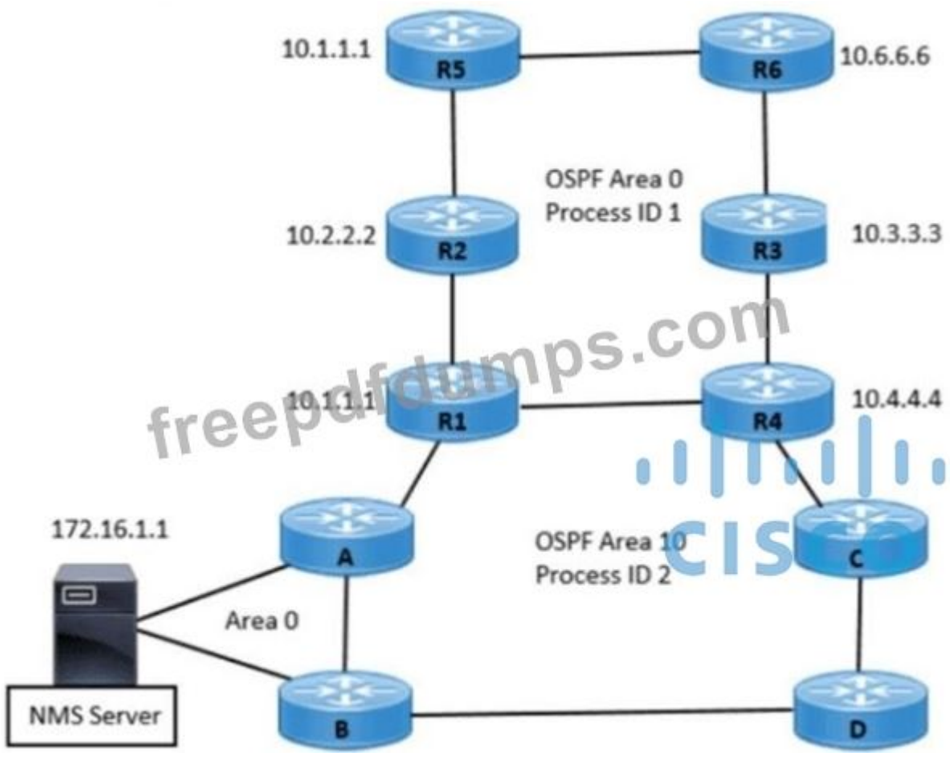
Cisco.300-420.v2026-09-21.q157

Exam Code:	300-420
Exam Name:	Designing Cisco Enterprise Networks
Certification Provider:	Cisco
Free Question Number:	157
Version:	v2026-09-21
# of views:	107
# of Questions views:	1570
https://www.freepdfdumps.com/Cisco.300-420.v2026-09-21.q157.html	

NEW QUESTION: 1



Refer to the exhibit An engineer is designing an OSPF solution with these requirements:

* NMS server will manage R5 and R6.

* Upon failure of R1. all NMS traffic should be routed through R4.

* Upon failure of the link between R5 and R6. all traffic destined for 10.6.6.6 should be routed through R4 Which solution must the engineer choose?

A. Advertise 172.16.1.1 into OSPF process 1 with high cost on R1.

B. Apply static routes on R2 and R3 with IP SLA tracking toward R5 and R6.

C. Enable the default-Information originate command with a higher metric on R2 to R1.

D. Redistribute OSPF process 1 into process 2 on R1 and R4.

Answer: D (LEAVE A REPLY)

Redistributing OSPF process 1 into process 2 on both R1 and R4 provides the required backup routing behavior between the management network and the target routers. The design requirements are failure driven:

if R1 fails, NMS traffic must use R4, and if the link between R5 and R6 fails, traffic to 10.6.6.6 must still route through R4. When separate OSPF processes or domains exist, controlled redistribution at both exit points allows reachability to be maintained through either redistribution point. Advertising a single address into OSPF with high cost on R1 does not solve both failure cases. Static routes with IP SLA on R2 and R3 would be more manual and less aligned with the dynamic OSPF design. Default-information originate would provide a default path but would not accurately advertise the specific internal reachability required for the NMS destinations. The redistribution design must also include route filtering or tagging in production to prevent feedback loops. Reference topics: OSPF redistribution, redundant redistribution points, management reachability, route control, failure-path design.

NEW QUESTION: 2



Refer to the exhibit. A customer needs to apply QoS to the network management traffic passing through the GigabitEthernet0/2 interface. All eight queuing classes are in use, so the new requirement must be integrated into the existing policy. Which solution must the customer choose?

A. Mark traffic to DSCP CS5 and assign it to the SIGNALLING class. Then, baseline existing queue sizes to determine if additional bandwidth can be provisioned to the SIGNALLING class.

B. Mark the traffic to DSCP CS4 and assign it to the SIGNALLING class. Then, prioritize traffic within the class.

C. Mark the traffic to DSCP CS6 and assign it to the ROUTING class Then, prioritize traffic within the class.

D. Mark the traffic to DSCP CS2 and assign it to the ROUTING class Then, baseline existing queue sizes to determine if additional bandwidth can be provisioned to the ROUTING class

Answer: D (LEAVE A REPLY)

Network management traffic should be marked and treated as a dedicated operations class, not mixed into voice signaling or real-time queues. Cisco enterprise QoS design guidance commonly reserves DSCP CS2 for operations, administration, and management traffic such as SNMP, SSH, syslog, NTP, and other management flows. The question states that all eight queuing classes are already used, so the design should integrate management traffic into the closest existing class and validate that the queue has enough bandwidth. Option D is the best fit because it uses CS2 and requires baselining before allocating more bandwidth. CS6 is normally used for network control and routing protocols; using it for ordinary management traffic would compete with routing adjacencies and control-plane stability. CS5 and CS4 are also unsuitable because they are normally associated with video/signaling-related classes in many enterprise QoS models. The professional design action is to classify management traffic explicitly, mark it CS2 at the trust boundary, map it to the existing class, and then measure queue utilization before changing bandwidth allocations. Reference topics: Cisco QoS baseline, DSCP CS2, network management traffic, queue bandwidth design.

NEW QUESTION: 3

What is the purpose of a Cisco SD-Access underlay network?
A. to abstract IP-based connectivity from physical connectivity
B. to emulate LAN segments to transport Layer 2 frames over a Layer 3 network
C. to establish physical connectivity between switches and routers
D. to provide virtualization by encapsulating network traffic over IP tunnels
Answer: C (LEAVE A REPLY)
The Cisco SD-Access underlay provides the physical and IP transport foundation that allows fabric nodes to reach each other. In plain terms, it is the routed connectivity built from the switches, routers, links, and loopback reachability that the overlay depends on. Cisco SD-Access then uses the overlay to provide virtualization, segmentation, policy, and endpoint mobility on top of that transport. This is why the underlay is best described as establishing the connectivity between the switches and routers that form the fabric. Option A describes the abstraction performed by the overlay, not the underlay. Option B is closer to a Layer 2 tunneling description and is not the main purpose of the SD-Access underlay. Option D describes encapsulated overlay behavior, especially VXLAN-based data-plane transport, not the physical/routed foundation itself. A good SD- Access design validates underlay reachability, routing adjacencies, MTU, loopback advertisement, and redundancy before onboarding the overlay services. Without a stable underlay, LISP control-plane mappings and VXLAN data-plane forwarding cannot operate predictably. Reference topics: SD-Access underlay, routed campus foundation, fabric node reachability, overlay dependency.

NEW QUESTION: 4
An engineer uses Postman and YANG to configure a router with:
* OSPF process ID 400
* network 192.168.128.128/25 enabled for Area 0
Which get-config reply verifies that the model set was designed correctly?

A.

```
<rpc-reply message-id="urn:uuid:1b3d05cd-8118-3e6a-6c05-021345678aaf" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:
 <data>
  <native xmlns="http://cisco.com/ns/yang/ned/ios">
   <router>
    <ospf>
     <id>400</id>
     <network>
      <ip>1192.168.128.128</ip>
      <mask>0.0.0.128</mask>
      <area>0</area>
     </network>
    </ospf>
   </router>
  </native>
 </data>
</rpc-reply>
```

cleardfdumps.com

<rpc-reply message-id="urn:uuid:1b3d05cd-8118-3e6a-6c05-403478311aaf" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:nc="urn:ietf:para
<data>
<native xmlns="http://cisco.com/ns/yang/ned/ios">
<router>
<ospf>
<id>400</id>
<network>
<ip>192.168.128.128</ip>
<mask>0.0.0.127</mask>
<area>0</area>
</network>
</ospf>
</router>
</native>
</data>
</rpc-reply>

B.

cisco
freepdfdumps.com

C.

<rpc-reply message-id="urn:uuid:1b3d05cd-8118-3e6a-6c05-012354678aaf" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:nc="u
<data>
<native json="http://cisco.com/ns/yang/ned/ios">
<router>
<ospf>
<id>400</id>
<network>
<ip>192.168.128.128</ip>
<mask>0.0.0.127</mask>
<area>0</area>
</network>
</ospf>
</router>
</native>
</data>
</rpc-reply>

cisco
freepdfdumps.com

D.

```
<rpc-reply message-id="urn:uuid:1b3d05cd-8118-3e6a-6c05-012435678aaf" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
 <data>
  <native xmlns="http://cisco.com/ns/yang/ned/ios">
   <router>
    <ospf>
     <id>400</id>
     <network>
      <ip>192.168.128.128</ip>
      <mask>255.255.255.128</mask>
      <area>0</area>
     </network>
    </ospf>
   </router>
  </native>
 </data>
</rpc-reply>
```

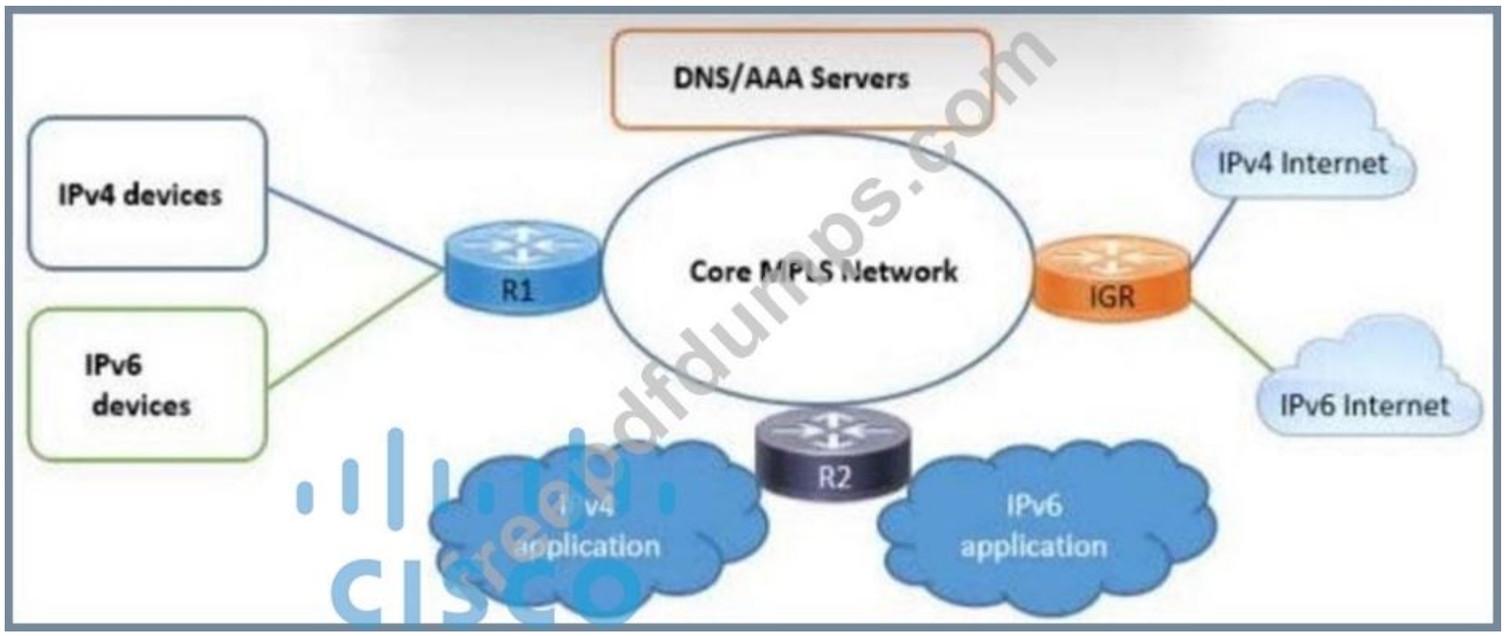
Answer: (SHOW ANSWER)

The valid get-config reply must show OSPF process ID 400 with network 192.168.128.128/25 enabled for Area 0. A /25 prefix has a block size of 128 addresses, so 192.168.128.128/25 is a valid network boundary and covers addresses from 192.168.128.128 through 192.168.128.255. Any reply that shows the wrong OSPF process, an incorrect network boundary, the wrong prefix length or wildcard equivalent, or a nonzero area does not verify the requested model set. In Cisco IOS XE model-driven configuration, the readback has to match the schema path and values used by the model, not simply produce an approximate CLI interpretation.

Option B is the reply that confirms the intended OSPF object. This verification method is important because YANG automation can fail silently in operational workflows when the payload is syntactically valid but targets the wrong container or key. A get-config verification step confirms the running datastore contains the exact OSPF process and area membership required by the design. Reference topics: YANG OSPF configuration, NETCONF get-config, prefix validation, IOS XE native models.

NEW QUESTION: 5

Refer to the exhibit.



An architect must design an IPv6 migration solution for an enterprise customer to support these requirements:

* Clients will transition to the new IPv6 network, which provides NAT64 and IPv6 DNS resolution services, using the same DNS name that points to the IPv4 address.

* The service provider will create a client-facing IPv6 interface with a new IPv6 virtual address that points to the same IPv4 DNS server.
* The service provider will support clients that use global IPv6 addresses and encapsulate IPv4 packets into IPv6 tunnels.
Which two migration solutions must the architect choose? (Choose two.)
A. Use dual-stack lite from the MPLS network to the IGR.
B. Use IPv6 tunneling from the devices to the core MPLS network.
C. Use dual-stack lite from the devices to the core MPLS network.
D. Use NAT44/64 from the MPLS network to the IGR.
E. Use NAT44/64 from the devices to the core MPLS network.
Answer: (SHOW ANSWER)
The design combines Dual-Stack Lite from the devices toward the core MPLS network and NAT44/64 from the MPLS network toward the Internet gateway router. The requirements describe IPv6-facing clients that still need access to IPv4 services and DNS behavior that allows the same service name to be reached during migration. Dual-Stack Lite carries IPv4 customer traffic across an IPv6 access or provider network by tunneling IPv4 packets over IPv6 toward a provider translation point. NAT44/64 services then provide translation or address-family interworking for IPv4-only resources. This combination is used when customer- facing access moves to IPv6 while legacy IPv4 services remain reachable during transition. A simple IPv6 tunnel from devices to the MPLS core does not provide the required IPv4 service translation. Using only NAT44/64 from end devices also shifts complexity to clients rather than placing translation at the network edge. The selected choices match the staged migration model: IPv6 transport with encapsulation in the access
/core side and translation at the provider or gateway boundary. Reference topics: IPv6 transition, DS-Lite, NAT64, DNS64, IPv4-over-IPv6 tunneling.

NEW QUESTION: 6
Drag and drop the characteristics from the left onto the YANG modules they describe on the right. Not all options are used.



freepdfdumps.com
Answer:



Explanation:
The YANG module mapping follows the distinction between standards-based, vendor-neutral, and vendor- native data models. IETF YANG models are standards-driven and are appropriate when the design goal is protocol-level interoperability across vendors for common network functions. OpenConfig models are also vendor-neutral, but they are developed by the OpenConfig community and commonly use a structured separation between intended configuration and operational state. Cisco native models expose platform- specific Cisco IOS XE, IOS XR, or NX-OS capabilities and are often the most complete option when a feature is unique to a Cisco operating system. The selected drag-and-drop answer aligns each characteristic with the model family that owns that behavior. In a professional automation design, the engineer should not choose a model only because it is familiar. The correct choice depends on feature coverage, device support, platform release, schema stability, and whether the workflow must operate consistently across vendors. For multivendor intent, choose IETF or OpenConfig first; for Cisco-specific feature depth, use Cisco native models. Reference topics: YANG, IETF models, OpenConfig, Cisco native models, NETCONF and RESTCONF automation.

NEW QUESTION: 7
An engineer is designing a QoS solution for a campus. The design must guarantee real-time traffic delivery during congestion, minimize the bandwidth consumption for possible virus or worm attacks, and reduce flooding of excessive traffic during times of congestion. Which two solutions must the engineer select?
(Choose two.)
A. Create a shaping policy to drop excessive traffic and a strict queue for real-time traffic.
B. Apply queuing on the distribution to core links
C. Create a policing policy to drop excessive traffic and a strict queue for real-time traffic.
D. Create a scavenger queue for excessive traffic and a strict queue for real-time traffic
E. Apply queuing on the access to distribution links.
Answer: D,E (LEAVE A REPLY)
The design should use a scavenger queue for excessive or suspicious traffic and apply queuing on the access- to-distribution uplinks where congestion can occur. Cisco QoS campus design commonly uses a less-than- best-effort scavenger class to constrain traffic associated with worms, peer-to-peer applications, bulk background flows, or other nonbusiness traffic. Placing this traffic in a scavenger queue minimizes the bandwidth it can consume during congestion without harming real-time applications. Real-time voice or video traffic should be placed in a strict-priority queue so it receives low delay and low jitter treatment when links are congested. Applying queuing on access-to-distribution links is important because those uplinks are common campus congestion points, especially during bursts or abnormal traffic events. A shaping policy is not the right tool for campus LAN uplinks, and indiscriminate policing of real-time traffic can cause drops that damage voice and video quality. Reference topics: campus QoS, scavenger class, strict-priority queue, access uplink queuing, congestion management, worm mitigation.

NEW QUESTION: 8

A company requires a private WAN design that allows remote sites to connect to HQ. The design must ensure that:
* traffic is always encrypted
* forwarding overhead is reduced
* management of security is centralized
* multicast traffic is supported
Which technology must the company select?
A. IPsec P2P
B. mGRE
C. DMVPN Phase 3
D. GET VPN
Answer: (SHOW ANSWER)
GET VPN is the correct technology for this private WAN requirement. Cisco Group Encrypted Transport VPN secures unicast and multicast traffic over a private WAN using IPsec encryption and group key management. Its key advantage is that it preserves the original packet header instead of building permanent point-to-point tunnels, which reduces forwarding overhead compared with GRE/IPsec or DMVPN overlays.
Cisco describes GET VPN as using a centralized key server to distribute group keys and security policy to group members, satisfying the requirement for centralized security management. It also supports multicast efficiently because the core or provider network can perform multicast replication while the traffic remains encrypted. IPsec point-to-point would not scale well across many remote sites and lacks group key management. mGRE alone does not encrypt traffic. DMVPN Phase 3 supports dynamic encrypted tunnels and spoke-to-spoke communication, but it still relies on tunnel encapsulation and does not reduce overhead the way GET VPN does on a private WAN. Reference topics: GET VPN, GDOI, IPsec group encryption, private WAN, multicast encryption.

NEW QUESTION: 9
Which protocol is the Cisco SD-Access data plane based on?
A. OMP
B. VXLAN
C. NHRP
D. LISP
Answer: B (LEAVE A REPLY)
The Cisco SD-Access data plane is based on VXLAN. Cisco SD-Access separates the fabric into functional planes: the underlay provides IP reachability, the control plane uses LISP to map endpoints to locations, and the data plane uses VXLAN encapsulation to carry user traffic across the routed underlay. VXLAN allows the fabric edge node to encapsulate endpoint traffic and forward it through the fabric while preserving virtual network and segmentation information. This is what enables SD-Access to provide Layer 2 and Layer 3 overlay services over a Layer 3 routed transport. OMP is the control-plane protocol for Cisco SD-WAN, not SD-Access. NHRP is associated with DMVPN tunnel resolution. LISP is essential in SD-Access, but its role is endpoint ID to routing locator mapping, not packet encapsulation in the data plane. Therefore, VXLAN is the correct protocol for the SD-Access data plane. Reference topics: SD-Access fabric, VXLAN data plane, LISP control plane, overlay encapsulation, routed underlay.

NEW QUESTION: 10
A company plans to transition to IPv6. They will link their IPv4 addresses to the lowest significant bits of the new Ipv6 addresses. A network administrator with an employee id: 4264:42:116 is preparing a mapping schema for the new IPv6 addresses. Which address does the 172.16.10.0/24 network translate to?
A. 2001:db8:abcd::ac10:a00/120
B. 2001:db8:abcd:172:16:10::/96
C. 2001:db8:abcd:11d8:a00/120
D. 2001:db8:ac10:0a00::/64
Answer: A (LEAVE A REPLY)

The IPv4 network 172.16.10.0/24 maps to 2001:db8:abcd::ac10:a00/120 when the IPv4 address is placed into the lowest significant bits of the IPv6 address. The calculation converts each IPv4 octet into hexadecimal: 172 is AC, 16 is 10, 10 is 0A, and 0 is 00. Combining those bytes gives AC10:0A00, displayed in IPv6 notation as ac10:a00. Because a /24 IPv4 network has 8 host bits, the mapped IPv6 network leaves the same 8 host bits, resulting in a /120 IPv6 prefix. Option B incorrectly places decimal IPv4 octets into separate IPv6 hextets and uses /96, which represents a much larger mapping block than the /24 network. Option C does not represent
172.16.10.0 in hexadecimal. Option D changes the higher-order IPv6 prefix and does not match the supplied address plan. The correct address therefore preserves the planned IPv6 prefix and embeds the IPv4 network in the low-order bits. Reference topics: IPv4-to-IPv6 address mapping, hexadecimal conversion, IPv6 prefix length, migration addressing plans.

NEW QUESTION: 11
Which type of rendezvous point deployment is standards-based and support dynamic RP discovery?
A. Auto-RP
B. Anycast-RP
C. bootstrap router
D. static RP
Answer: C (LEAVE A REPLY)
Bootstrap Router is the standards-based mechanism for dynamic Rendezvous Point discovery in PIM Sparse Mode. Cisco supports several RP discovery methods, but they are not equivalent from a standards perspective.
Auto-RP is Cisco-developed and widely used in Cisco environments, but it is not the standards-based choice.
Static RP configuration is simple and deterministic, but it does not provide dynamic discovery and requires every router to be configured consistently. Anycast RP is a redundancy and load-sharing design in which multiple RPs share the same RP address, usually with MSDP or another synchronization method, but it is not the RP discovery mechanism itself. BSR allows candidate RPs and candidate BSRs to advertise RP-set information through the PIM domain so routers can learn RP mappings dynamically. This makes BSR suitable for large PIM domains where manual static RP configuration becomes operationally risky. In design, the architect should also protect BSR/RP placement, scope group ranges, and validate convergence behavior during RP or BSR failure. Reference topics: PIM-SM, Bootstrap Router, RP discovery, Auto-RP comparison, multicast control-plane design.

NEW QUESTION: 12
An architect is designing a network for an enterprise site. The design must use an active/backup design for the WAN. It must guarantee the SLA for several applications regardless of which connection is used. Which deployment model should the architect choose?
A. MPLS WAN from two separate ISPs
B. hybrid WAN using MPLS VPN and internet VPN from a single ISP
C. hybrid WAN using MPLS VPN and internet VPN from two separate ISPs
D. internet WAN from two separate ISPs
Answer: (SHOW ANSWER)
An MPLS WAN from two separate ISPs is the strongest fit when the enterprise requires an active/backup WAN design with guaranteed SLAs for applications on either connection. MPLS VPN services are typically delivered with provider-backed service-level commitments for latency, jitter, packet loss, and availability, which is why they are selected for predictable enterprise application transport. Using two separate MPLS providers improves resiliency against a provider failure while preserving SLA-backed behavior on the backup path. A hybrid MPLS and Internet VPN model can be cost effective, but the Internet path generally does not provide the same deterministic SLA unless enhanced by a specific managed service contract. Internet from two providers improves reachability diversity but does not inherently guarantee application SLA. A single ISP hybrid design creates a provider dependency and weakens availability. Therefore, dual-provider MPLS is the conservative design for guaranteed active/backup WAN service. Reference topics: MPLS WAN design, service-level agreements, active/backup WAN, provider diversity, application performance assurance.

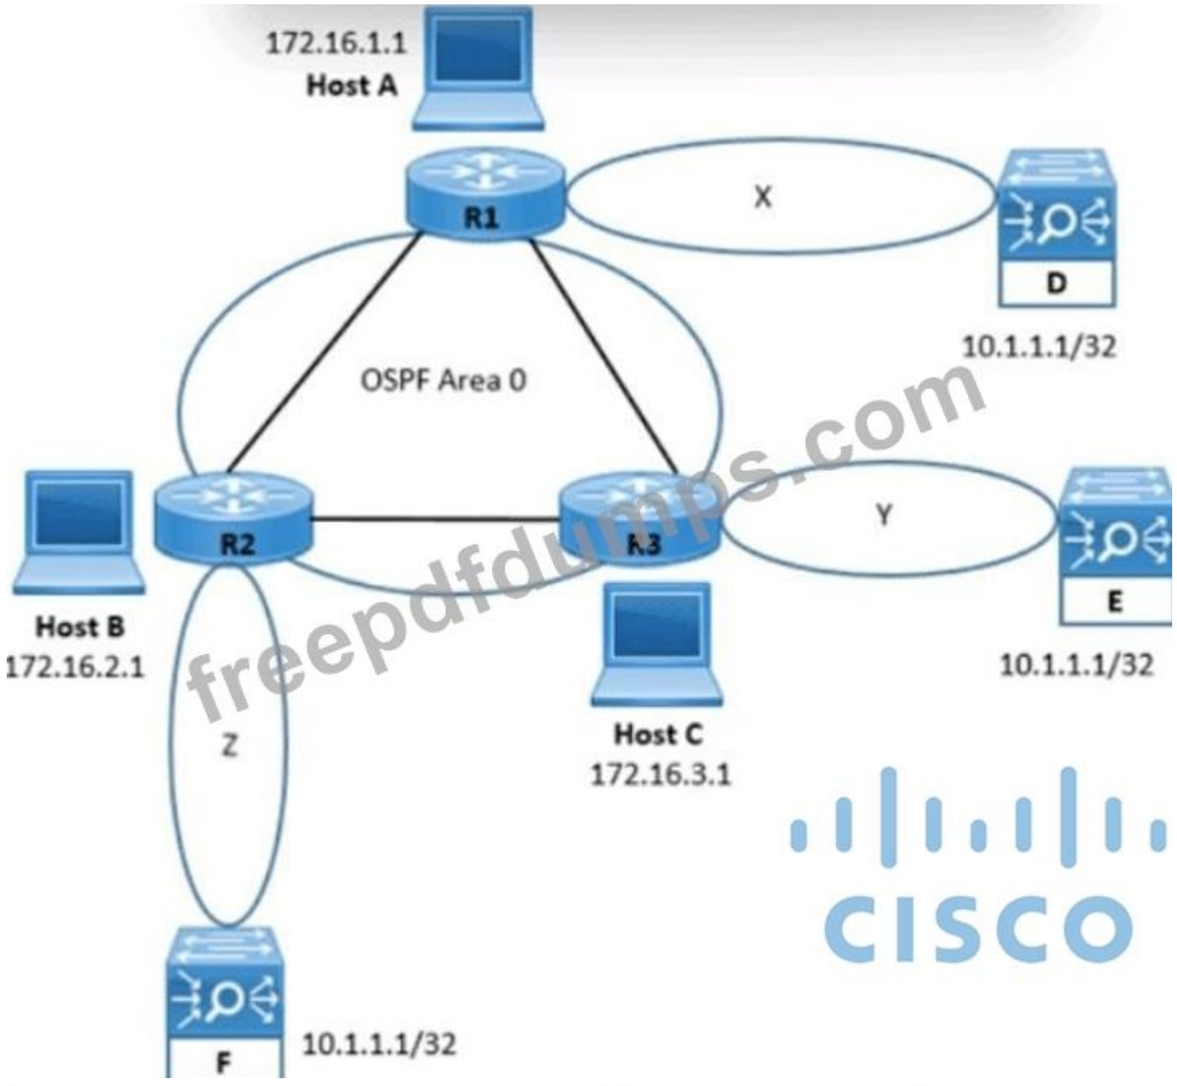
NEW QUESTION: 13
In the SD-WAN underlay network, which WAN Edge VPN ID is defined as the transport VPN and is used to carry control traffic?
A. VPN 0

B. VPN 512
C. VPN 128
D. VPN 256
Answer: A (LEAVE A REPLY)
VPN 0 is the transport VPN in Cisco SD-WAN. It carries control traffic and provides connectivity from WAN Edge routers toward the transport networks, such as MPLS, Internet, LTE, or other WAN services. Control connections to Cisco SD-WAN controllers are established through transport interfaces that belong to VPN 0.
Data VPNs carry service-side user traffic and are separate from the transport VPN. VPN 512 is commonly used for out-of-band management, not for carrying overlay control connections. VPN 128 and VPN 256 are not the reserved transport VPN identifiers. The design importance is that VPN 0 must have correct addressing, color assignment, routing reachability, NAT traversal considerations, and certificate-based controller connectivity. If VPN 0 is misdesigned, the WAN Edge cannot establish stable control connections, cannot receive OMP routes and policies, and cannot build the secure overlay properly. Therefore, VPN 0 is the correct transport VPN for underlay control traffic. Reference topics: Cisco SD-WAN VPN 0, transport VPN, control connections, WAN Edge onboarding, overlay formation.

NEW QUESTION: 14
An existing network solution is using BFD in echo mode. Several of the network devices are experiencing high CPU utilization which an engineer has determined is related to the BFD feature. Which solution should the engineer leverage to reduce the CPU load?
A. Implement slow timers between peers with low CPU resources.
B. Implement BED asynchronous mode between peers with low CPU resources.
C. Enable BFD multi-hop on the devices with low CPU resources.
D. Utilize carrier delay on all routers in the network.
Answer: A (LEAVE A REPLY)
When BFD echo mode contributes to high CPU utilization on lower-resource devices, the practical design response is to use slower BFD timers for those peers rather than continuing to run an overly aggressive failure- detection profile everywhere. Cisco BFD provides fast failure detection, but timer values must match platform capability, interface scale, and the operational value of subsecond detection. Echo mode can generate frequent packets, and in some environments the processing or punt behavior can create measurable CPU pressure.
Slowing the timers reduces packet rate and control-plane stress while keeping BFD available for faster detection than traditional routing protocol dead timers. Moving to asynchronous mode may be appropriate in some designs, but the listed option says "BED" and does not directly address the requirement as cleanly as reducing timer aggressiveness. BFD multihop is used for non-directly connected peers and is not a CPU-relief mechanism. Carrier delay changes interface state notification behavior, but it does not solve BFD packet processing load. Therefore, implement slower timers on peers with limited resources. Reference topics: BFD design, echo mode, failure-detection timers, CPU utilization.

