DUMPSARENA

Nokia Scalable IP Networks

Alcatel-Lucent 4A0-100

Version Demo

Total Demo Questions: 35

Total Premium Questions: 359

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Volume A	137
Topic 2, Volume B	202
Total	359

QUESTION NO: 2

With a subnet mask of 255.255.255.224, which of the IP addresses below are valid host addresses? (Choose all that apply)

A. 17.23.119.63

B. 87.99.12.159

C. 95.11.22.93

D. 192.11.25.87

ANSWER: C D

Explanation:

A mask of 255.255.255.224 is a /27 prefix: its final octet has 32 addresses per subnet. Therefore, subnet ranges advance in blocks of 32 in the last octet: 0–31, 32–63, 64–95, 96–127, and so on. Within each block, the first address is the subnet address and the last address is the directed broadcast address; all addresses strictly between those two values are usable unicast host addresses.

95.11.22.93 belongs to the 64–95 block. The subnet address is 95.11.22.64 and the broadcast address is 95.11.22.95, making .93 a valid host address. **192.11.25.87** is also in the 64–95 block, where .64 and .95 retain the reserved subnet and broadcast roles. Since .87 is between those boundaries, it is likewise a valid host address. This calculation follows standard IPv4 CIDR subnetting behavior and the conventional reservation of all-zero host bits for the network and all-one host bits for the directed broadcast address. See RFC 4632 and RFC 950.

QUESTION NO: 3

Which of the following is required if devices on different VLANs wish to communicate with each other?

A. They can communicate if they are connected to the same switch.

B. They can communicate if they are in the same broadcast domain.

C. They can communicate if there is a router to connect them.

D. They cannot communicate.

ANSWER: C

Explanation:

They can communicate if there is a router to connect them. A VLAN creates a distinct Layer 2 broadcast domain, so traffic from one VLAN is not bridged directly into another VLAN. For a host in one VLAN to reach a host in another, the packet must be forwarded by a Layer 3 device that has an interface, subinterface, or routed virtual interface associated with each relevant VLAN. The sending host uses its configured default gateway for off-subnet traffic; the router then examines the destination IP address, applies its routing table and any relevant policy, and forwards the packet toward the destination VLAN. This function is commonly called inter-VLAN routing. It may be provided by a dedicated router using an IEEE 802.1Q trunk with logical subinterfaces, or by a multilayer switch using switched virtual interfaces. In either implementation, the essential requirement is Layer 3 routing between the separate IP networks/VLANs, together with appropriate addressing, gateway configuration, and permitted routing policy. VLAN membership alone does not supply a forwarding path between the VLANs. Nokia's service-routing documentation describes routed interfaces and IP forwarding behavior in its router operating system: Nokia SR OS Router Interfaces. The IEEE VLAN architecture and tagged-frame context are defined by IEEE 802.1Q.

QUESTION NO: 4

Which of the following is TRUE about a PE router within a service provider's MPLS network?

A. It resides on the customer premises.

B. It is not aware of the VPN services provided by the service provider.

C. It should connect to at least one other PE router.

D. It should connect to at least one CE device.

ANSWER: D

Explanation:

"It should connect to at least one CE device." is correct because a provider-edge router is the service-provider router at the boundary between the provider MPLS backbone and a customer site. Its essential service role is to provide the attachment point for customer-edge devices and to maintain the customer-facing routing context required to deliver VPN services. For Layer 3 VPNs, the provider-edge router associates customer-facing interfaces with the appropriate VPN routing and forwarding instance, learns or exchanges customer routes with the customer-edge device, and distributes VPN reachability across the provider network using Multiprotocol BGP and MPLS forwarding information. A customer-edge device does not need to run MPLS or participate in the provider's internal routing protocols; the provider-edge router performs the provider-side VPN and label-processing functions. Thus, a router acting as a provider edge must have at least one customer-facing attachment in order to serve that role. RFC 4364 defines the customer-edge/provider-edge relationship and describes the provider-edge router as the device that attaches directly to customer-edge routers: RFC 4364, BGP/MPLS IP Virtual Private Networks. MPLS VPN architectural terminology is also specified by the IETF in RFC 4026, Provider-Provisioned Virtual Private Network Terminology.

QUESTION NO: 6

Which of the following are NOT examples of a VPN service offered on the Alcatel-Lucent 7750 SR? Choose two.

