DUMPSBOSS.

Cisco Certified Design Expert (CCDE v3.0)

Cisco 400-007

Version Demo

Total Demo Questions: 27

Total Premium Questions: 277

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

Which two protocols are used bv SDN controllers to communicate with switches and routers? (Choose two)

A. OpenFlash

B. OpenFlow

C. NetFlash

D. Open vSwitch Database

E. NetFlow

ANSWER: B D

Explanation:

SDN architectures commonly separate the control plane (controller) from the data plane (switch/router) and rely on southbound interfaces to program forwarding behavior and exchange operational state. OpenFlow is a well-known southbound protocol that allows an SDN controller to install, modify, and delete flow entries in a device's flow tables, effectively programming how traffic is forwarded. This is a canonical example of controller-to-switch communication in SDN and is widely cited in SDN literature and vendor implementations.

Open vSwitch Database (OVSDB) is also used between controllers and switches—particularly with Open vSwitch—to perform management-plane configuration such as creating bridges, ports, VLANs, tunnels, and retrieving state. In many SDN solutions, OpenFlow handles flow programming while OVSDB handles device configuration and operational parameters, making them complementary southbound protocols used by controllers to communicate with forwarding devices.

References: Open Networking Foundation (OpenFlow), RFC 7047 - The Open vSwitch Database Management Protocol.

QUESTION NO: 2

SDWAN networks capitalize the usage of broadband Internet links over traditional MPLS links to offer more cost benefits to enterprise customers. However, due to the insecure nature of the public Internet, it is mandatory to use encryption of traffic between any two SDWAN edge devices installed behind NAT gateways. Which overlay method can provide optimal transport over unreliable underlay networks that are behind NAT gateways?

A. TLS

B. DTLS

C. IPsec

D. GRE

ANSWER: B

Explanation:

DTLS is the best fit here because it provides encryption while using UDP as the transport, which is generally more tolerant of loss and latency variation on “unreliable” Internet underlays than TCP-based approaches. In SD-WAN designs, DTLS is commonly used to secure control-plane and/or data-plane communications across the public Internet, and it is well-suited to NAT traversal because it can operate over a single UDP flow that NAT devices can translate and maintain with keepalives. This aligns with the requirement for encrypted traffic between SD-WAN edge devices located behind NAT gateways, while still enabling efficient overlay transport over broadband links where packet loss and jitter are expected. DTLS essentially brings TLS-like security properties (authentication, confidentiality, integrity) to datagram transport, which is why it is widely adopted in SD-WAN overlays that must function across NAT and variable-quality Internet paths. References: Cisco SD-WAN Design Guide, RFC 6347 (Datagram Transport Layer Security).

QUESTION NO: 3

You need to redesign your NMS system so that it can collect information without causing adverse effects in the network, such as high CPU utilization on network devices and network instability. Which two options will minimize the impact of the trusted NMS polling your network in this situation? (Choose two.)

A. Implement SNMP community restrictions that are associated with an ACL

B. Disable unused OIDs and MIBs on the NMS systems.

C. Unload unused MIBs from the network devices

D. Prevent polling of large tables through the use of SNMP OID restrictions.

ANSWER: B D

Explanation:

To minimize the operational impact of a trusted NMS, the most effective approach is to reduce the amount and “cost” of SNMP work the devices must perform per polling cycle. Preventing polling of large tables through the use of SNMP OID restrictions is correct because large table walks (for example, interface tables, routing tables, or ARP/neighbor tables) can be CPU-intensive and can generate bursts of management-plane traffic. Constraining what can be queried (or redesigning what is queried) directly reduces device processing and the volume of responses, which helps avoid high CPU and instability.

Disabling unused OIDs and MIBs on the NMS systems is also correct because it reduces unnecessary polling requests at the source. In practice, this means removing or disabling pollers/collectors for metrics you do not operationally need, avoiding broad “walks” when targeted GETs will do, and tuning polling frequency and scope. This lowers both device load and network overhead while still allowing required telemetry collection.

These recommendations align with Cisco guidance to protect the management plane and to avoid expensive SNMP operations (especially repeated table retrieval) by limiting what is polled and how often. See Cisco SNMP Best Practices and Cisco SAFE/Management Plane Protection concepts.

QUESTION NO: 4

What are two common approaches to analyzing and designing networks? (Choose two.)

A. bottom-up approach

B. high-low security approach

C. top-down approach

D. left-right approach

E. three-tier approach

ANSWER: A C

Explanation:

Two widely used methodologies for network analysis and design are the top-down approach and the bottom-up approach. The top-down approach starts from business goals, user/application requirements, and constraints, then translates those into logical architecture (services, segmentation, routing domains, security policy) and finally into physical design (devices, links, addressing, and implementation details). This is commonly recommended because it ensures the network is built to meet application performance, availability, and operational requirements rather than being driven by specific hardware choices.

The bottom-up approach begins with the existing physical infrastructure and technologies (cabling, switching, routing platforms, WAN circuits, and current configurations) and works upward to determine what capabilities and limitations exist, then maps those findings to application and business needs. This approach is common in troubleshooting, assessments, migrations, and redesigns where an installed base must be evaluated and evolved. In practice, many real-world design efforts blend both: top-down to define intent and success criteria, and bottom-up to validate feasibility and constraints of the current environment.

References: Cisco Press – Top-Down Network Design (concepts), Cisco Learning Network – CCDE exam topics (design methodology emphasis)

QUESTION NO: 5

A legacy enterprise is using a Service Provider MPLS network to connect its head office and branches. Recently, they added a new branch to their network. Due to physical security concerns, they want to extend their existing IP CCTV network of the head office to the new branch, without any routing changes in the network. They are also under some time constraints. What is the best approach to extend the existing IP CCTV network to the new branch, without incurring any IP address changes?

A. GRE

B. L2TPv3

C. VXLAN

D. EoMPLS

ANSWER: B

Explanation:

The best approach is to extend Layer 2 connectivity between the head office and the new branch so the CCTV devices can remain in the same IP subnet and no routing or IP addressing changes are required. L2TPv3 is well-suited for this because it can pseudowire Ethernet frames (or VLANs) across an IP/MPLS core, effectively creating a point-to-point Layer 2 circuit between sites. This allows the enterprise to “stretch” the existing CCTV VLAN/subnet to the new branch quickly, meeting the time constraint while keeping the service provider MPLS underlay unchanged from the enterprise routing perspective. In practice, L2TPv3 can be deployed as a CE-based solution (no dependency on the provider offering a specific L2VPN service), which is often attractive for rapid turn-up when the enterprise controls the edge routers but cannot change the provider core. This design preserves the existing IP plan for the CCTV network and avoids introducing new routed segments

or renumbering. For background on L2TPv3 pseudowires and their use to transport Layer 2 over an IP network, see RFC 3931 (L2TPv3) and Cisco's L2TPv3 configuration guidance such as Cisco L2TPv3 configuration example.

