PassQueen
a.k.a dumpsqueen.com

Nokia Multicast Protocols

Alcatel-Lucent 4A0-108

Version Demo

Total Demo Questions: 26

Total Premium Questions: 266

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Volume A	58
Topic 2, Volume B	60
Topic 3, Volume C	60
Topic 4, Volume D	68
Total	266

QUESTION NO: 1

Which of the following data delivery models are most applicable to multicast data delivery? (Choose two)

A. One to one

B. Many to all

C. One to many

D. Many to many

E. One to all

ANSWER: C D

Explanation:

One to many and **Many to many** are the multicast delivery models. In one-to-many delivery, a single source transmits packets to a multicast group address, and the network replicates the traffic only where needed so that all interested receivers in that group receive the stream. This is the classic model for a live video feed, market-data publisher, or a single source distributing content to many listeners. In many-to-many delivery, multiple sources can send traffic to the same multicast group while multiple receivers join that group. Each participating receiver can receive traffic from the active sources, subject to any configured source-filtering policy. This model supports collaborative or shared-group applications such as multiparty conferencing and distributed information exchange. IP multicast is defined around host groups: senders address a datagram to a group, and the multicast infrastructure delivers it to members of that group rather than establishing separate unicast sessions per receiver. Protocol Independent Multicast supports both shared-tree and source-based forwarding approaches used to construct these multicast distribution paths. See RFC 1112, Host Extensions for IP Multicasting and RFC 4601, Protocol Independent Multicast—Sparse Mode.

QUESTION NO: 2

Which of the following PIM messages are sent towards RP? (Choose two)

A. (*, G) Join/Prune

B. (S.G) Join/Prune

C. Register Stop

D. (S, G, rpt) Join/Prune

ANSWER: A D

Explanation:

`(*, G) Join/Prune` is sent toward the Rendezvous Point because it establishes or removes multicast distribution-tree state on the shared tree. A receiver's designated router sends the Join/Prune message to its reverse-path-forwarding neighbor in the direction of the Rendezvous Point, and each upstream router processes and propagates that shared-tree request as required. This creates the receiver-to-Rendezvous-Point branch of the bidirectional shared-tree path used by sparse-mode multicast.

`(S, G, rpt) Join/Prune` is also sent toward the Rendezvous Point. In practice, its important use is the `(S, G, rpt)` prune generated when a last-hop router changes from the shared tree to a shortest-path tree. The router retains or builds source-specific state toward the source while sending the rendezvous-point-tree prune upstream toward the Rendezvous Point. This prevents duplicate traffic from continuing to arrive through the shared tree for that source and group while preserving multicast forwarding through the shortest-path tree.

These direction rules follow the PIM Join/Prune processing and sparse-mode shortest-path-tree transition procedures specified in RFC 7761, PIM Join/Prune messages and RFC 7761, RPT-to-SPT transition.

QUESTION NO: 3

Which of the following best describes the impact of an RP failure in a PIM ASM network?

A. All existing multicast traffic is stopped immediately, as the PIM ASM network must have an RP to transmit the multicast traffiC.

B. The multicast traffic will be impacted, but not immediately, as multicast routers will not discover the RP failure until the keep alive timer expires.

C. No impact at all, as an RP is not required in a PIM ASM network.

D. The traffic is not impacted if the multicast traffic flows have already been established, and the traffic flow does not involve the RP in the patH.

ANSWER: D

Explanation:

The traffic is not impacted if the multicast traffic flows have already been established, and the traffic flow does not involve the RP in the path. In PIM Sparse Mode Any-Source Multicast, the Rendezvous Point is essential to the initial receiver-to-source discovery process: receivers join the shared multicast distribution tree rooted at the RP, and sources initially register traffic with it. However, once the last-hop router has learned the active source and switched from the shared tree to the source-specific shortest-path tree, multicast packets travel directly from the source toward receivers. The RP is then no longer part of that packet-forwarding path. Consequently, failure of the RP does not interrupt such an already-established source-tree flow. This behavior is an important operational distinction between the RP's control-plane role in ASM group discovery and its possible role in the data-plane shared tree. Existing flows that have moved to shortest-path-tree forwarding can therefore continue while RP redundancy or recovery is addressed. PIM-SM procedures for shared trees, source registration, and shortest-path-tree transitions are specified in RFC 7761. The IETF multicast working-group material also provides background on PIM and ASM operation at IETF PIM Working Group.