A. VPWS Virtual Private Wire Service

B. VPLS Virtual Private LAN Service.

C. VPRN Virtual Private Routed Networks.

D. MLPPP Multi-link Point-to-Point protocol

E. VPDFS Virtual Private Dark Fiber Service.

ANSWER: D E

Explanation:

The correct selections are **MLPPP Multi-link Point-to-Point protocol** and **VPDFS Virtual Private Dark Fiber Service.** MLPPP is a link-layer protocol that bundles multiple PPP links into one logical connection. It can be used as an access or transport encapsulation feature, but it is not itself a customer VPN service. Likewise, "Virtual Private Dark Fiber Service" is not a defined 7750 SR VPN service offering. Dark fiber ordinarily describes unlit, physical optical-fiber infrastructure made available by a provider, rather than an emulated Layer 2 or Layer 3 VPN service delivered by the router. The 7750 SR service portfolio does include standardized VPN service families such as VPWS, VPLS, and VPRN. VPWS supplies point-to-point Ethernet connectivity, VPLS provides multipoint Ethernet LAN connectivity, and VPRN provides routed IP VPN connectivity. These services use the router's service architecture and MPLS/IP capabilities to isolate customer traffic while sharing provider infrastructure. Nokia describes the 7750 SR as a service router supporting business and mobile networking services, including Layer 2 and Layer 3 VPN capabilities; see the Nokia 7750 Service Router product page and the 7750 SR Services Overview documentation.

QUESTION NO: 7

Which of the following statements about OSPF routing protocol is FALSE?

A. It is a link state protocol.

B. It supports hierarchy using multiple areas.

C. It supports VLSM and address aggregation.

D. It uses hop-count for route selection.

ANSWER: D

Explanation:

It uses hop-count for route selection. is the false statement. OSPF is a link-state interior gateway protocol that selects routes using a cost metric, not a hop-count metric. Each OSPF interface is assigned a cost, commonly derived from its bandwidth or configured explicitly by the network operator. Routers advertise these link costs in link-state advertisements, build a synchronized link-state database for the area, and run the Shortest Path First algorithm to calculate the lowest-total-cost paths to destinations. The selected path therefore reflects the cumulative OSPF cost of its constituent links. A path with more routed hops can be preferred when its total cost is lower than that of a path with fewer hops. This metric behavior enables administrators to influence routing decisions according to link capacity, engineering requirements, or policy through interface-cost configuration. RFC 2328 specifies that OSPF associates a cost with every router interface and uses the sum of interface costs as the routing metric; it also describes the Shortest Path First calculation used to derive routes. See RFC 2328: OSPF Version 2, especially its metric and shortest-path sections, and the route calculation procedure.

QUESTION NO: 9

Which of the following statements are TRUE regarding SONET/SDH? (Choose 2)

A. They are deployed over point-to-point physical topologies to allow sub-50ms convergence.

B. An ADM (Add Drop Multiplexer) is used to connect various sites to the infrastructure.

C. IP datagrams are encapsulated in a PPP frame for transmission over an SDH network,

D. IP datagrams are encapsulated in an Ethernet frame for transmission over a SONET network.

ANSWER: B C

Explanation:

An ADM (Add Drop Multiplexer) is used to connect various sites to the infrastructure. SONET/SDH transport networks commonly use add/drop multiplexers at network nodes to insert (“add”) local tributary traffic into a high-speed optical signal and extract (“drop”) traffic destined for that site, while allowing the remaining traffic to continue through the network. This makes ADMs a fundamental building block for connecting sites on SONET/SDH transport infrastructure.

IP datagrams are encapsulated in a PPP frame for transmission over an SDH network. This describes Packet over SONET/SDH (POS): IP packets are carried by PPP, with PPP transported over a SONET/SDH payload. PPP supplies a well-defined point-to-point encapsulation and link-control mechanism for IP traffic, while SONET/SDH provides the synchronous optical transport layer. RFC 2615 specifies PPP encapsulation over SONET/SDH and identifies support for IP as a primary application. This model was widely used for high-speed router interconnection across carrier optical transport networks. See RFC 2615: PPP over SONET/SDH and the ITU-T G.707 recommendation for the SDH interface and multiplexing framework.

QUESTION NO: 10

You are working from a particular Command Line Interface (CLI) context, and want to see the commands available from your current context. What command can you issue to view this information?

A. view tree

B. tree

C. info detail

D. info

ANSWER: B

Explanation:

The `tree` command is used to display the CLI command hierarchy that is available from the current context. In SR OS, the command-line interface is organized as a hierarchical tree: each context contains commands and may contain additional subcontexts. Issuing `tree` lets an operator inspect that hierarchy from where they are currently positioned, making it useful for discovering valid configuration or operational commands without first navigating away from the active context.

This is particularly valuable when working in a large SR OS configuration, because command availability depends on the current branch of the CLI. The displayed tree provides a structured view of the applicable command paths, helping the operator identify both directly usable commands and related subordinate contexts. Nokia's SR OS documentation describes the CLI's hierarchical structure and its command-discovery behavior in the CLI usage material. See the Nokia SR OS CLI documentation and the SR OS CLI command reference for command syntax and CLI navigation details.

QUESTION NO: 11

What are the tasks of the Routing Table Manager (RTM)? Choose two

A. RTM selects the best route from among multiple routing protocols.

B. RTM selects the best route using the metric as criteria

C. RTM sends its best route to all its directly connected neighbors.

D. RTM installs its best route in the routing table.

ANSWER: A D

Explanation:

The Routing Table Manager is the SR OS control-plane component that consolidates route information submitted by routing protocols and other route sources. "RTM selects the best route from among multiple routing protocols." is correct because RTM arbitrates between competing routes to the same destination that originate from different sources. This selection is based primarily on the configured route preference, allowing SR OS to choose the active route source for a prefix consistently. Protocol-specific route calculation occurs before a route is submitted to RTM; RTM then performs the inter-protocol selection required for the system routing table.