QUESTION NO: 6

In the case of outsourced IT services, the RTO is defined within the SLA. Which two support terms are often included in the SLA by IT and other service providers? (Choose two.)

A. network size and cost

B. support availability

C. network sustainability

D. network reliability

E. resolution time

ANSWER: B E

Explanation:

In outsourced IT services, SLAs commonly define operational support commitments that directly affect service restoration and continuity outcomes. "support availability" is a typical SLA term because providers must specify when and how customers can access support (for example, 24x7x365 vs. business hours, follow-the-sun coverage, and which channels are supported). This is foundational for incident handling and is frequently paired with severity definitions and escalation paths so both parties understand how support is delivered.

"resolution time" is also a common SLA term because it sets expectations for how quickly incidents or service requests will be restored or resolved, often expressed as targets by priority/severity (for example, P1 restored within X hours). This aligns closely with continuity objectives such as RTO, since RTO is effectively a time-based restoration requirement and providers typically map their incident management targets to those business requirements.

These two terms are widely used across IT service management practices and are consistent with ITIL-style service level management and incident management constructs used by many providers. References: AXELOS ITIL guidance, Cisco Technical Support Services overview.

QUESTION NO: 7

Refer to the exhibit.

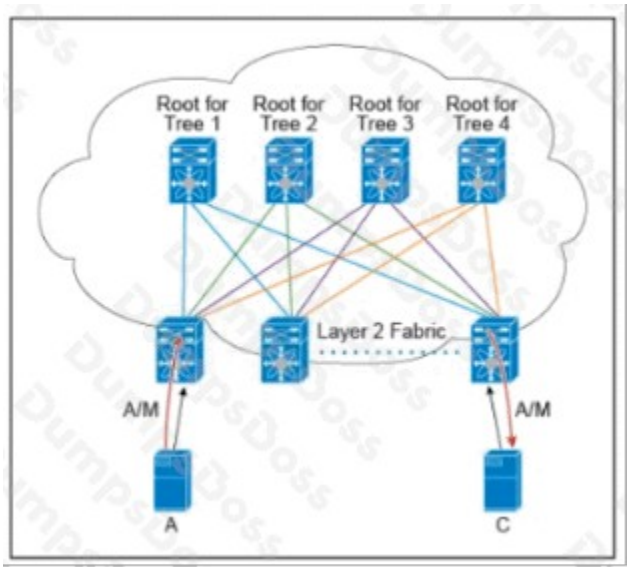

There are multiple trees in the Cisco FabricPath All switches in the Layer 2 fabric share the same view of each tree. Which two concepts describe how the multicast traffic is load-balanced across this topology? (Choose two)

A. A specific (S.G) traffic is not load-balanced

B. All trees are utilized at the same level of the traffic rate

C. Every leaf node assigns the specific (S.G) to the same tree.

D. A specific (S.G) multicast traffic is load-balanced across all trees due to better link utilization efficiency.

E. The multicast traffic is generally load-balanced across all trees

ANSWER: A C

Explanation:

In Cisco FabricPath, multicast forwarding uses multiple equal-cost trees (often referred to as multiple “topologies” or “trees”) so that multicast distribution can be spread across the fabric rather than pinned to a single spanning-tree instance. The key idea is that multicast traffic is mapped to a specific FabricPath tree based on the multicast flow (source/group), and that mapping is consistent across the fabric so all switches make the same forwarding decision for that flow. This provides deterministic forwarding (no loops) while also enabling load distribution: different multicast flows can be placed onto different trees, which improves overall link utilization and reduces the chance that one tree becomes a hotspot. As a result, a given (S,G) flow is carried on one selected tree, but across many (S,G) flows the fabric can leverage all available trees, effectively load-balancing multicast traffic across them for better efficiency. This behavior is part of FabricPath’s multi-tree forwarding model described in Cisco’s FabricPath documentation.

References: Cisco FabricPath configuration/operation overview, Cisco NX-OS FabricPath configuration guide

QUESTION NO: 8

Which main IoT migration aspect should be reviewed for a manufacturing plant?

A. Sensors

B. Security

C. Applications

D. Wi-Fi Infrastructure

E. Ethernet Switches

ANSWER: B

Explanation:

Security is the main IoT migration aspect that should be reviewed for a manufacturing plant because introducing or expanding IoT in an industrial environment materially changes the cyber-risk profile of the operation. Manufacturing plants commonly include OT/ICS assets (PLCs, HMIs, SCADA, drives, sensors, and gateways) that were not originally designed for Internet connectivity, frequent patching, or strong authentication. IoT migration often increases connectivity between IT and OT domains, adds new device identities, expands remote access needs, and introduces new protocols and management planes—all of which must be governed by a clear security architecture.

From a CCDE design perspective, security review spans segmentation (zones/conduits), least-privilege access, secure remote operations, device onboarding and lifecycle management, monitoring/telemetry, and incident response readiness. It also includes aligning the design to industrial security standards and guidance (for example, IEC 62443 concepts) and Cisco's validated approaches for industrial networks. Treating security as a primary migration aspect helps ensure availability and safety requirements are preserved while enabling new IoT capabilities.

References: Cisco Industrial Security, CISA Industrial Control Systems (ICS)

QUESTION NO: 9

Which two types of planning approaches are used to develop business-driven network designs and to facilitate the design decisions? (Choose two)

A. cost optimization approach

B. strategic planning approach

C. modular approach

D. tactical planning approach

E. business optimization approach

ANSWER: B D

Explanation:

Business-driven network design typically starts by translating business goals into a roadmap and then into actionable design choices. The strategic planning approach is used to align the network architecture with long-term business direction (for example, growth, mergers, cloud adoption, security posture, and service enablement). It helps define target-state capabilities, guiding principles, and high-level architecture decisions so the design supports business outcomes rather than just technical preferences.

