PassQueen
a.k.a dumpsqueen.com

Ericsson Certified Associate - IP Networking

Ericsson ECP-206

Version Demo

Total Demo Questions: 7

Total Premium Questions: 75

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

Which two statements are true about OSPF operation on a broadcast Ethernet network? (Choose two.)

A. The DR and BDR are elected using interface priority and router ID.

B. All routers on the segment must form full adjacencies with every other router.

C. A DROTHER router normally forms full adjacencies with the DR and BDR.

D. The router with the lowest OSPF cost is always elected as the DR.

ANSWER: A C

Explanation:

On a broadcast multi-access network, OSPF elects a designated router and a backup designated router to reduce the number of full adjacencies required between routers. The election uses the interface priority first; if priorities are equal, the highest router ID is selected. Routers that are neither DR nor BDR normally remain in the 2-Way state with one another, while forming full adjacencies with the DR and BDR.

QUESTION NO: 2

How do peers recognize each other's ability to use MP-BGP extensions?

A. They advertise the ability using the capabilities field during the session establishment.

B. They advertise the ability using the capabilities field in the End-of-RIB marker.

C. They advertise the ability using the capabilities field in the hello packet.

D. They advertise the ability using the capabilities field in the update packet.

ANSWER: A

Explanation:

They advertise the ability using the capabilities field during the session establishment. BGP peers negotiate support for Multiprotocol BGP by exchanging capabilities as part of the BGP OPEN-message processing that occurs when a session is established. The Multiprotocol Extensions capability identifies an Address Family Identifier and Subsequent Address Family Identifier that the sender is able to receive. This lets each peer determine the protocol families, such as IPv6 unicast or VPN routes, for which multiprotocol reachability information may be exchanged after the session becomes operational. Capability advertisement is therefore a session-establishment function rather than ordinary route advertisement. The capability mechanism is defined as an optional parameter of the BGP OPEN message, and it is designed to allow peers to announce supported BGP extensions without requiring a separate BGP version for every feature. RFC 5492 specifies the BGP capability advertisement mechanism and the OPEN-message encoding, while RFC 4760 defines the multiprotocol capability and the address-family information used by Multiprotocol BGP. See RFC 5492: Capabilities Advertisement with BGP-4 and RFC 4760: Multiprotocol Extensions for BGP-4.

QUESTION NO: 3

Which router function advertises external routes in OSPF?

A. ABR

B. ASBR

C. designated router

D. backbone router

ANSWER: B

Explanation:

ASBR is correct because an Autonomous System Boundary Router connects an OSPF routing domain to an external routing domain or another source of routing information. Its defining OSPF function is route redistribution: it imports non-OSPF routes, such as routes learned through static configuration, connected interfaces, or another routing protocol, and advertises them into the OSPF autonomous system as external routes. In normal OSPF areas, those redistributed routes are described with AS-external LSAs (type 5). When redistribution occurs in a not-so-stubby area, the ASBR originates type 7 LSAs, which can be translated to type 5 LSAs by an area border router when appropriate. OSPF external metrics can be advertised as external type 1, where the internal cost to reach the ASBR is included in route selection, or external type 2, where the external metric is the primary value. This role lets OSPF routers learn how to reach destinations outside their own OSPF domain while maintaining OSPF's link-state topology information internally. The IETF OSPFv2 specification defines an AS boundary router as a router that exchanges routing information with routers belonging to other autonomous systems and originates AS-external advertisements. See RFC 2328, section 2.4 and RFC 2328, AS-external-LSA format.

QUESTION NO: 4

What is a reason for using VLANs in an IP network?