"RTM installs its best route in the routing table." is also correct. After determining the active route, RTM maintains that selected route in the appropriate routing table and makes the route available for forwarding resolution. This centralized role ensures that the route used by the router reflects the selected source and associated next-hop information. Nokia's SR OS routing documentation describes the routing table as the common destination for routes learned through supported routing protocols and managed through route preference and route-selection processing. See the Nokia SR OS documentation portal and Nokia's Service Router Operating System overview.

QUESTION NO: 12

Which of the following statements are TRUE about the Alcatel-Lucent SR 7750? (Choose 2)

A. It can boot without a compact flash in the CF3 slot.

B. It cannot boot without a compact flash in the CF3 slot.

C. It can operate without a compact flash in the CF3 slot once it has booted.

D. It cannot operate without a compact flash in the CF3 slot once it has booted.

ANSWER: B C

Explanation:

The SR 7750 requires a CompactFlash card installed in the CF3 slot during the boot process. The CF3 card provides the boot media from which the system accesses the required boot image and system software information. Consequently, the chassis cannot complete its normal startup sequence when the required CompactFlash boot media is absent from CF3.

After the router has successfully loaded and is running, its operation is no longer dependent on continual access to the CompactFlash card in CF3. The operating system and active software execute from system memory, so the card can be removed after boot without immediately stopping forwarding or normal router operation. Operationally, administrators should still treat CompactFlash media carefully because it can hold boot software, configuration files, logs, and other persistent data needed for future restarts or maintenance. Before removing or replacing media, follow the applicable hardware procedures and ensure that required files are safely available elsewhere.

Nokia's SR documentation describes the boot-media and hardware procedures for the 7750 SR family in the 7750 SR Installation Guide and the broader platform documentation is available through the Nokia Service Router documentation portal.

QUESTION NO: 13

Which of the following best describes a repeater?

A. A passive device simply used to connect two or more cables. Does not generate or amplify any signals.

B. A device that receives and retransmits a signal out its ports, but does not do any Layer 2 analysis of the data.

C. A device that receives a signal and based on the Layer 2 destination address, makes a decision on which ports the signal should be retransmitted.

D. A device that receives a signal and based on the Layer 3 destination address, makes a decision on which ports the signal should be retransmitted.

ANSWER: B

Explanation:

A repeater is a Physical Layer device whose purpose is to receive an incoming signal, regenerate or reshape it as necessary, and transmit it onward. This regeneration counteracts signal attenuation and degradation over a transmission medium, allowing a network segment to extend farther than the original signal characteristics would otherwise permit. The defining characteristic is that it operates on the signal itself rather than interpreting frames: it does not inspect Ethernet MAC addresses, build a forwarding table, or make forwarding decisions based on Layer 2 information. Therefore, "A device that receives and retransmits a signal out its ports, but does not do any Layer 2 analysis of the data." is the best description. In classic Ethernet terminology, a multiport repeater is commonly called a hub; it repeats received traffic to its other ports without selectively switching frames. This is consistent with the OSI model's Physical Layer role, which concerns transmission and regeneration of bits and signals rather than data-link addressing. Nokia's networking documentation is available through the Nokia documentation portal, and Cisco's overview of the OSI model describes the Physical Layer's signaling and transmission functions: Cisco OSI reference.

QUESTION NO: 15

What types of addresses are part of encapsulation at the transport layer?

A. Source and destination MAC addresses

B. Source and destination IP addresses

C. Source and destination port numbers

D. Source and destination email addresses

ANSWER: C

Explanation:

Source and destination port numbers are part of transport-layer encapsulation. The transport layer, represented by protocols such as TCP and UDP, adds a header to application data to form a segment (TCP) or datagram (UDP). This header includes a source port and a destination port. Port numbers identify the communicating application processes or services on the sending and receiving hosts, allowing the receiving system to deliver incoming data to the appropriate application. This function is commonly called multiplexing and demultiplexing: many application sessions can share one host and IP address while remaining distinguishable through their transport-layer port values. For example, a client may use a temporary source port while sending traffic to a server's well-known service port. TCP defines source and destination port fields as 16-bit values in its header, and UDP uses the same fundamental addressing mechanism. The Internet Assigned Numbers Authority maintains the registry of assigned service names and port numbers at IANA Service Name and Port Number Registry. The TCP header format, including its source and destination port fields, is specified in RFC 9293.

QUESTION NO: 16

Which of the following are considered Virtual Private Networks? Choose all that apply.