The tactical planning approach complements this by focusing on near- to mid-term execution: prioritizing initiatives, selecting technologies and features, sequencing migrations, and defining concrete implementation steps and constraints (budget

cycles, operational readiness, risk, and timelines). Together, strategic and tactical planning provide a structured way to make design decisions that are traceable to business requirements and feasible to deliver.

This strategic/tactical split is consistent with Cisco's enterprise architecture and design methodologies that emphasize aligning technology strategy to business strategy and then operationalizing it through programs and projects. See Cisco's Enterprise Architecture resources at https://www.cisco.com/c/en/us/solutions/enterprise-networks/enterprise-architecture.html and Cisco Validated Design guidance at https://www.cisco.com/c/en/us/solutions/design-zone.html.

QUESTION NO: 10

The Service Provider you work for wants to deploy CoS functionality on the P routers of the MPLS core, to provide a complete CoS solution to all customers that purchase services such as dedicated Internet access, MPLS Layer 3 VPN, and Layer 2 VPN (pseudowire). The design requirements are - The network supports four service queues with equal treatment for delay, jitter, and packet loss. - Queues are numbered 03, where 0 is the default queue. - Three queues have one treatment, whereas the other queue has either one or two treatments. - The Service Provider manages control traffic, whereas the customers manage business critical and best effort. Which two statements could you recommend to allow for the appropriate level of bandwidth allocation? (Choose two.)

A. Control plane 80%. Real Time 10%, Business Critical 5%, Best Effort 5%

B. Control plane 20%, Real Time 40%, Business Critical 30%, Best Effort 10%

C. Control plane 90%, Real Time 5%, Business Critical 5%, Best Effort 0%

D. Control plane 10%, Real Time 30%, Business Critical 40%, Best Effort 20%

ANSWER: B D

Explanation:

"Control plane 20%, Real Time 40%, Business Critical 30%, Best Effort 10%" and "Control plane 10%, Real Time 30%, Business Critical 40%, Best Effort 20%" are appropriate recommendations because they allocate bandwidth across four queues in a way that matches common SP DiffServ/MPLS core design goals: protect provider-controlled traffic (routing/control/management) while still leaving the majority of capacity to customer traffic classes, and ensure every queue has an explicit, non-zero share so congestion management remains predictable. In an MPLS core, CoS is typically implemented with a small number of traffic classes mapped to a small number of hardware queues, using bandwidth guarantees (CBWFQ) and, where needed, strict priority for true real-time traffic. The requirement that three queues have one treatment and one queue has one or two treatments aligns with a model where most classes use standard bandwidth queuing while one class (often real-time) may additionally use priority/LLQ behavior. Both recommended splits keep control-plane bandwidth bounded (so it cannot starve customer services) while still reserving enough to maintain routing stability under stress, and they provide meaningful differentiation between real-time, business-critical, and best-effort. See Cisco's QoS design guidance for DiffServ and MPLS QoS models: Cisco DiffServ Overview and Cisco MPLS QoS Configuration Guide.

QUESTION NO: 11

Router R1 is a BGP speaker with one peering neighbor over link "A". When the R1 link/interface "A" fails, routing announcements are terminated, which results in the tearing down of the state for all BGP routes at each end of the link. What is this a good example of?

A. fault isolation

B. resiliency

C. redundancy

D. fate sharing

ANSWER: D

Explanation:

This is a good example of fate sharing. Fate sharing describes a design condition where multiple services, control-plane relationships, or forwarding outcomes depend on the same underlying component, so they fail together when that shared component fails. In this scenario, the BGP peering session and all routes learned over that session are dependent on the single physical/logical interface and link “A”. When that interface goes down, the BGP session is torn down (loss of TCP connectivity/keepalives), which immediately removes the associated BGP adjacency state and causes all routes learned via that neighbor to be withdrawn. The key point is that the control-plane relationship (the BGP session) and the reachability information (the BGP routes) share the same fate as the single link; there is no alternate path or independent mechanism keeping the session up. This is exactly the kind of shared-risk dependency CCDE designs try to identify and mitigate by introducing path diversity, alternate peerings, or other mechanisms that reduce shared failure domains. For additional background on BGP session behavior and how loss of connectivity impacts route advertisement/withdrawal, see Cisco BGP Troubleshooting Overview and RFC 4271 (BGP-4).

QUESTION NO: 12

You have been tasked with designing a data center interconnect as part of business continuity You want to use FCoE over this DCI to support synchronous replication. Which two technologies allow for FCoE via lossless Ethernet or data center bridging? (Choose two.)

A. DWDM

B. EoMPLS

C. SONET/SDH

D. Multichassis EtherChannel over Pseudowire

E. VPLS

ANSWER: B D

Explanation:

FCoE requires a lossless Ethernet service (Data Center Bridging features such as Priority Flow Control and ETS) because Fibre Channel traffic cannot tolerate drops the way typical TCP/IP applications can. Therefore, the DCI must preserve an Ethernet framing/service end-to-end so that DCB can be applied across the interconnect and the FCoE VLAN can be extended between sites. An Ethernet-over-MPLS service provides a Layer 2 Ethernet circuit (EVC) between data centers, which can be used to extend the DCB-capable Ethernet domain and carry FCoE frames transparently. Similarly, a multichassis EtherChannel over a pseudowire is a design pattern used to extend an Ethernet port-channel across a transport pseudowire, effectively presenting an Ethernet link/aggregation to the data center edge and enabling the extension of VLANs used for FCoE. These approaches align with Cisco guidance that FCoE is carried over lossless Ethernet and that L2 transport services (such as Ethernet pseudowires) are used when extending Ethernet-based storage traffic between sites. For background on FCoE/DCB requirements and Ethernet transport concepts, see Cisco FCoE overview and Cisco Pseudowire (PWE3) configuration guide.

QUESTION NO: 13

Refer to the exhibit.