QUESTION NO: 4

Which of the following are characteristics of the one-to-many model of multicast data delivery? (Choose two)

A. Receivers are generally unknown by the sourcE.

B. Receivers are generally known by each other.

C. It is well suited for applications in which the receivers provide feedback to each other and the sourcE.

D. There is a single data flow for each group on the source segment.

E. All devices are potentially capable of generating multicast datA.

F. It is well suited for applications such as video on demand, where each source sends a separate data stream to each of the many different receivers.

ANSWER: A D

Explanation:

Multicast uses a one-to-many delivery model in which a sender transmits packets to a multicast group address and the network replicates the traffic only where distribution paths diverge toward interested receivers. Consequently, **"Receivers are generally unknown by the sourcE."** is a defining characteristic: the sender addresses the group rather than maintaining individual receiver identities or separate unicast sessions. Group membership and forwarding-state mechanisms in the network determine where the traffic is delivered.

"There is a single data flow for each group on the source segment." is also correct. The source emits one multicast stream for the group onto its attached network segment. That stream is not duplicated once per receiver at the source; efficient replication occurs downstream as required by the multicast distribution tree. This behavior conserves source-side bandwidth and makes multicast appropriate for distributing the same content to many recipients, such as live video or market-data feeds. The Internet multicast host model defines sending to a group independently of knowledge of its individual members, while IGMP supports receiver membership signaling to the local multicast router. See [RFC 1112, Host Extensions for IP Multicasting] and [RFC 2236, Internet Group Management Protocol].

QUESTION NO: 5

Which of the following commands will verify, either directly or indirectly, if PIM is configured on an interface on the Alcatel-Lucent 7750 SR? (Choose two)

A. show router pim interface

B. show router pim status detail

C. show router pim status

D. show router interface detail

ANSWER: A D

Explanation:

`show router pim interface` is the direct verification command because it displays Protocol Independent Multicast interface information for the router instance. Its output identifies the IP interfaces on which PIM is enabled and provides operational PIM data associated with those interfaces, such as the configured mode and designated-router-related information. This makes it the standard command to confirm that an interface participates in PIM.

`show router interface detail` provides indirect verification. The detailed output for an IP interface includes the multicast-routing protocol status and interface attributes, allowing an operator to see whether PIM has been enabled for that specific router interface while also validating the interface's administrative and operational context. This is useful when troubleshooting because PIM configuration alone is not sufficient for multicast forwarding; the underlying routed interface must also be present and operational.

These commands therefore complement each other: one examines PIM's interface table directly, and the other examines the detailed routed-interface state where PIM enablement is reported. Nokia's SR OS multicast documentation describes PIM interface configuration and operational verification in the [SR OS PIM documentation]; the broader command and router-interface context is available in the [SR OS router interface documentation].

QUESTION NO: 6

Which of the following PIM messages are sent via unicast? (Choose two)

A. (*.G) Join/Prune

B. (S.G) Join/Prune

C. Register

D. Candidate-RP-Advertisement

ANSWER: C D

Explanation:

Register and **Candidate-RP-Advertisement** are sent using unicast in PIM Sparse Mode operation. When a designated router receives multicast traffic from a directly connected source but does not yet have a source-tree forwarding state toward the rendezvous point, it encapsulates the multicast packet in a PIM Register message and unicasts that message directly to the rendezvous point. This lets the rendezvous point learn about the active source and begin forwarding traffic through the shared tree or initiate source-tree procedures. PIM Register behavior and its unicast IP delivery are specified in [RFC 7761, section 4.9].

Candidate-RP-Advertisement is also a unicast control message. A candidate rendezvous point sends its advertisement directly to the elected bootstrap router so that the bootstrap router can include the candidate rendezvous point and its group-range information in the bootstrap information it distributes throughout the PIM domain. The PIM Bootstrap Router mechanism defines this unicast transmission in RFC 5059, section 3.1. These two messages therefore have specific unicast destinations rather than being transmitted as link-local PIM multicast control traffic.