NEW QUESTION: 15
In a multicast network, which condition must be met for an RPF check to be performed on the RP address ' ?
A. The PIM DM device receives a multicast packet and has no directly connected members
B. The PIM router or multilayer switch has a shared-tree state
C. The PIM router or multilayer switch has a source-tree state
D. The PIM DM device receives a multicast packet and has no directly connected PIM neighbor
Answer: (SHOW ANSWER)
The RPF check is performed against the RP address when the PIM router has shared-tree state. In PIM sparse mode, multicast traffic initially flows through the rendezvous point using the shared tree, represented as (*,G) state. For shared-tree traffic, the router verifies that packets arrive from the direction of the RP, because the RP is the root of the shared distribution tree. When the router has source-tree state, represented as (S,G), the RPF check is performed against the multicast source address instead. Dense mode behavior and directly connected membership are not the determining conditions in this question. The key distinction is whether forwarding state is shared-tree or source-tree. Cisco multicast design relies on this behavior to keep PIM forwarding loop free as traffic transitions from shared-tree delivery to shortest-path tree delivery where applicable. Selecting the shared-tree state option correctly identifies when the RP address is used for the RPF calculation. Reference topics: PIM sparse mode, shared tree, source tree, RPF toward RP, (*,G) and (S,G) state.

NEW QUESTION: 16



freepdfdumps.com

Refer to the exhibit. An engineer is designing an OSPF solution for a customer. The design must take into consideration:

* Application load balancers D. E. and F are in different geographical locations and are OSPF-enabled.
* Hosts A, B. and C connect to an application through the load balancers using IP address 10.1.1.1/32.
* In the event of a failure of one of the load balancers, hosts must still have access to the application.

Which solution must the engineer choose?

A. All load balancers to be co-located in area 0.
B. X, Y, and Z to be configured as different areas
C. At least one load balancer to be in area 0.
D. X, Y and Z to be configured as the same area

Answer: A (LEAVE A REPLY)

The load balancers should be co-located in area 0 because the design is using the same /32 application address from multiple OSPF-enabled devices. This is an anycast-style design: several load balancers advertise the same host route, and the routing protocol sends clients to the topologically best available instance. In OSPF, route preference is strongly affected by route type. Intra-area routes are

preferred over interarea routes, and interarea summarization can also hide useful path details. If identical anycast host routes are originated from different nonbackbone areas, hosts may not always select the closest or most resilient load balancer, and failure behavior can become inconsistent. Placing the load-balancer anycast advertisements in area 0 keeps the routing decision consistent and avoids unnecessary interarea preference complications. Configuring X, Y, and Z as different areas would make the problem worse, and placing only one load balancer in area 0 does not provide uniform resilience. Reference topics: OSPF area design, anycast host routes, intra-area preference, backbone area, application load-balancer resilience.

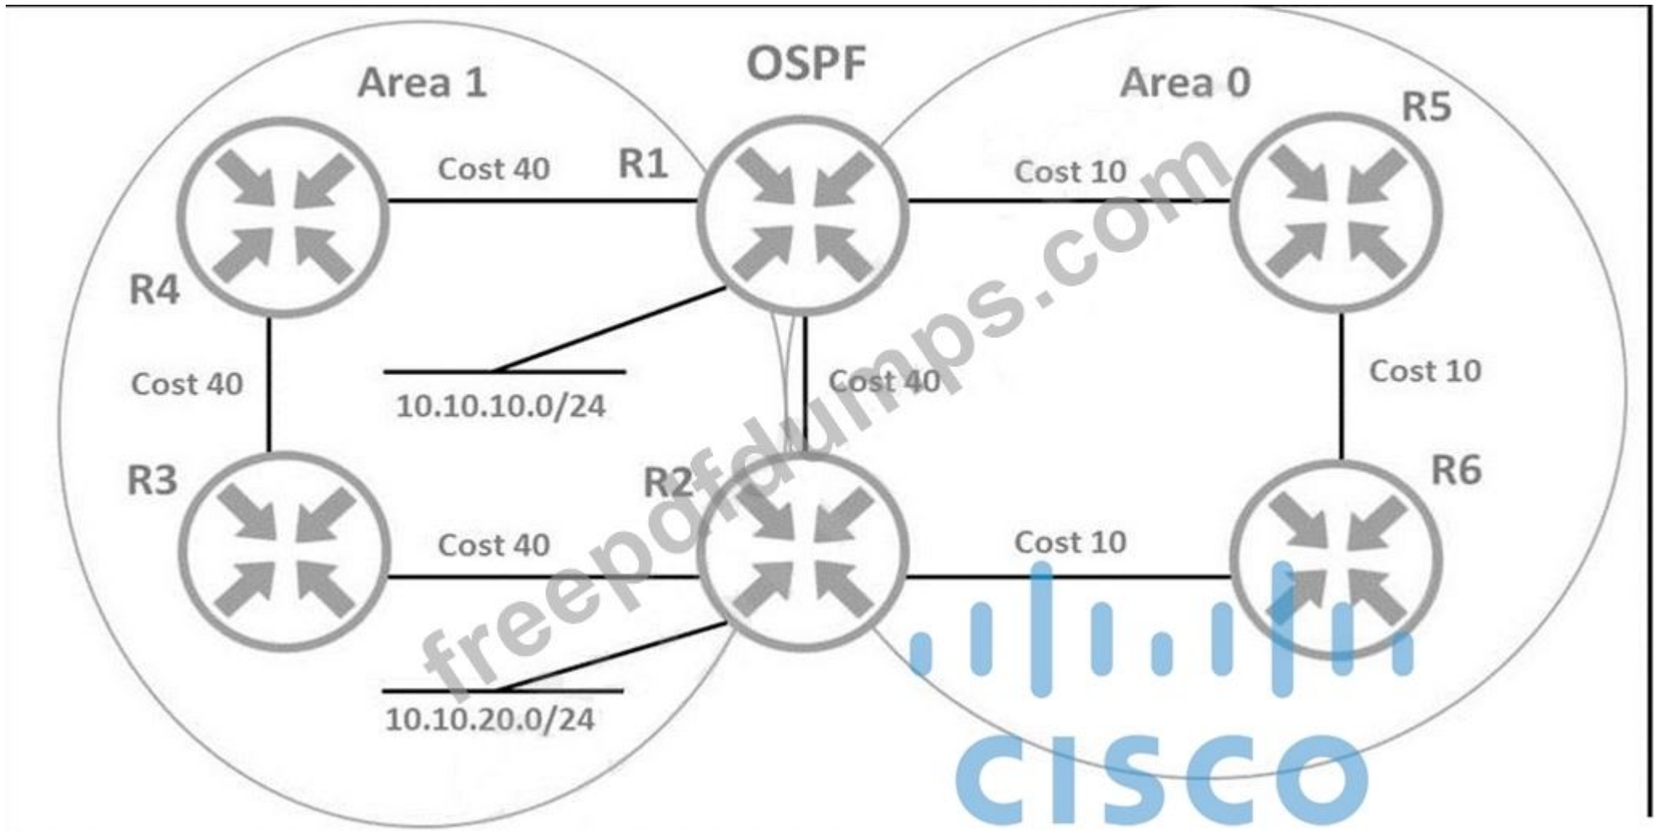
Valid 300-420 Dumps shared by Actual4test.com for Helping Passing 300-420 Exam! Actual4test.com now offer the **newest 300-420 exam dumps**, the Actual4test.com 300-420 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-420 dumps with Test Engine here: https://www.actual4test.com/300-420_examcollection.html (**457** Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 17
Which two best practices must be followed when designing an out-of-band management network? (Choose two.)
A. Enforce access control
B. Facilitate network integration
C. Back up data using the management network
D. Ensure that the management network is a backup to the data network
E. Ensure network isolation
Answer: (SHOW ANSWER)
An out-of-band management network should be isolated from the production data network and protected with strict access control. Cisco SAFE design principles treat the management plane as a sensitive function because it provides administrative access to infrastructure devices. Isolation ensures that production outages, routing problems, or compromised user segments do not automatically remove the operator's ability to manage routers, switches, firewalls, and controllers. Access control ensures that only authorized administrators and management systems can reach the management interfaces and protocols. The management network should not be treated as a backup data path, because that undermines separation and can expose management devices to production traffic risks. It should also not be used as a general-purpose data backup network. "Facilitate network integration" is too vague and can conflict with the isolation requirement. Therefore, the best practices are to enforce access control and ensure network isolation. In a mature design, this normally includes dedicated management interfaces or VRFs, jump hosts, AAA, logging, and tightly controlled management protocol access.

NEW QUESTION: 18
What is a benefit of using VRRPv3 as compared to VRRPv2?
A. VRRPv3 supports IPv4 and IPv6
B. VRRPv3 supports authentication
C. VRRPv3 supports preemption
D. VRRPv3 supports stateful switchover
Answer: A (LEAVE A REPLY)
VRRPv3 improves protocol applicability by supporting both IPv4 and IPv6 address families. VRRPv2 is associated with IPv4 virtual router redundancy, while VRRPv3 extends the protocol so the same first-hop redundancy concept can be used for IPv6 default gateways as well. In campus and branch designs, this matters because IPv6 deployments should not require a separate redundancy design from IPv4 when a standards-based FHRP is preferred. Preemption is not a unique VRRPv3 advantage because VRRP has long supported priority-based master election and preemption behavior. Authentication is also not the correct differentiator in this comparison; VRRPv3 removed the simple protocol authentication field from the base protocol and expects security to be provided by the underlying network or IPsec where required. Stateful switchover is a platform high-availability capability, not a VRRPv3 protocol feature. Therefore, the key benefit of VRRPv3 over VRRPv2 for enterprise design is IPv4 and IPv6 support under the same redundancy protocol model.

NEW QUESTION: 19

Refer to the exhibit.
C0FD9 F48C9ACDC725EA850EC2476EE1E



An architect must design a solution that uses the direct link between R1 and R2 for traffic from 10.10.10.0/24 toward network 10.10.20.0/24. Which solution should the architect include in the design?

A. Configure the OSPF cost of the link to a value lower than 30.

B. Lower the Administrative Distance for OSPF area 0.

C. Place the link into area 2 and install a new link between R1 and R2 in area 0.

D. Configure the link to provide multiarea adjacency.

Answer: A (LEAVE A REPLY)

Lowering the OSPF cost on the direct R1-to-R2 link is the cleanest way to force traffic for 10.10.20.0/24 to prefer that link. OSPF calculates the shortest path using cumulative interface cost, and Cisco designs normally tune OSPF path selection by changing interface cost rather than administrative distance. Administrative distance determines preference between different routing protocols, not between OSPF paths inside the same OSPF process. Moving the link into another area or creating a multiarea adjacency adds design complexity and does not directly solve the requirement unless the resulting cost still becomes lower. The exhibit states that the architect must use the direct link for traffic from 10.10.10.0/24 toward 10.10.20.0/24. Therefore, the direct link must have a lower total OSPF metric than the alternate path. Configuring the link cost to a value below 30 makes OSPF choose it as the preferred intra-area route. This is deterministic, easy to validate with the OSPF database and routing table, and consistent with Cisco hierarchical IGP design practice. Reference topics: OSPF cost, SPF path selection, metric tuning, intra-area routing.

NEW QUESTION: 20

An engineer uses Postman and YANG to configure a router with:

* OSPF process ID 200
* network 172.16.10.128/26 enabled for Area 0

Which get-config reply verifies that the model set was designed correctly?

A.

```
<rpc-reply message-id="urn:uuid:1b3d05cd-8118-3e6a-6c05-411157936aaf" xmlns="urn.ietf:params:
xml:ns:netconf:base:1.0" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
 <data>
  <native xmlns="http://cisco.com/ns/yang/ned/ios">
   <router>
    <ospf>
     <id>200</id>
     <network>
      <ip>172.16.10.128</ip>
      <mask>0.0.0.63</mask>
      <area>0</area>
     </network>
    </ospf>
   </router>
  </native>
 </data>
</rpc-reply>
```

B.

```
<rpc-reply message-id="urn.uuid:1b3d05cd-8118-3e6a-6c05-012435678aaf“
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:nc="urn.ietf.params:xml:ns:netconf:base:1.0">
 <data>
  <native xmlns="http://cisco.com/ns/yang/ned/ios">
   <router>
    <ospf>
     <id>200</id>
     <network>
      <ip>172.16.10.128</ip>
      <mask>255.255.255.192</mask>
      <area>0</area>
     </network>
    </ospf>
   </router>
  </native>
 </data>
</rpc-reply>
```

freepdfdumps.com

```
<rpc-reply message-id="urn:uuid:1b3d05cd-8118-3e6a-6c05-021345678aaf" xmlns="urn:ietf:params:
xml:ns:netconf:base:1.0" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data>
   <native xmlns="http://cisco.com/ns/yang/ned/ios">
    <router>
     <ospf>
      <id>200</id>
      <network>
       <ip>172.16.10.128</ip>
       <mask>0.0.0.192</mask>
       <area>0</area>
      </network>
     </ospf>
    </router>
   </native>
  </data>
</rpc-reply>
```


C.


```
<rpc-reply message-id="urn:uuid:1b3d05cd-8118-3e6a-6c05-0123546678aaf" xmlns="urn:ietf:params:
xml:ns:netconf:base:1.0" xmlns:nc="urn:ietf:params:json:ns:netconf:base:1.0">
  <data>
   <native json="http://cisco.com/ns/yang/ned/ios">
    <router>
     <ospf>
      <id>200</id>
      <network>
       <ip>172.16.10.128</ip>
       <mask>0.0.0.63</mask>
       <area>0</area>
      </network>
     </ospf>
    </router>
   </native>
  </data>
</rpc-reply>
```


D.

Answer: (SHOW ANSWER)

The correct get-config reply is the one that shows OSPF process ID 200 and the network 172.16.10.128/26 assigned to Area 0 in the same YANG model hierarchy that was used for configuration. In model-driven configuration, verification is not based on visual similarity to CLI output; it is based on whether the operational or running datastore contains the intended structured objects at the correct schema

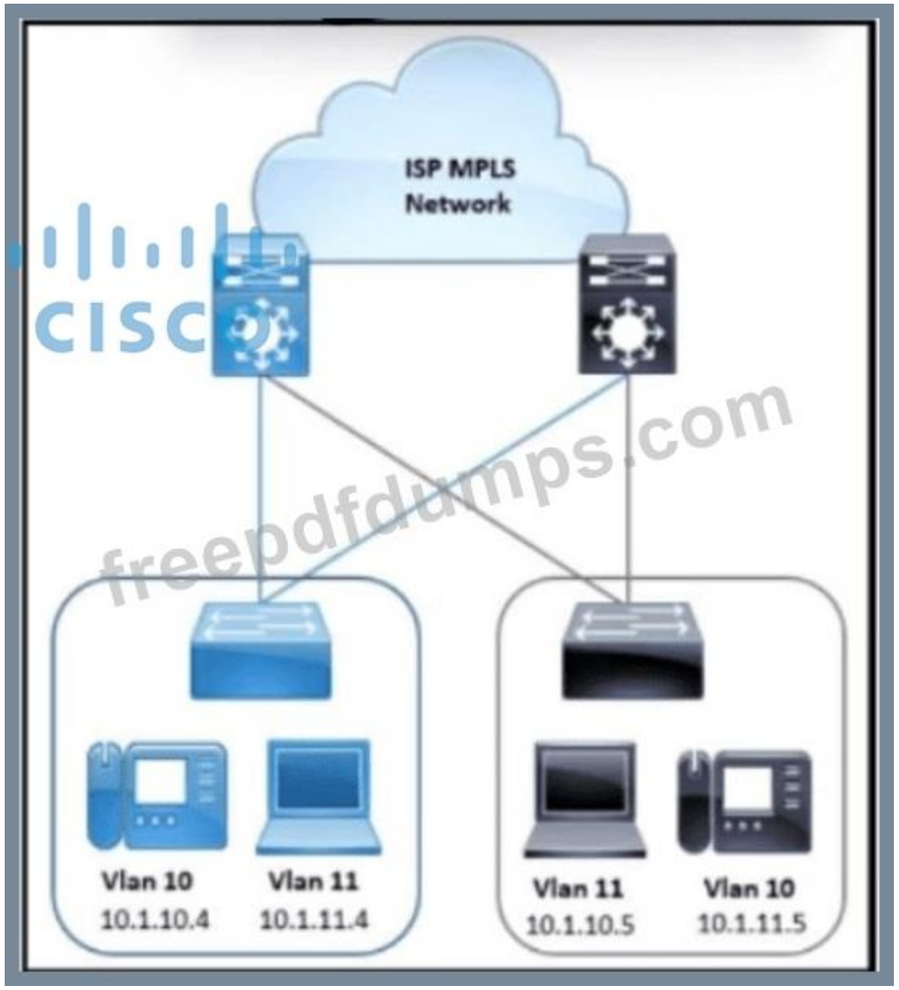
freepdfdumps.com

paths. For IOS XE native OSPF configuration, the process identifier, network statement, wildcard or prefix representation, and area association must be represented exactly as the model requires. A /26 network beginning at
172.16.10.128 covers addresses 172.16.10.128 through 172.16.10.191, so replies showing a different base network, wrong process ID, wrong area, or invalid mask cannot verify the design. Option A is the matching reply. The professional validation workflow is to send the configuration, read back the intended datastore using get-config, and confirm the model fields rather than relying only on CLI screen output. Reference topics: NETCONF get-config, IOS XE native OSPF YANG, datastore verification, OSPF area configuration.

NEW QUESTION: 21
In a Cisco SD-Access fabric, which node facilities connectivity between the fabric and networks external to the fabric?
A. intermediate
B. edge
C. control plane
D. border
Answer: D (LEAVE A REPLY)
The border node provides connectivity between the SD-Access fabric and networks external to the fabric. In Cisco SD-Access, fabric edge nodes connect endpoints, intermediate nodes provide underlay transport, and control-plane nodes maintain endpoint-to-location mappings. Border nodes sit at the boundary of the fabric and connect the overlay to external Layer 3 networks such as WAN, data center, Internet edge, shared services, or other fabrics. They de-encapsulate VXLAN traffic leaving the fabric and advertise or import reachability using VRFs and routing protocols. A border can be internal, external, or anywhere border depending on whether it connects to shared services, outside networks, or both. Intermediate nodes do not provide external fabric exit; they only transport traffic within the underlay. Edge nodes connect endpoints but are not normally the external gateway for the whole fabric. Control-plane nodes provide mapping services, not external connectivity. Reference topics: SD-Access border node, fabric exit, VRF mapping, external connectivity, overlay-to-underlay boundary.

NEW QUESTION: 22
A company wants to enable several third-party video conferencing networks based on multicast services. The video conferencing platform can accommodate numerous and dispersed senders and receivers in third-party networks The interior routing protocol is OSPF and the exterior routing protocol is BGP Which multicast mode must be selected to achieve this goal?
A. Enable MP-BGP
B. Enable BIDIR-PIM
C. Configure MSDP
D. Set IGMPv2.
Answer: B (LEAVE A REPLY)
BIDIR-PIM is the correct multicast mode for a video conferencing environment with numerous dispersed senders and receivers. Bidirectional PIM is designed for many-to-many multicast applications where receivers may also transmit and where maintaining per-source trees would create unnecessary state and control-plane overhead. Instead of building a shortest-path tree for every source, BIDIR-PIM uses a shared tree rooted at the rendezvous point and forwards traffic bidirectionally. This makes it attractive for collaborative applications and conferencing scenarios with many active participants. MP-BGP can carry multicast VPN or multicast routing information in service-provider designs, but it is not itself the multicast forwarding mode for this requirement. MSDP is used to exchange source information between PIM-SM domains, not to optimize many- to-many forwarding within the multicast design. IGMPv2 handles host membership signaling on local segments and does not provide the routed multicast mode. Reference topics: BIDIR-PIM, many-to-many multicast, shared tree, video conferencing multicast, PIM design.

NEW QUESTION: 23



freepdfdumps.com

Refer to the exhibit. An architect must design a resilient gateway solution based on these requirements:

* VLAN 10 and VLAN 11 support voice and video applications.
* Link and node failures must have minimal impact on traffic.
* Provide protection against false hello packets.
* Support IPv6.

Which solution must the architect choose?

A. GLBP with IP SLA tracking

B. VRRP version 2 with authentication

C. HSRP version 2 with MD5 authentication

D. VRRP version 2 with object tracking

Answer: (SHOW ANSWER)

HSRP version 2 with MD5 authentication is the correct resilient gateway choice because the design requires IPv6 support and protection against false hello packets. HSRPv2 supports both IPv4 and IPv6 gateway redundancy and improves group scaling compared with the original version. MD5 authentication protects the HSRP control exchange by ensuring that only authorized routers participate in the redundancy group, reducing the risk of a rogue or misconfigured device sending false hello messages and influencing active gateway selection. GLBP provides load balancing but does not satisfy the same

IPv6 and authenticated hello requirement in this option set. VRRP version 2 is IPv4-focused and therefore does not meet the IPv6 requirement; VRRPv3 would be the relevant standards-based IPv6-capable version, but it is not offered.
Object tracking helps adjust gateway priority after link failure, but it does not protect against false hello packets by itself. Reference topics: HSRPv2, first-hop redundancy, IPv6 gateway redundancy, MD5 authentication, campus high availability.

NEW QUESTION: 24
An engineer is tasked with designing a dual BGP peering solution with a service provider. The design must meet these conditions:
* The routers will not learn any prefix with a subnet mask greater than /24.
* The routers will determine the routes to include in the routing table based on the length of the mask alone.
* The routers will make this selection regardless of the service provider configuration.
Which solution should the engineer include in the design?
A. Use a route map and access list to block the desired networks, and apply the route map to BGP neighbors inbound.
B. Use a route map and prefix list to block the desired networks, and apply the route map to BGP neighbors outbound.
C. Use an IP prefix list to block the desired networks and apply the IP prefix list to BGP neighbors outbound.
D. Use an IP prefix list to block the desired networks and apply the IP prefix list to BGP neighbors inbound.
Answer: D (LEAVE A REPLY)
An inbound IP prefix list is the correct tool because the routers must decide which prefixes to accept from the service provider based only on prefix length. Prefix lists are designed to match network prefixes and mask- length ranges using operators such as le and ge. Applying the prefix list inbound to the BGP neighbors prevents routes with masks longer than /24 from entering the local BGP table, regardless of what the service provider advertises. This conserves memory and ensures the customer routing policy is enforced locally. An access list is a poor fit because it does not natively express prefix-length logic as cleanly as a prefix list.
Applying the policy outbound would filter routes the customer sends to the provider, which is the wrong direction; the requirement is about routes the customer learns. A route map could reference a prefix list, but the simplest and most exact answer is to use an IP prefix list inbound. Therefore, the design should apply an inbound prefix list to both BGP peers to block prefixes greater than /24. Reference topics: BGP route filtering, prefix lists, inbound policy, maximum accepted prefix length.

NEW QUESTION: 25
How is sub-second failure of a transport link detected in a Cisco SD-WAN network?
A. Hellos are sent between the WAN Edge routers and the vSmart controller.
B. BFD runs on the IPsec tunnels between WAN Edge routers.
C. BGP is used between WAN Edge routers and the vSmart controller.
D. Link state change messages are sent between vSmart controllers.
Answer: (SHOW ANSWER)
Cisco SD-WAN detects subsecond transport failure by running BFD across the IPsec tunnels between WAN Edge routers. The WAN Edge devices form encrypted data-plane tunnels over each transport, and BFD probes those tunnels to measure liveliness, loss, latency, jitter, and path availability. When BFD detects a failed or degraded tunnel, the SD-WAN fabric can withdraw or avoid that path quickly, enabling application- aware routing and fast failover. Hellos between WAN Edge routers and vSmart controllers are control-plane mechanisms and do not directly detect data-plane transport failure between edge routers. BGP is not the SD- WAN overlay control protocol for this function; OMP is used with vSmart for overlay route and policy exchange. Link-state messages between vSmart controllers do not measure transport-path liveliness.
Therefore, BFD on IPsec tunnels is the correct failure-detection mechanism. The design should ensure that BFD timers, tunnel scale, controller policy, and SLA classes are chosen to balance fast failover with platform resources and transport stability. Reference topics: Cisco SD-WAN BFD, IPsec tunnels, transport failure detection, application-aware routing, data-plane liveliness.

NEW QUESTION: 26

A network engineer must optimize a campus OSPF deployment Currently each time a type 1 or type 2 LSA is generated within an area, the OSPF process must recompute the entire SPT Which solution improves the recomputation process?

A. iSPF

B. BFD

C. SPF

D. PRC

Answer: A (LEAVE A REPLY)

Incremental SPF is the correct solution because the problem specifically states that Type 1 or Type 2 LSAs force the OSPF process to recompute the entire shortest path tree. Cisco OSPF documentation states that when changes to a Type 1 or Type 2 LSA occur, the entire SPT is normally recomputed, and that incremental SPF allows the system to recompute only the affected part of the tree. This improves convergence and reduces CPU use because the router does not have to run a full Dijkstra calculation for every topology change. BFD detects failures quickly, but it does not change how OSPF recalculates the SPT after the LSA arrives.

Standard SPF is the default full calculation behavior and therefore does not solve the optimization requirement. PRC is associated with partial route calculation behavior in some link-state contexts, but for this OSPF question the Cisco feature name is iSPF. Reference topics: OSPF incremental SPF, Type 1 LSA, Type

2 LSA, SPF optimization, campus OSPF scaling.

NEW QUESTION: 27

Company A recently acquired another company. Users of the newly acquired company must be able to access a server that exists on Company A's network, both companies use overlapping IP address ranges. Which action conserves IP address space and provides access to the server?

A. Use a single IP address to create overload NAT

B. Use a single IP address to create a static NAT entry

C. Build one-to-one NAT translation for every user that needs access

D. Re-IP overlapping address space in the acquired company

Answer: B (LEAVE A REPLY)

The best answer is to use a single IP address for a static NAT entry that represents the server with a nonoverlapping address. When two companies use overlapping private address space, routing directly between the networks can fail because each side may already have a local route for the same prefix. The requirement is narrow: users in the acquired company need access to one server in Company A, and the design must conserve address space. A static NAT mapping for the server solves the destination overlap problem with one translated address, giving remote users a unique reachable address for that server without renumbering either company. Overload NAT is normally used to translate many inside clients behind one address for outbound access; it is not the most precise solution for publishing one overlapping server. One-to- one NAT for every user wastes addresses and adds operational overhead. Readdressing the acquired company would solve the overlap permanently, but it does not conserve effort or address space. Therefore, a single static NAT entry for the server is the most efficient design.

NEW QUESTION: 28

A company is working with a service provider to design a BGP policy. The company is dual-homed with the provider and wants to control which link inbound traffic transits. Also, the company will advertise several networks to the provider and needs propagation to go no further. Which BGP attribute meet these requirements?

A. AS-path

B. MED

C. community

D. local preference

Answer: C (LEAVE A REPLY)

BGP community is the correct attribute because it can support both policy requirements through coordination with the service provider. Communities are route tags that an upstream provider can match to apply routing policy, such as adjusting local preference to influence which dual-homed provider link is preferred for inbound traffic. Communities also include well-known values used to constrain route propagation. Cisco documents no-export as a community that prevents a route from being advertised outside the receiving AS, and no-advertise as a community that prevents advertisement to any peer.

This directly matches the requirement that several customer-advertised networks must propagate no further. MED can influence inbound path selection in some dual-link designs, but it does not provide the propagation-control function. AS- path prepending can influence inbound traffic but is less precise and again does not solve the no-further- advertisement requirement. Local preference is normally applied inside an AS, not directly by the customer to external peers. Reference topics: BGP communities, no-export, no-advertise, provider policy, inbound traffic engineering.

NEW QUESTION: 29
Which design consideration must be made when dual WAN Edge routers are deployed at a branch site?
A. Use BGP AS-path prepending to influence egress traffic and use MED to influence ingress traffic from the branch.
B. HSRP priorities must match the OMP routing policy to prefer one WAN Edge over the other.
C. Traffic must be symmetrical as it egresses the WAN Edges and returns from remote sites for DPI to function properly.
D. Configure BFD between WAN Edge routers to detect sub-second link failures.
Answer: (SHOW ANSWER)
When dual WAN Edge routers are deployed at a branch, first-hop redundancy and SD-WAN overlay path preference must be aligned. The correct design consideration is that HSRP priorities on the service side should match the OMP routing policy that determines which WAN Edge is preferred. If the LAN gateway points hosts to one WAN Edge while OMP policy prefers the other edge for return or remote-site traffic, traffic can become asymmetric or hairpinned through the wrong device. Proper alignment gives predictable active/standby or active/active behavior and avoids blackholing during failover. Option A describes traditional BGP path manipulation with AS-path prepending and MED; those are not the primary SD-WAN branch dual-edge design controls. Option C overstates symmetry as a universal requirement for DPI; Cisco SD-WAN can handle many traffic patterns, but gateway and OMP alignment remain the core branch design issue. Option D is wrong because SD-WAN already uses BFD across data-plane tunnels, not as a special direct adjacency between the two local WAN Edges for this purpose. Reference topics: Cisco SD-WAN branch redundancy, HSRP, OMP policy, WAN Edge preference, service-side high availability.

NEW QUESTION: 30
An engineer is designing an IPv4 addressing plan for an enterprise with 1000 branches. Each branch requires a prefix for data and a prefix for voice. Each prefix must accommodate up to 128 hosts, and prefixes must facilitate summarization at aggregation points in the network. The security team requires a simple method for identifying voce prefixes. Which allocation does the engineer recommend from the RFC1918 address space?
A. /24 prefixes for data from 10.0.0.0/15 and /24 prefixes for voice from 172.16.0.0/15
B. /24 prefixes for data from 10.0.0.0/8 and /24 prefixes for voice from the next contiguous /24 prefix per site
C. /25 prefixes for data from 10.0.0.0/8 end /25 prefixes for voice from the next contiguous /25 prefix per branch
D. /24 prefixes for data from 10.0.0.0/8 and /24 prefixes for voice from 172.16.0.0/12
Answer: B (LEAVE A REPLY)
The best allocation is to give each branch adjacent /24 prefixes from the 10.0.0.0/8 private range, using one
/24 for data and the next contiguous /24 for voice. Each /24 provides 254 usable host addresses, which satisfies the requirement for up to 128 hosts. More importantly, pairing two contiguous /24s creates a clean
/23 summary per branch. For example, 10.0.0.0/24 for data and 10.0.1.0/24 for voice summarize as 10.0.0.0
/23; the next branch can use 10.0.2.0/24 and 10.0.3.0/24 summarized as 10.0.2.0/23. This pattern supports aggregation and keeps the security team's identification simple because voice prefixes can consistently use the odd-numbered third octet. Option C uses /25 prefixes, but /25 provides only 126 usable addresses, which fails a strict requirement for up to 128 hosts if the subnet and broadcast addresses are excluded. Separate
172.16 ranges reduce per-site summarization clarity. Reference topics: RFC 1918 addressing, VLSM, route summarization, branch addressing design, data and voice prefix allocation.

NEW QUESTION: 31
When expanding an existing Cisco SD-Access network, in addition to the control plane, which two device roles are needed to create an additional fabric site? (Choose two.)
A. leaf
B. cEdge

C. WLC
D. edge
E. border
Answer: D,E (LEAVE A REPLY)
An additional Cisco SD-Access fabric site requires edge and border roles in addition to the control-plane function. Edge nodes connect wired endpoints, access-layer devices, and sometimes access points into the fabric. They perform endpoint detection, anycast gateway functions, and fabric encapsulation for user traffic.
Border nodes provide connectivity between the fabric site and external networks such as the WAN, data center, internet edge, or shared services. Without edge nodes, the site has no practical access-layer fabric attachment for endpoints. Without border nodes, the site cannot properly exchange traffic with resources outside the local fabric. A wireless LAN controller may be needed when fabric wireless is part of the design, but it is not one of the mandatory fabric roles for every additional site. Leaf and cEdge are not Cisco SD- Access fabric roles in this context. Reference topics: Cisco SD-Access fabric roles, edge node, border node, control-plane node, fabric site expansion.