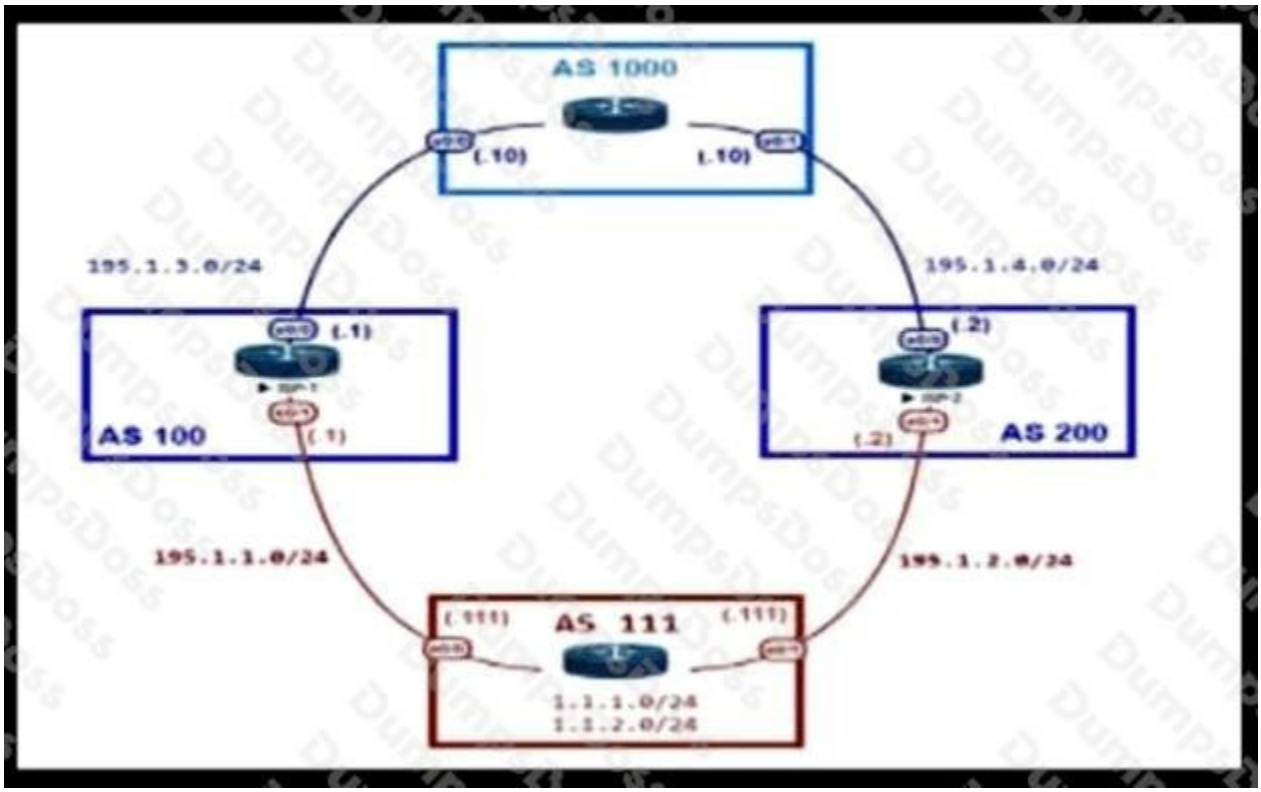

An engineer is designing the network for a multihomed customer running in AS 111 does not have any other Ass connected to it. Which technology is more comprehensive to use in the design to make sure that the AS is not being used as a transit AS?

A. Configure the AS-set attribute to allow only routes from AS 111 to be propagated to the neighbor ASs.

B. Use the local preference attribute to configure your AS as a non-transit" AS.

C. include an AS path access list to send routes to the neighboring ASs that only have AS 111 in the AS path field.

D. Include a prefix list to only receive routes from neighboring ASs.

ANSWER: C

Explanation:

include an AS path access list to send routes to the neighboring ASs that only have AS 111 in the AS path field is the most comprehensive approach because it directly enforces the core requirement for a multihomed “stub” AS: only advertise the customer’s own prefixes (originated by that AS) to upstream neighbors, and never advertise routes learned from one upstream to the other. In BGP, preventing transit is fundamentally an outbound policy problem—controlling what you export. By matching the AS_PATH so that the only AS present is the customer’s own AS, you ensure that any route learned from an external neighbor (which would include other ASNs in the AS_PATH) is denied from being advertised to the other neighbor. This is a robust design pattern for stub/multihomed customers because it scales beyond individual prefixes and continues to work even if the customer adds more internal prefixes, as long as they are originated within the same AS. Cisco documents using AS-path filters (often combined with prefix-lists/route-maps) as a standard BGP policy tool for controlling route

advertisement and preventing unintended transit behavior. See: Cisco BGP Filtering and Policy Control and Cisco BGP Configuration Guide and Technical Notes.

QUESTION NO: 14

Refer to the exhibit.

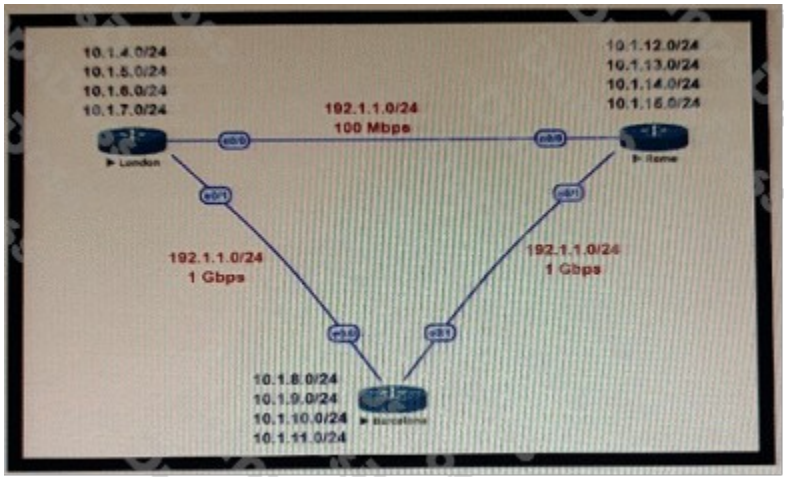

This network is running OSPF as the routing protocol. The internal networks are being advertised in OSPF London and Rome are using the direct link to reach each other although the transfer rates are better via Barcelona Which OSPF design change allows OSPF to calculate the proper costs?