A. IES

B. VPWS

C. VPLS

D. VPRN

ANSWER: B C D

Explanation:

VPWS, VPLS, and VPRN are VPN service types in the Alcatel-Lucent/Nokia service model. VPWS (Virtual Private Wire Service) provides a point-to-point Layer 2 private connection, emulating a dedicated wire between two customer attachment points. VPLS (Virtual Private LAN Service) provides multipoint Layer 2 connectivity, allowing geographically separated customer sites to operate as though they share one Ethernet LAN. VPRN (Virtual Private Routed Network) provides Layer 3 IP VPN connectivity using separate customer routing instances, commonly with MPLS transport in the provider network.

These services all deliver logically isolated connectivity over shared service-provider infrastructure. Their distinction is the VPN layer and topology offered: VPWS is point-to-point Layer 2, VPLS is multipoint Layer 2, and VPRN is Layer 3 routed VPN service. This classification aligns with the VPN service framework defined in RFC 4664, which identifies virtual private wire and virtual private LAN service types, and with the IP VPN architecture in RFC 4364. Nokia SR OS documentation is available through the Nokia Service Router documentation portal.

QUESTION NO: 17

Which of the following protocols belongs to the OSI suite of protocols? (Choose 2)

A. OSPF

B. BGP

C. X.500

D. IS-IS

E. Ethernet

ANSWER: C D

Explanation:

X.500 and IS-IS belong to the OSI protocol suite. X.500 is the family of OSI Directory standards, defining a distributed directory service and its information model. Its purpose is to make directory information—such as names, organizations, and related attributes—available across interconnected systems. The International Telecommunication Union maintains the X.500-series directory recommendations, including the overview of concepts, models, and services.

IS-IS, meaning Intermediate System to Intermediate System, was originally specified by ISO as a connectionless network-layer routing protocol for OSI environments. It enables intermediate systems (routers) to exchange topology information and calculate paths through an OSI network. Although IS-IS was later adapted for IP routing and is widely deployed in modern IP/MPLS service-provider networks, its origin and base specification are firmly in the OSI suite: ISO/IEC 10589. This OSI heritage is also reflected in the IETF's Integrated IS-IS specification, which describes extensions enabling IS-IS to carry IP routing information. See the ITU-T X.500 recommendation and RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments.

QUESTION NO: 18

Which of the following is a function of the Link layer?

A. It is responsible for encapsulating packets into frames for transmission on physical media.

B. It is responsible for encapsulating packets into IP datagrams and routing them.

C. It is responsible for encapsulating application data into TCP/UDP messages.

D. It is responsible for the timing of the signals on physical media.

ANSWER: A

Explanation:

It is responsible for encapsulating packets into frames for transmission on physical media. The Link layer provides the interface between the network-layer protocol, such as IP, and the particular network technology used to carry traffic on a local link. After IP produces a packet, the Link layer prepares it for delivery across that link by placing the packet into a technology-specific frame. For example, an Ethernet implementation adds an Ethernet header and frame check sequence, allowing the receiving interface to recognize the frame, identify its intended local destination, and check for transmission errors. The frame is then delivered to the Physical layer for actual signaling over copper, fiber, or another medium. This link-local framing role is fundamental in routed IP networks: at every router hop, the incoming Layer 2 frame is removed and the IP packet is encapsulated again into a new frame appropriate to the next outgoing link. RFC 1122 describes the link layer as handling communication with the physical network and notes that IP datagrams are transmitted through a link-layer protocol. See RFC 1122, Section 2.3 and the ISO OSI reference model overview.

QUESTION NO: 19

Choose two true statements that characterize Distance Vector Routing?

A. Routers send a copy of their routing table to their neighbors periodically.

B. Routers flood link information throughout the entire area.

C. Network converges quickly (within several seconds) after a topology change.

D. Routers do not have precise knowledge of the entire network topology.

ANSWER: A D

Explanation:

Distance-vector routing is characterized by each router maintaining route information based on the distance (metric) and direction (next hop) to a destination. "Routers send a copy of their routing table to their neighbors periodically" is correct because classic distance-vector protocols such as RIP distribute routing updates at regular intervals to directly connected neighbors. These updates allow neighbors to calculate or refresh their own best routes using the advertised metrics. RIP version 2, for example, specifies periodic response messages containing the router's routing information; see RFC 2453.

"Routers do not have precise knowledge of the entire network topology" is also correct. A distance-vector router learns reachability and metrics from adjacent routers rather than collecting a complete map of all routers and links. Its forwarding decisions therefore rely on the best next hop learned for each destination, not on a full topology database. This model

follows the distributed Bellman-Ford approach described in RIP's route-processing behavior. The routing-information exchange and metric-based route selection are detailed in RFC 1058. These characteristics distinguish distance-vector operation from protocols that construct a complete link-state view of the routing domain.

QUESTION NO: 20

What is the function of Layer 2?