QUESTION NO: 7

Which of the following events could trigger a PIM router sending an (S,G) Join message? (Choose three)

A. Receiver performing IGMP version 3 (S, G) report

B. Receiver performing enhanced IGMP version 2 signaling

C. Router performing PIM Source Path Tree signaling

D. Router performing translation to convert a standard IGMP (*, G) Join into a PIM (S, G) Join

E. RP performing translation to convert a PIM (*, G) Join into a PIM (S, G) Join

ANSWER: A

Explanation:

Receiver performing IGMP version 3 (S, G) report is the correct selection because IGMPv3 allows an attached receiver to express interest in multicast traffic from a particular source for a particular multicast group. The receiver's source-and-group interest is represented as an (S,G) state on the router's downstream interface. Where the router must obtain that source-specific traffic from an upstream PIM neighbor, it creates corresponding PIM (S,G) forwarding state and sends a PIM (S,G) Join toward the selected source or the appropriate reverse-path-forwarding neighbor. This is the fundamental control-plane relationship between IGMPv3 source filtering at the receiver-facing edge and PIM source-specific multicast forwarding in the routed network. The Join message builds the distribution tree in the upstream direction, while multicast packets subsequently flow downstream from the source to the interested receiver. RFC 3376 defines IGMPv3 source-filtering reports, including source-specific membership information, at RFC 3376. The PIM specification describes (S,G) Join/Prune state and the use of Join messages to construct multicast forwarding state at RFC 7761.

QUESTION NO: 8

Which of the following multicast packets can be forwarded on the Internet?

A. A multicast packet destined to FF70:: 1234

B. A multicast packet destined to FF1E:: 1234

C. A multicast packet destined to FF02:: 1234

D. A multicast packet destined to FF3A:: 1234

ANSWER: B

Explanation:

A multicast packet destined to FF1E:: 1234 can be forwarded on the Internet because it uses the IPv6 multicast address format in which the final hexadecimal digit of the second octet specifies multicast scope. In FF1E::1234, the scope value is E, which denotes global scope. Global-scope multicast addresses are intended to be routable across the Internet, subject to multicast routing policy, administrative controls, and the availability of multicast routing between networks. The preceding flag value of 1 identifies a transient multicast group: the group is not permanently assigned by IANA, but that does not limit its routing scope. IPv6 multicast forwarding decisions must honor the destination address's scope; a global-scope destination is therefore eligible for forwarding beyond a local link, site, or organization. RFC 4291 defines the IPv6 multicast address structure and identifies scope value E as global scope. IANA also maintains the authoritative IPv6 multicast address and scope registry. See RFC 4291, Section 2.7 and the IANA IPv6 Multicast Address Space Registry.

QUESTION NO: 9

An IGMP version 3 receiver signals INCLUDE mode report with an empty INCLUDE list. What does this indicate?

A. It received the query, but it does not wish to join any group.

B. It is willing to receive traffic for the group from any sourcE.

C. This is the initial join message; a selection of source will follow.

D. It indicates that it is leaving the group.

ANSWER: D

Explanation:

It indicates that it is leaving the group. In IGMPv3, a receiver's membership state is represented by a filter mode and a source list. INCLUDE mode means the receiver wants multicast packets for the group only from the sources explicitly named in its list. Consequently, an INCLUDE record with an empty source list expresses that there are no sources from which the receiver wishes to receive traffic for that multicast group. This is the IGMPv3 state used to withdraw group membership, and it has the operational effect of a leave. Routers processing IGMPv3 reports use this state transition to update their multicast forwarding state and, where appropriate, begin the group-and-source-specific query procedure to determine whether any other receivers remain interested. RFC 3376 explicitly defines an empty INCLUDE list as the state in which a host is not willing to receive any multicast data sent to the group, and specifies its use when leaving the group. This source-filter-based signaling replaces the separate, simple leave indication used by earlier IGMP behavior while retaining the required membership-withdrawal semantics. See the IGMPv3 host-state definition in RFC 3376, section 3 and the report-record processing rules in RFC 3376, section 4.2.12.

QUESTION NO: 10

Which of the following about the signaling of multicast group information in an MVPN service is TRUE when MPLS P2MP tunnels are used?

A. BGP Auto-Discovery is used for the signaling of the PMSIs, while either mLDP or RSVP-TE P2MP is used for the signaling of customer multicast group membership.

B. Either BGP Auto-Discovery or PIM is used for the signaling of the PMSIs, while either mLDP or RSVP-TE P2MP is used for the signaling of customer multicast group membership.

C. BGP Auto-Discovery is used for the signaling of the PMSIs, as well as for the signaling of customer multicast group membership.

D. BGP Auto-Discovery is used for the signaling of the PMSIs, while either BGP Auto-Discovery or PIM is used for the signaling of customer multicast group membership.