A. Change the OSPF reference bandwidth to accommodate faster links.

B. Filter the routes on the link between London and Rome

C. Change the interface bandwidth on all the links.

D. Implement OSPF summarisation to fix the issue

ANSWER: A

Explanation:

Change the OSPF reference bandwidth to accommodate faster links is the correct design change because OSPF interface cost is derived from a simple formula: reference bandwidth divided by the interface bandwidth. By default, many OSPF implementations use a 100 Mbps reference bandwidth, which causes all links at 100 Mbps and above (FastEthernet, Gigabit, 10G, etc.) to collapse to the same cost value. When multiple high-speed paths end up with identical (or insufficiently differentiated) costs, OSPF may prefer a topologically “shorter” path (fewer hops) even if the actual throughput is better on an alternate path with higher-capacity links. Increasing the reference bandwidth (consistently across all routers in the OSPF domain) restores cost granularity so that faster links receive lower costs and OSPF can select the path that better reflects real link speeds. This is a design-level fix because it preserves OSPF’s metric behavior without manipulating routing with filters or relying on potentially misleading interface bandwidth settings for non-routing purposes.

References: Cisco OSPF Cost Calculation and Reference Bandwidth, Cisco IOS XE OSPF Configuration Guide (auto-cost reference-bandwidth)

QUESTION NO: 15

Company XYZ Is running a redundant private WAN network using OSPF as the underlay protocol The current design accommodates for redundancy In the network, but it Is taking over 30 seconds for the network to reconverge upon failure Which technique can be Implemented In the design to detect such a failure in a subsecond?

A. STP

B. fate sharing

C. OSPF LFA

D. BFD

E. flex links

ANSWER: D

Explanation:

Bidirectional Forwarding Detection (BFD) is the correct technique to detect failures in subsecond timeframes in an OSPF-based WAN. OSPF's default failure detection is typically driven by hello/dead timers, which commonly results in multi-second (or longer) detection and thus slower convergence. BFD provides a lightweight, protocol-independent mechanism to rapidly detect forwarding-path failures between two adjacent devices by sending frequent control packets and declaring the session down after a small number of missed packets. When integrated with OSPF, BFD can immediately signal neighbor loss to the routing process, triggering fast reconvergence without relying on aggressive OSPF hello timers (which can increase CPU load and risk instability at scale). This makes BFD a common design best practice for subsecond failure detection on WAN links, including routed Ethernet, MPLS L3VPN PE-CE adjacencies, and other routed underlays where rapid detection is required.

References: Cisco - Bidirectional Forwarding Detection (BFD) Overview, Cisco IOS XE - Configuring BFD for OSPF

QUESTION NO: 16

A customer has a functional requirement that states HR systems within a data center should be segmented from other systems that reside in the same data center and same VLAN. The systems run legacy applications by using hard-coded IP addresses. Which segmentation method is suitable and scalable for the customer?

A. data center perimeter firewalling

B. VACLs on data center switches

C. transparent firewalling

D. routed firewalls

ANSWER: B

Explanation:

VACLs on data center switches is the most suitable and scalable approach when you must enforce segmentation between endpoints that remain in the same VLAN and cannot tolerate IP addressing changes due to hard-coded IP dependencies. VLAN access control lists are applied to an entire VLAN and can filter traffic between hosts within that VLAN (intra-VLAN), which is exactly the problem statement: HR systems and non-HR systems share the same VLAN but still require isolation.

Because VACL enforcement happens in the switching infrastructure, it avoids redesigning the L3 boundary, avoids introducing new routed hops, and does not require renumbering or moving workloads to different subnets—key constraints for legacy applications. From a scalability perspective, VACLs can be centrally managed per VLAN and can be combined with object-group style constructs (platform-dependent) and operational processes to consistently enforce policy across many hosts without inserting additional inline devices per segment. Cisco documents VACLs as a mechanism to control traffic within a VLAN and between VLANs, making them a practical tool for intra-VLAN segmentation in campus/data center switching designs. See Cisco - Configuring ACLs (includes VLAN ACL concepts) and Cisco Nexus - Security Command Reference (VACL/ACL features).

QUESTION NO: 17

Which two control plane policer designs must be considered to achieve high availability? (Choose two.)

A. Control plane policers are enforced in hardware to protect the software path, but they are hardware platform dependent in terms of classification ability.

B. Control plane policers are really needed only on externally facing devices.

C. Control plane policers can cause the network management systems to create false alarms.

D. Control plane policers must be processed before a forwarding decision is made.

E. Control plane policers require that adequate protocols overhead are factored in to allow protocol convergence.

ANSWER: A E

Explanation:

For high availability, control-plane policing (CoPP/CPPr) must be designed so it protects the CPU without unintentionally degrading routing stability and operations. It is correct that control plane policers are enforced in hardware to protect the software path, but they are hardware platform dependent in terms of classification ability. In practice, CoPP is implemented using platform-specific capabilities (different match granularity, queueing, and punt-path behavior), so an HA design must account for what the specific hardware can classify and police consistently—especially across redundant supervisors/line cards or mixed platforms.

It is also correct that control plane policers require that adequate protocols overhead are factored in to allow protocol convergence. HA depends on keeping essential control-plane traffic (routing protocols, adjacency/keepalives, ARP/ND, etc.) within policer rates during bursts (e.g., reconvergence, link flaps, ISSU/SSO events). If rates are set too tightly without headroom, CoPP can drop legitimate control traffic, slowing or preventing convergence and creating instability during failures—the exact moments HA is most critical.

References: Cisco – Control Plane Policing (CoPP) Overview, Cisco – Control Plane Policing Configuration Guide

QUESTION NO: 18

You have been asked to design a high-density wireless network for a university campus. Which two principles would you apply in order to maximize the wireless network capacity? (Choose two.)

A. Implement a four-channel design on 2.4 GHz to increase the number of available channels

B. Choose a high minimum data rate to reduce the duty cycle.

C. increases the number of SSIDs to load-balance the client traffic.

D. Make use of the 5-GHz band to reduce the spectrum utilization on 2.4 GHz when dual-band clients are used.

E. Enable 802.11n channel bonding on both 2.4 GHz and 5 GHz to increase the maximum aggregated cell throughput.

ANSWER: B D

Explanation:

To maximize capacity in a high-density campus WLAN, the most impactful design principles are to reduce airtime consumption and to increase usable spectrum. Choosing a high minimum data rate is a classic capacity technique because it discourages (or disassociates) very low-rate clients and reduces the amount of airtime required per frame. Since 802.11 is a shared medium, airtime is the limiting resource; higher basic/minimum rates shorten transmission time for management/control and client traffic, improving overall cell efficiency and allowing more users to be served with acceptable performance.