A. Layer 2 is responsible for encapsulating packets into frames for transmission on physical media.

B. Layer 2 is responsible for encapsulating packets into IP datagrams and routing them.

C. Layer 2 is responsible for encapsulating application data into TCP/UDP messages.

D. Layer 2 is responsible for the timing of the signals on physical media.

ANSWER: A

Explanation:

Layer 2, the OSI Data Link layer, provides node-to-node delivery over a particular local network link. Its fundamental function is to take a Layer 3 packet, such as an IP packet, and encapsulate it in a Layer 2 frame that can be transmitted across the physical medium. The frame adds link-specific information, most notably source and destination media access control (MAC) addresses, and typically includes a frame check sequence (FCS) so the receiving interface can detect transmission errors. On Ethernet networks, this framing allows switches and network interface cards to identify the intended local recipient and handle the packet correctly on that LAN segment. The physical transmission itself—electrical, optical, radio signaling, bit timing, connectors, and media characteristics—is a Layer 1 responsibility; Layer 2 uses the service supplied by Layer 1 to send complete frames. This packet-to-frame encapsulation is therefore the essential Layer 2 operation described by the correct statement. The Internet architecture similarly distinguishes the link layer as the layer responsible for communication across a local network link, below IP. See RFC 1122, Section 2.3 and the IEEE 802.1D overview.

QUESTION NO: 21

Which of the following are considered Virtual Private Network (VPN) technologies? (Choose 3)

A. BGP

B. VPWS

C. VPRN

D. ISP

E. VPLS

ANSWER: B C E

Explanation:

VPWS, VPRN, and VPLS are VPN service technologies implemented on Alcatel-Lucent/Nokia Service Routers. VPWS (Virtual Private Wire Service) provides a point-to-point Layer 2 VPN, emulating a private circuit between two customer attachment points across a provider packet network. VPRN (Virtual Private Routed Network) provides a Layer 3 IP VPN: customer sites connect to separate virtual routing instances, allowing isolated IP routing and forwarding for each customer. VPLS (Virtual Private LAN Service) provides a multipoint Layer 2 VPN that makes geographically separated customer sites appear to share one Ethernet LAN broadcast domain. Together, these services represent the principal Layer 2 and Layer 3 VPN models commonly delivered by an MPLS/IP service-provider network. Nokia's service documentation identifies VPWS, VPLS, and VPRN as supported service types and describes their respective Layer 2 or Layer 3 service behavior. The IETF VPN framework likewise distinguishes provider-provisioned Layer 2 VPNs, including point-to-point and multipoint Ethernet services, from Layer 3 VPNs. See the Nokia SR OS Services Overview and RFC 4664: Framework for Layer 2 VPNs.

QUESTION NO: 22

Which two protocols are used for the dynamic signaling of MPLS labels (Choose two)?

A. RSVP-TE

B. CSPF

C. PNNI

D. LDP

ANSWER: A D

Explanation:

RSVP-TE and LDP are the MPLS control-plane protocols used to signal and distribute label information dynamically. RSVP-TE extends the Resource Reservation Protocol for traffic-engineered MPLS label-switched paths. It uses RSVP PATH and RESV signaling messages to establish an explicitly routed LSP and carries MPLS label information as part of that signaling process. This enables an ingress router to request a path with particular traffic-engineering constraints while intermediate routers allocate and signal labels for the LSP.

LDP, the Label Distribution Protocol, provides label distribution for forwarding equivalence classes between LSRs. Neighboring routers establish LDP sessions, advertise label mappings, and use those bindings to build MPLS forwarding entries. In a typical IP/MPLS network, LDP can therefore create dynamically signaled transport LSPs that follow the underlying IGP shortest paths, while RSVP-TE is used when the network needs explicitly signaled traffic-engineered paths. Both protocols consequently perform dynamic MPLS label signaling. See RFC 3209: RSVP-TE and RFC 5036: LDP Specification.

QUESTION NO: 23

What is the preference value in the routing table used for?

A. To select from routes with different next-hops.

B. To select from routes with different costs.

C. To select from routes learned from different routing protocols.

D. To select from routes with different prefixes.

ANSWER: C

Explanation:

“To select from routes learned from different routing protocols.” is correct. In Alcatel-Lucent/Nokia SR OS routing, route preference is the administrative ranking assigned to a route source or protocol. When the router has multiple candidate routes to the same destination prefix that were learned through different protocols, it uses preference to determine which source is considered more trustworthy. The route with the preferred value is eligible to become the active route and be installed in the forwarding table, subject to the normal route-selection process. This mechanism lets an operator define a deliberate hierarchy between sources such as static routes, OSPF, IS-IS, BGP, and other routing information, so that route selection remains predictable when equivalent destination information is received from more than one protocol. Preference is therefore a property used for comparing route sources, rather than a substitute for destination-prefix matching. Further route attributes, including protocol-specific metrics, can be considered after the applicable route source is selected. Nokia’s SR OS documentation provides the routing-protocol and route-selection context in its Unicast Routing Protocols documentation; the general forwarding and routing-table model is also described in RFC 1812.

QUESTION NO: 24

How does address summarization reduce the routing table size?