ANSWER: C

Explanation:

BGP Auto-Discovery is used for the signaling of the PMSIs, as well as for the signaling of customer multicast group membership. In an RFC 6514 BGP/MPLS IP VPN multicast deployment, BGP is the inter-provider-edge control-plane mechanism used to distribute the multicast VPN information needed to construct and use provider multicast service interfaces. The inclusive provider multicast service interface is advertised with BGP auto-discovery routes, allowing participating provider edges to discover the multicast distribution tunnel and the VPN context in which it applies. Customer multicast traffic and receiver interest are then represented across the provider network by BGP C-multicast routes. These routes carry the customer multicast source and group information, along with the route target and provider multicast service interface information needed for the receiving provider edges to join the appropriate distribution state. This BGP-based distribution is independent of whether the underlying provider tunnel is instantiated using MPLS point-to-multipoint technology such as mLDP or RSVP-TE P2MP. Those protocols establish or maintain the transport tunnel, whereas BGP distributes the MVPN membership and service-interface control information. RFC 6514 defines these BGP route types and their role in MVPN multicast state distribution: RFC 6514. The related MVPN architecture and provider multicast service interface model are described in RFC 6513.

QUESTION NO: 11

What are the first two actions a router performs after PIM is enabled?

A. BSR election followed by DR election.

B. DR election followed by PIM Hello messages.

C. DR election followed by BSR election.

D. BSR election followed by PIM Hello messages.

E. PIM Hello messages followed by BSR election.

F. PIM Hello messages followed by DR election.

ANSWER: F

Explanation:

PIM Hello messages followed by DR election. is correct. When PIM is enabled on a multicast-capable interface, the router begins PIM neighbor discovery by transmitting PIM Hello messages. These messages advertise the router's presence and PIM parameters to other PIM routers on the shared network segment. Hello processing establishes and maintains PIM neighbor relationships; the configured Hello period is normally 30 seconds, with startup Hellos sent more quickly to speed initial convergence.

Once routers on the multiaccess segment have exchanged Hellos and learned each other's DR priorities and addresses, they can determine the Designated Router. The DR is elected independently for each PIM interface and provides the key local role of handling multicast receiver and source-facing functions on that LAN. The election uses the PIM DR Priority Hello option when present; otherwise, the highest interface IP address is used as the tie-breaker. Thus, neighbor discovery through Hello exchange necessarily precedes a meaningful DR election. This behavior is specified in the PIM sparse-mode protocol definition, including the Hello and DR-election procedures. See RFC 7761, Protocol Independent Multicast—Sparse Mode, especially its PIM Hello and Designated Router election sections, and RFC 4601 for the foundational PIM-SM behavior.

QUESTION NO: 12

Refer to the following exhibit.