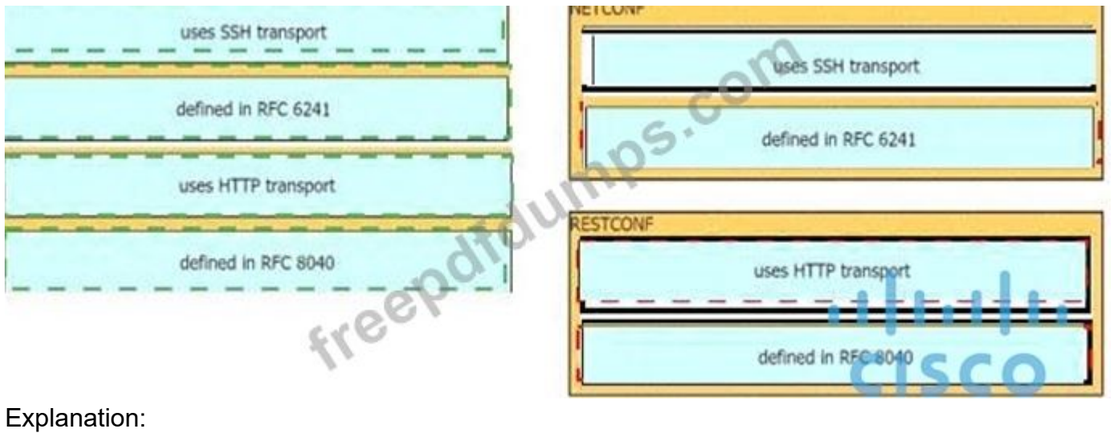
Valid 300-420 Dumps shared by Actual4test.com for Helping Passing 300-420 Exam! Actual4test.com now offer the **newest 300-420 exam dumps**, the Actual4test.com 300-420 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-420 dumps with Test Engine here: https://www.actual4test.com/300-420_examcollection.html (**457** Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 32
When vEdge router redundancy is designed, which FHRP is supported?
A. HSRP
B. OMP
C. GLBP
D. VRRP
Answer: D (LEAVE A REPLY)
When Cisco SD-WAN edge router redundancy is designed, VRRP is the supported first-hop redundancy protocol in the SD-WAN edge context. VRRP provides a virtual default gateway for LAN-side devices, allowing one WAN Edge router to act as the active gateway while another can take over if the active router or tracked condition fails. Cisco SD-WAN designs use VRRP with tracking and policy alignment so branch LAN traffic can fail over to the appropriate edge router while the SD-WAN overlay handles transport and tunnel selection. HSRP and GLBP are traditional campus first-hop redundancy protocols, but they are not the supported FHRP answer for vEdge router redundancy in this design context. OMP is the SD-WAN control- plane routing protocol used between WAN Edge routers and vSmart controllers; it is not a first-hop redundancy protocol for LAN hosts. Therefore, the correct FHRP for vEdge router redundancy is VRRP. The design should also ensure that VRRP priorities, tracking, and SD-WAN routing preferences align so traffic exits the intended edge during normal and failure conditions.

NEW QUESTION: 33
Drag and drop the characteristics from the left onto the configuration protocols they describe on the right.



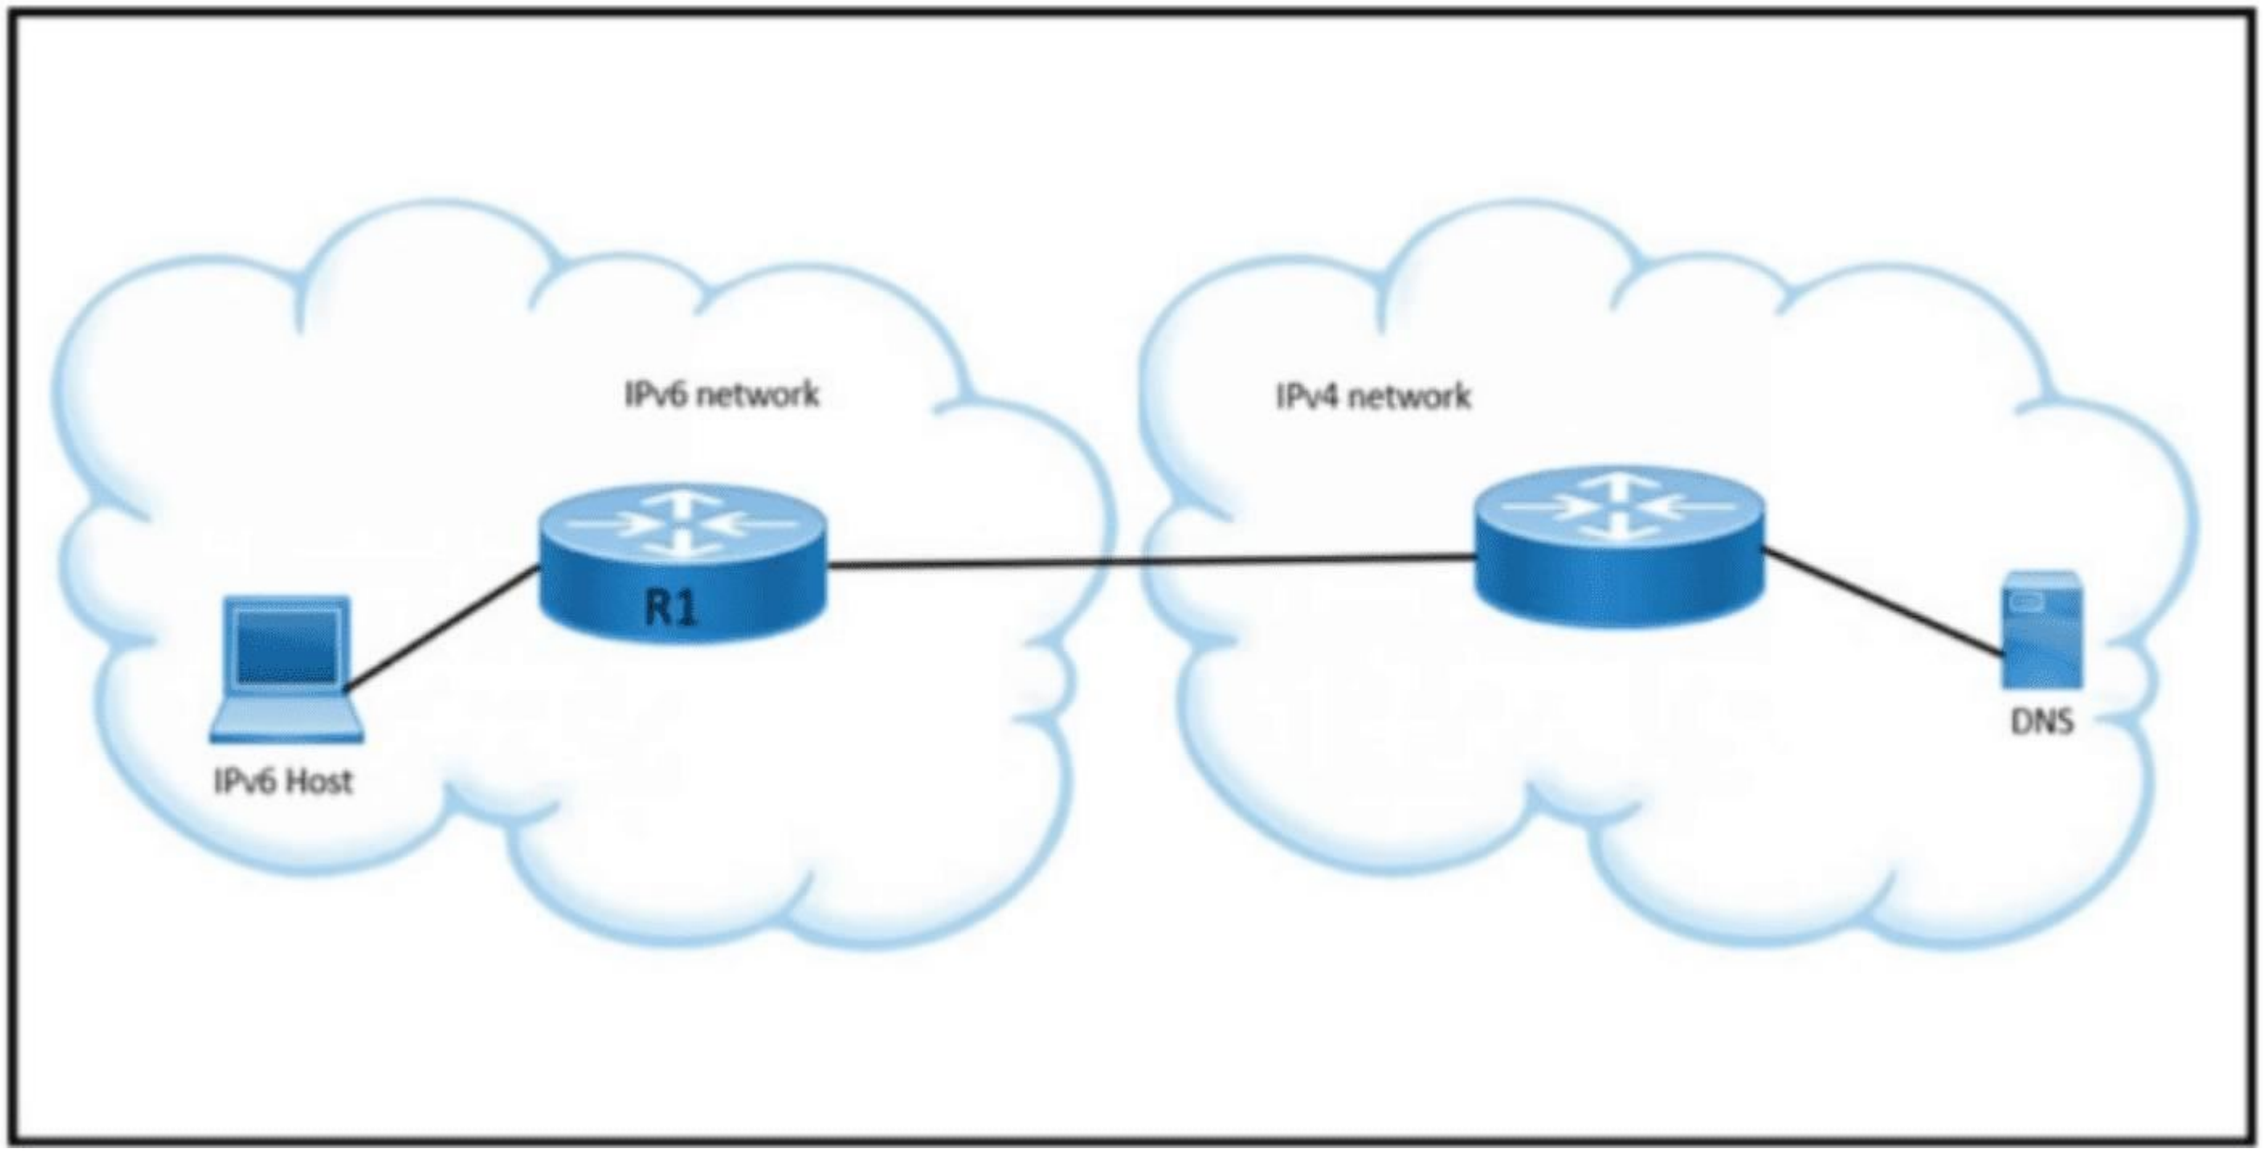
Answer:



Explanation:

The configuration-protocol mapping distinguishes NETCONF and RESTCONF behavior in model-driven programmability. NETCONF is an XML-based network configuration protocol that runs over a secure transport such as SSH and supports structured operations such as get-config, edit-config, lock, unlock, commit, and datastore manipulation. It is well suited to transactional configuration workflows where candidate and running datastores are supported. RESTCONF exposes YANG-modeled data through a REST- style interface over HTTP or HTTPS, using familiar methods such as GET, POST, PUT, PATCH, and DELETE. It can encode data in XML or JSON depending on headers and implementation. Both protocols use YANG models to define the structure and semantics of configuration and operational data, but they differ in transport style, operation model, and tooling. The selected mapping therefore places NETCONF characteristics with the protocol that provides RPC-style operations and datastore handling, while RESTCONF receives HTTP method and URI-based characteristics. In design, the choice depends on controller support, device operating system, required transaction semantics, security policy, and developer tooling. Reference topics: NETCONF, RESTCONF, YANG, datastores, XML and JSON encoding.

NEW QUESTION: 34



Refer to the exhibit. An engineer must connect the IPv6 island to the IPv4-only network to provide IPv6 hosts access to file servers and DNS services in the IPv4 network. Which NAT should the engineer choose?

A. stateless NAT66

B. stateful NAT66

C. static NAT-PT

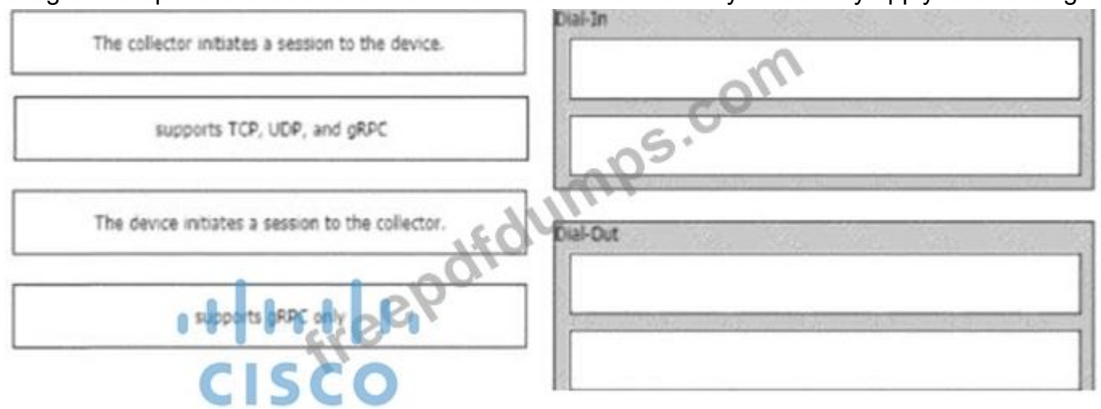
D. dynamic NAT-PT

Answer: D (LEAVE A REPLY)

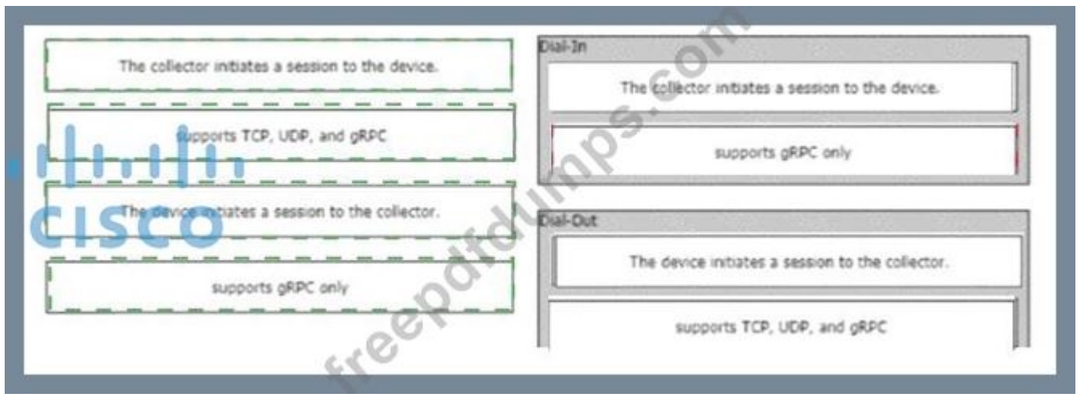
Dynamic NAT-PT is the correct selection when IPv6 hosts in an IPv6 island must access IPv4-only servers and services. NAT-PT translates between IPv6 and IPv4 packet headers and allows IPv6-only nodes to communicate with IPv4-only destinations by dynamically creating translation state. In this scenario, IPv6 hosts need access to file servers and DNS services in an IPv4 network, so a dynamic translation mechanism is required. Stateless NAT66 and stateful NAT66 are IPv6-to-IPv6 translation mechanisms and do not connect IPv6 hosts to IPv4-only resources. Static NAT-PT would be appropriate

for fixed one-to-one mappings where specific hosts are permanently mapped, but the question describes general access from an IPv6 island to IPv4 services, making dynamic NAT-PT the better match. In modern designs, NAT64 and DNS64 are commonly preferred successors, but within the answer choices NAT-PT is the IPv6-to-IPv4 transition mechanism.
Reference topics: IPv6 transition, NAT-PT, IPv6-to-IPv4 translation, DNS access, dual-protocol migration.

NEW QUESTION: 35
Drag and drop the characteristics from the left onto the telemetry mode they apply to on the right.



Answer:

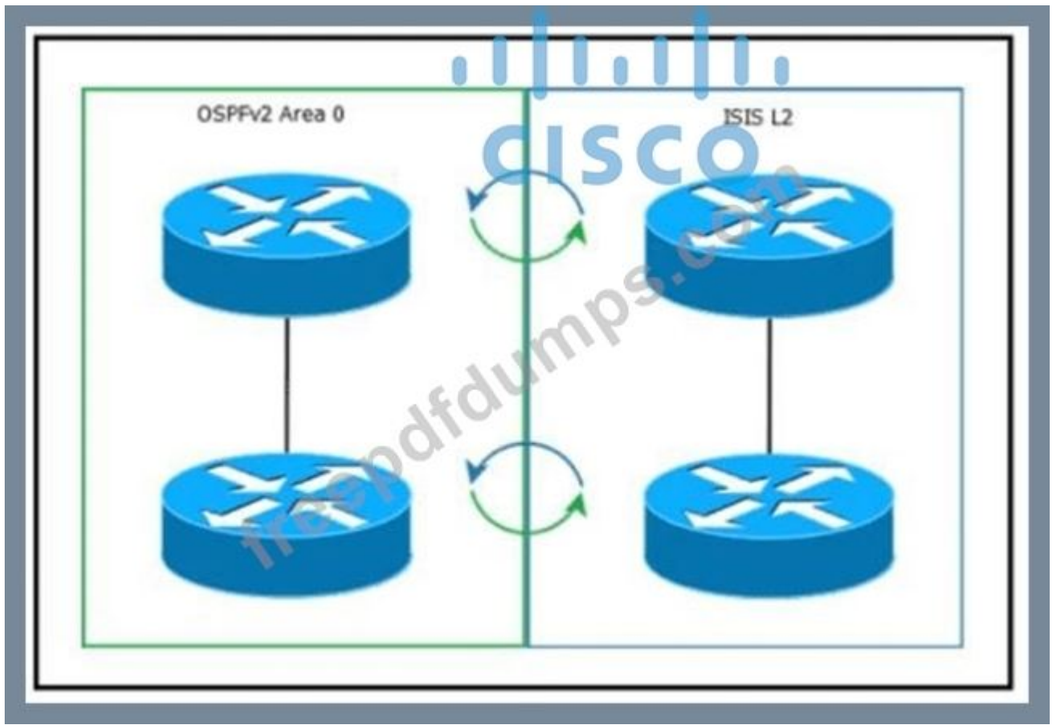


Explanation:
Dial-In: collector initiates a session to the device; supports gRPC only. Dial-Out: device initiates a session to the collector; supports TCP, UDP, and gRPC.
Model-driven telemetry can be deployed in dial-in or dial-out mode, and the difference is which side initiates the session and where the subscription is controlled. In dial-in mode, the collector initiates the session to the network device and requests the telemetry subscription. This makes the collector responsible for establishing the session and asking for the data stream. In the context of the provided options, dial-in is paired with support for gRPC only. In dial-out mode, the network device initiates the session toward the collector according to a configured subscription. This is useful when devices sit behind NAT or when the operational model requires devices to push telemetry to a known collector. In the provided choices, dial-out is paired with support for TCP, UDP, and gRPC. The mapping is not arbitrary; it follows Cisco telemetry terminology, where dial-in is collector-initiated and dial-out is device-initiated. This distinction is important in network automation designs because firewall policy, collector scaling, and subscription lifecycle management differ between the two modes.

NEW QUESTION: 36
An engineer is upgrading a company's main site to include a connection to a second ISP. The company will receive full Internet routing tables from both ISPs via BGP. The engineer must ensure that the company does not become a transit autonomous system. Which solution should be included in this design?
A. Tag incoming routes from both ISPs with BGP community no-export.

B. Lower the MED for updates sent to the secondary ISP.
C. Use a route-map to prevent all prefixes from being advertised to either ISP.
D. Modify the local-preference for routes incoming from the primary ISP.
Answer: C (LEAVE A REPLY)
The company must prevent itself from advertising routes learned from one ISP to the other ISP. If it receives full Internet tables from both providers and then re-advertises those learned prefixes, it can unintentionally become a transit autonomous system. The proper BGP design is to apply outbound route policy so only the company's own originated prefixes are advertised to the ISPs, while routes learned from ISP1 are never advertised to ISP2 and routes learned from ISP2 are never advertised to ISP1. A route map or equivalent prefix filtering policy applied outbound to the ISP neighbors achieves that goal. Local preference influences outbound path selection inside the company AS; it does not prevent transit advertisements. MED influences how a neighboring AS may enter the company network, but it does not stop the company from advertising third-party routes. Tagging incoming routes with no-export is not as deterministic as simply filtering what is advertised to each provider and may still allow undesirable routing behavior within the directly connected provider. Therefore, the design should include outbound route filtering with a route map.

NEW QUESTION: 37

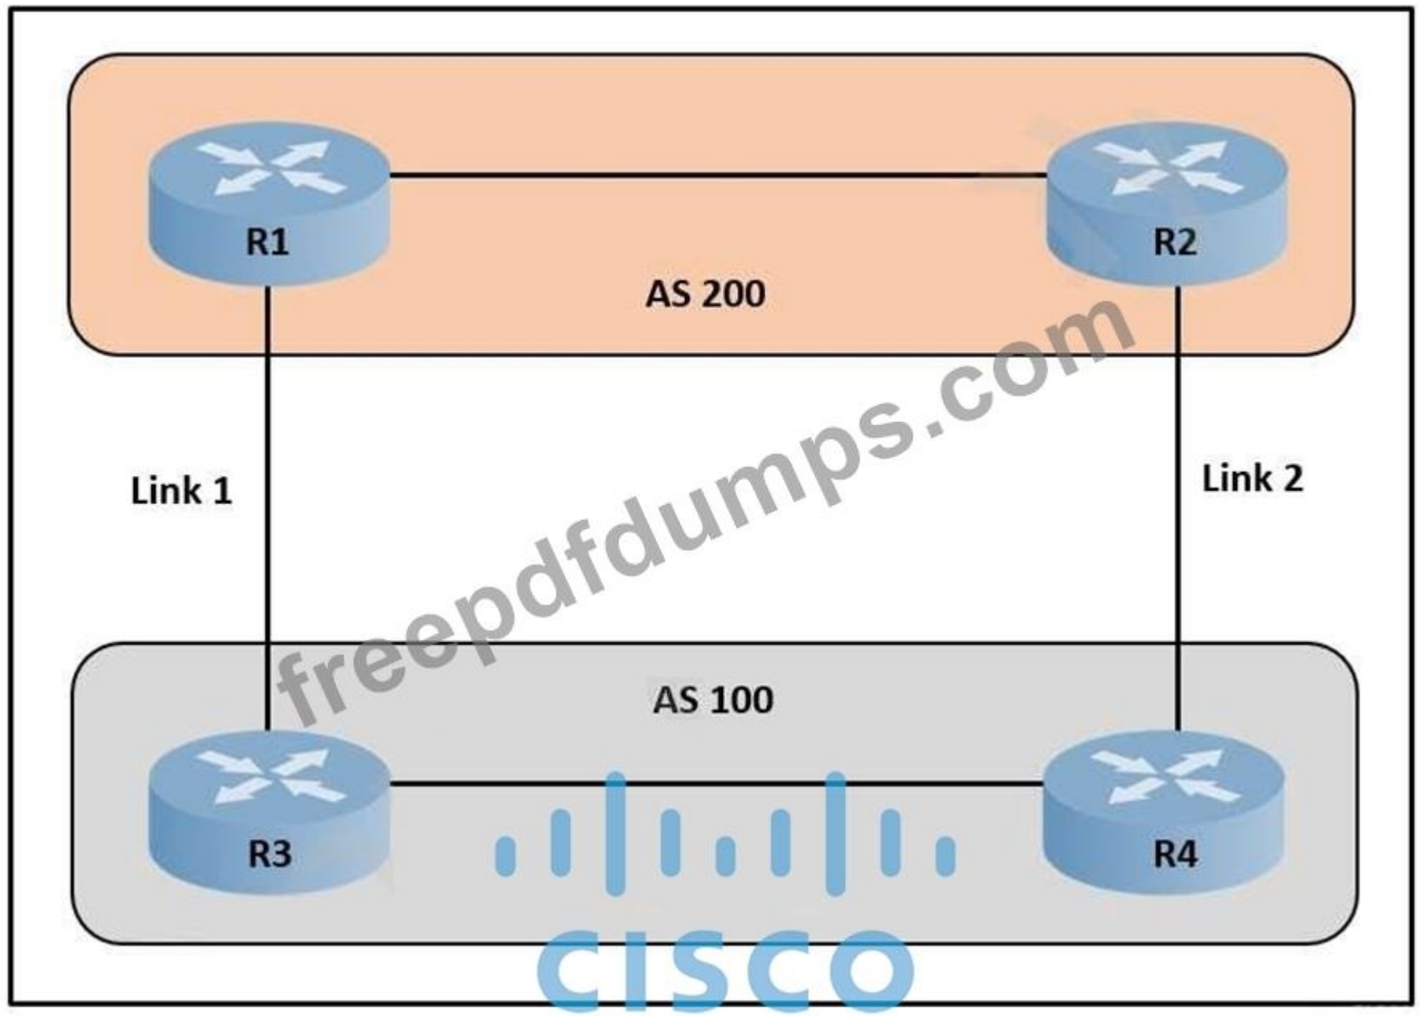


Refer to the exhibit. An architect is designing a network that requires route redistribution. The design must prevent route feedback and the creation of routing loops. The OSPF domain is using default metrics, and the IS-IS domain is using narrow metrics. Which solution must the architect select?
A. Change the IS-IS administrative distance to 105.
B. Change the OSPF area to a nonbackbone stub area
C. Use route filtering with an ACL or prefix list.
D. Use route tagging with a route map.
Answer: D (LEAVE A REPLY)
Route tagging with a route map is the correct redistribution loop-prevention method. In a design with mutual redistribution between OSPF and IS-IS, a route learned from one domain can be redistributed into the other domain and then fed back into the original domain at another redistribution point. This creates route feedback, suboptimal routing, or loops. Cisco routing design commonly prevents this by tagging routes at the point of redistribution and then filtering or denying routes with that tag when they are seen again at another redistribution boundary. A route map can set the tag during redistribution and

match the tag on the reverse redistribution direction. Changing administrative distance may influence local route selection, but it does not reliably prevent feedback throughout the network. Stub areas reduce external LSA scope in OSPF, but they do not solve a multipoint redistribution loop between OSPF and IS-IS. Prefix filtering can work, but tagging scales better and preserves route identity across redistribution points. Reference topics: route redistribution, route tagging, route maps, OSPF, IS-IS, loop prevention.

NEW QUESTION: 38
How is end-to-end microsegmentation enforced in a Cisco SD-Access architecture?
A. VLANs are used to segment traffic at Layer 2.
B. 5-tuples and ACLs are used to permit or deny traffic.
C. SGTs and SGTACLs are used to control access to various resources.
D. VRFs are used to segment traffic at Layer 3.
Answer: C (LEAVE A REPLY)
End-to-end microsegmentation in Cisco SD-Access is enforced with Security Group Tags and Security Group ACLs. Virtual networks provide macrosegmentation by separating routing tables and address spaces, but SGTs provide more granular group-based policy inside or across those virtual networks. Cisco SD-Access segmentation guidance explains that users or endpoints are assigned to security groups through SGTs and that SGACLs control communication between those groups. This avoids building large numbers of VLANs or IP ACLs tied to topology. VLANs alone provide Layer 2 separation but do not deliver identity-based end-to-end microsegmentation. Five-tuple ACLs can filter traffic, but they are more operationally complex and less policy-oriented than group-based access control. VRFs are useful for macrosegmentation, not microsegmentation within a virtual network. Therefore, SGTs and SGACLs are the correct enforcement method. In a professional design, the architect should define scalable groups according to business roles and application access needs, then use Cisco DNA Center and ISE integration to automate consistent policy deployment across the fabric.

NEW QUESTION: 39
Refer to the exhibit.
C0FD9F48 C9ACDC725EA850EC2476EE1E



freepdfdumps.com
CISCO

A network engineer is designing a network for AS100. The design should ensure that all traffic enters AS100 via link 1 unless there is a network failure. In the event of a failure, link 2 should function as the path for incoming traffic. Which solution should the design include?

A. Modify the next-hop attribute on R3.

B. Use AS-Path prepending on R3.

C. Modify the next-hop attribute on R4.

D. Use AS-Path prepending on R4.

Answer: D (LEAVE A REPLY)

AS-path prepending is the correct inbound path-control tool for this design. The objective is for traffic entering AS100 to use link 1 during normal operation and link 2 only after a failure. In BGP, the AS_PATH attribute is a primary path-selection attribute used by external autonomous systems; a shorter AS path is normally preferred over a longer path. By prepending AS100 multiple times on advertisements sent from R4 across link 2, AS100 makes the backup path look less attractive to upstream networks while leaving it available for failover. Modifying next-hop on R3 or R4 does not influence how external networks choose between two entry points into AS100. Prepending on R3 would make the preferred link less attractive, which is the opposite of the requirement. The design should also document provider policy because some providers local-preference routes before AS path, but within the choices, prepending on R4 is the appropriate inbound traffic-engineering method. Reference topics: BGP AS_PATH, inbound path control, multihoming, backup Internet edge design.

NEW QUESTION: 40
An engineer must propose a QoS architecture model that allows an application to inform the network of its traffic profile and to request a particular type of service to support its bandwidth and delay requirements. The application requires consistent and dedicated bandwidth end to end. Which QoS architecture model meets these requirements?
A. DiffServ
B. LLQ
C. WRED
D. IntServ
Answer: (SHOW ANSWER)
Integrated Services is the QoS model that allows an application to signal its traffic requirements to the network and request specific service treatment. Cisco QoS guidance distinguishes IntServ from DiffServ by explaining that IntServ follows a signaled-QoS model, while DiffServ uses provisioned per-hop behavior based on markings. IntServ uses RSVP to reserve resources along the end-to-end path. That reservation process allows the network to admit or deny the flow based on available bandwidth and policy, which is why IntServ aligns with applications that require consistent and dedicated bandwidth. DiffServ is much more scalable and widely deployed in enterprise networks, but it does not provide per-flow resource reservation by itself. LLQ is a queuing mechanism commonly used inside a DiffServ design for delay-sensitive traffic such as voice. WRED is a congestion-avoidance technique, not an end-to-end reservation architecture. Because the question explicitly says that the application informs the network of its traffic profile and requests a service level, the correct QoS architecture is IntServ.

NEW QUESTION: 41
A client is moving to Model-Driven Telemetry and requires periodic updates. What must the network architect consider with this design?
A. Updates that contain changes within the data are sent only when changes occur.
B. Empty data subscriptions do not generate empty update notifications.
C. Periodic updates include a full copy of the data that is subscribed to.
D. The primary push update is sent immediately and cannot be delayed.
Answer: C (LEAVE A REPLY)
For model-driven telemetry with periodic updates, the architect must understand that each periodic update sends the subscribed data at the configured interval. Cisco telemetry guidance distinguishes periodic subscriptions from on-change subscriptions. A periodic subscription streams a full copy of the subscribed data element or table at each interval, even if the underlying state has not changed. That behavior is useful when the collector needs regular time-series samples, trending, or baseline monitoring, but it creates predictable bandwidth and collector-load requirements. On-change subscriptions are different: they publish updates when the subscribed state changes and are more efficient for state changes such as interface status, memory thresholds, or protocol adjacency m ovement. Empty subscription behavior and initial update timing may matter in specific implementations, but they are not the main design consideration in this question. Since the customer requires periodic updates, the design must size the transport, collector, retention system, and subscription intervals around repeated full data delivery. Reference topics: model-driven telemetry, periodic subscription, on-change subscription, YANG-modeled operational data, collector sizing.