Making use of the 5-GHz band is also fundamental for high density because it offers substantially more non-overlapping channels than 2.4 GHz, enabling better frequency reuse and lower co-channel contention. Steering dual-band clients to 5 GHz reduces congestion on 2.4 GHz and increases total system capacity by spreading clients across more spectrum. These principles align with Cisco high-density WLAN guidance: optimize for airtime efficiency and leverage 5-GHz channel availability for capacity and reuse.

References: Cisco Wireless LAN Design Guide for High Density Client Environments, Cisco WLAN Radio Frequency (RF) Design Basics

QUESTION NO: 19

You were tasked to enhance the security of a network with these characteristics:

• A pool of servers is accessed by numerous data centers and remote sites

• The servers are accessed via a cluster of firewalls

• The firewalls are configured properly and are not dropping traffic

• The firewalls occasionally cause asymmetric routing of traffic within the server data center.

Which technology should you recommend to enhance security by limiting traffic that could originate from a hacker compromising a workstation and redirecting flows at the servers?

A. Poison certain subnets by adding static routes to Null0 on the core switches connected to the pool of servers.

B. Deploy uRPF strict mode.

C. Limit sources of traffic that exit the server-facing interface of the firewall cluster with ACLs.

D. Deploy uRPF loose mode

ANSWER: D

Explanation:

Deploy uRPF loose mode is the best fit because it provides source-address validation while tolerating the asymmetric routing that is explicitly present in the server data center. The threat described is a compromised workstation attempting to redirect flows at the servers (for example, by spoofing source addresses to appear as a trusted host/subnet). Unicast Reverse Path

Forwarding helps mitigate source spoofing by checking whether the router has a route to the packet's source address in the routing table. In loose mode, the check is simply "is there any route to the source," rather than requiring the return path to use the same interface the packet arrived on. That makes it appropriate in environments with multiple paths, firewall clusters, and occasional asymmetry, where strict mode would create false drops. Using uRPF at the appropriate ingress points (often at the data center edge or aggregation) reduces the ability of an attacker to inject traffic with illegitimate source addresses toward the server pool, improving overall security without relying on perfectly symmetric forwarding. See Cisco's uRPF overview and configuration guidance: Cisco uRPF configuration and concepts and Cisco uRPF feature documentation.

QUESTION NO: 20 - (DRAG DROP)

The network team in XYZ Corp wants to modernize their infrastructure and is evaluating an implementation and migration plan to allow integration MPLS-based, Layer 2 Ethernet services managed by a service provider to connect branches and remote offices. To decrease OpEx and improve

response times when network components fail, XYZ Corp decided to acquire and deploy new routers. The network currently is operated over E1 leased lines (2 Mbps) with a managed CE service provided by the telco.

Drag and drop the implementation steps from the left onto the corresponding targets on the right in the correct order.

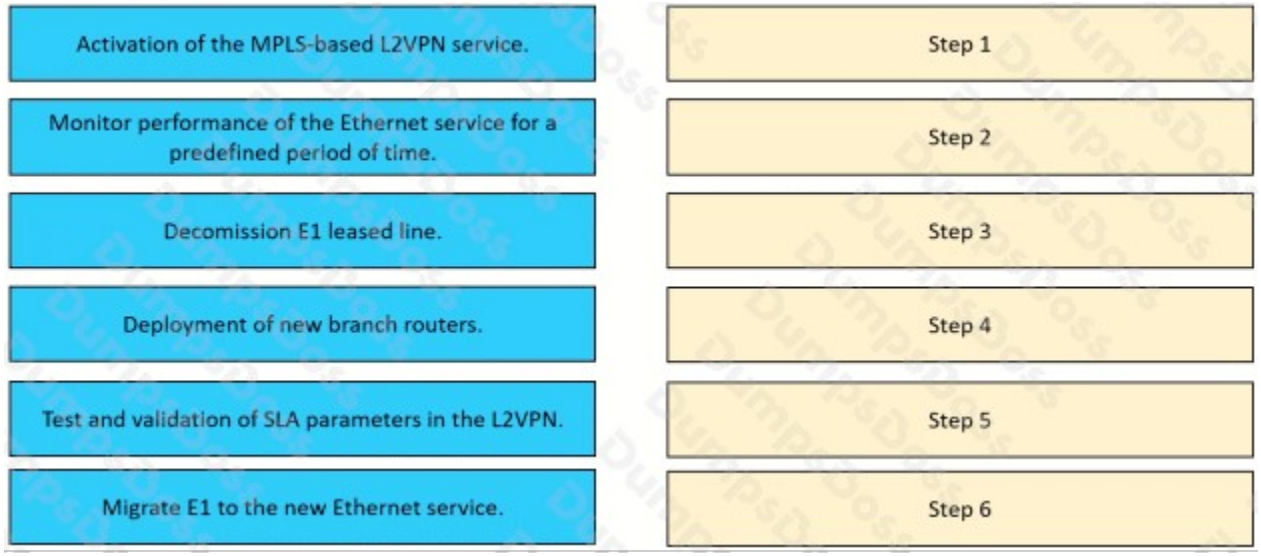

ANSWER:

Explanation:

The correct order follows a typical “build → turn up → accept → migrate → stabilize → retire” migration lifecycle used in enterprise WAN refresh projects. You start by deploying the new branch routers so the customer edge is physically installed, powered, cabled, and ready for cutover activities. With the CE in place, the service provider can activate the MPLS-based L2VPN service and hand off the Ethernet access, giving you an operational underlay/transport to work with.

Once the service is up, you validate it against the contracted service levels (latency, jitter, loss, availability, and any class-of-service behavior) before moving production traffic. This acceptance step is important because it confirms the provider service is delivering what was purchased and that the new CE configuration (interfaces, QoS, routing/bridging as applicable) behaves as designed. After acceptance, you migrate the existing E1-based connectivity to the new Ethernet/L2VPN service, typically via a planned cutover window, ensuring routing adjacencies and application reachability are restored on the new path.

After migration, you monitor performance for a predefined stabilization period to catch intermittent issues (errors, microbursts, QoS mis-marking, MTU problems, or provider-side impairments) that may not appear during initial testing. Only when the new service is proven stable do you decommission the E1 leased line, avoiding premature teardown that could remove your fallback option. This sequence aligns with common provider service turn-up and acceptance practices and with phased migration guidance for WAN transitions. References: Cisco WAN and Branch Solutions, Cisco Layer 2 VPN Overview.