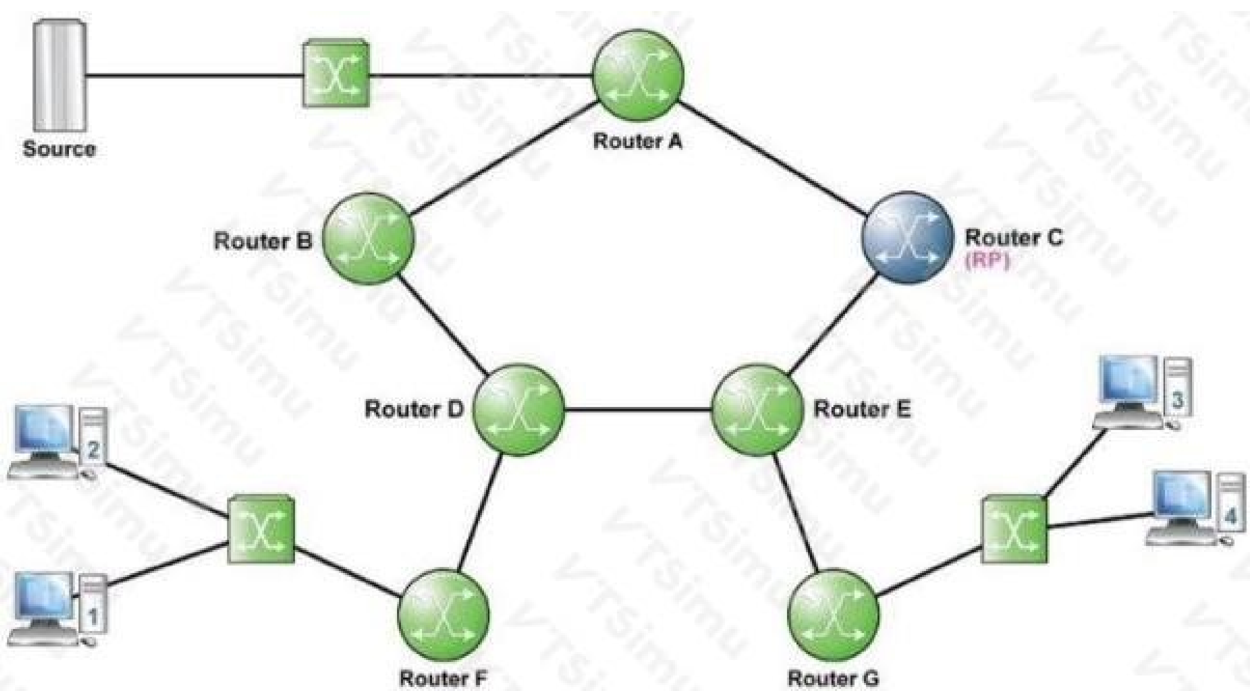

Assuming all links have an equal metric in the IGP, where will Source Path Tree have been created if Receiver 2 and Receiver 3 have joined the multicast group and the source is sending and the switchover has occurred? (Choose two)

A. Router F-D-B-A

B. Router G-E-C-A

C. Router C-E-D-B-A

D. Router G-E-C

E. Router F-D-E-C

ANSWER: A B

Explanation:

After SPT switchover, the last-hop routers serving Receiver 2 and Receiver 3 join the source directly by sending PIM (S,G) joins toward the source's unicast reverse-path-forwarding neighbor. The resulting multicast forwarding state follows the IGP shortest path from each receiver-side router back to the source, rather than continuing to use the shared tree toward the rendezvous point. With equal IGP metrics, the topology shown produces the source-specific branch **Router F-D-B-A** for the receiver reached through Router F and the source-specific branch **Router G-E-C-A** for the receiver reached through Router G. These are the two SPT branches created after the receivers have joined, traffic is active, and switchover has completed. PIM uses the route to the source to perform the RPF check and establish the upstream direction for source-specific multicast state; this ensures traffic is accepted only from the expected path toward the source. Nokia's multicast-routing documentation describes PIM shared-tree to shortest-path-tree behavior and source-specific forwarding state. The PIM-SM protocol behavior is also specified in RFC 7761; see Nokia's multicast routing documentation for platform implementation context.

QUESTION NO: 13

Which of the following about the use of BGP A-D in a MVPN service is TRUE?

A. If the I-PMSI is signaled by BGP A-D, then the S-PMSI must be signaled by BGP A-D.

B. If the I-PMSI is signaled by PIM, then the S-PMSI must be signaled by BGP A-D.

C. Customer group membership must be signaled by BGP A-D.

D. Customer group membership must be signaled by BGP A-D or PIM.

ANSWER: D

Explanation:

Customer group membership must be signaled by BGP A-D or PIM. An MVPN needs a multicast control-plane mechanism to convey receiver interest for a customer multicast group so that multicast forwarding state can be built toward the relevant provider-edge routers. Depending on the MVPN control-plane model and deployment, this function is conveyed through BGP-based MVPN signaling or through PIM signaling. BGP-based MVPN uses BGP route families to distribute multicast-related state between provider-edge routers, while PIM-based procedures provide the corresponding multicast join/prune signaling in designs that use PIM for this purpose.

This flexibility is fundamental to MVPN interoperation and migration scenarios: the service architecture separates customer multicast-state signaling from the provider multicast tunnel that transports the traffic. The provider therefore has a defined mechanism to learn which provider-edge routers require delivery for a customer group, regardless of whether the applicable control-plane signaling is BGP-based or PIM-based. RFC 6514 defines the BGP procedures and route types used for BGP/MPLS IP VPN multicast, including the distribution of multicast state between PEs. RFC 6513 describes the overall BGP-based MVPN architecture and its multicast control-plane functions. See RFC 6514 and RFC 6513.

QUESTION NO: 14

Which two of the following addresses should not be used in a multicast network at the same time?

A. 239.99.12.35

B. 237.87.12.35

C. 231.57.12.35

D. 229.227.12.35

E. 228.139.12.35

ANSWER: A D

Explanation:

239.99.12.35 and **229.227.12.35** should not be used simultaneously because they map to the same Ethernet multicast MAC address. IPv4 multicast-to-Ethernet mapping uses the multicast OUI `01:00:5E` together with only the low-order 23 bits of the IPv4 multicast group address. Consequently, the upper five bits of the multicast group identifier are not represented in the resulting MAC address. This is a deliberately non-unique mapping: multiple IPv4 multicast addresses can therefore resolve to one Layer 2 multicast destination.