NEW QUESTION: 42
An organization is designing a detailed QoS plan that limits bandwidth to specific rates. Which two parameters are supported be the traffic policing feature? (Choose two.)
A. violating
B. marking
C. shaping
D. bursting
E. conforming
Answer: A,E (LEAVE A REPLY)
Cisco traffic policing evaluates packets against a configured traffic contract and classifies them into action states such as conforming, exceeding, and violating. The policy then applies configured actions to those states, such as transmit, drop, or remark. In a two-rate or three-color policing design, conforming traffic is within the committed rate, exceeding traffic is above the committed rate but may still be within an excess burst allowance, and violating traffic is outside the permitted profile. Therefore, conforming and violating are valid policing parameters or state/action categories in the context of the question.

Shaping is not a policing parameter; it is a different QoS mechanism that buffers excess traffic and sends it later. Marking can be an action applied by policing, but the listed policing state that directly corresponds to Cisco class-based policing syntax is conform or violate. Bursting is related to token bucket sizing, but the answer choices specifically align to policing state categories. The correct pair is violating and conforming.

NEW QUESTION: 43
An engineer must design an in-band management solution for a customer with branch sites. The solution must allow remote management of the branch sites using management protocols over an MPLS WAN. Queueing is implemented at the remote sites using these classes:

- Class1 equals voice traffic
- Class2 equals mission-critical traffic
- Class3 equals default traffic

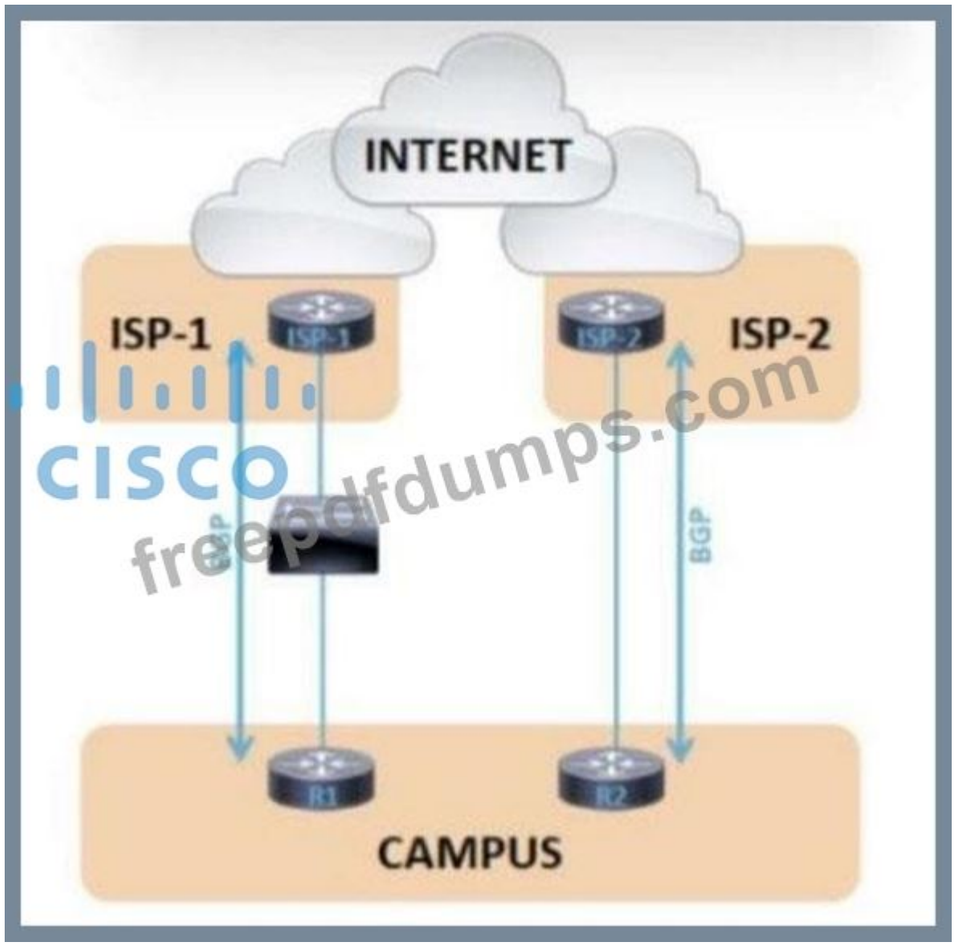
How must the solution prioritize the management traffic over the WAN?
A. Mark the traffic with DSCP CS1 and map into Class2 with a minimum bandwidth assigned by reducing the bandwidth available to Class3.
B. Mark the traffic with DSCP CS6 and map into Class1 with a minimum bandwidth assigned by reducing the bandwidth available to Class2
C. Mark the traffic with DSCP EF and map into Class1 with a minimum bandwidth assigned by reducing the bandwidth available to Class2.
D. Mark the traffic with DSCP CS2 and map into Class2 with a minimum bandwidth assigned by reducing the bandwidth available to Class3
Answer: D (LEAVE A REPLY)
For in-band management over the production WAN, the proper marking is DSCP CS2 mapped into the management or operations class with a minimum bandwidth guarantee. Cisco enterprise QoS guidance separates network control traffic from network management traffic. Routing protocol and control-plane traffic should not be diluted by ordinary management flows. CS6 is typically reserved for routing and network- control traffic; EF is reserved for low-latency voice bearer traffic; CS1 is commonly used for scavenger or low-priority traffic. Network management flows such as SSH, SNMP, NTP, FTP-based maintenance, and syslog require reliable delivery, but they should not starve business traffic or real-time services. Therefore, marking with CS2 and mapping the traffic into Class2 is the correct design choice. The bandwidth adjustment should come from a lower-priority class, not from the routing/control or real-time class. Because this is in- band management, the architect must also ensure access control, source restrictions, AAA, logging, and management-plane protection. QoS classification alone does not secure management access; it only gives the traffic an appropriate forwarding treatment over the MPLS WAN. Reference topics: in-band management, DSCP CS2, enterprise QoS classes, management-plane traffic.

NEW QUESTION: 44
What are the two purpose of the RPF check in multicast routing?
A. to ensure that multicast packets are forwarded if they arrived on the interface used to route traffic back to the source address
B. to ensure that multicast packets are forwarded if they arrived on the interface used to route traffic to the destination address
C. to ensure that multicast packets, no matter the interface they arrived on. are forwarded out all interfaces
D. to ensure that multicast packets are dropped if they arrived on the interface used to route traffic to the destination address
E. to ensure that multicast packets are dropped if they arrived on the interface used to route traffic back to the source address
Answer: (SHOW ANSWER)
The RPF check verifies that multicast traffic arrives on the interface the router would use to reach the source or RP, and it prevents forwarding when traffic arrives from the wrong direction. Cisco multicast forwarding is intentionally different from unicast forwarding because a multicast packet may need to be replicated toward many receivers. Without Reverse Path Forwarding, multicast loops and duplicate packets could occur. Option A captures the forwarding condition: the packet is accepted when it arrives on the interface used to route back toward the source address. Option D is the closest available drop condition in the answer set, expressing that packets not aligned with the correct reverse path are discarded rather than forwarded. Option B incorrectly checks the route toward the destination group, because multicast groups are not reached through ordinary unicast destination routing. Option C would flood traffic and create loops. Option E reverses the correct logic by dropping packets that arrive on the proper reverse-path interface. Reference topics: multicast RPF, loop prevention, source-tree forwarding, shared-tree forwarding, PIM.

NEW QUESTION: 45
An architect must design a topology for a WAN network that satisfies these requirements:

* Devices must be able to make informed decisions.
* Suboptimal paths are allowed only In case of a failure.
* Backup paths must always be available.
Which topology must the architect select?
A. partial mesh
B. hub and spoke
C. full mesh
D. Clos
Answer: (SHOW ANSWER)
A full-mesh topology best satisfies the requirement that devices make informed forwarding decisions, backup paths always exist, and suboptimal paths are used only after a failure. In a full mesh, each site or WAN node has direct connectivity to every other relevant node. This gives the routing protocol complete path visibility and allows traffic to use direct optimal paths under normal conditions. If a direct link fails, alternate paths remain available through other nodes, so the network can reroute without depending on a single hub. A hub- and-spoke design is simpler and less expensive, but spoke-to-spoke traffic must traverse the hub, which is suboptimal even during normal operation. A partial mesh can reduce cost but may still force some traffic over indirect paths when no direct connection exists. Clos is a data center fabric topology concept and is not the typical WAN topology answer here. The full mesh is the most resource-intensive design, but it best matches the stated forwarding and backup requirements.
Reference topics: WAN topology, full mesh, path optimality, backup paths, routing convergence.

NEW QUESTION: 46
Refer to the exhibit.



freepdfdumps.com

The failover time of ISP-2 is significantly shorter than ISP-1 when an interface on the ISP router toward the campus network fails. Which solution minimizes the downtime to the sub-second?

A. Aggressive timers

B. Next-hop address tracking

C. Graceful-restart

D. BFD

Answer: D (LEAVE A REPLY)

BFD is the correct mechanism for sub-second failure detection. The question describes unequal failover behavior when an interface on an ISP router toward the campus fails. In many WAN and Internet edge designs, BGP timers or next-hop tracking may not react fast enough when the physical failure is not directly visible to the customer router. Cisco defines Bidirectional Forwarding Detection as a lightweight protocol that provides rapid forwarding-path failure detection independent of the routing protocol timers. When BFD is associated with BGP, OSPF, EIGRP, or static routing, the routing process receives failure notification quickly and can withdraw or replace the route without waiting for longer protocol hold timers. Aggressive timers can improve reaction time but increase session churn and CPU load. Next-hop address tracking improves BGP response when next-hop reachability changes, but it is not the most precise sub-second detection tool.

Graceful restart is intended to preserve forwarding during control-plane restart, not detect failures faster.

Reference topics: BFD, WAN edge resiliency, BGP failure detection, sub-second convergence.

Valid 300-420 Dumps shared by Actual4test.com for Helping Passing 300-420 Exam! Actual4test.com now offer the **newest 300-420 exam dumps**, the Actual4test.com 300-420 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-420 dumps with Test Engine here: https://www.actual4test.com/300-420_examcollection.html (**457** Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

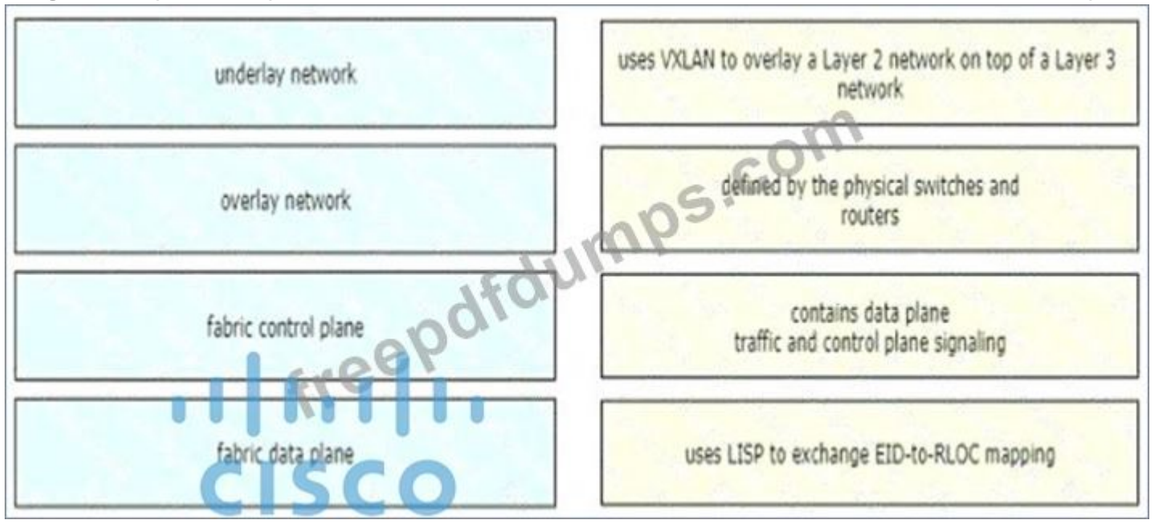
NEW QUESTION: 47
An engineer is designing a QoS solution for a customer The customer ' s internet connection has a bandwidth of 10 Mbps. The design must ensure that traffic bursts of data do not exceed the bandwidth of the connection and that received traffic does not starve out business-critical traffic Which solution must the engineer choose?
A. Configure the queuing default queue for shaping inbound and policing outbound.
B. Configure the queuing default queue for shaping inbound and policing inbound.
C. Configure the queuing default queue for shaping outbound and policing inbound.
D. Configure the queuing default queue for shaping outbound and policing outbound.
Answer: C (LEAVE A REPLY)
The correct QoS design is to shape outbound traffic and police inbound traffic. Traffic shaping buffers and schedules outbound packets so the customer does not send bursts beyond the 10 Mbps Internet connection rate. This is useful on egress because the local router controls the transmit queue and can smooth bursts before they reach the provider. Cisco QoS design distinguishes shaping from policing: shaping delays excess traffic, while policing drops or remarks excess traffic immediately. Inbound shaping is generally not effective because the traffic has already arrived on the interface; the device cannot delay packets before they consume the inbound circuit. Inbound policing can protect business-critical traffic by limiting or dropping lower- priority excess traffic as it enters the network. Policing outbound would enforce the rate but would cause avoidable drops during bursts, which is less desirable when the router can shape before transmission.
Therefore, outbound shaping and inbound policing satisfy both rate-control and protection requirements.
Reference topics: traffic shaping, traffic policing, QoS queuing, Internet edge QoS, congestion management.

NEW QUESTION: 48
A router running ISIS is showing high CPU and bandwidth utilization. An engineer discovers that the router is configured as L1/L2 and has L1 and L2 neighbors. Which step optimizes the design to address the issue?
A. Make this router a DIS for each of the interfaces
B. Disable the default behavior of advertising the default route on the L1/L2 router
C. Configure the router to be either L1 or L2
D. Configure each interface as either L1 or L2 circuit type
Answer: (SHOW ANSWER)
Cisco IS-IS supports Level 1, Level 2, and Level 1/Level 2 operation, and the design should restrict each link to the level it actually needs. Cisco documentation for the isis circuit-type command states that it "configures the type of adjacency that is required for neighbors on the specified interface." In this case, the router is configured as L1/L2 and has both L1 and L2 neighbors, which causes unnecessary LSP flooding, SPF processing, and database maintenance on links that may not need both levels. Making the entire router only L1 or only L2 would be too blunt, because the router may legitimately need to participate in both domains on different interfaces. Making the router a DIS also does not solve the underlying issue; it only controls designated intermediate system behavior on broadcast networks. The best optimization is to configure each interface with the correct circuit type, such as level-1 for intra-area links and level-2-only for backbone links.
This reduces control-plane load while preserving the intended two-level hierarchy.

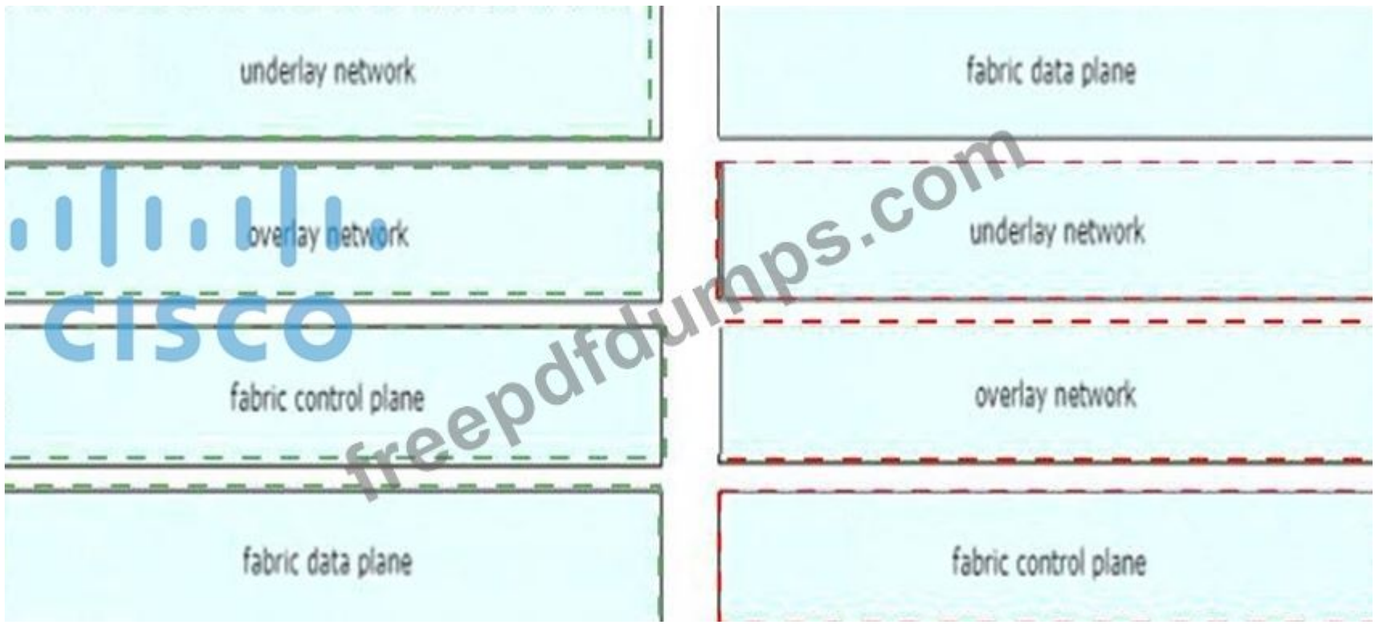
NEW QUESTION: 49
Refer to the exhibit. An engineer must ensure that the QoS design guarantees bandwidth for the applications, and an application can request a particular type of service to support its delay requirements. Which solution must the engineer select?
A. IntServ with DSCP
B. DiffServ with DSCP

C. IntServ with RSVP
D. DiffServ with RSVP
Answer: C (LEAVE A REPLY)
The correct solution is IntServ with RSVP. The requirement says the application must be able to request a specific service level and receive guaranteed bandwidth with support for delay requirements. Cisco describes Integrated Services as a QoS model that uses Resource Reservation Protocol to signal per-flow resource requirements across the network. RSVP allows applications or endpoints to request bandwidth and delay treatment; each participating network device can admit or reject the reservation based on available resources.
This is exactly the behavior needed when an application must explicitly request a service profile rather than simply inherit a class marking. DiffServ with DSCP is more scalable, but it only marks traffic and applies class-based per-hop behavior. It does not provide per-flow reservation or allow the application to explicitly reserve resources. DSCP by itself is a marking mechanism, not a reservation protocol.
RSVP belongs with IntServ in the classic QoS model. Reference topics: Integrated Services, RSVP, resource reservation, QoS guarantees, delay-sensitive applications.

NEW QUESTION: 50
Drag and drop the components in a Cisco SD-Access architecture from the left onto their descriptions on the right.



Answer:

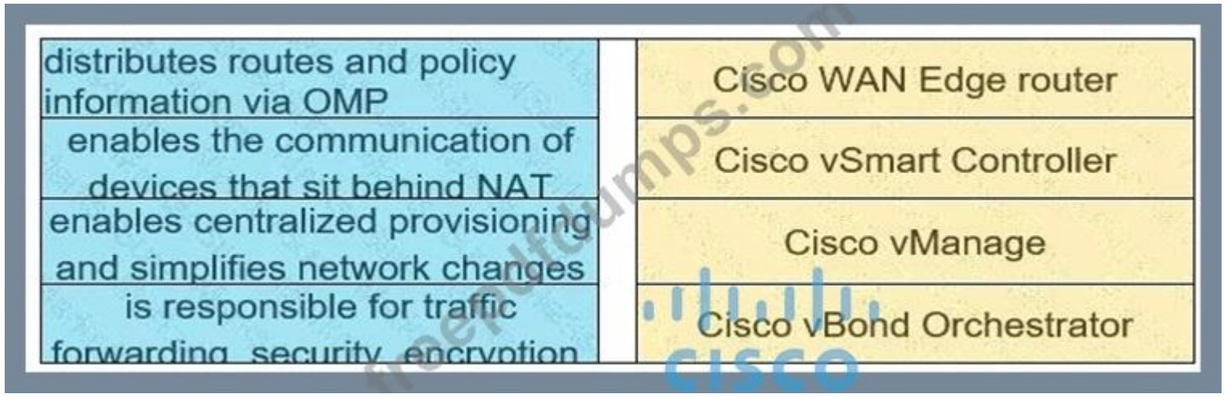


freepdfdumps.com

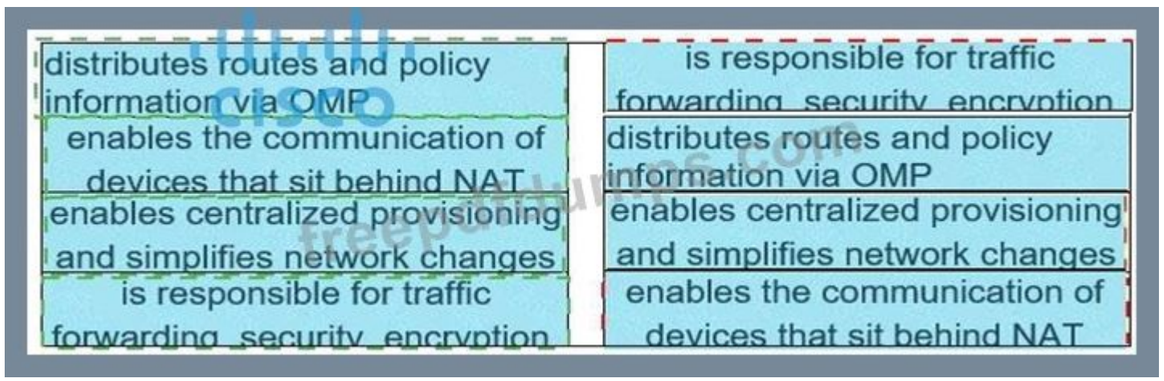
Explanation:
The Cisco SD-Access component mapping is based on the normal fabric roles. A fabric edge node is the access-facing fabric device that connects wired endpoints and fabric access points, acts as the anycast Layer 3 gateway, encapsulates endpoint traffic into VXLAN, and registers endpoint reachability with the control plane. A control-plane node supplies the LISP map-server and map-resolver functions, maintaining endpoint- to-location mappings. A border node connects the fabric to external Layer 3 domains, shared services, WAN, Internet, or another fabric/transit domain, and it de-encapsulates traffic as it leaves the fabric. Intermediate nodes provide routed IP transport inside the underlay and move packets between edge, border, and control- plane nodes without owning endpoint policy or registration. This role separation is important because a design failure often comes from assigning the wrong function to the wrong node. Access-facing policy and endpoint onboarding belong at the edge; external connectivity belongs at the border; location resolution belongs to the control plane; transport belongs to intermediate nodes. Reference topics: SD-Access fabric roles, edge node, border node, control-plane node, intermediate node.

NEW QUESTION: 51
Which two overlay network design considerations must be made for a Cisco SD-Access network? (Choose two.)
A. Dedicated IGP process for the fabric
B. LAN automation for deployment
C. Reduce subnets and simplify DHCP management
D. Avoid overlapping IP subnets
E. Layer 3 to the access design
Answer: C,D (LEAVE A REPLY)

NEW QUESTION: 52
Drag and drop the descriptions from the left onto the Cisco SD-WAN component they describe on the right.



Answer:



Explanation:

WAN Edge: traffic forwarding, security, encryption, QoS, and routing protocols; vSmart: distributes routes and policy information via OMP; vManage: centralized provisioning and simplified network changes; vBond:

communication for devices behind NAT.

The Cisco SD-WAN component mapping follows the normal separation of management, control, orchestration, and data-plane roles. WAN Edge routers are the data-plane devices at branches, data centers, campuses, or cloud edges. They forward user traffic and apply routing, QoS, encryption, segmentation, and security services. vSmart controllers provide the centralized control plane and distribute reachability and policy information with OMP, including prefixes, TLOCs, and service routes. vManage is the management- plane system; it supplies centralized provisioning, device templates, policy configuration, software lifecycle functions, monitoring, and operational visibility. vBond operates as the orchestration-plane component, authenticating devices during onboarding and helping devices behind NAT discover and establish the correct secure control connections. This separation is critical in enterprise design because each component has different redundancy, reachability, and scale requirements. The architect should not blur vManage and vSmart roles: vManage configures and monitors, while vSmart computes and distributes control policy. Reference topics: Cisco Catalyst SD-WAN architecture, WAN Edge data plane, vSmart OMP, vManage management plane, vBond orchestration.

is responsible for traffic forwarding, security, encryption, QoS,and routing protocols distributes routes and policy information via OMP enables centralized provisioning and simplifies network changes enables the communication of devices that sit behind NAT

NEW QUESTION: 53

An engineer is designing a networking solution to allow two hosts to communicate-one host located within the company A network and the other within the company B network. The two companies have no other plans for future additional connections. Both companies want to use a single secure and encrypted internet connection, and the configuration must be as simple as possible. Which network solution must the engineer choose?

A. single DMVPN with EIGRP routing

B. routed IPsec tunnel with OSPF routing

C. policy-based IPsec tunnel with static routing
D. MPLS VPN provided service with BGP routing
Answer: C (LEAVE A REPLY)
A policy-based IPsec tunnel with static routing is the simplest secure design for connecting two hosts in two companies when no additional future connectivity is planned. The requirement is narrow: one host in Company A must communicate with one host in Company B over a single encrypted Internet connection.
Policy-based IPsec is straightforward for this type of point-to-point encryption because crypto ACLs define exactly which interesting traffic should be protected. Static routing is sufficient because there is no complex topology, no dynamic path selection requirement, and no anticipated expansion. A DMVPN design would add unnecessary NHRP, multipoint GRE, and dynamic routing complexity. A routed IPsec tunnel with OSPF is useful when many networks or dynamic route changes are expected, but it is more complex than needed.
MPLS VPN with BGP would require provider service and does not meet the simple encrypted Internet connection requirement. Reference topics: site-to-site IPsec, policy-based VPN, static routing, point-to-point secure connectivity, WAN simplicity.

NEW QUESTION: 54
Which function does the Cisco SD-Access intermediate node perform?
A. Act as LISP proxy tunnel router.
B. Route and transport IP traffic.
C. Act as an anycast Layer 3 gateway.
D. Map users to a virtual network.
Answer: B (LEAVE A REPLY)
An SD-Access intermediate node routes and transports IP traffic between fabric nodes. It is part of the underlay transport path, comparable to a routed distribution or core device that provides IP reachability between edge, border, and control-plane nodes. Intermediate nodes do not perform endpoint registration, VXLAN encapsulation for attached users, virtual-network mapping, or anycast gateway services. Those functions belong primarily to fabric edge nodes. They also do not act as LISP proxy tunnel routers, which is a border-node function used for communication between the fabric and external networks. The intermediate node's role is deliberately simple: maintain underlay routing, provide resilient transport, and forward encapsulated traffic across the fabric infrastructure. This separation of roles is what lets Cisco SD-Access scale by keeping endpoint intelligence at the edge and mapping intelligence in the control plane while the intermediate layer remains a stable IP transport network. Reference topics: SD-Access intermediate node, routed underlay, IP transport, fabric roles, edge and border node separation.

NEW QUESTION: 55
A customer plans to adopt distributed QoS in their enterprise WAN. The policy must allow for individual packet marking according to the type of treatment required and for forwarding based on hop-by-hop treatment locally defined on each device. Which technology must the customer select?
A. CBWFQ
B. LLQ
C. Diffserv
D. IntServ
Answer: (SHOW ANSWER)
DiffServ is the correct technology because the customer wants distributed QoS based on packet marking and per-hop treatment. Cisco differentiated services designs classify and mark traffic at the network edge, normally using DSCP values, and each device along the path applies a locally defined per-hop behavior to that traffic class. This model scales well in enterprise WANs because routers do not need to maintain per-flow state across the entire network. CBWFQ and LLQ are queuing mechanisms that can implement DiffServ treatment, but they are not the overall QoS architecture requested by the question. IntServ uses RSVP and per- flow resource reservation, which provides stronger admission-style guarantees but does not scale as well for large distributed WAN environments. The question specifically mentions individual packet marking and forwarding based on hop-by-hop treatment, which is the language of DiffServ. A complete design should define trust boundaries, DSCP markings, class maps, queue policy, and provider handoff behavior. Reference topics: DiffServ, DSCP marking, per-hop behavior, distributed QoS, WAN QoS architecture.

NEW QUESTION: 56

When differentiating between IETF. OpenConfig. and Cisco native YANG models, how does the use of containers differ?

A. OpenConfig uses one container for operational data and another container for configuration data, and IETF and Cisco native models use a single container for operational data and configuration data.

B. IETF and Cisco native models use a single container for operational data and configuration data, and OpenConfig uses one container for operational data and another container for configuration data.