A. to implement virtual routing

B. to enable MAC address learning on a router port

C. to isolate hosts within the same IP subnet

D. to isolate hosts across multiple IP subnets

ANSWER: C

Explanation:

"to isolate hosts within the same IP subnet" is correct because a VLAN creates a separate Layer 2 broadcast domain, independently of the physical switch location of connected devices. Switch ports assigned to different VLANs do not forward ordinary Ethernet broadcasts, unknown unicasts, or Layer 2 traffic between those VLANs. This segmentation lets an administrator separate groups of hosts logically on shared switching infrastructure, rather than requiring a separate physical switch network for every group. In an IP network, VLANs are commonly aligned with IP subnet boundaries, but the fundamental function of a VLAN is Layer 2 isolation; IP addressing itself does not determine VLAN membership. Using VLAN separation can constrain broadcast propagation, support administrative separation of user or service groups, and establish clear boundaries where inter-VLAN communication can be controlled through a Layer 3 gateway and policy mechanisms. Cisco's VLAN overview describes VLANs as logical broadcast domains created on switching infrastructure, while IEEE 802.1Q defines the VLAN tagging mechanism used to identify traffic belonging to a VLAN across Ethernet links. See Cisco: What Is a VLAN? and IEEE 802.1Q standard information.

QUESTION NO: 5

Which two statements are true about IPv6 Duplicate Address Detection (DAD)? (Choose two.)

A. A node sends a Neighbor Solicitation with the unspecified address as the source address for its tentative address.

B. A Neighbor Advertisement received in response indicates that the tentative address is already in use.

C. A node uses an ARP broadcast to detect whether its IPv6 address is duplicated.

D. DAD is required only when assigning a global unicast address.

ANSWER: A B

Explanation:

Before assigning a tentative IPv6 unicast address for normal use, a node performs DAD by sending a Neighbor Solicitation for that address. The solicitation uses the unspecified address (::) as its source and is sent to the address's solicited-node multicast address. If another node responds with a Neighbor Advertisement, the address is already in use and must not be assigned to the interface. DAD is part of IPv6 Neighbor Discovery and does not use IPv4 ARP broadcasts.

QUESTION NO: 6

Which two protocols apply to both IPv4 and IPv6? (Choose two.)

A. SNMP

B. ARP

C. DNS

D. MD

ANSWER: A C

Explanation:

SNMP and DNS both apply in IPv4 and IPv6 networks because they provide application-layer services independently of the IP address family used for transport or addressing. SNMP is used to monitor and manage network elements through operations such as retrieving management information, receiving notifications, and, where authorized, changing configuration. SNMP implementations can communicate using either IPv4 or IPv6 transport, allowing the same management architecture to support dual-stack and IPv6-only devices. The IETF defines transport mappings for SNMP that explicitly include UDP over IPv6.

DNS also operates across both address families. DNS resolves names to resource records appropriate to the requested service and network protocol. In particular, an IPv4 host address is represented by an A record, while an IPv6 host address is represented by an AAAA record. A DNS resolver and authoritative server can themselves communicate over IPv4 or IPv6, and DNS data can provide either or both address types. This makes DNS fundamental to dual-stack deployments, where clients may use name resolution to obtain IPv4 and IPv6 connectivity information. See the IETF SNMP transport mapping in RFC 3417 and IPv6 DNS extensions in RFC 3596.

QUESTION NO: 7

In a company network, a host sends an Ethernet frame destined to the address FF:FF:FF:FF:FF:FF.

What will an Ethernet switch do with this frame?

A. It will forward the frame to only one port.

B. It will forward the frame to all ports.

C. It will drop it.

D. It will send the frame only to management port.

ANSWER: B

Explanation:

It will forward the frame to all ports. FF:FF:FF:FF:FF:FF is the reserved Ethernet broadcast MAC address. A broadcast frame is intended for every station within the same Layer 2 broadcast domain, so a switch floods it through each applicable forwarding port in that VLAN. In practical switch operation, this means the frame is transmitted out all eligible ports except the ingress port on which it was received; sending it back through that same port would be unnecessary. The frame is not forwarded beyond the local broadcast domain by a router, unless a specific higher-layer broadcast-relay mechanism is configured for the protocol involved. Broadcast delivery is fundamental to common LAN operations, including address-

resolution and certain discovery processes. Ethernet MAC addressing defines the all-ones address as the broadcast address, and switching behavior ensures that all hosts on the relevant LAN can receive such traffic. Cisco’s Ethernet switching overview describes broadcast frames as being flooded to all ports in the VLAN, while the IEEE Registration Authority identifies the broadcast MAC address as the all-ones value. See Cisco LAN switching technologies and IEEE MAC address tutorial.