For both listed groups, the retained 23-bit value is `0x630C23`. Each maps to Ethernet MAC address `01:00:5E:63:0C:23`. Hosts and switches using Layer 2 multicast forwarding cannot distinguish these two groups by destination MAC address alone. Using both groups on the same multicast-capable LAN can cause receivers and network devices to process traffic for the unintended group, producing unnecessary traffic delivery and potentially affecting multicast filtering behavior. The mapping procedure and its intentional 32-to-1 address ambiguity are specified in RFC 1112, Section 6.4; multicast address allocation principles are also documented in RFC 5771.

QUESTION NO: 15

Which of the following are NOT characteristics of the many-to-many model of multicast data delivery? (Choose two)

A. Receivers are generally unknown by the source.

B. Receivers are generally known by each other, since any receiver may be a source.

C. It is well suited for applications in which the receivers provide feedback to each other and the source.

D. All devices are potentially capable of generating multicast data.

E. It is well suited for applications such as video on demand, where each source sends a separate data stream to each of the many different receivers.

ANSWER: B E

Explanation:

In the many-to-many multicast model, membership in a multicast group does not imply that members know one another. A receiver can join a group dynamically and receive traffic sent to the group address, while the identity and number of other receivers are normally not exposed to it. Although any group member may potentially transmit, that capability does not make receivers generally known by each other. Therefore, "Receivers are generally known by each other, since any receiver may be a source." is not a characteristic of this model.

"It is well suited for applications such as video on demand, where each source sends a separate data stream to each of the many different receivers." is also not a many-to-many multicast characteristic. Multicast distribution is based on sending a packet once to a group address and having the network replicate it only where delivery paths diverge; it is not based on a source maintaining a distinct stream for every receiver. Many-to-many multicast is intended for group communication in which multiple members can be senders and receivers. The multicast host model and group-address delivery behavior are defined in RFC 1112; multicast service concepts are also described by RFC 3170.

QUESTION NO: 16

Which of the following are mandatory functions of a multicast receiver on a multicast network using IGMP version 2? (Choose two)

A. Report their multicast group membership to their local router

B. Signal their local router as they join and leave multicast groups

C. Generate feedback to the network or source

D. Listen for announcements advertising available sessions from the source

E. Generate multicast data to the group

ANSWER: A B

Explanation:

IGMP version 2 provides the host-to-router signaling required for routers to learn which multicast traffic must be forwarded onto a LAN. A receiver must report its multicast group membership to its local router. Hosts send an IGMP Membership Report when they join a group and also respond to router Membership Queries, subject to IGMP's report-suppression behavior. These reports allow the local multicast router to maintain group membership state for the connected network segment.

A receiver must also signal its local router as it joins and leaves multicast groups. IGMPv2 added the Leave Group message, enabling a host that is leaving a group to notify the router promptly when it is the appropriate reporting host. The router can then issue group-specific queries to determine whether any remaining receivers still require the stream. Together, membership reporting and join/leave signaling let multicast forwarding follow active receiver interest rather than requiring multicast data to be sent unnecessarily to every LAN.

The normative behavior and message processing are specified in RFC 2236: Internet Group Management Protocol, Version 2. General IETF multicast protocol resources are available from the IETF Standards site.

QUESTION NO: 17

Which features were introduced in IGMP version 3? (Choose two)

A. Leave group message

B. Report suppression was eliminated

C. IGMP snooping

D. TTL of 1 for IGMP messages

E. A modified membership report to include Source address

ANSWER: B E

Explanation:

IGMPv3 introduced source filtering, allowing a host to state not only its interest in a multicast group but also the multicast sources from which it wishes to receive traffic. This capability is carried in the revised Version 3 Membership Report format, which contains group records and can list source addresses. The feature is the foundation for source-specific multicast behavior and enables routers to maintain source-aware multicast membership state.

IGMPv3 also eliminated report suppression for its membership reports. Instead of relying on hosts hearing another host's report and suppressing their own response, the querier can obtain membership state from individual hosts. RFC 3376 explicitly specifies that Version 3 reports are not subject to the usual report-suppression mechanism and defines report records that convey the host's current filter mode and source list. These two changes work together: source-address reporting provides the detailed membership information, while the removal of suppression ensures that this host-specific state can be learned reliably. See [RFC 3376: Internet Group Management Protocol, Version 3], especially the membership-report and report-suppression sections, and the IETF's [RFC 3376 specification page].