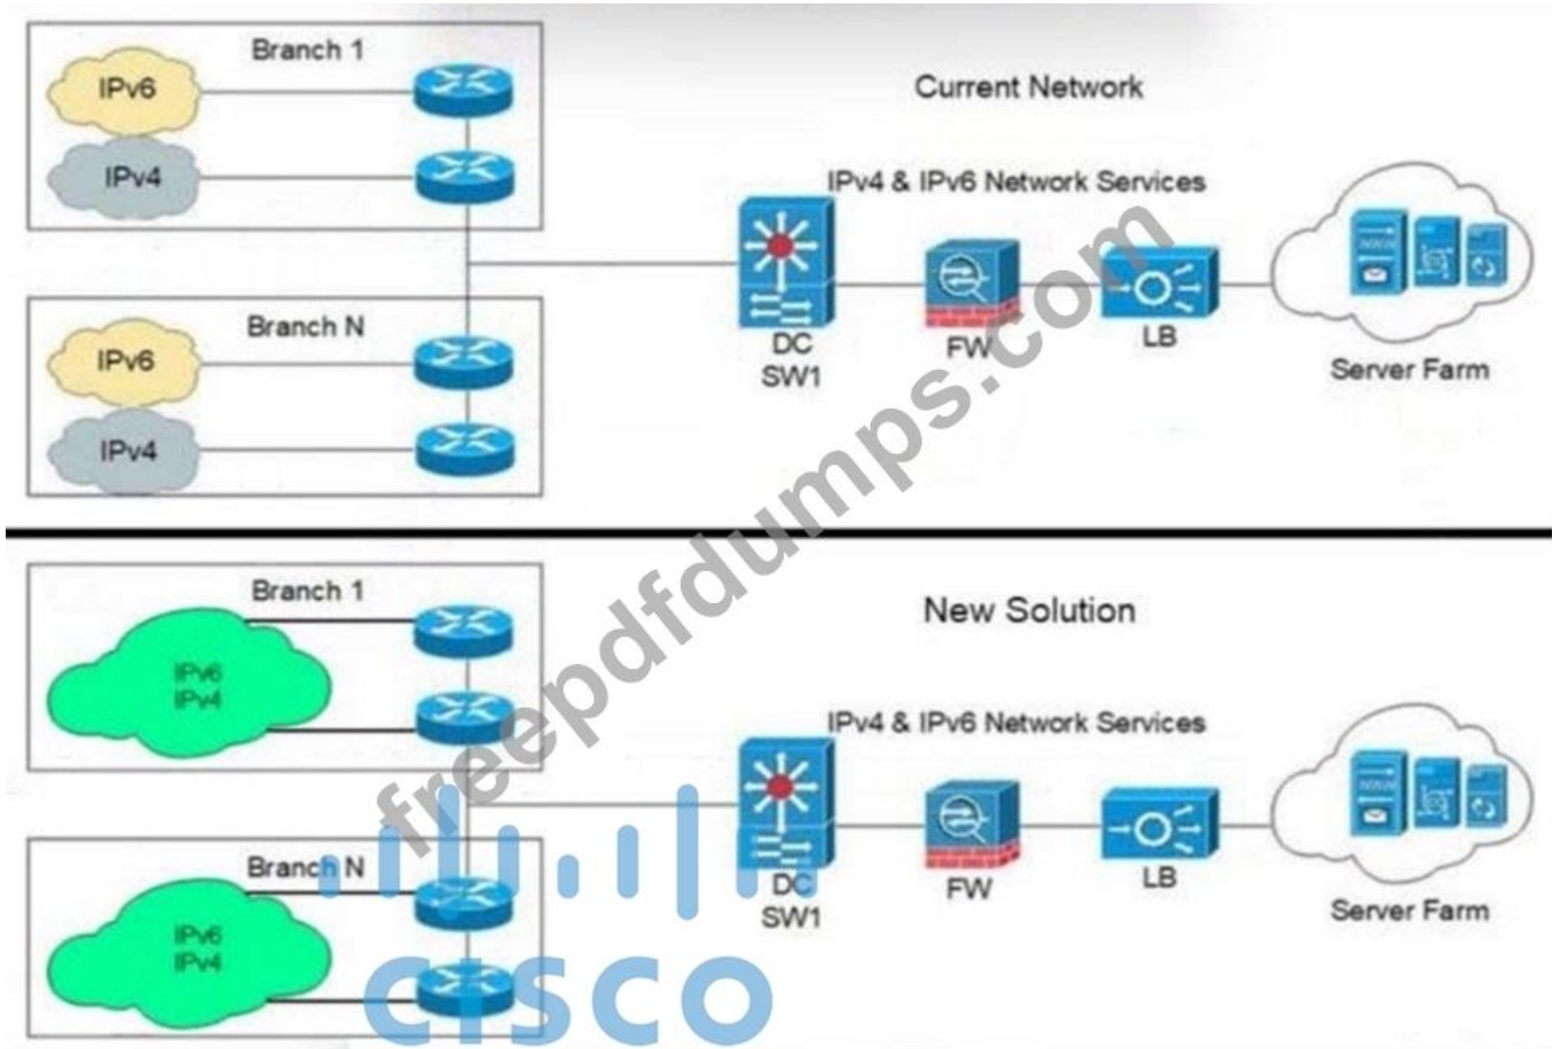
C. IETF and Cisco native models use one container for operational data and another container for configuration data, and OpenConfig uses a single container for operational data and configuration data.

D. Cisco native models use one container for operational data and another container for configuration data, and OpenConfig and IETF use a single container for operational data and configuration data.

Answer: B (LEAVE A REPLY)

OpenConfig separates intended configuration from operational state by using separate config and state containers, while many IETF and Cisco native models use a single hierarchy for configuration and operational data or expose operational data differently. This distinction is important in automation because the client must know whether it is writing desired configuration or reading actual device state. OpenConfig's structure was designed to provide consistent multivendor modeling, with configuration data placed under config containers and observed state under state containers. Cisco native models tend to mirror the Cisco device CLI and feature hierarchy more closely, while IETF models focus on standards-based interoperability. Option B correctly states the contrast presented in the question. Option A reverses the relationship. Option C incorrectly assigns split containers to IETF and native models. Option D incorrectly isolates Cisco native models as the only split- container style. Reference topics: YANG model structure, OpenConfig config and state containers, IETF YANG, Cisco native models, network automation data modeling.

NEW QUESTION: 57



freepdfdumps.com

Refer to the exhibit. An architect is developing a solution to consolidate networks while retaining device redundancy. The routing protocol for the WAN routers must be open standard, ensure high availability, and provide the fastest convergence time. Which solution must the design include?

A. both routers running EIGRP

B. one router running OSPFv2 and other OSPF v3

C. one router running ISIS and other OSPF v3

D. both routers running OSPFv2

Answer: (SHOW ANSWER)

Both WAN routers should run OSPFv2. The requirement explicitly says that the routing protocol must be open standard, provide high availability, and offer fast convergence. OSPF is an open-standard link-state routing protocol and is broadly supported across vendors. It builds a link-state database, calculates best paths with SPF, supports equal-cost multipath, and can be tuned with BFD, SPF throttling, LSA throttling, and hierarchical areas for resilient enterprise WAN designs. EIGRP is fast and operationally efficient in Cisco- only environments, but it is not the best answer when the requirement emphasizes an open-standard routing protocol. Running OSPFv2 on one router and OSPFv3 on another is invalid for a single IPv4 routing design because they would not form the intended common routing domain. Mixing IS-IS and OSPFv3 also fails the requirement for a unified design between the redundant routers. Therefore, both routers should run OSPFv2.

Reference topics: OSPFv2, open-standard routing, WAN routing high availability, link-state convergence, ECMP and BFD integration.

NEW QUESTION: 58

An organization plans to deploy multicast across two different autonomous systems. Their solution must allow RPs to:

*discover active sources outside their domain

*use the underlying routing information for connectivity with other RPs

*announce sources joining the group

Which solution supports these requirements?

A. MSDP

B. SSM

C. PIM-SM

D. PIM-DM

Answer: A (LEAVE A REPLY)

MSDP is the correct solution for multicast across different autonomous systems when RPs must discover active sources outside their local domain. In PIM Sparse Mode, each multicast domain can have its own rendezvous point. Without a mechanism to exchange source information, receivers in one domain may not know about active sources registered with an RP in another domain. Multicast Source Discovery Protocol solves that problem by allowing RPs to establish peer relationships and exchange Source-Active information.

Cisco multicast design uses MSDP with PIM-SM when separate domains need to share multicast source reachability while relying on the unicast routing table for connectivity between RPs. SSM does not use an RP and depends on receivers joining a specific source and group. PIM-SM by itself works inside a domain but does not provide interdomain source discovery between RPs. PIM-DM is flood-and-prune and is not suitable for this sparse, interdomain design. Therefore, the organization should deploy MSDP between the RP devices or RP domains. Reference topics: MSDP, PIM-SM, interdomain multicast, Source-Active messages, RP-to- RP source discovery.

NEW QUESTION: 59

Which method will filter routes between EIGRP neighbors within the same autonomous system?

A. distribute-list

B. policy-based routing

C. leak-map

D. route tagging

Answer: A (LEAVE A REPLY)

A distribute list is the correct method to filter routes exchanged between EIGRP neighbors within the same autonomous system. In Cisco EIGRP designs, distribute lists can be applied inbound or outbound under the EIGRP process to control which routes are accepted from or advertised to neighbors. The filter can reference an access list, prefix list, or route map depending on platform and configuration style. This allows the designer to limit route visibility without changing the forwarding behavior of packets that are already routed.
Policy-based routing is a data-plane forwarding feature that changes next-hop selection for matching traffic; it does not filter EIGRP route advertisements between neighbors. Leak maps are used with route summarization to selectively advertise more specific routes through a summary, not as the general route-filtering method between EIGRP neighbors. Route tagging is valuable in redistribution designs to identify route origin and prevent feedback loops, but the tag alone does not filter routes unless a route map or distribute list uses it.
Therefore, the correct route-filtering mechanism between EIGRP neighbors is a distribute list.

NEW QUESTION: 60
How is internet access provided to a WAN edge router that is connected to a MPLS transport link?
A. OMP advertises a default route from a WAN Edge router that is connected to the MPLS and internet transport networks
B. Internet access must be provided at the WAN Edge router through either a 4G/5G link or local Internet circuit
C. An extranet must be provided in the MPLS transport network to allow private traffic to reach the public internet
D. TLOC extensions are used to route traffic to a WAN Edge router that is connected to the Internet transport network
Answer: D (LEAVE A REPLY)
TLOC extension is the Cisco SD-WAN mechanism that allows a WAN Edge router without direct access to a transport to reach that transport through another WAN Edge router at the same site. Cisco describes TLOC extension as a redundancy feature in SD-WAN networks. In the question, the WAN Edge router is connected to an MPLS transport link, but Internet access is available through another WAN Edge device. With TLOC extension, the MPLS-only edge can extend traffic through the peer router toward the Internet transport, allowing the site to use multiple transports without requiring every edge router to have a physical connection to every circuit. OMP can advertise routes and TLOCs, but OMP alone does not create physical Internet access for a router attached only to MPLS. Requiring a local 4G/5G or Internet circuit on the same router is not necessary when TLOC extension is part of the design. An MPLS extranet is also not the normal SD-WAN solution for Internet transport reachability.
Therefore, TLOC extension is the correct design method.

NEW QUESTION: 61

```
SW2#show spanning-tree vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    24577
              Address     0011.0022.0033
              Cost        4
              Port        3 (FastEthernet0/3)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
              Address     0011.00bb.00cc
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 15

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Fa0/3               Root FWD 4         128.3    P2p
Gig0/2              Altn BLK 4         129.4    P2p

SW2#show run interface Gig0/2
Building configuration...
!
Interface GigabitEthernet0/2
 switchport mode dynamic desirable

SW2#show run interface Fa0/3
Building configuration...

Current configuration : 102 bytes
!
Interface FastEthernet0/3
 switchport mode dynamic desirable
 spanning-tree cost 4
```



Refer to the exhibits. An engineer is troubleshooting an issue in which the Gig0/2 interface on a Cisco switch named SW2 fails to become the root port. Which two commands must be run on SW2 to resolve this issue?

(Choose two.)

A.

```
SW2(config)# interface Fa0/3
```

B.
```
SW2(config-if)# no spanning-tree cost 4
```
C.
```
SW2(config)# interface Gig0/2
```
D.
```
SW2(config-if)# spanning-tree cost 4
```
E.
```
SW2(config-if)# spanning-tree cost 5
```

Answer: C,D (LEAVE A REPLY)

The selected command pair is appropriate because the issue is that GigabitEthernet0/2 on SW2 is not becoming the root port as expected. In spanning-tree design, the root port is selected based on the lowest root path cost, then upstream bridge ID, then port ID when necessary. If the wrong port is selected, the engineer must adjust the local spanning-tree parameters on SW2 so that the intended interface has the best path toward the root bridge. Typical corrective commands lower the spanning-tree cost or set the port priority on the desired interface or VLAN instance. The correct answer pair is therefore the combination that changes the STP decision variables for Gig0/2 rather than changing unrelated features such as PortFast, BPDU Guard, or VLAN trunking. PortFast is unsafe on switch-to-switch links, and protection features do not make an interface more attractive as a root port. The fix must influence STP path selection while keeping the topology loop-free.

Reference topics: STP root port election, path cost, port priority, root bridge path selection, Catalyst Layer 2 troubleshooting.

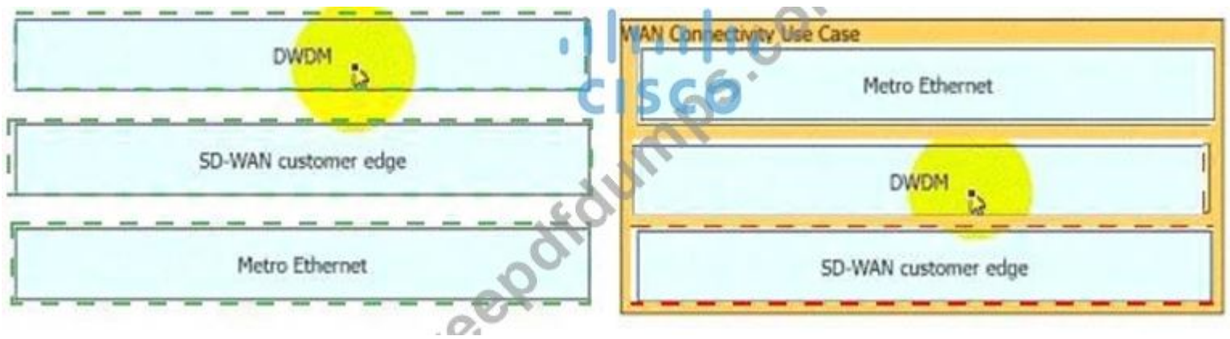
Valid 300-420 Dumps shared by Actual4test.com for Helping Passing 300-420 Exam! Actual4test.com now offer the **newest 300-420 exam dumps**, the Actual4test.com 300-420 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-420 dumps with Test Engine here: https://www.actual4test.com/300-420_examcollection.html (**457** Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 62

Drag and drop the types of WAN connectivity from the left onto the connectivity use cases on the right.



Answer:



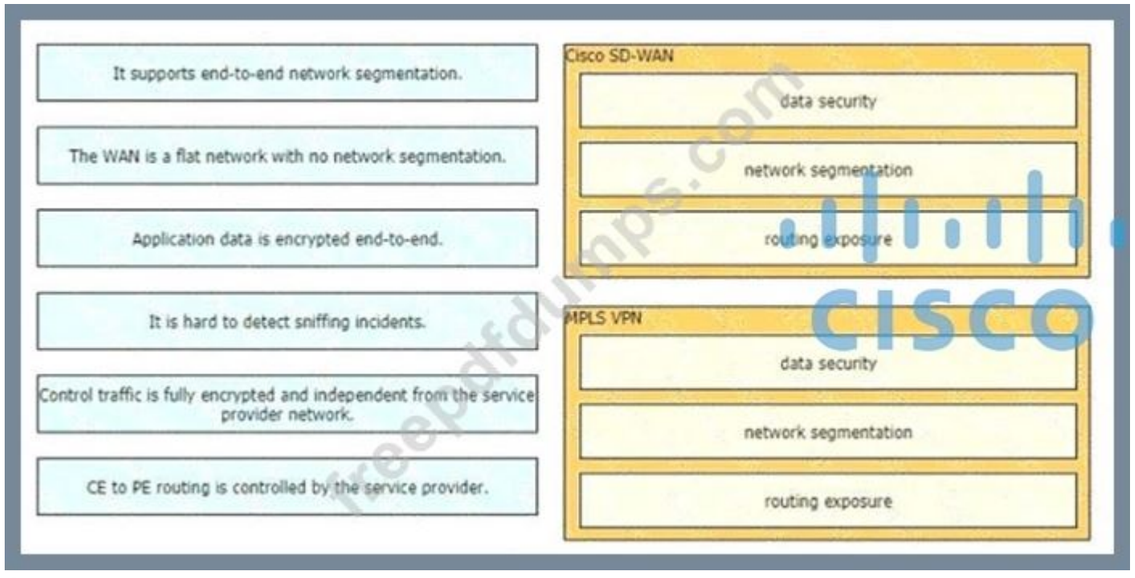
Explanation:
The WAN connectivity mapping distinguishes private, public, and overlay connectivity use cases. Enterprise WAN design normally selects the transport according to the required security, availability, latency, cost, and operational control. MPLS Layer 3 VPN is commonly used when the enterprise needs private routed WAN service with provider-managed segmentation and predictable QoS behavior. Internet VPN designs are used when cost, rapid deployment, or broadband availability is more important, with IPsec or SD-WAN providing encryption and overlay control. Layer 2 services such as VPWS or VPLS are selected when the customer requires Ethernet adjacency, VLAN transport, or transparent LAN-like connectivity across sites. The drag- and-drop answer preserves those distinctions by matching each WAN type to the use case it is designed to support. A professional WAN architecture can combine these transports in a hybrid model, but the design must still identify which service carries private routing, which carries encrypted overlay traffic, and which provides Layer 2 extension. Reference topics: WAN transport selection, MPLS L3VPN, Internet VPN, Metro Ethernet, VPWS, VPLS, hybrid WAN design.

NEW QUESTION: 63
A large chain of stores currently uses MPLS-based T1 lines to connect their stores to their data center. An architect must design a new solution to improve availability and reduce costs while keeping these considerations in mind:
* The company uses multicast to deliver training to the stores.
* The company uses dynamic routing protocols and has implemented QoS.
* To simplify deployments, tunnels should be created dynamically on the hub when additional stores open.
Which solution should be included in this design?
A. VPLS
B. GET VPN
C. DMVPN
D. IPsec
Answer: (SHOW ANSWER)
DMVPN is the correct design because it satisfies all of the stated branch requirements without requiring manually built static tunnels for every store. Cisco DMVPN uses multipoint GRE with NHRP and IPsec to create a scalable hub-and-spoke overlay, and later phases allow dynamic spoke-to-spoke tunnels when branch- to-branch traffic is required. GRE supports multicast and dynamic routing protocols, while IPsec supplies encryption across the lower-cost WAN transport. This fits the requirement to deliver multicast training, use dynamic routing, preserve QoS markings and policy behavior, and simplify future branch additions because new spokes can register dynamically with the hub. VPLS would provide Layer 2 multipoint service, but it is typically a provider-delivered service and does not meet the secure dynamic VPN requirement as directly.
GET VPN is efficient for private MPLS environments, but it does not create dynamic GRE tunnels from the hub for new stores. Static IPsec is secure but operationally poor at scale and does not natively carry multicast without an overlay. Therefore, the design should include DMVPN. Reference topics: DMVPN, mGRE, NHRP, IPsec protection, multicast over VPN.

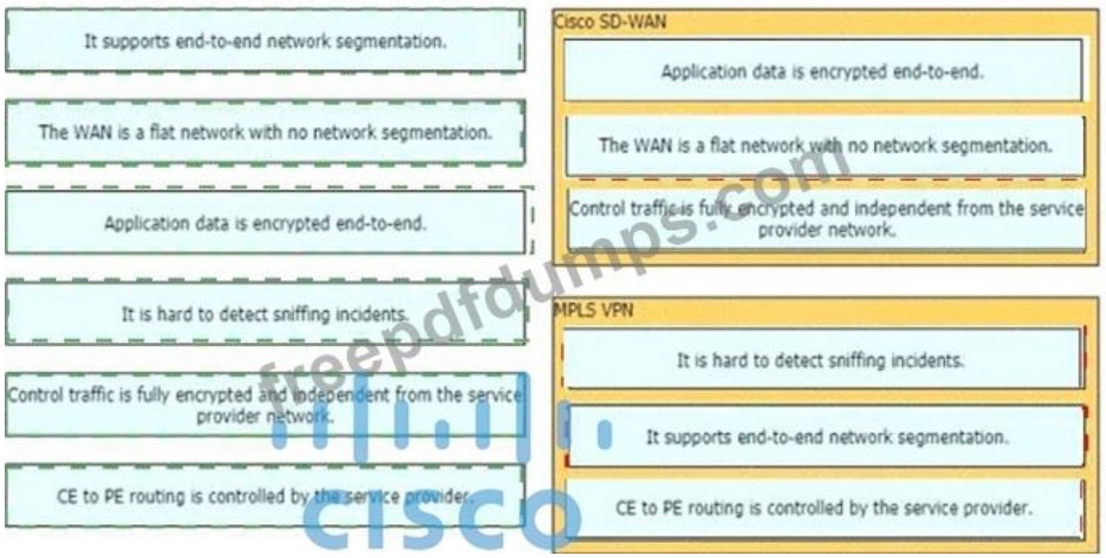
NEW QUESTION: 64
A network engineer is redesigning a company ' s QoS solution. The company is currently using IP Precedence, but the engineer plans to move to DiffServ. It is important that the new solution provide backward compatibility with the current solution. Which technology should the design include?

A. expedited forwarding
B. assured forwarding
C. class selector code points
D. default per hop behavior
Answer: C (LEAVE A REPLY)
Class selector code points provide backward compatibility between IP Precedence and DiffServ. IP Precedence uses the three most significant bits of the former Type of Service field. DiffServ expands QoS marking by using the six-bit DSCP field, but class selector values preserve compatibility by setting the three high-order bits to match the corresponding IP Precedence value while leaving the lower bits as zeros. For example, CS3 corresponds to the old precedence value 3, and CS6 corresponds to precedence 6. This allows devices that still interpret IP Precedence to provide broadly consistent treatment while newer DiffServ-aware devices can use DSCP-based policy. Expedited forwarding is used for low-loss, low-latency traffic such as voice. Assured forwarding provides drop precedence within classes. Default per-hop behavior represents best- effort treatment. None of those are specifically the backward-compatibility mechanism for migration from IP Precedence. Therefore, the design should include class selector code points during the transition to DiffServ.
Reference topics: DiffServ, DSCP, class selector, IP Precedence compatibility, QoS marking migration.

NEW QUESTION: 65
Drag and drop the description from the left onto the corresponding WAN connectivity types and categories on the right.



Answer:



Explanation:

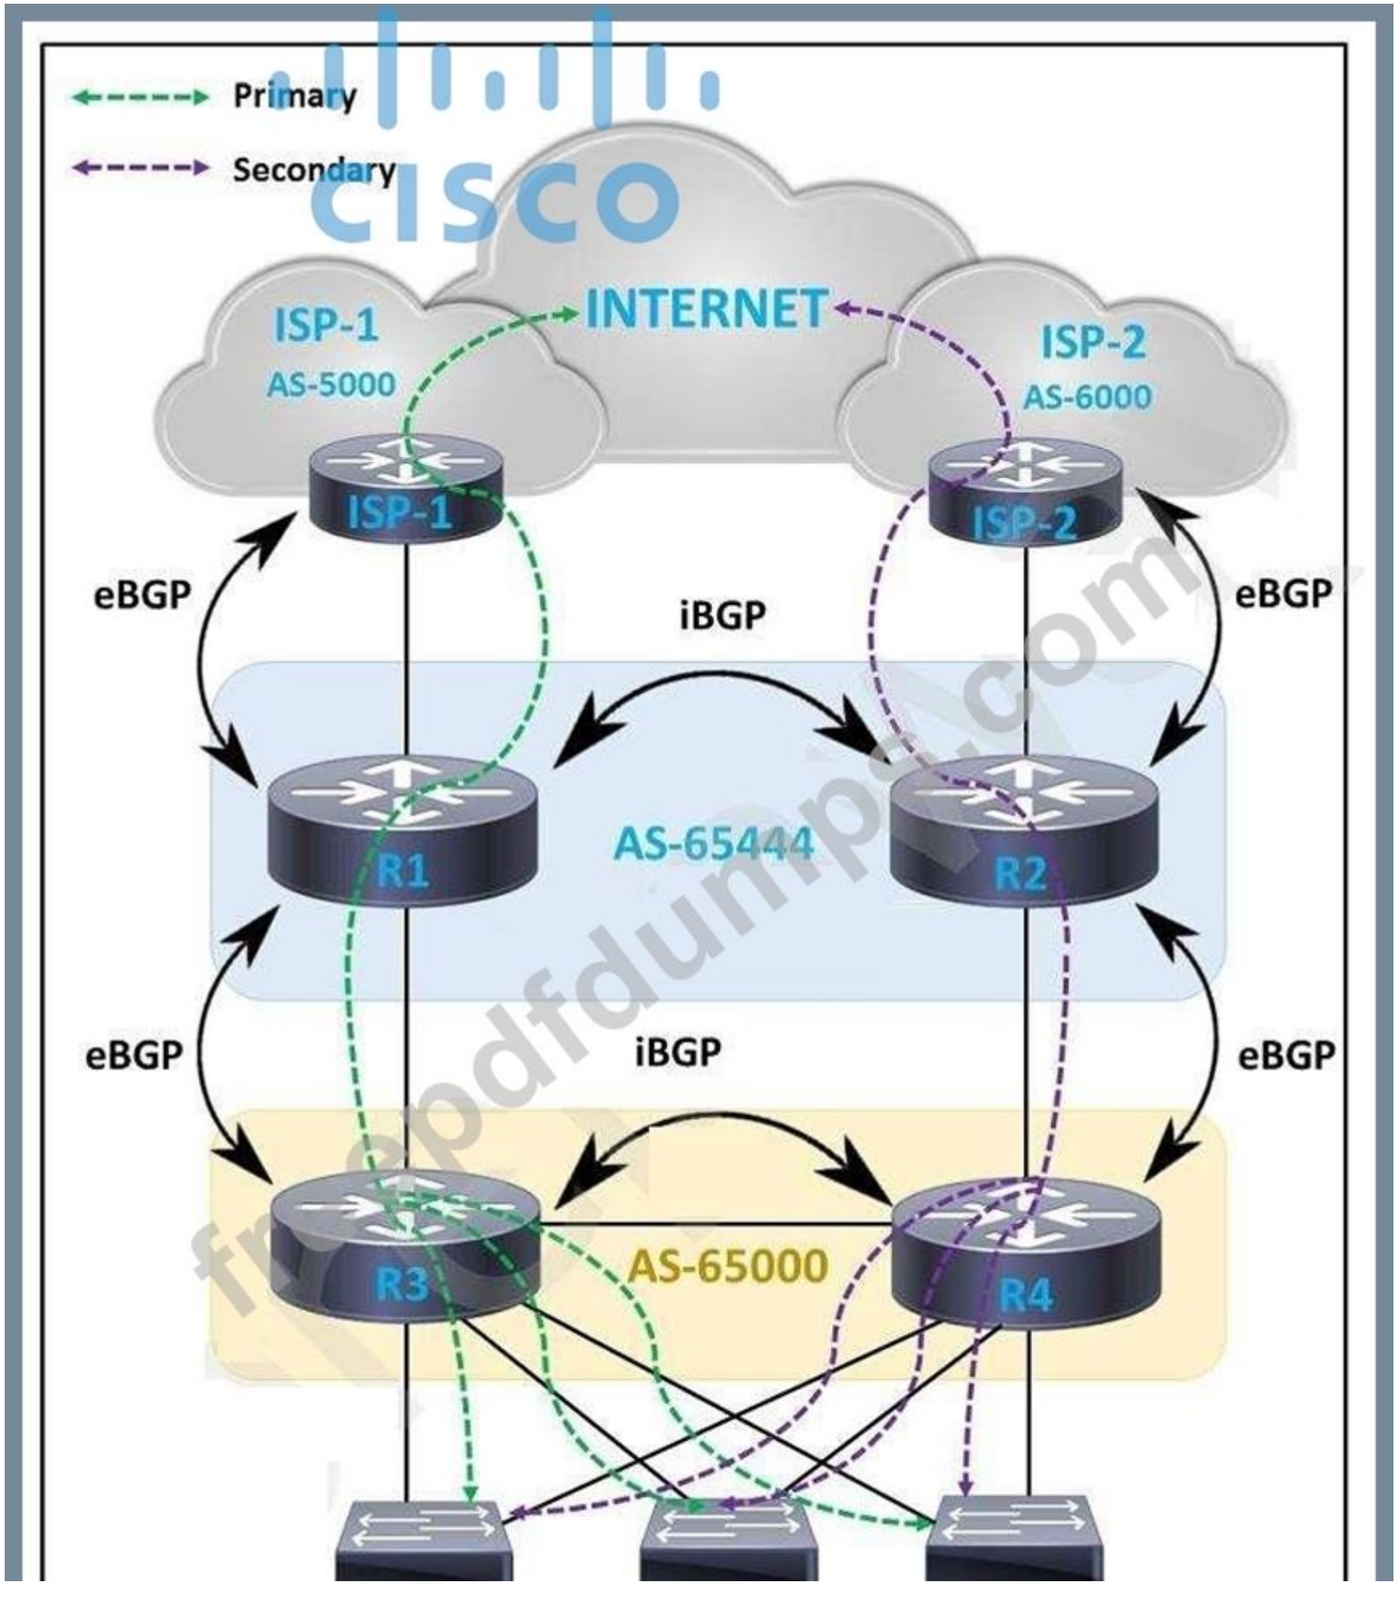
The drag-and-drop mapping separates WAN connectivity technologies by the service characteristics they deliver. Cisco enterprise WAN designs distinguish private WAN options, such as MPLS Layer 3 VPN, Ethernet WAN, VPLS, and VPWS, from Internet-based options such as broadband, LTE, and VPN overlays.

Layer 3 MPLS VPN gives the provider responsibility for routing between customer sites and is commonly used when private any-to-any routing, QoS treatment, and segmentation are required. Layer 2 services such as VPWS and VPLS extend customer Layer 2 adjacency across a provider network; VPWS is point-to-point, while VPLS is multipoint. Internet and cellular links are typically lower-cost and fast to deploy, but they require overlay security, such as IPsec, DMVPN, or SD-WAN encryption, when used for enterprise private traffic. The selected mapping therefore aligns each description with the WAN category that actually provides that behavior. In design work, the architect must compare latency, resiliency, SLA, QoS preservation, encryption requirements, routing control, and operational ownership before selecting the WAN type.

Reference topics: enterprise WAN connectivity, MPLS VPN, VPLS, VPWS, Internet VPN, cellular backup.

NEW QUESTION: 66

Refer to the exhibit.


Primary
Secondary
INTERNET
ISP-1
AS-5000
ISP-2
AS-6000
ISP-1
ISP-2
eBGP
iBGP
eBGP
R1
R2
AS-65444
eBGP
iBGP
eBGP
R3
R4
AS-65000



An engineer must design a WAN solution so that ISP-1 is always preferred over ISP-2. The path via ISP-2 is considered as a backup and must be used only when the path to ISP-1 is down. Which solution must the engineer choose?

A. R1:
- Routes advertised to ISP-1: 0x AS-path prepend
- Routes received from ISP-1: HIGH local-preference
- Routes advertised to R2: no action
- Routes received from R2: community NO-EXPORT

R2:
- Routes advertised to ISP-2:5x AS-path prepend
- Routes received from ISP-2: LOW local-preference
- Routes advertised to R1: community NO-ADVERTISE
- Routes received from R1: no action

B. R1:
- Routes advertised to ISP-1: 0x AS-path prepend
- Routes received from ISP-1: HIGH local-preference
- Routes advertised to R2: community NO-EXPORT
- Routes received from R2: no action

R2:
- Routes advertised to ISP-2: 5x AS-path prepend
- Routes received from ISP-2: LOW local-preference
- Routes advertised to R1: no action
- Routes received from R1: no action

C. R1:
- Routes advertised to ISP-1: 0x AS-path prepend
- Routes received from ISP-1: LOW local-preference
- Routes advertised to R2: community NO-ADVERTISE
- Routes received from R2: no action

R2:
- Routes advertised to ISP-2: 5x AS-path prepend
- Routes received from ISP-2: HIGH local-preference
- Routes advertised to R1: no action
- Routes received from R1: community NO-ADVERTISE

D. R1:
- Routes advertised to ISP-1: 5x AS-path prepend

- Routes received from ISP-1: LOW local-preference
- Routes advertised to R2: community NO-ADVERTISE
- Routes received from R2: no action

R2:

- Routes advertised to ISP-2: 0x AS-path prepend
- Routes received from ISP-2: HIGH local-preference
- Routes advertised to R1: community NO-EXPORT
- Routes received from R1: no action

Answer: B (LEAVE A REPLY)

The best design is the option that prefers ISP-1 for both outbound and inbound traffic while keeping ISP-2 available as a backup. For outbound traffic from the enterprise, local preference is the correct BGP attribute because it is evaluated inside the local AS and higher local preference wins. R1 should assign a high local preference to routes learned from ISP-1, while R2 should assign a low local preference to routes learned from ISP-2. For inbound traffic from the Internet toward the enterprise, AS-path prepending is the practical signal sent to external networks. Advertising routes to ISP-1 with no prepending and routes to ISP-2 with five prepends makes the ISP-1 path look shorter and therefore more attractive. The R1-to-R2 internal advertisements should not be blocked by NO-ADVERTISE because each edge router must still share reachability internally for failover. Option B matches these controls without unnecessarily suppressing routes between the edge routers. Reference topics: BGP local preference, AS-path prepending, active/backup Internet design, inbound and outbound path control.