A. By allowing a block of routes to be represented by one route.

B. By only keeping best routes in the routing table.

C. By rejecting duplicate route advertisements.

D. By not advertising directly connected routes.

ANSWER: A

Explanation:

By allowing a block of routes to be represented by one route is correct. Address summarization, also called route aggregation, combines multiple contiguous networks that share a common prefix into a single, less-specific prefix. Rather than installing and advertising an individual route for every constituent subnet, a router can use one summary route that covers the entire address block. For example, several contiguous /24 networks may be represented by one appropriately aligned /22 prefix. This reduces the number of route entries that must be stored, processed during route lookups, and exchanged in routing updates. It also improves scalability by limiting the propagation of detailed topology information beyond the part of the network where that detail is needed. For a summary to accurately represent its component networks, the networks must be contiguous and align with the summary prefix boundary. Traffic matching the aggregate is forwarded toward the summarizing router, which retains or can reach the more-specific routes needed to deliver traffic within the summarized block. This hierarchical aggregation principle is fundamental to CIDR-based IP routing and is especially valuable in large service-provider and enterprise networks. See the IETF's CIDR specification in RFC 4632 and Nokia's routing documentation at Route Summarization.

QUESTION NO: 25

Which ATM adaptation layer is commonly used for transporting IP datagrams or non-real time data?

A. AAL0

B. AAL1

C. AAL2

D. AAL3/4

E. AAL5

ANSWER: E

Explanation:

AAL5 is the ATM Adaptation Layer commonly used to carry IP datagrams and other non-real-time, variable-length data over ATM networks. It was designed as a streamlined data-service adaptation layer: the sender places a higher-layer packet into a Common Part Convergence Sublayer protocol data unit, appends an 8-byte trailer, and segments the resulting payload into ATM cells. The trailer includes a length field and a cyclic redundancy check, allowing the receiving endpoint to reassemble the cells, identify the original packet length, and validate the completed payload. This relatively low-overhead design made AAL5 the normal encapsulation choice for IP-over-ATM and LAN-emulation-era data networking. RFC 1483 specifies multiprotocol encapsulation over AAL5, including the encapsulation of IP datagrams, while the AAL5 specification defines the adaptation layer's message-mode data-transfer and trailer behavior. Because its service is optimized for packet data rather than timing-sensitive circuit emulation, AAL5 directly matches the requirement for IP datagrams or non-real-time data. See RFC 1483: Multiprotocol Encapsulation over ATM Adaptation Layer 5 and ITU-T Recommendation I.363.5.

QUESTION NO: 26

Which of the following are examples of Time Division Multiplexing (Choose two)?

A. Ethernet.

B. Token-Ring.

C. SONET/SDH.

D. E1 carrier.

E. PPP

ANSWER: C D

Explanation:

SONET/SDH and E1 carrier are examples of Time Division Multiplexing because each carries multiple lower-rate channels by assigning them defined time positions within a repeating higher-rate transmission structure. SONET and SDH are synchronous transport hierarchies: tributary signals are mapped into structured frames, and the available transport capacity is organized into time slots and payload containers that recur at precise intervals. This synchronous framing and allocation of bandwidth is a core application of TDM in carrier transport networks. The ITU-T SDH framework is specified in Recommendation G.707, which defines the network-node interface rates and synchronous transport structure.

An E1 carrier is the classic digital TDM system. Its 2.048 Mb/s frame contains 32 time slots, each normally transmitting 8 bits per frame. Time slot 0 provides framing and related functions, while the remaining slots can be allocated for voice or data channels, including channel-associated signaling arrangements. Because every channel receives a recurring, fixed temporal position in the E1 frame, E1 directly illustrates synchronous TDM. The E1 primary-rate frame and time-slot structure are defined by the ITU-T in Recommendation G.704.

QUESTION NO: 27

Which of the following statements describes the purpose of the Transport Layer in the TCP/IP stack?

A. Provides a data transport service to higher protocol layers.

B. Provides a data transport service for routing and control protocols such as OSPF and ICMP.

C. Provides a universal address plan to uniquely identify every device in the network.

D. Defines a standard method for framing data for transmission on the physical network medium.

ANSWER: A

Explanation:

Provides a data transport service to higher protocol layers. This correctly describes the core role of the TCP/IP transport layer: it supplies end-to-end communication services to application and other upper-layer protocols. The transport layer accepts data from an application, identifies the relevant communicating processes by using port numbers, divides data into transport units as needed, and delivers the data to the appropriate process at the destination. Depending on the protocol selected, this service can include reliable, ordered delivery, acknowledgements, retransmission, flow control, and connection management through TCP, or lightweight connectionless datagram delivery through UDP.

This layer operates above Internet Protocol and allows applications to communicate independently of the underlying network technologies and individual network paths. In the Internet protocol suite's architectural description, TCP and UDP are specifically identified as transport-layer protocols, while applications use their services for host-to-host process communication. RFC 1122 describes the transport layer and its relationship to the application and internet layers: RFC 1122: Requirements for Internet Hosts—Communication Layers. The Internet Engineering Task Force's TCP specification further defines the reliable transport service provided by TCP: RFC 9293: Transmission Control Protocol.