QUESTION NO: 18

What is the default version of IGMP on the Alcatel-Lucent 7750 SR?

A. Version 1

B. Version 2

C. Version 3

D. There is no default setting, you must configure the version when you enable IGMP on 7750 SR

ANSWER: C

Explanation:

The default IGMP version on an Alcatel-Lucent 7750 SR is **Version 3**. SR OS uses IGMPv3 by default when IGMP is enabled for an IPv4 multicast interface. IGMPv3 is the current IGMP protocol version specified for IPv4 multicast membership management and adds source-filtering capability. This capability lets a receiver express not only interest in a multicast group, but also whether it wants traffic from particular multicast sources or wants to exclude particular sources. That information is important for source-specific multicast services and enables the router to maintain more precise multicast forwarding state. The configured IGMP version can be changed when interoperability with an environment using an earlier IGMP version is required, but no version configuration is necessary merely to enable IGMP because the default operating version is already IGMPv3. Nokia's SR OS multicast documentation describes the IGMP configuration and operational behavior for supported protocol versions, while the IETF specification defines IGMPv3 membership reports and source filtering. See the [Nokia SR OS IGMP documentation] and [RFC 3376: Internet Group Management Protocol, Version 3].

QUESTION NO: 19

Which features are introduced in IGMP version 2? (Choose two)

A. Leave group message

B. Report suppression is eliminated

C. IGMP snooping

D. TTL of 1 for IGMP messages

E. Formalized IGMP Querier election

ANSWER: A E

Explanation:

IGMP version 2 introduces the **Leave group message**, which enables a host to explicitly notify local multicast routers that it is leaving a multicast group. This significantly improves the ability of a router to stop forwarding traffic for a group when there are no remaining receivers on that subnet. After receiving a leave, the router can issue group-specific membership queries to determine whether any other hosts still require the multicast stream, rather than waiting only for the normal membership timer to expire.

IGMP version 2 also formalizes **IGMP Querier election**. When multiple multicast routers are present on the same LAN, the router with the numerically lowest IP address becomes the querier. The elected querier sends the periodic general membership queries, while the other routers monitor those queries and can take over if the current querier disappears. This provides deterministic, interoperable querier selection and avoids unnecessary duplicate querying.

These mechanisms are specified in the IGMPv2 protocol definition, RFC 2236. See [RFC 2236: Internet Group Management Protocol, Version 2] and the IETF's [RFC 2236 protocol text].

QUESTION NO: 20

In IPv6 multicast ASM mode, which of the following methods cannot be used to obtain the RP address?

A. Static

B. Router solicitation message

C. Bootstrap message

D. Embedded RP

ANSWER: B

Explanation:

Router solicitation message is correct because it is an IPv6 Neighbor Discovery Protocol control message used by a host to request Router Advertisements from local IPv6 routers. Its purpose is to help a node promptly learn on-link configuration information, such as default-router and prefix information; it does not distribute Protocol Independent Multicast rendezvous point mappings or provide an RP address for an ASM multicast group. IPv6 ASM receivers and sources require an RP mapping so that (*,G) multicast state and initial source registration can be directed toward the appropriate rendezvous point. Supported RP-address acquisition mechanisms include manual/static configuration, the PIM Bootstrap Router mechanism, and the IPv6 Embedded-RP mechanism. Embedded-RP encodes RP information in an IPv6 multicast group address for the defined embedded-RP range, allowing routers to derive the relevant RP directly from that group address. The Bootstrap Router mechanism instead disseminates RP-set information throughout the PIM domain. These mechanisms operate as part of multicast/PIM control-plane behavior, whereas Router Solicitation is strictly a local-link Neighbor Discovery exchange. The definition and processing rules for Router Solicitations are specified in RFC 4861. IPv6 Embedded-RP operation and its group-address encoding are specified in RFC 3956.

QUESTION NO: 23

With IGMP snooping enabled, how does a LAN switch identify ports with routers attached?