NEW QUESTION: 67

An architect must design a QoS model for a business-critical application that Is delay-sensitive and requires high bandwidth. The company ' s head office hosts the application, and DMVPN tunnels protected with IPsec provide connectivity between the head office and branches. Which solution must the architect choose?

A. RSVP
B. IntServ
C. WRED
D. DiffServ

Answer: (SHOW ANSWER)

DiffServ is the correct QoS model for a business-critical, delay-sensitive, high-bandwidth application across DMVPN tunnels protected with IPsec. Differentiated Services classifies and marks traffic, usually with DSCP, and then applies per-hop behavior at each network device. This model scales well across enterprise WANs because routers do not need to maintain per-flow reservation state. DMVPN and IPsec overlays can carry marked traffic, and the WAN design can apply queuing, shaping, and congestion avoidance based on those markings. IntServ with RSVP can provide explicit resource reservation, but it is stateful and does not scale well across many branch tunnels. RSVP also becomes operationally difficult across encrypted overlay environments. WRED is only a congestion avoidance mechanism, not a complete QoS architecture. Because the application needs differentiated treatment but the WAN must remain scalable, DiffServ is the practical Cisco enterprise design model. Reference topics: DiffServ, DSCP marking, DMVPN QoS, IPsec QoS pre- classify, per-hop behavior, enterprise WAN QoS.

NEW QUESTION: 68

A network engineer must design a multicast solution based on these requirements:

* interactive communication
* must not use source trees
* users must register
* 100 multicast sources

Which solution must the company choose?

A. MSDP
B. PIM-DM
C. any-source multicast

D. BIDIR PIM

Answer: D (LEAVE A REPLY)

BIDIR-PIM is the correct multicast solution for interactive many-to-many communication with many sources when source trees must not be used. Bidirectional PIM is specifically designed for applications where receivers can also be senders and where the number of sources can grow without creating per-source multicast state throughout the network. Instead of building shortest-path trees for each source, BIDIR-PIM uses a shared tree rooted at the rendezvous point and forwards traffic bidirectionally along that tree. This reduces multicast routing-table growth and control-plane churn in environments with many active sources. PIM dense mode is not appropriate because it floods and prunes traffic and is inefficient at scale. Any-source multicast with PIM-SM may create source trees, which violates the requirement. MSDP is used to exchange source information between PIM-SM domains and is not itself the forwarding mode for this campus or WAN multicast design. Reference topics: BIDIR-PIM, many-to-many multicast, shared tree, multicast source scaling, PIM design.

NEW QUESTION: 69

What is the purpose of a control plane node in a Cisco SD-Access network fabric?

A. to maintain the endpoint database and mapping between endpoints and edge nodes

B. to detect endpoints in the fabric and inform the host tracking database of EID-to-fabric-edge node bindings

C. to identify and authenticate endpoints within the network fabric

D. to act as the network gateway between the network fabric and outside networks

Answer: A (LEAVE A REPLY)

The SD-Access control-plane node maintains the endpoint database and resolves mappings between endpoints and the fabric edge nodes where those endpoints are located. Cisco SD-Access uses LISP in the fabric control plane. Edge nodes register endpoint identifiers with the control-plane node, and other fabric nodes query the control plane when they need to locate a destination endpoint. This control-plane function is commonly described through LISP Map-Server and Map-Resolver behavior. Option B is not correct because endpoint detection occurs at the edge node, not at the control-plane node. The edge is the device connected to users, access points, or endpoint-facing networks, so it learns local endpoint presence and registers that information.

Option C refers to authentication and identity, which is integrated with Cisco ISE. Option D describes the border-node role, because border nodes connect the fabric to external networks. The corrected answer is therefore A. A stable SD-Access design must provide redundant control-plane reachability and scale the control-plane role according to fabric size, mobility, and endpoint churn. Reference topics: SD-Access control plane, LISP Map-Server, LISP Map-Resolver, endpoint-to-location mapping.

NEW QUESTION: 70

A company requires a private WAN design that allows remote sites to connect to HQ. The design must ensure that:

* traffic is always encrypted
* forwarding overhead is reduced
* management of security Is centralized
* multicast traffic is supported

Which technology must the company select?

A. iPiac P2P

B. GET VPN

C. DMVPN Phase 3

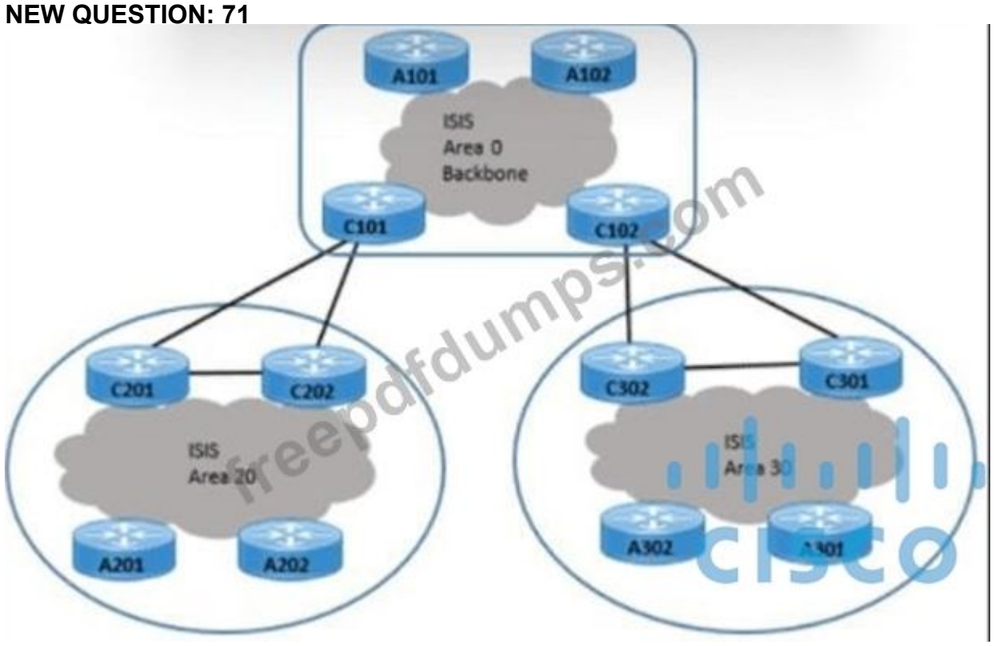
D. mGRE

Answer: (SHOW ANSWER)

GET VPN is the correct private WAN encryption technology when traffic must always be encrypted, multicast must be supported, forwarding overhead must be reduced, and security management must be centralized. Cisco Group Encrypted Transport VPN protects traffic over a private WAN without building traditional point-to-point or multipoint tunnel overlays. Because it preserves the original IP header, the provider or enterprise routing core can still make normal forwarding decisions, and multicast, QoS, and routing behavior are easier to preserve than with heavy tunnel overlays. GET VPN uses a key server model to centralize group policy and distribute encryption keys to group members. This reduces the operational burden of maintaining many permanent tunnels. DMVPN supports multicast through GRE and provides dynamic tunnels, but it adds tunnel encapsulation overhead and does not meet the reduced-overhead private WAN requirement as cleanly. Plain mGRE lacks the full centralized group encryption

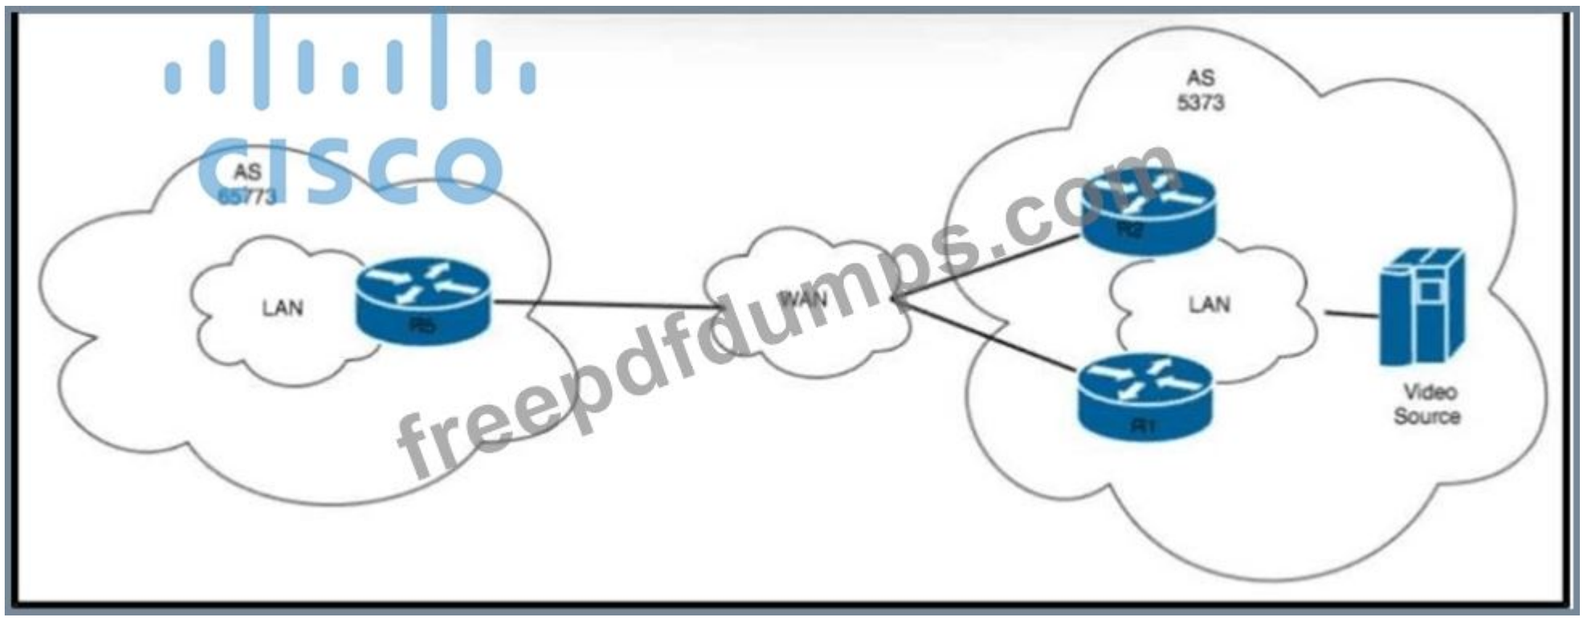
model. Reference topics: GET VPN, group encryption, private WAN security, multicast support, key server, tunnel-less IPsec.

NEW QUESTION: 71



Refer to the exhibit. An architect is designing a hierarchical ISIS solution for a customer with these requirements:
* Routers will double In all areas within the next 24 months.
* Link flaps within areas 20 and 30 must not impact the backbone area.
* Traffic originating from A201 and A302 routers must connect to application servers in the backbone.
Which design must the architect select?
A. C201 Level 1/2, A301 Level 1/2 and A102 Level 1/2
B. C101 Level 1/2. A201 Level 1, and A101 Level 2
C. C102 Level 2. A202 Level 2, and A102 Level 1
D. C302 Level 2. A302 Level 1/2. and A101 Level 2
Answer: A (LEAVE A REPLY)
The selected IS-IS design must place the correct routers as Level 1/Level 2 boundary routers so that area failures stay local while traffic can still reach backbone application services. In a hierarchical IS-IS design, Level 1 routers exchange detailed topology only inside their own area. Level 2 routers form the backbone used for interarea connectivity. L1/L2 routers connect the local area to the backbone and advertise reachability between the two levels. Areas 20 and 30 are expected to double in size, and link flaps inside those areas must not impact the backbone. Therefore, the design should keep internal area routers at Level 1 and use only the boundary devices that connect those areas to the backbone as L1/L2. Option A matches that principle by selecting the appropriate boundary routers for interlevel connectivity while avoiding unnecessary L1/L2 adjacencies throughout the topology. The other options either place the wrong routers in Level 2 only or fail to provide a clean transition between local areas and the backbone. Reference topics: IS-IS hierarchy, L1/L2 boundary design, area containment, backbone stability, route leaking.

NEW QUESTION: 72



freepdfdumps.com

Refer to the exhibit. A company specializing in VoD content creation has two offices in a separate multicast domain connected by a WAN link. BGP communication has been established between the offices. Clients are inside the LAN in each office. In AS5373. R2 has been selected as RP. What must the network architect design to deliver VoD content to clients in AS65773?

A. MSDP

B. PIM ASM with Auto-RP

C. PIM SSM

D. PIM ASM with BSR

Answer: A (LEAVE A REPLY)

MSDP is required because the design connects separate PIM-SM multicast domains and needs multicast source discovery across domain boundaries. In each autonomous system, receivers join a multicast group through the local rendezvous point. Without MSDP, an RP in one domain does not automatically learn about active sources registered to an RP in another domain. MSDP allows RPs to exchange Source-Active information so that receivers in AS65773 can learn about VoD sources in AS5373 and build the required multicast forwarding state. Auto-RP and BSR are mechanisms for discovering or advertising rendezvous point information inside a PIM domain; they do not solve interdomain source discovery between separate multicast domains. PIM SSM can work for source-specific applications, but the question states that R2 has been selected as RP, which indicates an ASM/PIM-SM design. With BGP already established between offices, MSDP complements the interdomain routing design by exchanging multicast source reachability information.

Reference topics: MSDP, PIM-SM, interdomain multicast, Source-Active messages, rendezvous point resilience.

NEW QUESTION: 73

An engineer must use YANG with an XML representation to configure a Cisco IOS XE switch with these specifications:

* IP address 10.10.10.10/27 configured on the interface GigabitEthernet2/1/0

* connectivity from a directly connected host 10.10.10.1/27

Which YANG data model set must the engineer choose?

A.

```
<interfaces xmlns="urn:ietf:params:xml:ns:yang:ietf-interfaces">
 <interface>
  <name>GigabitEthernet2/1/0</name>
  <type xmlns:ianaift="urn:ietf:params:xml:ns:yang:iana-if-type">ianaift:ethernetCsmacd</type>
  <enabled>false</enabled>
  <ipv4 xmlns="urn:ietf:params:xml:ns:yang:ietf-ip">
   <address>
    <ip>10.10.10.10</ip>
    <netmask>255.255.255.224</netmask>
   </address>
  </ipv4>
 </interface>
</interfaces>
```

B.

```
<interfaces YANG="urn:ietf:params:xml:ns:yang:ietf-interfaces">
 <interface>
  <name>GigabitEthernet2/1/0</name>
  <type YANG:ianaift="urn:ietf:params:xml:ns:yang:iana-if-type">ianaift:ethernetCsmacd</type>
  <enabled>true</enabled>
  <ipv4 YANG="urn:ietf:params:xml:ns:yang:ietf-ip">
   <address>
    <ip>10.10.10.10</ip>
    <netmask>255.255.255.224</netmask>
   <address>
  </ipv4>
 </interface>
</interfaces>
```

C.

```
<interfaces json="urn:ietf:params:json:ns:yang:ietf-interfaces">
 <interface>
  <name>GigabitEthermet2/1/0</name>
  <type json:ianaift="urn:ietf:params:json:ns:yang:iana-if-type">ianaift:ethernetCsmacd</type>
  <enabled>true</enabled>
  <ipv4 json="urn:ietf:params:json:ns:yang:ietf-ip">
   <address>
    <ip>10.10.10.10</ip>
    <netmask>255.255.255.224</netmask>
   </address>
  </ipv4>
 </interface>
```

```
<interfaces xmlns="urn:ietf:params:xml:ns:yang:ietf-interfaces">
  <interface>
    <name>GigabitEthernet2/1/0</name>
    <type xmlns:ianaift="urn:ietf:params:xml:ns:yang:iana-if-type">ianaift:ethernetCsmacd</type>
    <enabled>true</enabled>
    <ipv4 xmlns="urn:ietf:params:xml:ns:yang:ietf-ip">
      <address>
        <ip>10.10.10.10</ip>
        <netmask>255.255.255.224</netmask>
      </address>
    </ipv4>
  </interface>
</interfaces>
```

freepdfdumps.com

D.

Answer: D (LEAVE A REPLY)

The correct model set is the one that configures the IOS XE interface GigabitEthernet2/1/0 with IPv4 address 10.10.10.10 and prefix length 27 using valid XML according to the selected YANG schema. The directly connected host address 10.10.10.1/27 confirms the subnet relationship; both addresses belong to 10.10.10.0 /27, so the interface configuration must use that prefix length rather than a host mask or an unrelated network.

YANG-based configuration is strict: the namespace, container hierarchy, interface type, interface name, and address fields must match the device model. If an option places GigabitEthernet under the wrong container, represents the mask incorrectly, or uses a schema from a different model family, the device will reject the payload or configure the wrong object. Option D correctly represents the IOS XE XML structure and the required addressing. In real operations, the engineer should follow this with a get-config readback against the running datastore to prove that the structured configuration committed as intended. Reference topics: IOS XE interface YANG, XML payload structure, NETCONF configuration, IPv4 subnet validation.

NEW QUESTION: 74

What is the purpose of a border node in a Cisco SD-Access fabric?

A. connect devices to a network

B. perform traffic encapsulation and de-encapsulation

C. perform network virtualization

D. expand a network

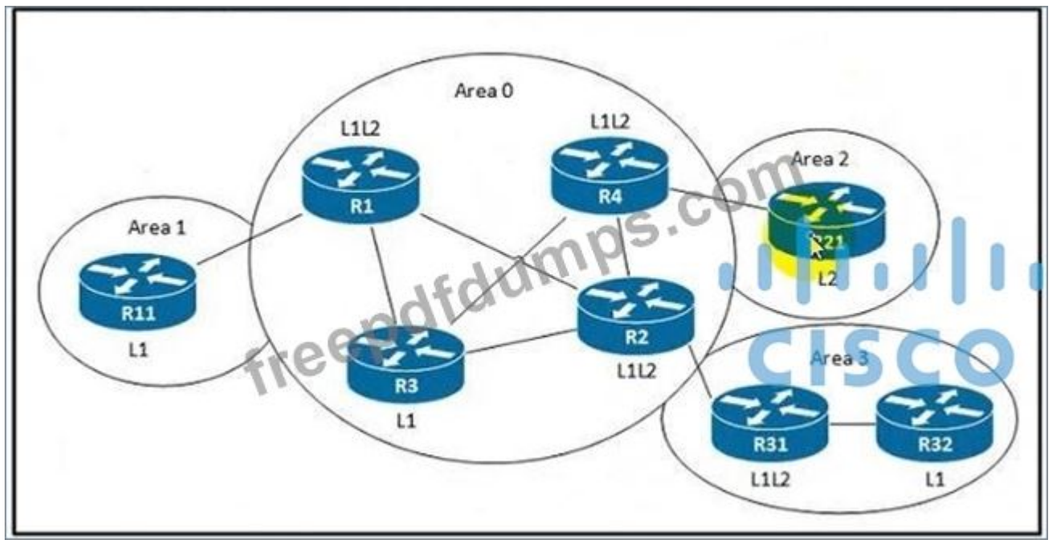
Answer: A (LEAVE A REPLY)

The border node provides connectivity between the SD-Access fabric and networks outside the fabric. In Cisco SD-Access, edge nodes connect endpoints to the fabric, intermediate nodes transport routed IP traffic through the underlay, control-plane nodes maintain endpoint location mappings, and border nodes connect the fabric to external Layer 3 domains such as the data center, WAN, internet edge, or shared services network. A border node may perform proxy tunnel router functions and can de-encapsulate fabric traffic that is destined outside the overlay, but its design purpose is external connectivity for the fabric site. Network virtualization is primarily delivered through virtual networks, VRF separation, LISP control-plane mappings, and VXLAN encapsulation. Expanding the network is too vague and does not identify a Cisco SD-Access role. Therefore, among the available options, the border node is the component used to connect the fabric to the required external network resources. Reference topics: Cisco SD-Access border node, external connectivity, fabric roles, LISP proxy tunnel router, VN-to-VRF mapping.

NEW QUESTION: 75
What is the purpose of a TLOC extension in a Cisco SD-WAN network fabric?
A. to facilitate WAN Edge router redundancy within a site
B. to identify the physical interface where a WAN Edge router connects to the WAN transport network
C. to expand the number of colors that are potentially applied to a network transport interface
D. to aggregate multiple physical interfaces into a single logical Interface
Answer: (SHOW ANSWER)
A TLOC extension facilitates WAN Edge router redundancy within a site. In Cisco SD-WAN, a Transport Location identifies a WAN transport attachment using attributes such as system IP, color, and encapsulation.
A TLOC extension allows one WAN Edge router to use a transport circuit physically connected to another WAN Edge router at the same site. This is valuable when a branch has two edge routers and two transports, but each transport is cabled to only one router. With TLOC extension, each router can still gain logical access to both transports, improving redundancy and allowing better path diversity without adding more service- provider circuits or handoffs. The feature does not simply identify the physical interface; that is part of TLOC definition. It does not increase the number of colors or create a port-channel-style aggregate. A sound design also accounts for failure domains, service-side routing, and whether the inter-router link has enough bandwidth for extended transport traffic. Reference topics: Cisco SD-WAN TLOC, TLOC extension, WAN Edge redundancy, transport color, branch resiliency.

NEW QUESTION: 76
Which nonproprietary mechanism can be used to automate rendezvous point distribution in a large PIM domain?
A. Embedded RP
B. BSR
C. Auto-RP
D. Static RP
Answer: B (LEAVE A REPLY)
The standards-based mechanism for dynamic rendezvous point distribution in a large PIM domain is Bootstrap Router. BSR allows candidate RPs and candidate BSRs to participate in a standards-based election and advertisement process so PIM routers can learn the RP set dynamically. This avoids configuring every router manually with static RP information and avoids depending on a Cisco-proprietary mechanism. Auto- RP can also distribute RP information dynamically, but Auto-RP is Cisco proprietary and therefore does not meet the nonproprietary requirement. Embedded RP is associated with IPv6 multicast group addressing and is not the general-purpose answer for automating RP discovery in a large IPv4 PIM domain. Static RP configuration is simple and deterministic, but it does not automate RP distribution and becomes operationally fragile as the multicast domain grows. In Cisco multicast design, BSR is the correct choice when the network requires dynamic RP discovery, vendor-neutral behavior, and scalable PIM-SM operation across a large domain.

Valid 300-420 Dumps shared by Actual4test.com for Helping Passing 300-420 Exam! Actual4test.com now offer the **newest 300-420 exam dumps**, the Actual4test.com 300-420 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-420 dumps with Test Engine here: https://www.actual4test.com/300-420_examcollection.html (**457** Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)



Refer to the exhibit. A customer experienced an unexpected network outage when the link between R1 and R2 went down. An architect must design a solution to ensure network continuity in the event the link fails again.

Which solution should the design include?

A. Make R31 an L1 router.

B. Make R3 an L1L2 router

C. Make Area 0 L2-only

D. Make R11 an L2 router.

Answer: B (LEAVE A REPLY)

Making R3 an L1/L2 router restores the required IS-IS hierarchy and prevents the outage caused by the failed R1-to-R2 link. In IS-IS, Level 1 routers have detailed topology only inside their own area, while Level 2 routers provide the interarea backbone. Routers that connect a Level 1 area to the Level 2 backbone must operate as Level 1/Level 2 devices so they can participate in both databases and provide an exit path for area traffic. If a path between areas depends on a router that is only Level 1 or only Level 2 in the wrong location, traffic can be blackholed when a key link fails. Configuring R3 as L1/L2 gives the topology an additional valid transition point between the local area and the backbone, improving continuity after the R1-to-R2 failure. Making unrelated routers L1 or L2 does not necessarily create the required interarea attachment.

Reference topics: IS-IS Level 1, Level 2 backbone, L1/L2 routers, interarea reachability, hierarchical IGP design.

NEW QUESTION: 78

Which queuing structure is used on SD-WAN Edge routers?

A. FIFO

B. LLQ+WFQ

C. 1P-4Q-2T

D. Priority

Answer: B (LEAVE A REPLY)

Cisco SD-WAN edge routers use a queuing model based on LLQ and WFQ concepts. Low Latency Queuing provides strict priority treatment for delay-sensitive traffic such as voice and real-time media, while Weighted Fair Queuing gives proportional service to other traffic classes based on bandwidth allocation. This combination aligns with SD-WAN requirements where business-critical and real-time traffic must be protected without starving normal data applications. FIFO is not appropriate for a multi-class WAN design because it does not distinguish application priority during congestion. A pure priority queue would risk starving lower-priority traffic if it were not controlled by policing or bandwidth limits. Hardware-specific campus queuing models such as 1P-4Q-2T describe certain switch egress queue structures, not the generic SD- WAN edge queuing architecture. In a professional SD-WAN QoS design, traffic is classified, marked, mapped to forwarding classes, and then scheduled using priority and

weighted queues according to business intent. Reference topics: Cisco SD-WAN QoS, LLQ, WFQ, forwarding classes, application-aware WAN policy.

NEW QUESTION: 79
A network engineer must design an MSDP multicast solution to provide RP resilience in a network with two separate domains. Also, multicast sources and receivers must register with the local RP. Which solution must the engineer choose?
A. Configure the RP has value to 0, and traffic will route to the closest RP
B. Configure the RP loopback interface with the same IP address/32, and traffic will route to the closest RP
C. Configure the RP group ranges to split the multicast traffic, and traffic will route to the longest match
D. Configure the RP priority with the same value, and traffic will route to the closest RP
Answer: B (LEAVE A REPLY)
The correct design is to use the same anycast RP address on loopback interfaces so receivers and sources register with their nearest available rendezvous point. Cisco Anycast RP designs use a common RP address advertised from multiple RP routers, allowing unicast routing to direct PIM joins and registers to the closest RP. MSDP is then used between the RP routers so each RP can learn about active sources registered to the other RP. This provides resilience and local registration while keeping RP reachability consistent for the multicast domain. Option B captures the key behavior: the RP loopback is configured with the same /32 address, and normal routing sends traffic to the closest RP. Configuring a hash value, group range splitting, or equal priority alone does not provide the Anycast RP behavior required here. A professional design should also use unique loopbacks for MSDP peer sessions while keeping the shared anycast loopback as the RP address. Reference topics: Anycast RP, MSDP, PIM-SM rendezvous points, source registration, multicast high availability.

NEW QUESTION: 80
Refer to the exhibit. A network engineer is designing an OSPF solution to connect a company ' s remote to a newly provisioned MPLS VPN backbone. Some of the branches have a direct dark fiber connection between each other. The engineer wants to ensure that the dark fibers are used only when the MPLS core is unavailable. Which solution must the engineer choose?
A. Stub area
B. Sham link
C. Virtual link
D. NSSA
Answer: (SHOW ANSWER)
An OSPF sham link is the correct MPLS VPN design when backdoor links exist between customer sites and should be used only if the MPLS core is unavailable. In MPLS Layer 3 VPN designs that run OSPF between CE and PE routers, the MPLS VPN path can appear as an interarea route, while a direct dark-fiber backdoor link may appear as an intra-area route. OSPF prefers intra-area routes over interarea routes, so traffic could incorrectly prefer the dark-fiber path during normal operation. Cisco sham links solve this by creating an intra- area logical link across the MPLS VPN backbone between PE routers, making the MPLS path competitive or preferred according to OSPF cost. A virtual link connects disconnected OSPF area 0 sections and is not the right MPLS VPN tool. Stub and NSSA area types control external LSA behavior, not MPLS-backdoor path preference. The sham link preserves the desired routing hierarchy while keeping the direct dark fiber as backup. Reference topics: OSPF sham link, MPLS L3VPN, PE-CE routing, backdoor links, route preference.

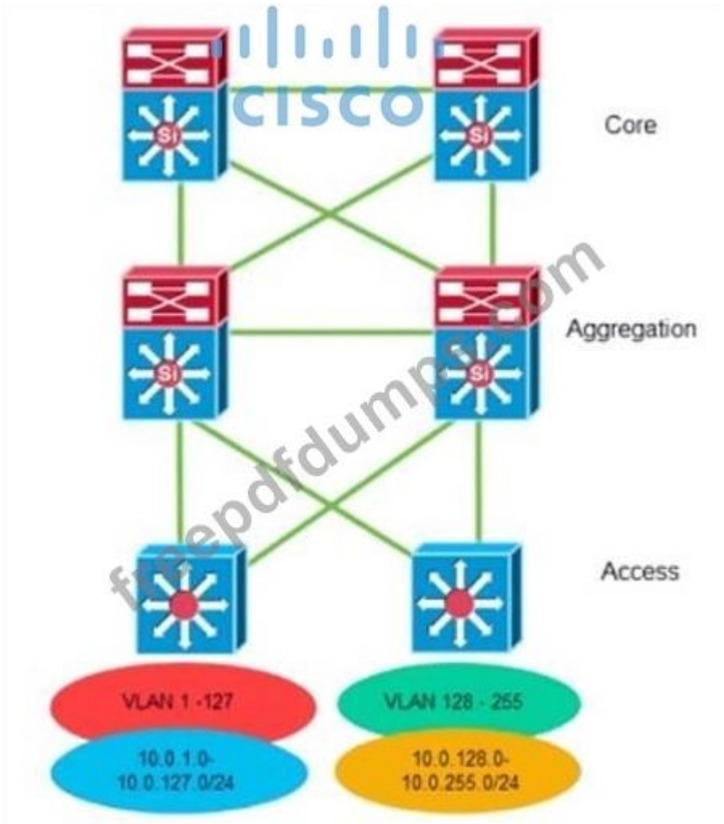
NEW QUESTION: 81
An engineer is designing a BGP solution supporting a VXLAN environment over a Layer 3 IPv4 network fabric with these requirements
* provide Layer 2 adjacency
* allow VM migration of workloads between sites
* IGP is OSPF
Which BGP address family must the engineer choose?
A. VPNv4
B. IPv4 unicast
C. L2VPN VPLS-VPWS

D. L2VPNEVPN
Answer: (SHOW ANSWER)
The required BGP address family is L2VPN EVPN. The design calls for VXLAN over a Layer 3 IPv4 fabric while preserving Layer 2 adjacency and allowing workload mobility between sites. EVPN is the BGP control plane used with VXLAN overlays to advertise MAC reachability, IP host information, and Ethernet segment information in a scalable way. It separates the Layer 2 and Layer 3 overlay control plane from the underlay IGP, which in this case is OSPF. VPNv4 is used for MPLS Layer 3 VPN routes, not VXLAN Layer 2 adjacency. IPv4 unicast advertises ordinary IP prefixes and does not carry MAC mobility or Ethernet VPN information. L2VPN VPLS or VPWS address families are associated with MPLS-based Layer 2 VPN services and do not provide the VXLAN EVPN control-plane model required here. EVPN is therefore the correct family for VXLAN fabrics that need host mobility and scalable Layer 2 extension. Reference topics:
BGP EVPN, VXLAN, L2VPN EVPN address family, MAC reachability, workload mobility.