QUESTION NO: 28

Which of the following statements are TRUE regarding IP-filters? (Choose 2)

A. IP-filters may be applied on interface ingress, egress, or both.

B. Multiple IP-filters can be applied in each direction.

C. An IP-filter may only be applied to a single interface if the scope is set to template.

D. An IP-filter may have multiple match criteria per entry.

E. The default-action of an IP-filter is always drop.

ANSWER: A D

Explanation:

IP-filters may be applied on interface ingress, egress, or both. In SR OS, an IP filter policy can be associated with the appropriate traffic direction on an interface, allowing the router to inspect and act on packets as they enter the interface, leave it, or in both directions when the design requires separate ingress and egress policies. This makes filters useful for traffic classification, access control, and protection at interface boundaries.

An IP-filter may have multiple match criteria per entry. A filter entry is designed to match a traffic flow using a combination of packet fields rather than a single attribute alone. Depending on the filter type and platform support, criteria can include such fields as source and destination prefixes, protocol, ports, and other packet attributes. The criteria within an entry collectively define the traffic to which that entry's action applies. This capability permits precise policies, for example matching a particular protocol between defined address ranges. Nokia's SR OS documentation describes filter-policy configuration, filter entries, match fields, actions, and attachment contexts in its SR OS documentation library and Router Configuration Guide filter-policy section.

QUESTION NO: 29

Given the following (Network address: 208.40.224.0, Network prefix: /28.) What is the address range, including the network and broadcast addresses, for the fourth subnetwork?

A. 208.40.224.0 208.40.224.255

B. 208.40.4.0 208.40.4.255

C. 208.40.224.16 208.40.224.31

D. 208.40.224.48 208.40.224.63

E. 208.40.227.0 - 208.40.224.255

ANSWER: D

Explanation:

The correct range is **208.40.224.48 208.40.224.63**. A /28 IPv4 prefix leaves four host bits, so each subnet contains 2^4, or 16, consecutive addresses. Consequently, subnet boundaries advance in blocks of 16 in the final octet. Starting from the stated network address, the first subnet spans 208.40.224.0 through 208.40.224.15; the next boundaries occur at .16, .32, and .48. Therefore, the fourth subnet begins at 208.40.224.48. Its 16-address block ends immediately before the following /28 boundary, which is 208.40.224.64, making 208.40.224.63 the broadcast address. The requested range explicitly includes both special addresses: .48 is the network address identifying this particular subnet, and .63 is the directed broadcast address for every host within it. This follows standard classless inter-domain routing prefix interpretation: the prefix specifies the fixed network portion and the remaining bits enumerate addresses within that prefix. See RFC 4632, Classless Inter-domain Routing and RFC 950, Internet Standard Subnetting Procedure.

QUESTION NO: 30

How many of the front access card slots of the Alcatel-Lucent 7750 SR-12 router are dedicated for redundant common equipment?

A. 1

B. 2

C. 7

D. 10

ANSWER: B

Explanation:

The correct answer is **2**. The 7750 SR-12 chassis provides twelve front-access card slots in total. Ten of these are service-card slots, intended for line-card/IOM-based service capacity, while two slots are reserved for the system's common equipment. The common-equipment positions accommodate the redundant switch-fabric and control/management functions required to maintain chassis operation and provide high availability. Reserving two dedicated slots permits an active and standby common-equipment arrangement, so that essential control-plane and switching resources can remain available if a common-equipment module fails or is replaced. This physical separation of service capacity from redundant system infrastructure is a key part of the SR-12 platform design. Nokia's 7750 SR documentation identifies the SR-12 as a twelve-slot chassis with ten slots for I/O modules and two slots for switch-fabric/control-processor modules. See the Nokia SR OS documentation portal and the Nokia 7750 Service Router product page for platform and hardware documentation.

QUESTION NO: 31

Which three of the following are important considerations when designing subnets? (Choose three.)

A. Number of subnetworks currently required.

B. Type of physical connectivity used by each segment.

C. Future network growth requirements.

D. Number of hosts each subnetwork will support.

E. Class of the network address to be used for the network.

ANSWER: A C D

Explanation:

"Number of subnetworks currently required," "Future network growth requirements," and "Number of hosts each subnetwork will support" are the core inputs to an IP subnetting plan. The present number of required subnetworks establishes how many distinct routed or broadcast domains must be addressed immediately. That requirement determines the quantity of subnet identifiers the addressing plan must provide. The anticipated number of hosts in each subnetwork determines the host-address capacity required per prefix, including sensible allowance for infrastructure addresses and operational expansion. These two sizing requirements must be evaluated together because allocating more bits to the subnet portion increases the number of available subnets while reducing the number of usable host addresses in each subnet. Future growth is equally important because a design that only fits current requirements can quickly lead to overlapping allocations, unnecessary renumbering, or fragmented address space as new sites, VLANs, services, and hosts are introduced. CIDR-based planning supports selecting prefixes according to actual subnet and host needs rather than relying on legacy address-class boundaries. RFC 950 defines the fundamental purpose of subnetting as partitioning an assigned network into multiple physical networks, while RFC 4632 documents CIDR's prefix-based allocation and aggregation model. See RFC 950 and RFC 4632.

