistically representative view of traffic, including applications, protocols, endpoints, and traffic direction, without requiring every packet to be exported. Its counter-sampling component polls and exports interface and other device counters at configured time intervals. Those periodic counter samples enable measurement of utilization, packet rates, errors, discards, and longer-term traffic trends.

This combination is particularly appropriate for a Brocade network where scalable, low-overhead visibility is required across high-speed interfaces. Random packet sampling supplies flow-level insight, while time-based counter sampling supplies accurate trend and performance data from device counters. The sFlow specification describes flow samples and counter samples as the core sampled-data types, and the sFlow overview explains their use for scalable network monitoring. See RFC 3176: sFlow and the sFlow Overview.

QUESTION NO: 15

Which two multicast routing protocols are supported on Brocade FastIron devices? (Choose two.)

A. DVMRP

B. PIM

C. MBGP

D. MOSPF

ANSWER: A B

Explanation:

DVMRP and **PIM** are supported multicast-routing protocol families on Brocade FastIron platforms. FastIron multicast configuration supports Distance Vector Multicast Routing Protocol operation for building multicast distribution trees using its multicast distance-vector behavior. DVMRP is a multicast routing protocol specified in RFC 1075, and FastIron documentation includes its multicast-routing configuration and operational support.

Protocol Independent Multicast is also supported by FastIron for multicast forwarding and distribution-tree construction. PIM operates independently of the specific unicast routing protocol used to establish reachability; its multicast modes use the existing unicast routing table for reverse-path forwarding decisions. This makes PIM a core FastIron multicast-routing feature for routed multicast deployments. The standards-based PIM behavior and terminology are defined in RFC 7761. Brocade FastIron product documentation and configuration guides are available through the Broadcom documentation portal, including the platform-specific Layer 3 and multicast configuration material.