QUESTION NO: 21

What is an architectural framework created by ETSI that defines standards to decouple network functions from proprietary hardware-based appliances and have them run in software on standard x86 servers?

A. NPIV

B. NFVIS

C. NFV

D. VNF

ANSWER: C

Explanation:

Network Functions Virtualization is the ETSI-defined architectural framework that standardizes how network functions (such as firewalls, load balancers, EPC/IMS components, etc.) are decoupled from dedicated, proprietary appliances and instead implemented as software running on industry-standard compute platforms (commonly x86) using virtualization technologies. ETSI NFV specifies the overall architecture and key building blocks—such as the NFV Infrastructure (NFVI), Virtual Network Functions, and the Management and Orchestration (MANO) framework—so that vendors and operators can build interoperable solutions. This is precisely the “framework created by ETSI” aspect in the question: it’s not a single product or a single virtualized function, but the standards-based reference architecture that enables those functions to run on standard servers and be lifecycle-managed consistently. This approach improves agility, scaling, and service deployment speed by treating network functions more like software workloads than fixed hardware appliances.

References: ETSI NFV, Cisco: Network Functions Virtualization (NFV)

QUESTION NO: 22

Which SDN architecture component is used by the application layer to communicate with the control plane layer to provide instructions about the resources required by applications?

A. southbound APIs

B. northbound APIs

C. orchestration layer

D. SDN controller

ANSWER: B

Explanation:

northbound APIs is correct because it is the interface exposed by the SDN control plane (typically the SDN controller) to the application layer, allowing applications and higher-level business logic to express intent and requirements to the controller. Through northbound APIs, applications can request network services (for example, connectivity, QoS, security policy, traffic engineering constraints, or segmentation) without needing to manage device-specific configurations. The controller translates these application-level requirements into network-wide policies and then programs the underlying infrastructure using southbound mechanisms. This separation is a core SDN principle: applications communicate “down” to the controller via northbound APIs, while the controller communicates “down” to network devices via southbound APIs. In practice, northbound APIs are often RESTful interfaces and/or intent-based abstractions that enable automation and integration with orchestration systems and network applications. This is consistently described in Cisco’s SDN architecture discussions where northbound interfaces connect applications to the controller and southbound interfaces connect the controller to forwarding devices.

References: Cisco — Software-Defined Networking (SDN) Overview, RFC 7426 — Software-Defined Networking (SDN): Layers and Architecture Terminology

QUESTION NO: 23

Which methodology is the leading lifecycle approach to network design and implementation?

A. PPDIOO

B. Waterfall model

C. Spiral model

D. V model

ANSWER: A

Explanation:

PPDIOO is the leading lifecycle approach to network design and implementation in Cisco-centric enterprise network projects. It is Cisco's structured methodology that spans the full lifecycle: Prepare (establish business case, strategy, and high-level requirements), Plan (detailed requirements and project planning), Design (logical and physical design), Implement (deployment and migration), Operate (day-to-day operations), and Optimize (continuous improvement and proactive tuning). This end-to-end lifecycle focus is what makes it particularly suitable for network design and implementation work, because it explicitly ties technical architecture decisions to business requirements and then carries those decisions through deployment and ongoing operations, rather than stopping at "design complete." It is also widely used as a best-practice framework in Cisco training and design guidance, and it aligns well with iterative improvement and operational feedback loops that are critical in real-world networks. For additional background, see Cisco Press's overview of PPDIOO and lifecycle design concepts at https://www.ciscopress.com/articles/article.asp?p=1608131&seqNum=3 and Cisco's Enterprise Architecture/PPDIOO discussion at https://www.cisco.com/c/en/us/about/press/internet-protocol-journal/back-issues/table-contents-34/123-ppdioo.html.

QUESTION NO: 24

Refer to the table.

CONNECTIVITY	CAPEX	OPEX ANNUAL	INSTALLATION FEE	TERM
DWDM over dark fiber	$250,000	$100,000	$30,000	60 months
CWDM over dark fiber	$150,000	$100,000	$25,000	60 months
MPLS	$50,000	$150,000	$75,000	12 months
Metro Ethernet	$45,000	$125,000	$5,000	36 months

A customer investigates connectivity options for a DCI between two production data centers. The solution must provide dual 10G connections between locations with no single points of failure for Day 1 operations. It must also include an option to scale for up to 20 resilient connections in the second year to accommodate isolated SAN over IP and isolated, dedicated replication IP circuits. All connectivity methods are duplex 10 Gbps. Which transport technology costs the least over two years, in the scenario?

A. Metro Ethernet

B. DWDM

C. CWDM

D. MPLS

ANSWER: B

Explanation:

DWDM is the least-cost option over two years in this scenario because the requirement to scale from an initial resilient pair of 10G links to as many as 20 resilient connections strongly favors a transport that can add many additional 10G circuits without repeatedly purchasing new physical access loops or paying per-circuit managed service premiums. With DWDM, once the fiber pair and optical line system are in place, incremental capacity is typically added by lighting additional wavelengths (lambdas) and inserting transponders/muxponders, which is generally more cost-effective at higher circuit counts than buying many discrete Ethernet or MPLS services. DWDM is also designed for high-capacity, long-haul/metro DCI and supports resilient designs (diverse paths, protection switching) without creating single points of failure. In contrast, service-based options often scale linearly in cost with each additional isolated circuit, which becomes expensive when expanding to many dedicated SAN/replication connections. This “high fixed cost, low marginal cost” scaling characteristic is why DWDM commonly wins total cost of ownership when the number of 10G circuits grows significantly over time.

References: Cisco Data Center Interconnect overview, Cisco Optical Networking (DWDM) solutions

QUESTION NO: 25

Which two statements describe the hierarchical LAN design model? (Choose two)

A. It is a well-understood architecture that provides scalability

B. It is the best design for modern data centers

C. It is the most optimal design but is highly complex

D. It provides a simplified design

E. Changes, upgrades, and new services can be introduced in a controlled and stagged manner

ANSWER: A E

Explanation:

The hierarchical LAN design model is a long-standing Cisco campus design approach that breaks the network into functional layers (typically access, distribution, and core). This structure is widely adopted and well understood in the industry, which makes it easier to scale the network over time by adding capacity or expanding at the appropriate layer without redesigning everything end-to-end. Because each layer has a clear role, the model supports modular growth and predictable behavior as the campus expands.