QUESTION NO: 32

What protocol is used by ping to verify IP network reachability?

A. DCHP

B. ICMP

C. ARP

D. NAT

ANSWER: B

Explanation:

ICMP is the protocol used by the ping utility to test IP reachability. Ping sends an ICMP Echo Request message to a destination IP address and, when the destination is reachable and permitted to respond, receives an ICMP Echo Reply. This exchange verifies that an IP-layer path exists between the source and destination and also allows ping to report round-trip delay and packet-loss statistics. ICMP is defined as an integral supporting protocol for IP operation: it carries control, diagnostic, and error-reporting messages rather than application data. In an IP network, a successful reply demonstrates bidirectional IP connectivity for the tested packets; however, a missing reply is not always conclusive evidence of a network failure, since routers, firewalls, hosts, or security policies may filter or suppress ICMP echo messages. The foundational ICMP specification defines the Echo and Echo Reply message types used by ping in IPv4. For IPv6, ping similarly uses ICMPv6 Echo Request and Echo Reply messages. See RFC 792, Internet Control Message Protocol and RFC 4443, ICMPv6.

QUESTION NO: 33

What problem was Spanning Tree Protocol (STP) primarily designed to solve?

A. Inability to perform switch to switch redundancy using multiple connections at Layer 2.

B. Missing protocol field in Ethernet II required to identify redundancy protocols at Layer 2.

C. Missing packet sequencing in Layer 2 Ethernet required to re-order packets on arrival.

D. Inability to forward multicast frames without a routing protocol.

ANSWER: A

Explanation:

Spanning Tree Protocol (STP) was designed to make redundant Layer 2 Ethernet connections usable without allowing those connections to create switching loops. Redundant physical links are valuable because they provide an alternate path if a switch or link fails. However, Ethernet frames have no Layer 2 time-to-live field, so a frame caught in a forwarding loop can circulate indefinitely. Such loops can result in broadcast storms, repeated delivery of frames, and instability in the switches' MAC-address forwarding tables. STP creates a loop-free logical topology from a physically redundant topology. It elects a root bridge, determines the lowest-cost paths toward that root, and places selected redundant ports into a non-forwarding state. The blocked path remains available as a backup and can transition to forwarding if the active path fails. Thus, the essential purpose is not to remove redundancy, but to prevent the Layer 2 loops that redundancy could otherwise introduce while retaining failover capability. This behavior is specified by IEEE bridging standards, including Rapid Spanning Tree Protocol in IEEE 802.1Q. See the IEEE 802.1Q standard page and Cisco's Spanning Tree Protocol overview.

QUESTION NO: 34

Which organization eventually became the standards body for IP and related protocols?

A. IEEE

B. IETF

C. NSFNET

D. ITU-T

E. APNIC

ANSWER: B

Explanation:

The Internet Engineering Task Force is the standards-development organization responsible for the technical specifications that define and evolve the Internet Protocol suite and many related Internet technologies. Its work is organized into focused working groups, whose output is published through the RFC (Request for Comments) series. Important IP-related standards, including IPv4 and IPv6 specifications, TCP, UDP, BGP, DNS-related protocol standards, and numerous routing, security, and management protocols, are developed or maintained through the IETF's open consensus-based process. The IETF grew from early U.S. government-sponsored Internet research activities and became the principal open standards forum for Internet protocols as the Internet expanded beyond its original research networks. Its standards process emphasizes interoperable implementations, open participation, technical review, and documented consensus. The RFC Editor maintains the archival RFC publication series, while the IETF publishes standards-track RFCs and other technical documents through that series. The IETF's role in creating and maintaining Internet protocol standards is described by the IETF About page and its standards process is documented in RFC 2026.

QUESTION NO: 35

Given the Mac Address 00-20-60-ce-2b:28, which part is the Organizationally Unique Identifier (OUI)?

A. 2b:28

B. 00-20

C. ce-2b:28

D. 00-20-60

ANSWER: D

Explanation:

The Organizationally Unique Identifier (OUI) is **00-20-60**. In the conventional 48-bit MAC address format, the address consists of six octets (48 bits). The first three octets, or first 24 bits, form the OUI and identify the organization to which IEEE assigned that prefix. Here, the address is intended to be read as `00-20-60-CE-2B-28`; therefore, its initial three octets are `00-20-60`. The remaining three octets, `CE-2B-28`, are the organization-defined portion, commonly called the extension identifier or device-specific identifier. Network devices use this division so vendors can assign distinct interface MAC addresses under an IEEE-registered organizational prefix. IEEE Registration Authority documentation describes an OUI as a 24-bit identifier assigned to an organization for use in constructing MAC addresses. The standard MAC-address structure and the purpose of the OUI are also summarized in Cisco's networking reference material. See the IEEE Registration Authority OUI information and Cisco MAC address overview.