NEW QUESTION: 82
A network engineer must design a multicast solution to prevent the spoofing of multicast streams and ensure efficient bandwidth utilization. The network will be merged with another multicast domain in the future, and the merge must require minimum effort. Which two solutions meet the customer requirements? (Choose two.)
A. PIM-SSM
B. IGMPv3
C. IGMPv2
D. PIM-SM
E. MSDP
Answer: (SHOW ANSWER)
PIM-SSM with IGMPv3 is the correct multicast design. Source-Specific Multicast requires receivers to identify both the multicast group and the permitted source, normally through IGMPv3 membership reports.
That source-specific join model prevents unauthorized or spoofed sources from injecting traffic into the group because receivers do not simply accept traffic from any source for a group address. It also improves bandwidth efficiency because the network builds forwarding state only for the requested source and group pair, without requiring shared trees or rendezvous point discovery. PIM-SM with MSDP can interconnect multicast domains, but it is built around Any-Source Multicast behavior and does not provide the same source- specific anti-spoofing property. IGMPv2 cannot signal source-specific joins, which is why IGMPv3 is required for SSM. Since the network may merge with another multicast domain later, SSM also reduces integration complexity by avoiding RP and MSDP dependencies. Therefore, PIM-SSM and IGMPv3 meet both the security and efficiency requirements. Reference topics: PIM-SSM, IGMPv3, source-specific joins, multicast security, bandwidth efficiency.

NEW QUESTION: 83
Which two functions does the control plane node provide in a Cisco SD-Access architecture? (Choose two.)
A. LISP proxy ETR
B. host tracking database
C. policy mapping
D. map server
E. endpoint registration
Answer: B,D (LEAVE A REPLY)
The Cisco SD-Access control-plane node provides the host tracking database and map-server function. SD- Access uses LISP-based control-plane behavior to separate endpoint identity from endpoint location. Fabric edge nodes register endpoints with the control-plane node, and the control-plane node maintains the endpoint- to-location database. When a fabric node needs to forward traffic to an endpoint, it queries the control-plane node so the destination endpoint can be resolved to the correct fabric location. In LISP terminology, this maps to the map-server and map-resolver role, while the SD-Access design discussion often refers to the endpoint database or host tracking database. The control-plane node does not enforce policy mapping by itself; policy is integrated with Cisco ISE and applied through scalable group policy in the fabric. It also does not function as a proxy ETR in this option set. Endpoint registration is initiated by the edge nodes, which detect local endpoints and inform the control plane. Therefore, host tracking database and map server are the correct functions. Reference topics: SD-Access control plane, LISP map server, host tracking database, EID-to- RLOC mapping.

NEW QUESTION: 84
Which information update is carried by OMP and enables the Cisco SD-WAN to build a secure overlay fabric on top of any public or private transport without regard for the actual link IP?
A. TLOC
B. RLOC
C. LISP PITR
D. DTLS
Answer: (SHOW ANSWER)
TLOC information is the OMP update that allows Cisco SD-WAN to build an overlay over any public or private transport without depending on the actual underlay link addressing. Cisco OMP advertises route reachability along with Transport Location mappings. A TLOC represents where and how a WAN Edge can be reached in the transport network, using attributes such as system IP, transport color, encapsulation, and next-hop information. Remote WAN Edge routers use this information to form secure data-plane tunnels to the correct transport attachment. RLOC is a LISP locator term associated with SD-Access and LISP designs, not the SD-WAN OMP construct. DTLS is the secure transport used for control connections, but it is not an OMP route information type. LISP PITR is unrelated to Cisco SD-WAN OMP. The secure overlay is built by combining OMP routes, TLOCs, policy, and encrypted tunnel establishment. Reference topics: Cisco SD- WAN OMP, TLOC routes, transport color, overlay routing, secure data-plane tunnels.

NEW QUESTION: 85


freepdfdumps.com
Refer to the exhibit. An architect is designing a Layer 3 campus network The design must hide network instability, reduce network overhead, and conserve critical device memory Which route summarization solution must the architect select?
A. * The VLAN subnets must be summarized into 10.0.0.0 ' 16 at the aggregation layer and advertised to the core layer
B. * The core layer must advertise a default route toward the aggregation layer

* The VLAN subnets must be summarized into 10 0 0 0/16 at the access layer and advertised to the aggregation layer

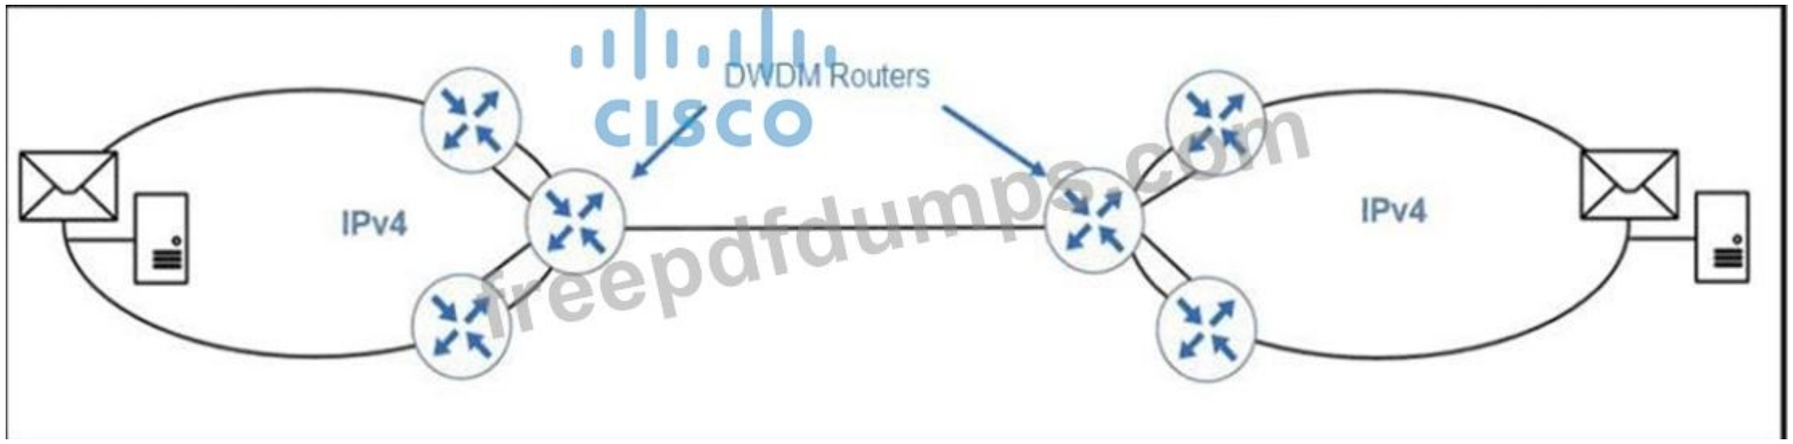
C. * The aggregation layer must advertise a default route toward the access layer.
* The VLAN subnets must be summarized into 10 0.0 0/16 at the aggregation layer and advertised to the core layer

D. * The aggregation layer must advertise a default route toward the core layer
* The VLAN subnets must be summarized into 10 0 0 0.116 at the aggregation layer and advertised to the access layer

Answer: (SHOW ANSWER)

The best summarization design is for the aggregation layer to advertise a default route toward the access layer and to summarize the VLAN subnets toward the core. Cisco hierarchical campus design uses summarization at natural layer boundaries so that access-layer instability is hidden from the rest of the network. If individual access VLANs flap, the core should not have to process every specific route change. Advertising a summary such as 10.0.0.0/16 from the aggregation layer toward the core reduces routing-table size, lowers CPU and memory use, and confines convergence events to the lower layer. At the same time, the access layer normally does not need full knowledge of the enterprise routing table; a default route from the aggregation layer is sufficient for northbound reachability. Summarizing at the access layer is normally less effective because the aggregation layer is the first stable boundary that sees multiple access networks. Advertising a default from the core toward aggregation would hide information in the wrong direction. Reference topics: hierarchical routing, route summarization, campus aggregation design, convergence containment, routing-table scaling.

NEW QUESTION: 86



Refer to the exhibit. An engineer is planning an IPv4 to IPv6 migration solution for a customer. The routers in the network can support IPv4 and IPv6, except for the DWDM routers. The DWDM routers provide a Layer 2 link in which the routers peer directly with each other across a DWDM circuit. The circuit also provides connectivity between the mail servers. Which IPv6 migration technique must the engineer deploy?

A. dual-stack

B. 6to4

C. ISATAP

D. 6rd

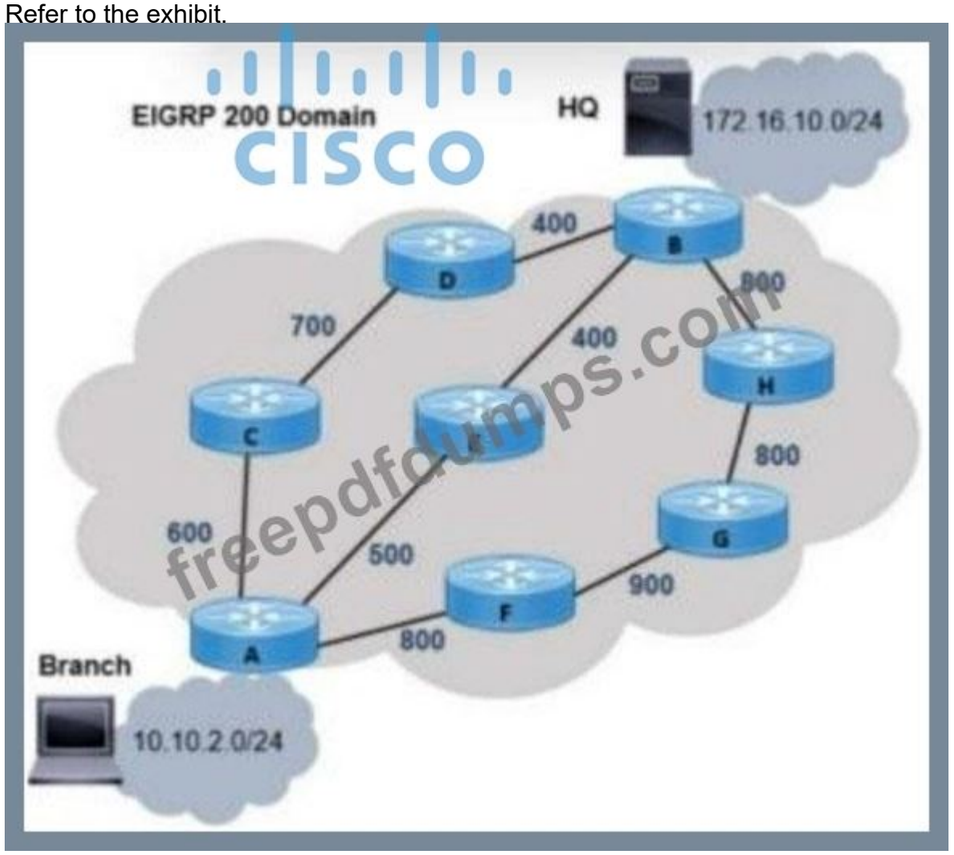
Answer: A (LEAVE A REPLY)

Dual-stack is the correct IPv6 migration technique because the DWDM devices provide Layer 2 transport between routers that already support both IPv4 and IPv6. Layer 2 transport is generally transparent to IPv6 because it forwards Ethernet frames rather than participating in Layer 3 routing. The routers at the ends of the DWDM circuit can peer directly over the Layer 2 service and run IPv6 in parallel with IPv4, while mail servers can also retain IPv4 reachability during the migration. Cisco dual-stack migration guidance uses simultaneous IPv4 and IPv6 operation to support phased application and host transition with minimal tunneling overhead. 6to4, ISATAP, and 6rd are tunneling or transition mechanisms used when IPv6 islands must cross IPv4-only routed infrastructure. In this case, the limiting devices are in the transport layer and do not need to route IPv6. Deploying dual-stack on the capable routers and servers is simpler, more transparent, and avoids unnecessary encapsulation. Reference topics: IPv6 dual-stack migration, Layer 2 transparency, IPv6 over Ethernet, tunneling avoidance, phased migration.

NEW QUESTION: 87

Which two functions are provided by the Cisco SD-WAN orchestration plane? (Choose two.)
A. centralized provisioning
B. primary authentication point
C. NAT traversal facilitation
D. Zero Touch Provisioning
E. troubleshooting and monitoring
Answer: B,C (LEAVE A REPLY)
Cisco Catalyst SD-WAN separates orchestration from management, control, and data forwarding. The orchestration-plane function is provided by vBond, which authenticates WAN Edge devices and helps establish secure connectivity between edges and controllers. In Cisco design terms, vBond is also the NAT traversal facilitator because it allows devices sitting behind NAT to discover reachable public and private address information and form the correct control connections. These two functions match the options for primary authentication point and NAT traversal facilitation. Centralized provisioning, troubleshooting, monitoring, and template-based configuration belong to vManage, the management plane. Route and policy distribution belongs to vSmart, the control plane, through OMP. Zero Touch Provisioning is part of the onboarding workflow, but the question asks for orchestration-plane functions, not the broader provisioning service. A sound SD-WAN design should deploy redundant vBond instances, ensure DNS reachability to them, and validate firewall rules so DTLS/TLS control connections can be established reliably across Internet and private transports. Reference topics: Cisco SD-WAN architecture, vBond orchestration, control connections, NAT traversal.

NEW QUESTION: 88
Refer to the exhibit.



freepdfdumps.com

An architect is designing an EIGRP solution based on these requirements:
* Traffic forwarding should use the best two paths while all links are available
* Single path failure must not impact traffic between branch and HQ

Which solution must the architect select?

A. Maximum-paths 2

B. Add-paths 2

C. Metric weights 010100

D. Variance 2

Answer: D (LEAVE A REPLY)

EIGRP variance is the correct design because the requirement is to use the best two paths while all links are available and maintain forwarding if one path fails. EIGRP supports unequal-cost load balancing when feasible successors exist and the variance multiplier allows additional loop-free paths to be installed in the routing table. Cisco EIGRP design uses the feasible distance and reported distance conditions to ensure that unequal-cost alternatives are loop free before they are used. The maximum-paths command controls how many equal or eligible paths can be installed, but it does not make an unequal path eligible. Add-paths is a BGP feature, not an EIGRP mechanism. Changing metric weights is rarely recommended because it affects metric calculation globally and can create inconsistent routing if not deployed everywhere. Variance 2 allows EIGRP to include backup paths whose metric is within two times the successor metric, assuming they satisfy the feasibility condition. Reference topics: EIGRP variance, unequal-cost load balancing, feasible successor, DUAL, path resiliency.

NEW QUESTION: 89

When is it advisable to provide dedicated control plane nodes within a Cisco SD-Access design?

A. in a small deployment where border nodes are not required

B. in a design where fabric edge nodes are unable to provide control plane functionality

C. in designs without Cisco DNA Center

D. when there is a requirement for frequent roaming of endpoints across fabric edge nodes

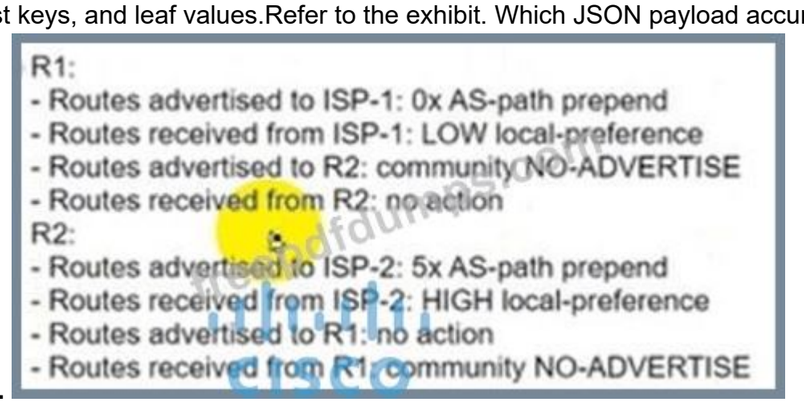
Answer: (SHOW ANSWER)

Dedicated control-plane nodes are advisable when frequent endpoint roaming across fabric edge nodes is expected. In Cisco SD-Access, the control-plane node maintains the host tracking database and resolves endpoint identifiers to routing locators. When endpoints move between fabric edge nodes, the fabric must update endpoint-to-location mappings quickly and consistently. In smaller deployments, the control-plane function can sometimes be colocated with border nodes, but larger or more mobile environments benefit from dedicated control-plane nodes to improve scale, resilience, and operational separation. The decision is not because Cisco DNA Center is absent; DNA Center is the management and automation platform for SD- Access. It is also not because edge nodes are unable to provide control-plane functionality, since that is not their fabric role. Frequent roaming creates more mapping updates and lookup activity, so dedicated control- plane resources help maintain stable endpoint registration and resolution performance. Reference topics:
Cisco SD-Access control-plane node, host tracking database, endpoint mobility, LISP map server, fabric scale.

NEW QUESTION: 90

An engineer is configuring an interface on a Cisco IOS XE device using a RESTCONF PATCH request. The target data model requires a strict hierarchy containing the correct namespace, container paths, list keys, and leaf values.Refer to the exhibit. Which JSON payload accurately represents the expected YANG data model structure to successfully commit the configuration?

A.



B.

R1:
- Routes advertised to ISP-1: 5x AS-path prepend
- Routes received from ISP-1: LOW local-preference
- Routes advertised to R2: community NO-ADVERTISE
- Routes received from R2: no action

R2:
- Routes advertised to ISP-2: 0x AS-path prepend
- Routes received from ISP-2: HIGH local-preference
- Routes advertised to R1: community NO-EXPORT
- Routes received from R1: no action

C.

R1:
Routes advertised to ISP-1: 0x AS-path prepend
Routes received from ISP-1: HIGH local-preference
Routes advertised to R2: no action
Routes received from R2: community NO-EXPORT

R2:
Routes advertised to ISP-2: 5x AS-path prepend
Routes received from ISP-2: LOW local-preference
Routes advertised to R1: community NO-ADVERTISE
Routes received from R1: no action

D.

R1:
- Routes advertised to ISP-1: 0x AS-path prepend
- Routes received from ISP-1: HIGH local-preference
- Routes advertised to R2: community NO-EXPORT
- Routes received from R2: no action

R2:
- Routes advertised to ISP-2: 5x AS-path prepend
- Routes received from ISP-2: LOW local-preference
- Routes advertised to R1: no action
- Routes received from R1: no action

Answer: D (LEAVE A REPLY)

The validated answer is Option D because this item depends on the exhibit containing configuration or model fragments. In Cisco programmability questions, the correct option is determined by whether the selected YANG or API payload accurately represents the intended configuration hierarchy, data types, and device- specific path. The answer must match the complete model structure, not merely contain the right-looking address, interface name, or protocol value. Option D should therefore be retained as the verified selection because it is the only option that aligns with the expected data model structure shown in the exhibit. For these items, the professional review must avoid inventing a replacement payload when the answer choices are embedded as images. The proper explanation is that YANG-based configuration validation requires the namespace, container path, list keys, and leaf values to match the model used by the target platform.

Reference topics: YANG configuration validation, RESTCONF or NETCONF payload structure, model hierarchy, Cisco IOS XE programmability, exhibit-based option verification.

NEW QUESTION: 91

A network engineer must segregate three interconnected campus networks using IS-IS routing. A two-layer hierarchy must be used to support large routing domains and to avoid more specific routes from each campus network being advertised to other campus network routers automatically. Which two actions does the engineer take to accomplish this segregation? (Choose two.)

A. Designate two IS-IS routers as BDR routers at the edge of each campus, and configure one BDR for all Level 1 routers and one BDR for all Level 2 routers.

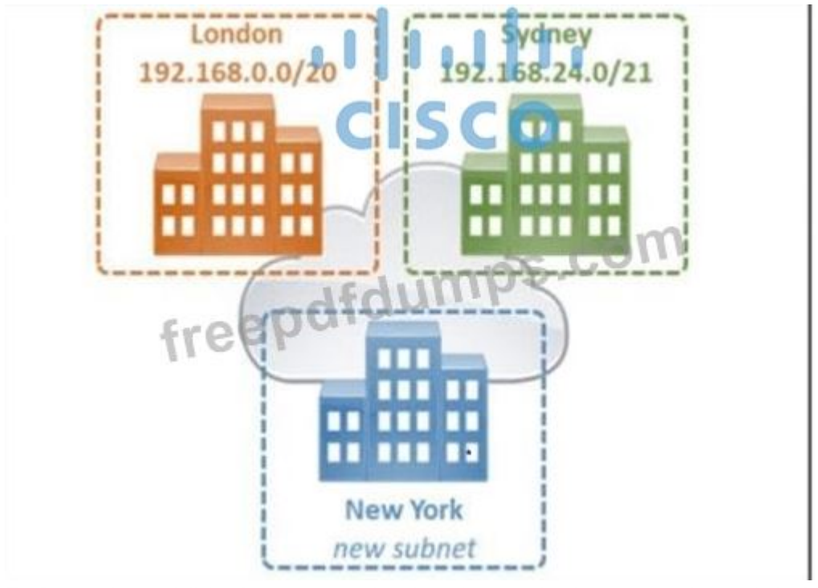
B. Designate two IS-IS routers from each campus to act as Level 1/Level 2 backbone routers at the edge of each campus network.

C. Assign the same IS-IS NET value for each campus, and configure internal campus routers with Level 1/ Level 2 routing.

D. Utilize different MTU values for each campus network segment. Level 2 backbone routers must utilize a larger MTU size of 9216.
E. Assign a unique IS-IS NET value for each campus, and configure internal campus routers with Level 1 routing.
Answer: (SHOW ANSWER)
A scalable two-level IS-IS design uses Level 1 areas for internal campus routing and Level 2 for the interarea backbone. Each campus should have a unique IS-IS NET area value so the internal routers belong to that campus's Level 1 area. Internal campus routers should be configured as Level 1 routers, which keeps detailed campus prefixes inside the local area and prevents unnecessary propagation of more-specific routes to other campuses. At the edge of each campus, routers that connect the campus to the backbone should operate as Level 1/Level 2 routers. These routers connect the local Level 1 area to the Level 2 backbone and can summarize or control routing information between levels. IS-IS does not use BDR routers; that is an OSPF term, so option A is invalid. Assigning the same NET value everywhere collapses the hierarchy and defeats the segregation goal. MTU differences are not a hierarchy or route-segmentation tool. Therefore, the design should use unique NET values per campus with Level 1 internal routers and Level 1/Level 2 routers at the campus edge.

Valid 300-420 Dumps shared by Actual4test.com for Helping Passing 300-420 Exam! Actual4test.com now offer the **newest 300-420 exam dumps**, the Actual4test.com 300-420 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-420 dumps with Test Engine here: https://www.actual4test.com/300-420_examcollection.html (**457** Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 92



Refer to the exhibit. A customer is planning to deploy a new branch in New York. The new office will not exceed 1024 users. Which subnet must be used to provide maximum number of host addresses while not providing more than necessary?
A. 192.168.8.0/21
B. 192.168.16.0/22
C. 192.168.16.0/21
D. 192.168.8.0/22
Answer: D (LEAVE A REPLY)
A /22 subnet is the smallest IPv4 block in the answer set that closely fits a branch expected not to exceed roughly one thousand users while avoiding unnecessary address waste. A /22 contains 1024 total addresses and 1022 usable host addresses after the network and broadcast addresses are reserved. That aligns with the requirement to provide the maximum number of host addresses without selecting a

larger block than necessary. A /21 provides 2048 total addresses and 2046 usable hosts, which is materially larger than the stated requirement and wastes address space. The specific 192.168.8.0/22 block also aligns on a valid /22 boundary because /22 networks increment by four in the third octet. A 192.168.16.0/22 is also mathematically valid, but the exhibit determines the available address range; the selected option uses the appropriate block from the displayed plan. Reference topics: IPv4 subnet sizing, CIDR, usable host calculation, branch address planning, VLSM.

NEW QUESTION: 93
Which topology within a network underlay eliminates the need for first hop redundancy protocols while improving fault tolerance, increasing resiliency, and simplifying the network?
A. virtualized topology
B. routed access topology
C. Layer 2 topology
D. logical fabric topology
Answer: (SHOW ANSWER)
A routed access topology is the underlay design that removes the normal need for first-hop redundancy protocols. In a traditional Layer 2 access design, end hosts usually rely on HSRP, VRRP, or GLBP at the distribution layer because the default gateway is shared across redundant distribution switches. In a routed access design, the Layer 3 boundary is pushed down to the access layer. The access switch routes toward the distribution layer using point-to-point routed links, so uplink failures are handled by the routing protocol instead of spanning tree and first-hop redundancy. Cisco campus design guidance highlights routed access as a way to improve convergence, avoid Layer 2 loop exposure across uplinks, and simplify resiliency by using equal-cost routed paths. Virtualized topology is too broad and does not inherently remove FHRP. A Layer 2 topology normally increases dependence on STP and gateway redundancy. A logical fabric topology is not the generic underlay topology described by the question. The correct choice is therefore routed access. Reference topics: routed access campus, Layer 3 access, FHRP elimination, routed underlay resiliency.

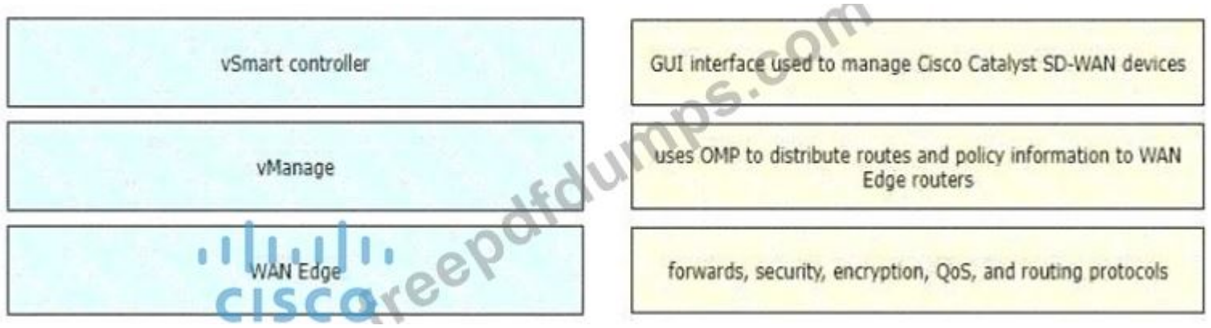
NEW QUESTION: 94
An engineer is working for a large cable TV provider that requires multiple sources streaming video on different channels using multicast with no rendezvous point. Which multicast protocol meets these requirements?
A. PIM-SM
B. PIM-SSM
C. any-source multicast
D. BIDIR-PIM
Answer: (SHOW ANSWER)
PIM Source-Specific Multicast is the correct multicast protocol when receivers can subscribe to traffic from known sources and the design must avoid rendezvous points. SSM uses the (S,G) model, where receivers explicitly join a multicast group from a specific source. Because the source is known at join time, the network does not need a shared tree or RP discovery process. This makes SSM attractive for broadcast-style streaming applications, such as video channels, where each stream source is known and receivers subscribe to specific channels. PIM Sparse Mode and Any-Source Multicast use an RP to allow receivers to discover active sources, so they do not meet the requirement of no rendezvous point. BIDIR-PIM also relies on an RP-rooted shared tree and is optimized for many-to-many applications rather than source-specific broadcast channels.
For a large cable TV provider streaming multiple video channels from defined sources, SSM simplifies the multicast control plane, eliminates RP dependency, and scales efficiently by creating only the source-specific state required by active receivers.

NEW QUESTION: 95
A network administrator is troubleshooting a DMVPN setup between the hub and the spoke. Which action should the administrator take before troubleshooting the IPsec configuration?
A. Verify the GRE tunnels.
B. Verify ISAKMP.
C. Verify NHRP.
D. Verify crypto maps.
Answer: A (LEAVE A REPLY)

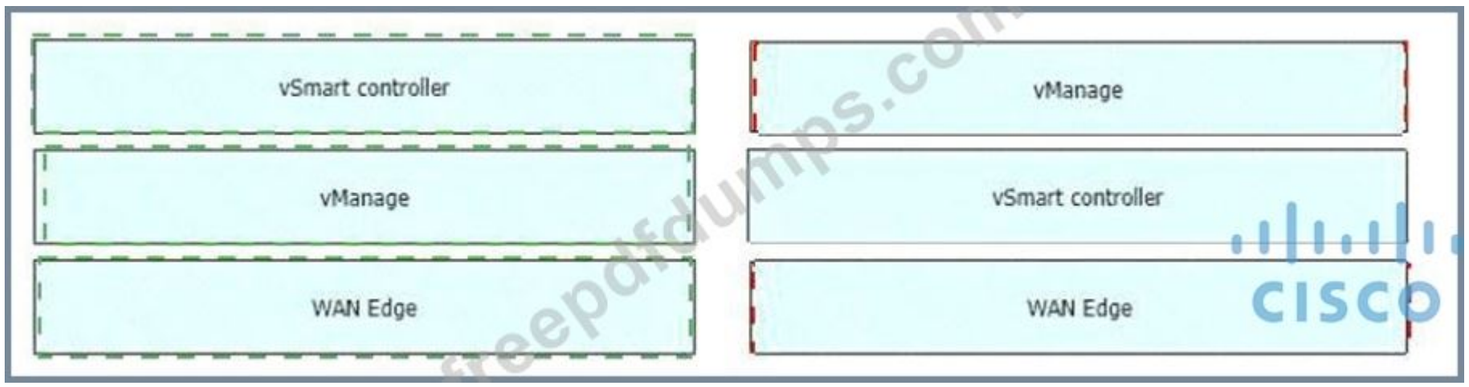
Before troubleshooting IPsec in a DMVPN deployment, the administrator should verify the GRE tunnel operation. DMVPN is built from multipoint GRE, NHRP, routing, and usually IPsec encryption. IPsec protects the tunnel traffic, but the logical tunnel and overlay reachability must be valid first. If the tunnel interface is down, incorrectly addressed, or unable to carry GRE traffic, troubleshooting crypto maps or ISAKMP will waste time because encryption cannot fix a broken overlay. After the GRE tunnel state is validated, the administrator can confirm NHRP registration and resolution, routing adjacencies, and then IPsec security associations. ISAKMP and crypto map troubleshooting are relevant only after basic tunnel operation is known to be correct. NHRP is also critical in DMVPN, but the question asks what to do before troubleshooting IPsec configuration; the first prerequisite is to verify that the GRE tunnel itself is functioning.
This layered troubleshooting order follows Cisco operational practice: confirm transport reachability, then tunnel interface behavior, then NHRP and routing, and finally cryptographic protection.