Another key characteristic is operational control: the layered separation allows changes, upgrades, and the introduction of new services to be rolled out in a staged, controlled way. For example, you can introduce new access-layer features (like updated edge security or PoE requirements) while keeping the core stable, or adjust distribution-layer policies (like routing, summarization, and filtering) with reduced blast radius. This aligns with Cisco’s guidance on modular, hierarchical enterprise campus design and lifecycle manageability. References: Cisco Design Zone for Campus, Cisco Campus Network Overview.

QUESTION NO: 26 - (DRAG DROP)

Drag and drop the multicast protocols from the left onto the current design situation on the right.

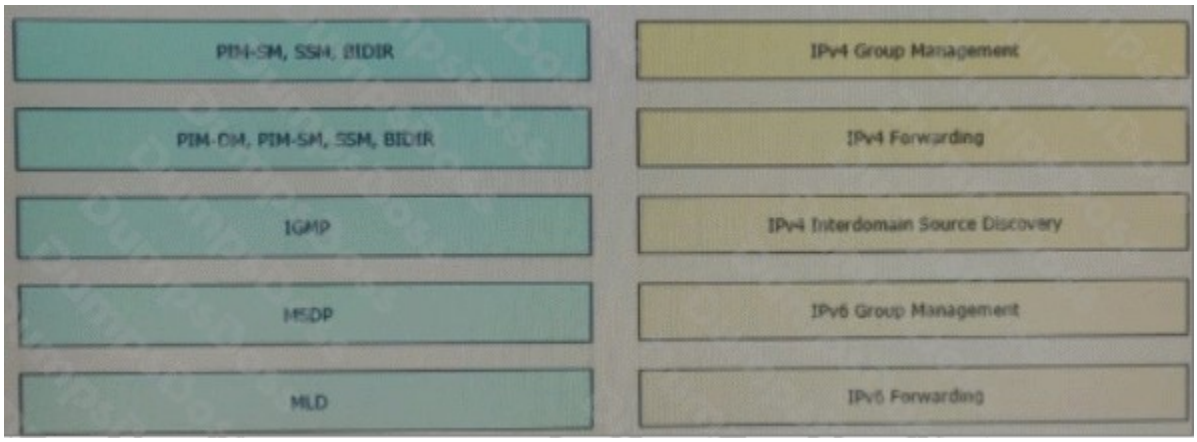

ANSWER:

Explanation:

This design question is really about placing each multicast-related protocol into the part of the multicast “control plane” where it operates, and separating IPv4 behavior from IPv6 behavior.

For IPv4 group management, hosts signal their interest in joining or leaving multicast groups to the local first-hop router using IGMP. That’s why IGMP belongs under “IPv4 Group Management.” For IPv6, the equivalent host-to-router membership protocol is MLD (Multicast Listener Discovery), so MLD belongs under “IPv6 Group Management.” These are explicitly defined as the host membership protocols for their respective IP versions and are not the routing/forwarding protocols themselves.

For multicast forwarding between routers inside a domain, Protocol Independent Multicast (PIM) is used to build distribution trees and drive multicast routing state. In IPv4 networks, you may encounter PIM Dense Mode (PIM-DM) as well as PIM Sparse Mode (PIM-SM), Source-Specific Multicast (SSM, typically using PIM-SM behavior with (S,G) joins), and Bidirectional PIM (BIDIR). Therefore, “IPv4 Forwarding” correctly maps to “PIM-DM, PIM-SM, SSM, BIDIR.” In IPv6, PIM-SM, SSM, and BIDIR are used, but PIM-DM is not generally used/standardized in the same way for IPv6 operational designs, so “IPv6 Forwarding” correctly maps to “PIM-SM, SSM, BIDIR.”

Finally, interdomain source discovery for Any-Source Multicast across different PIM-SM domains is historically handled by MSDP, which allows RPs in different domains to learn about active sources. That’s why MSDP belongs under “IPv4 Interdomain Source Discovery.”

References: Cisco IP Multicast Technology Overview, RFC 3376 (IGMPv3), RFC 3810 (MLDv2), RFC 3618 (MSDP).

QUESTION NO: 27

You are the lead IP network designer for a new service provider called XYZ, and you are working closely with the CTO to finalize design requirements. The CTO informs you that they want to transport IPv6 prefixes of customers through the XYZ network at this time; however, they need your advice on whether to deploy dual stack or MPLS 6PE/6VPE. Which two options do you recommend? (Choose two.)

A. Prepare the dual-stack infrastructure from the beginning, even if BGP prefixes would have to be announced via IPv4 in case you decide to maintain the BGPfree core

B. Use MPLS 6PE to simplify the operation and keep a BGP-free core. When the LDPv6 becomes available, change to 4PE or keep the core using both IPv4 and IPv6. The main goal is to keep the core BGP-free and ensure that IPv4, IPv6, VPNv4, and VPNv6 are all label-switched

C. Use MPLS 6VPE to simplify the operation and keep a BGP-free core. When the LDPv6 becomes available, change to 4PE or keep the core using both IPv4 and IPv6. The main goal is to keep the core BGP-free and ensure that IPv4, IPv6, VPNv4, and VPNv6 are all label-switched.

D. Build a dual-stack network. Enable BGP in the core. Redistribute EBGP routes into IGP.

ANSWER: A B C

Explanation:

For a service provider that needs to carry customer IPv6 while keeping the provider core operationally simple, the two broadly recommended approaches are to either build the infrastructure as dual stack from day one or to use MPLS-based IPv6 transport that avoids enabling IPv6 throughout the core. Preparing a dual-stack infrastructure early is a common best practice because it avoids repeated redesign later: core links, IGP, addressing plans, management/telemetry, and security controls can be made IPv6-capable even if the initial transport still relies on an IPv4 MPLS core and IPv4-based signaling/peering choices. In parallel, MPLS 6VPE is a well-established method to deliver IPv6 VPN services across an IPv4 MPLS backbone, leveraging MP-BGP to carry VPNv6 routes while keeping the P routers IPv4/MPLS-only (a “BGP-free core” from the transit perspective). This aligns with the goal of transporting customer IPv6 prefixes without requiring end-to-end IPv6 in the core. These recommendations match Cisco’s documented IPv6-over-MPLS service models and migration guidance for providers balancing feature delivery with operational risk.

References: Cisco MP-BGP/MPLS VPN Configuration Guide, Cisco IPv6 Design and Deployment