A. The same way as it identifies receivers

B. From the IGMP Host Membership Reports received

C. From the default gateway address of unicast frames

D. From the IGMP Queries received

E. The switch issues a message on all ports to which only routers respond

ANSWER: D

Explanation:

With IGMP snooping enabled, a LAN switch identifies a router-facing multicast port by inspecting received IGMP Query messages. An IGMP-capable multicast router, or the elected IGMP querier on that subnet, periodically transmits General Queries to determine which multicast groups have interested receivers. Because these queries originate from the multicast-routing side of the LAN, the snooping switch records the ingress interface as a multicast-router port (often called an *mrouter* port). The switch then forwards multicast control traffic and multicast streams, as required, toward that port so the router can route multicast traffic between network segments. This learning process is fundamental to IGMP snooping: the switch passively examines IGMP control packets rather than relying on ordinary unicast forwarding information. Router-port learning may also be supported through multicast-routing protocol control packets or static configuration on particular platforms, but IGMP Queries are the relevant mechanism in this question. Nokia's multicast documentation describes router ports as being dynamically learned from IGMP query messages, while RFC 2236 defines the IGMP query behavior used by multicast routers and queriers. See Nokia SR OS IGMP snooping documentation and RFC 2236: Internet Group Management Protocol, Version 2.

QUESTION NO: 24

What is the destination address of the MLDv2 general query?

A. FF02::1

B. FF02::2

C. FF02::16

D. The multicast group address of the multicast data that the router is receiving

ANSWER: A

Explanation:

FF02::1 is the destination address used by an MLDv2 general query. MLD, or Multicast Listener Discovery, is the IPv6 control protocol through which multicast routers discover whether nodes on a directly attached link have listeners for multicast traffic. A general query is intended to solicit current listener state from every multicast-capable node on the local link, rather than to query membership for one particular multicast group. It is therefore sent to the IPv6 link-local scope all-nodes multicast address, FF02::1. The FF02 scope value confines the query to the local link, which is essential because listener membership must be learned separately on each link and MLD messages are not routed beyond that link. In MLDv2, hosts receiving this query respond with reports describing their current multicast reception state, including source-filtering information where applicable. RFC 3810 specifies that general queries are sent to the all-nodes multicast address and distinguishes them from group-specific query behavior. See the normative MLDv2 specification in RFC 3810, especially the query-message and multicast-address definitions.

QUESTION NO: 25

What is the first action a router performs when it receives a multicast packet?

A. It determines if it is the DR for that segment.

B. It checks the PIM source group database for the matching group.

C. It sends a PIM Assert message to the RP.

D. It performs an RPF check on the packet.

E. It sends an IGMP join to the sourc

ANSWER: D

Explanation:

It performs an RPF check on the packet. Reverse Path Forwarding is the fundamental loop-prevention test used when a multicast router receives traffic. The router compares the interface on which the multicast packet arrived with the interface that its unicast routing table identifies as the best return path toward the multicast source. For source-specific multicast state, this is evaluated directly toward the source; for shared-tree traffic, the relevant RPF neighbor may be the rendezvous point. If the packet is received through the expected RPF interface, the router can process the applicable multicast forwarding state and replicate the packet to qualified outgoing interfaces. This early validation is essential because multicast distribution normally creates copies at branch routers, and accepting packets from arbitrary interfaces could create forwarding loops and duplicate traffic. PIM uses the unicast routing table for this RPF decision rather than requiring a separate multicast route calculation. The behavior and terminology are specified in the PIM-SM standard, which describes the RPF interface and RPF-neighbor checks used for multicast packet forwarding. See RFC 4601, Protocol Independent Multicast—Sparse Mode and Nokia 7750 SR multicast routing protocol documentation.

QUESTION NO: 26

An end station receiving a broadcast packet processes the data up to at least which layer of the OSI stack?

A. Layer 1

B. Layer 2

C. Layer 3

D. Layer 4

E. Layer 7

ANSWER: C

Explanation:

Layer 3 is correct because a broadcast packet is delivered across the local broadcast domain to every attached end station, and each receiving station must process the received Ethernet frame before the network-layer packet can be evaluated. After validating the Layer 2 frame and recognizing the broadcast destination, the station passes the encapsulated network-layer information to its IP implementation. At Layer 3, the host examines the IP destination and determines that the packet is addressed to a local or limited broadcast address and must be handled according to the applicable protocol. This network-layer processing is necessary even when no higher-layer application ultimately consumes the payload. In practical IPv4 environments, examples include limited broadcasts and subnet-directed broadcasts, whose delivery and host-handling behavior are defined at the IP layer. The use of broadcast therefore consumes processing resources on all end stations in the broadcast domain, which is one reason multicast is preferred when traffic is intended only for a defined group of receivers. RFC 919 describes the IP broadcast addressing model and specifies that broadcast datagrams are delivered to all hosts on the relevant network. RFC 1122 further defines host requirements for handling broadcast addresses and datagrams. See RFC 919 and RFC 1122.