NEW QUESTION: 96

Drag and drop the Cisco Catalyst SD-WAN components from the left to their definitions on the nght



Answer:

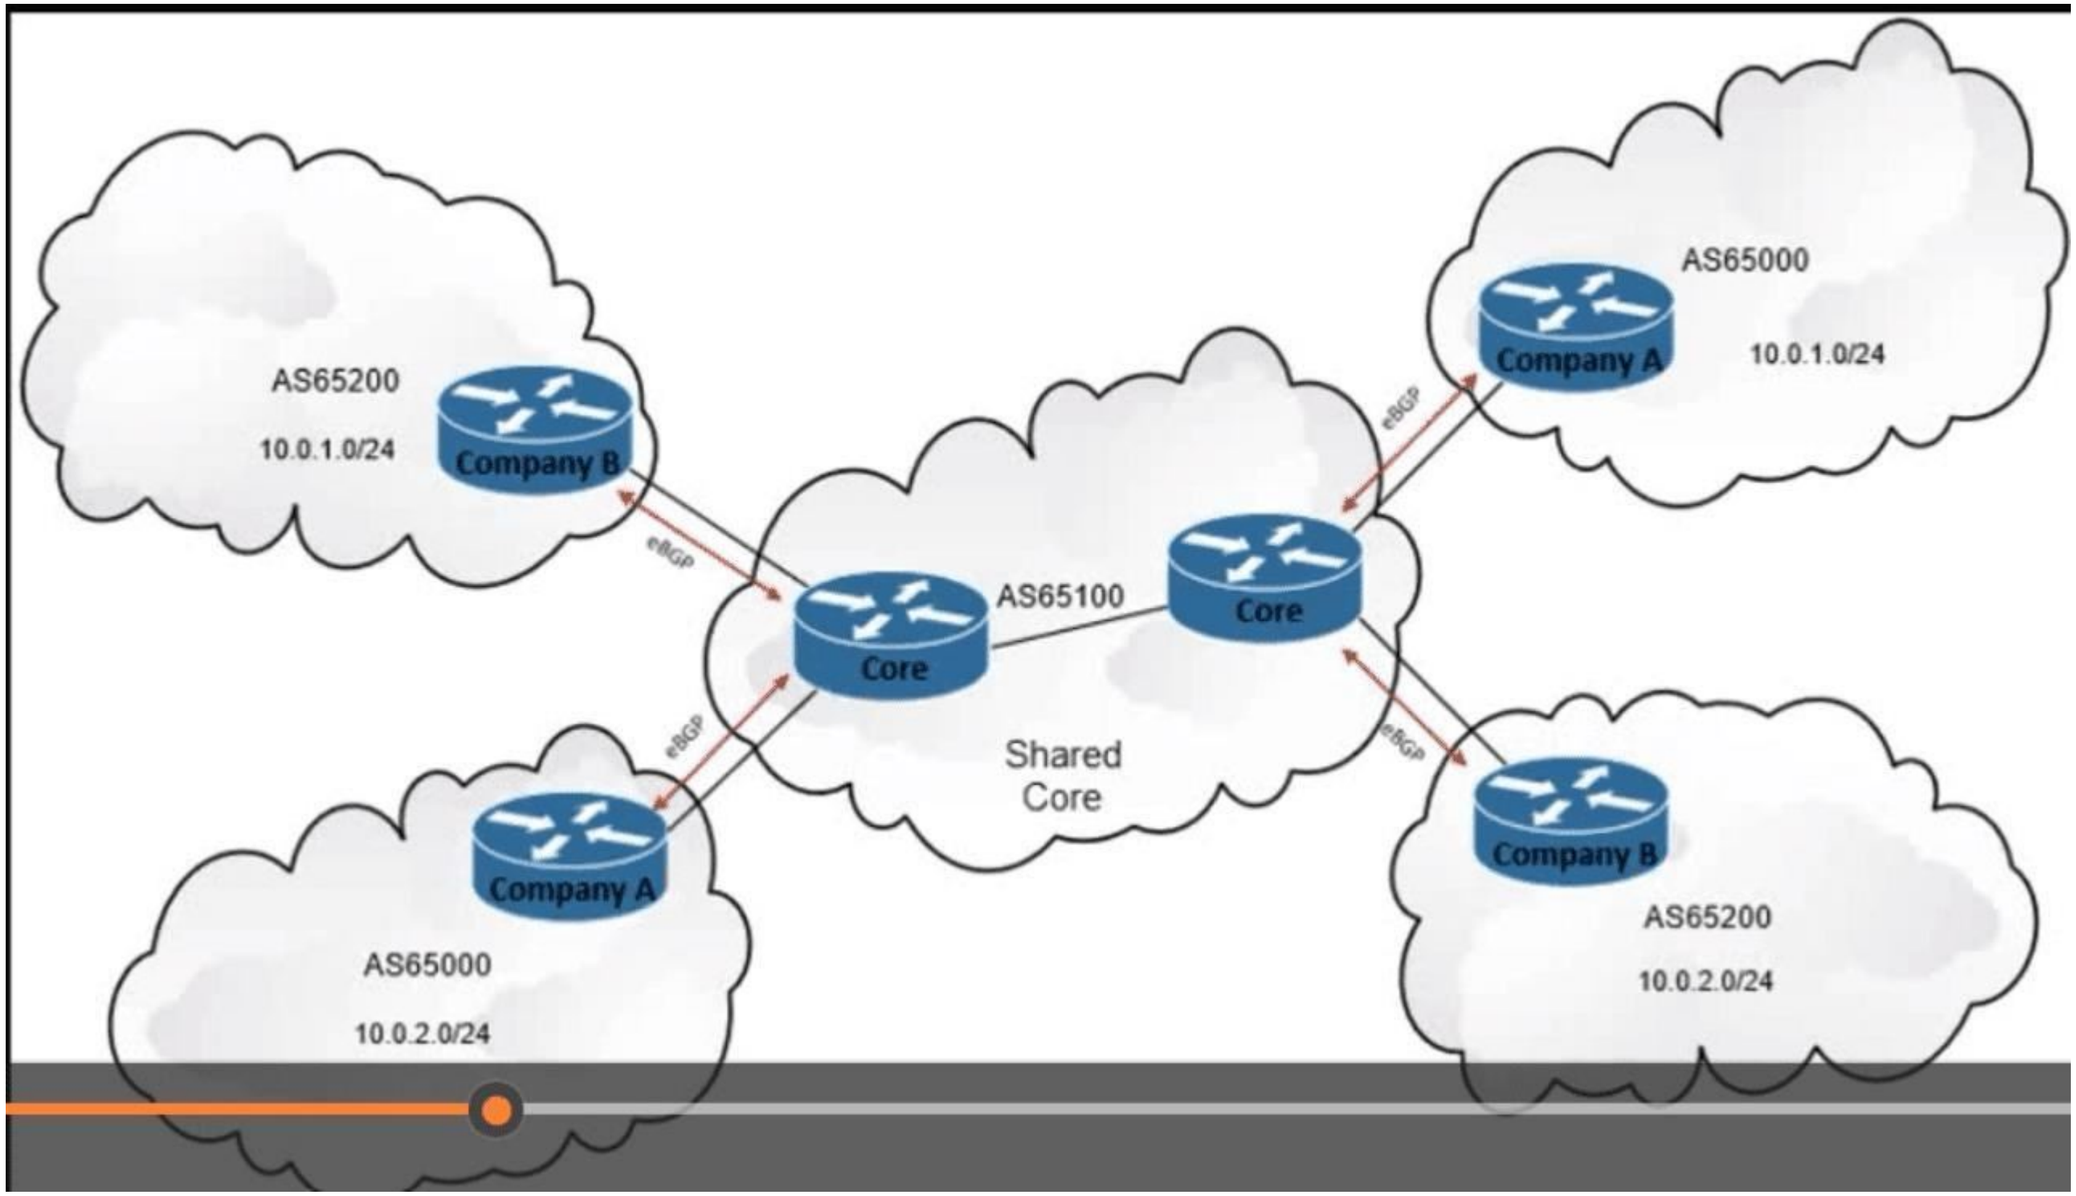


Explanation:

The drag-and-drop mapping is based on the standard Cisco Catalyst SD-WAN component roles. vManage, now Cisco Catalyst SD-WAN Manager, provides centralized management, configuration, monitoring, templates, and operational visibility. vSmart Controller provides the centralized control plane, distributes OMP routing information, and applies control policy to determine reachability and path selection. vBond Orchestrator supports initial authentication, orchestration, and NAT traversal so WAN Edge devices can join the overlay. WAN Edge routers provide the data plane at branches, campuses, data centers, and cloud edges; they forward user traffic, establish secure tunnels, and apply localized policies. Correctly separating these roles is critical in the design because management, control, orchestration, and forwarding scale independently.

Misidentifying vBond as the data-plane element or vManage as the route controller would lead to an incorrect SD-WAN architecture. Reference topics: Cisco Catalyst SD-WAN components, vManage, vSmart, vBond, WAN Edge, OMP control plane. This role separation is also what lets large deployments scale without mixing troubleshooting, provisioning, and forwarding responsibilities on one system.

NEW QUESTION: 97



Refer to the exhibit. Company A acquired Company B, and the companies want to use the same shared core.
However, the companies must not learn each other ' s prefixes. Which two address family solutions meet these requirements? Choose two.

A. IPv4 unicast for Company A and Company B

B. VPNv4 for the shared core and for Company A and Company B

C. IPv4 unicast and VPNv4 for Company A and Company B

D. IPv4 unicast and VPNv4 for shared core

Answer: A,D (LEAVE A REPLY)

The correct design uses IPv4 unicast for Company A and Company B and VPNv4 in the shared core. In an MPLS Layer 3 VPN architecture, customer edge routing toward the provider edge commonly uses the ordinary IPv4 unicast address family. The PE router then places the learned customer routes into the appropriate VRF. Across the shared provider core, MP-BGP carries those customer routes as VPNv4 prefixes by adding route distinguishers, which keep identical or overlapping customer prefixes unique. Route targets then control which VRFs import or export each route. Cisco describes this VRF

and VPNv4 model as the mechanism that prevents one VPN customer from learning another customer's prefixes while sharing the same provider backbone. VPNv4 is not normally used directly between the customer networks and their CE- facing sessions. Therefore, the customer sides use IPv4 unicast, and the shared core uses VPNv4 to preserve separation. Reference topics: MPLS Layer 3 VPN, MP-BGP VPNv4, VRF, route distinguisher, route targets, customer route isolation.

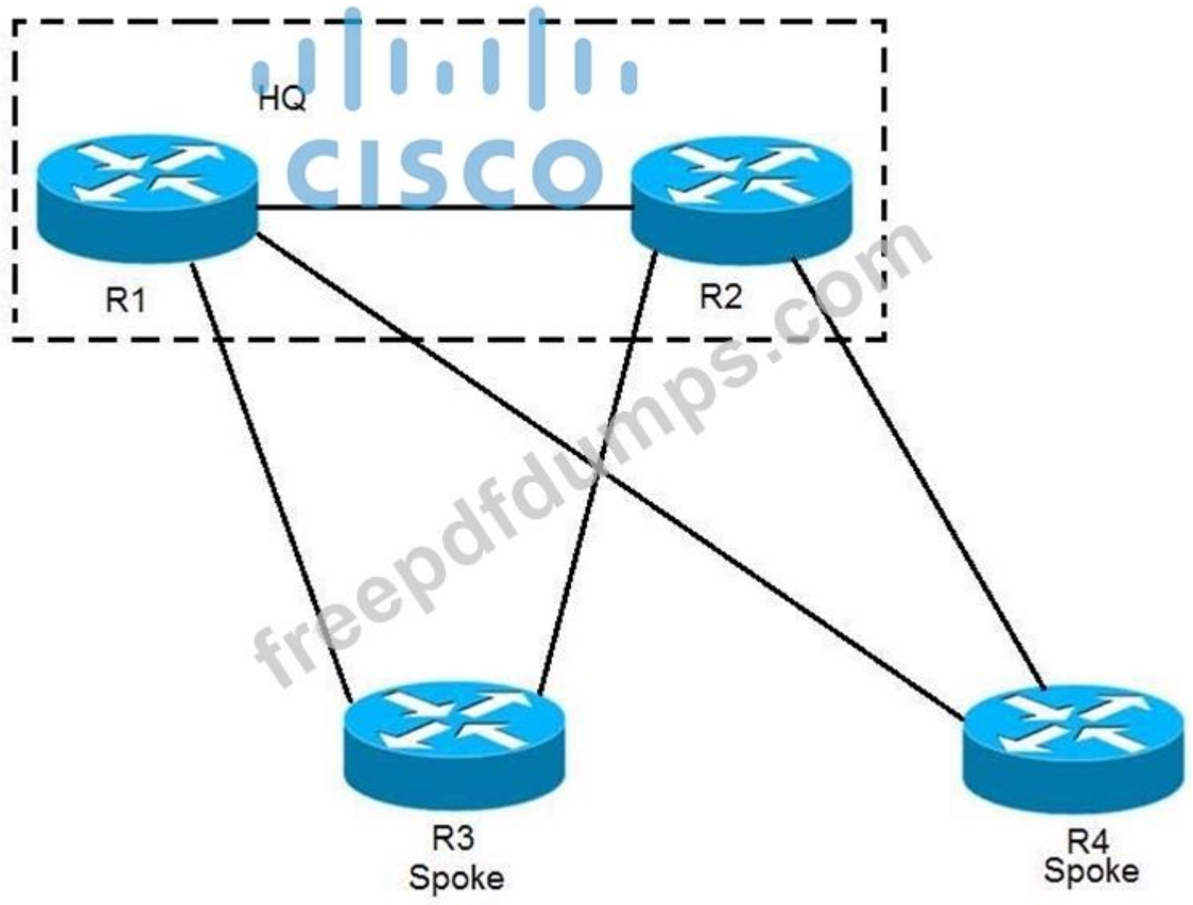
NEW QUESTION: 98
An engineer needs to design a management network for the company. The solution has these requirements:
* overlay network does not cause routing issues
* ease of troubleshooting for the operations team
* devices are accessed securely
Which solution meets these requirements?
A. VRF for management traffic and SSH keys for device access
B. Private VLANs for management traffic and TACACS+ for device access
C. Separate physical interfaces for management traffic and TACACS+ for device access
D. VLANs for management traffic and RADIUS for device access
Answer: A (LEAVE A REPLY)
The correct solution is again a management VRF with SSH keys for device access. This duplicate item tests the same design principle: management traffic should be logically isolated from the overlay and production forwarding paths. A VRF gives the management plane its own routing table, so changes or failures in the overlay do not normally alter management reachability. That makes the network easier to operate because engineers can troubleshoot device access, routing, and overlay behavior as separate concerns. SSH keys satisfy the secure-access requirement by providing encrypted management sessions with public-key authentication. Private VLANs and standard VLANs can segment Layer 2 traffic, but they do not solve routing separation by themselves. TACACS+ and RADIUS provide authentication, authorization, and accounting, but they do not prevent overlay routing from interacting with management reachability. A separate physical interface can help in some designs, but option A is the cleanest answer because it combines routing isolation with secure access. Reference topics: management VRF, SSH, secure management plane, route isolation, overlay troubleshooting.

NEW QUESTION: 99
Which control plane protocol is responsible for EID-to-RLOC mapping concerning SO-Access Architecture?
A. GBAC
B. LISP
C. CEF
D. VXLAN
Answer: (SHOW ANSWER)
LISP is the control-plane protocol responsible for Endpoint Identifier to Routing Locator mapping in Cisco SD-Access. In the SD-Access architecture, an endpoint address is treated as an EID, and the fabric node location is represented by an RLOC. The fabric control-plane node performs LISP map-server and map- resolver functions so edge nodes can register local endpoints and resolve the current location of remote endpoints. This separation allows the overlay to support host mobility, anycast gateway behavior, and segmentation without depending only on traditional subnet-bound forwarding. VXLAN is the fabric data- plane encapsulation used to carry traffic across the fabric and preserve policy information, but it does not provide the EID-to-RLOC mapping function. CEF is the local forwarding mechanism inside Cisco devices, and GBAC is not the SD-Access mapping protocol. Therefore, LISP is the correct answer because it provides the SD-Access fabric control plane. Reference topics: Cisco SD-Access control plane, LISP, EID, RLOC, map-server, map-resolver.

NEW QUESTION: 100
What are two benefits of designing an SD-WAN network fabric with direct Internet access implemented at every site? (Choose two.)
A. It decreases latency to applications hosted by public cloud service provider.
B. It decreases latency on Internet circuits.
C. It increases the speed of delivery of site deployments through zero-touch provisioning.

D. It increases the total available bandwidth on Internet circuits.
E. It alleviates network traffic on MPLS circuits.
Answer: (SHOW ANSWER)
Implementing direct Internet access at every Cisco SD-WAN site improves application experience and reduces unnecessary backhaul. Many modern business applications are hosted in SaaS or public cloud environments. If branch Internet traffic must hairpin through a central data center before reaching those applications, latency increases and central WAN or security stacks can become bottlenecks. SD-WAN DIA allows selected cloud and Internet-bound traffic to exit locally while still applying policy, segmentation, and security controls. This decreases latency to applications hosted by public cloud providers and improves user experience. It also alleviates traffic on MPLS circuits because Internet and cloud traffic no longer needs to consume private WAN bandwidth back to a centralized egress point. DIA does not inherently reduce the physical latency of an Internet circuit, and it does not increase the provisioned circuit bandwidth. Zero-touch provisioning may accelerate site deployment, but that is a separate SD-WAN onboarding benefit rather than a specific benefit of deploying DIA at every site. Therefore, the two correct benefits are reduced latency to public cloud applications and reduced load on MPLS circuits.

NEW QUESTION: 101
Refer to the exhibit.



EIGRP has been configured on all links. The spoke nodes have been configured as EIGRP stubs, and the WAN links to R3 have higher bandwidth and lower delay than the links to R4. When a link failure occurs at the R1-R2 link, what happens to traffic on R1 that is destined for a subnet attached to R2?
A. R1 has no route to R2 and drops the traffic
B. R1 load-balances across the paths through R3 and R4 to reach R2
C. R1 forwards the traffic to R3, but R3 drops the traffic

D. R1 forwards the traffic to R3 in order to reach R2

Answer: A (LEAVE A REPLY)

EIGRP stub routing is designed to prevent remote or spoke routers from being used as transit routers. In a hub- and-spoke or dual-homed remote design, a stub router advertises only the route types permitted by its stub configuration, such as connected or summary routes, and it is not queried for routes that it cannot reasonably provide. This reduces EIGRP query scope and helps avoid stuck-in-active conditions. The tradeoff is that the network must not depend on a stub site to carry transit traffic between other parts of the topology. In this scenario, the spoke nodes are configured as EIGRP stubs. When the direct R1-to-R2 link fails, R1 cannot use the alternate path through the other spoke-side topology as a transit path to reach a subnet attached to R2. The EIGRP stub design prevents the route from being advertised in a way that would make the spoke act as a transit router. As a result, R1 has no valid route to R2 ' s attached subnet and drops the traffic. This behavior is consistent with the purpose of EIGRP stub routing.

NEW QUESTION: 102

A company wants to switch from static to dynamic routing. The branches use DMVPN back to the hub using two internet connections. One internet connection speed is 10 Mbps, and the other is 100 Mbps. All locations use Cisco routers; however, the branch routers have limited memory and CPU resources. Which routing protocol and design solution must the company choose for optimal traffic forwarding during peak traffic times?

A. iBGP with the hub routers set up as route reflectors

B. OSPF deployed in area 0 with branch routers connected back via virtual links

C. EIGRP with branch routers as stub routers and variance enabled

D. ISIS with the hub and spoke routers configured in two different areas

Answer: C (LEAVE A REPLY)

EIGRP with stub routing and variance is the best fit for Cisco-only DMVPN branches with unequal Internet link speeds and limited branch resources. EIGRP works well over DMVPN because it supports hub-and- spoke designs, unequal-cost load balancing, and simple branch optimization through stub routing. Stub branches limit query propagation and reduce CPU, memory, and bandwidth consumption on small routers.

Variance allows EIGRP to use feasible unequal-cost paths, which is important when one Internet connection is 10 Mbps and the other is 100 Mbps. OSPF can run over DMVPN, but virtual links are not an appropriate branch design and OSPF is less efficient in this specific Cisco-only unequal-bandwidth scenario. IS-IS is not commonly selected for small DMVPN branch overlays. iBGP with route reflectors can scale, but it does not inherently use link utilization and unequal metric behavior as cleanly as EIGRP variance for this requirement.

Reference topics: EIGRP over DMVPN, stub routing, variance, unequal-cost load balancing, branch resource optimization.

NEW QUESTION: 103

An engineer needs to design a management network for the company. The solution has these requirements:

* overlay network does not cause routing issues

* ease of troubleshooting for the operations team

* devices are accessed securely

Which solution meets these requirements?

A. VRF for management traffic and SSH keys for device access

B. Private VLANs for management traffic and TACACS+ for device access

C. Separate physical interfaces for management traffic and TACACS+ for device access

D. VLANs for management traffic and RADIUS for device access

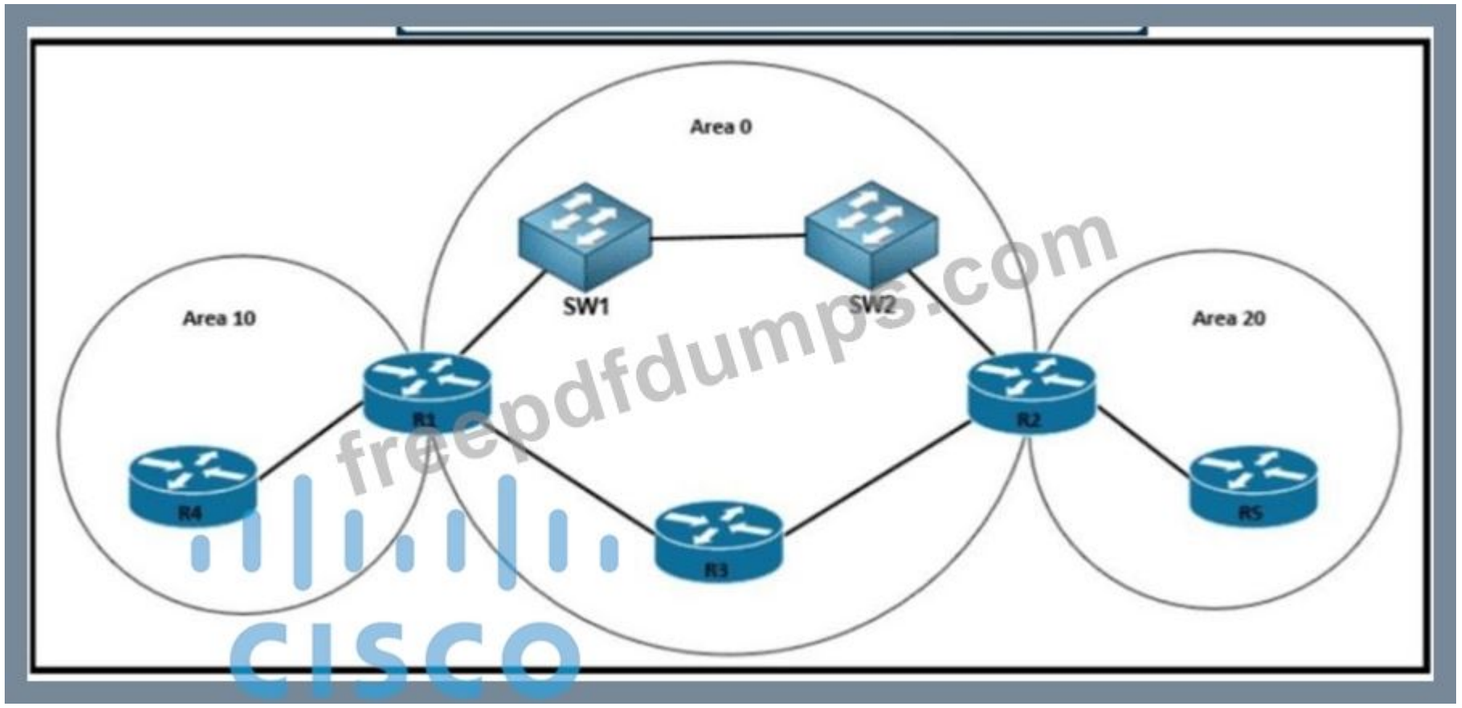
Answer: (SHOW ANSWER)

The correct solution is a management VRF with SSH keys for device access. A management VRF creates a separate routing and forwarding table for management traffic, keeping administrative reachability isolated from production and overlay routes. That directly satisfies the requirement that the overlay network must not cause routing issues and makes troubleshooting easier for operations teams because management reachability can be validated independently. Cisco VRF design uses separate routing tables to isolate traffic and can support overlapping addresses or independent routing policy. For secure device access, SSH with key-based authentication provides encrypted administrative sessions and stronger authentication than clear-text or password-only methods. Private VLANs and ordinary VLANs

provide Layer 2 separation but do not create independent routing tables. TACACS+ and RADIUS are AAA systems; they are useful but do not define the management transport architecture. Separate physical interfaces can be useful out-of-band, but the best match for overlay isolation and operational simplicity is management VRF plus SSH keys. Reference topics:
management VRF, secure device access, SSH public-key authentication, routing-table isolation, enterprise management design.

NEW QUESTION: 104
Which NETCONF operation creates filtering that is specific to the session notifications?
A. < create-subscription >
B. < commit >
C. < notification >
D. < logging >
Answer: A (LEAVE A REPLY)
NETCONF supports event notifications through a subscription model. The operation used to establish a notification subscription is create-subscription. When a client sends this operation, it can specify the stream and optional filters that determine which notifications are delivered to that NETCONF session. This is why create-subscription is the operation associated with session-specific notification filtering.
The commit operation is used with the candidate configuration datastore to apply configuration changes to the running datastore; it does not create a telemetry or event subscription. The notification element is the message sent by the server after a subscription exists, not the operation used to create one. Logging is not a NETCONF operation in the protocol operation set. In Cisco model-driven management designs, NETCONF provides structured configuration and operational access using YANG data models, while notifications allow the management station to consume specific events without relying on repeated screen scraping or broad polling.
Therefore, the correct operation for creating filtering specific to session notifications is create-subscription.

NEW QUESTION: 105



Refer to the exhibit. An architect must ensure a convergence time of 200 ms or less during a link failure within area 0. In addition, the solution must not impact the overall performance of the network. Which solution must the architect select?
A. UDLD

B. BFD
C. fast hellos
D. carrier delay
Answer: (SHOW ANSWER)
BFD is the correct design feature for detecting a link failure within 200 ms or less without materially increasing OSPF control-plane load. Cisco uses Bidirectional Forwarding Detection as a fast, lightweight failure-detection protocol that can be associated with routing protocols such as OSPF. Instead of relying on aggressive OSPF hello and dead timers, BFD performs rapid forwarding-path checks and notifies OSPF when the path fails. This allows fast convergence while avoiding the additional routing-protocol overhead and instability that can come from very low hello timers. UDLD detects unidirectional Layer 2 fiber failures, but it is not the routing adjacency failure-detection mechanism requested here. Carrier delay can adjust interface state transition timing, but it does not provide a routing-protocol independent liveness session. Fast hellos can reduce detection time, but they increase protocol traffic and can affect overall performance. BFD is therefore the proper design choice for fast and efficient OSPF convergence. Reference topics: BFD, OSPF convergence, fast failure detection, routed campus resiliency, control-plane efficiency.

NEW QUESTION: 106
Which type of rendezvous point deployment is standards-based and supports dynamic RP discovery?
A. bootstrap router
B. Anycast-RP
C. Auto-RP
D. static RP
Answer: A (LEAVE A REPLY)
Bootstrap Router is the standards-based mechanism that supports dynamic RP discovery in a PIM sparse- mode domain. Cisco documentation distinguishes Auto-RP, which is Cisco-proprietary, from BSR, which is defined for standards-based RP distribution. With BSR, candidate RPs advertise their availability, and the elected bootstrap router distributes RP-set information throughout the PIM domain. This removes the operational burden of manually configuring the same static RP information on every multicast router. Anycast- RP provides RP redundancy and load sharing but does not by itself provide standards-based dynamic RP discovery across the domain. Static RP is simple but not dynamic. Auto-RP can dynamically distribute RP information, but it is not the standards-based answer. The design value of BSR is most visible in large multicast domains, where routers need consistent RP information and manual changes are risky. Therefore, when the question asks for a standards-based RP deployment that supports dynamic RP discovery, the correct answer is bootstrap router. Reference topics: PIM sparse mode, Bootstrap Router, candidate RP, RP-set distribution, dynamic RP discovery.

Valid 300-420 Dumps shared by Actual4test.com for Helping Passing 300-420 Exam! Actual4test.com now offer the **newest 300-420 exam dumps**, the Actual4test.com 300-420 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-420 dumps with Test Engine here: https://www.actual4test.com/300-420_examcollection.html (**457** Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

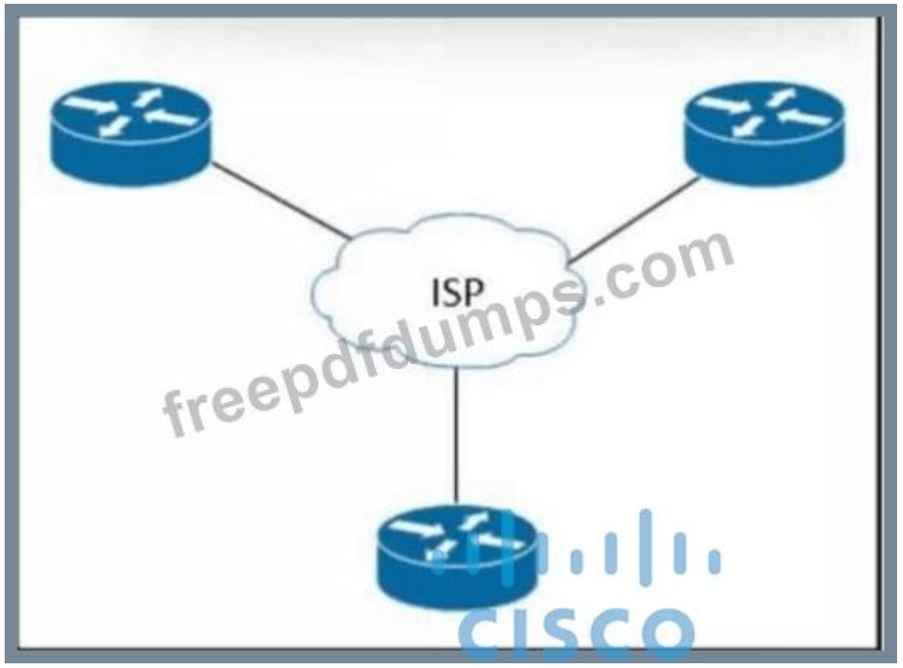
NEW QUESTION: 107
When a first hop redundancy solution is designed, which protocol ensures that load balancing occurs over multiple routers using a single virtual IP address and multiple virtual MAC addresses?
A. GLBP
B. IRDP
C. VRRP
D. HSRP
Answer: A (LEAVE A REPLY)
Gateway Load Balancing Protocol is the first-hop redundancy protocol that provides load balancing while presenting a single default gateway IP address to hosts. Cisco's GLBP documentation states that GLBP
"provides load balancing over multiple routers using a single virtual IP address and multiple virtual MAC addresses." That wording directly matches the requirement in the question. In a GLBP group, the

active virtual gateway answers ARP requests for the virtual IP address and can return different virtual MAC addresses associated with active virtual forwarders. As a result, multiple routers forward traffic for the same subnet rather than leaving one router idle during normal operation. HSRP and VRRP provide first-hop redundancy, but their basic operation uses one active or master forwarding device for a virtual IP at a time unless multiple groups are manually engineered per VLAN. IRDP is not the Cisco campus FHRP used for this design. Therefore, when the requirement is automatic load balancing across multiple routers using one virtual IP and multiple virtual MAC addresses, GLBP is the correct protocol.

NEW QUESTION: 108
An architect is designing a connectivity solution for a customer ' s two small branch offices. The customer wants a cost-effective design, no routing overload, and some down time during the year is acceptable. Which connectivity solution must the architect choose?
A. dual multihomed
B. single-homed
C. single multihomed
D. dual-homed
Answer: (SHOW ANSWER)
A single-homed WAN design is appropriate for two small branch offices when cost is the main concern, routing overload must be avoided, and occasional downtime is acceptable. Single-homed connectivity uses one WAN connection from the branch to the provider or headquarters. It is the simplest and least expensive model because it minimizes circuits, routers, routing adjacencies, and failover complexity. Dual-homed or multihomed designs improve availability, but they add cost and operational overhead through additional links, devices, routing policy, and monitoring. Dual multihomed designs are the most resilient but clearly exceed the requirement for a small cost-sensitive branch. A single multihomed design provides some provider or link diversity, but it is still more complex than necessary if the customer accepts some downtime. In Cisco WAN design, the correct topology must match the business availability requirement; not every site justifies redundant circuits. Reference topics: WAN topology selection, single-homed branch design, cost versus availability, routing simplicity, branch connectivity.

NEW QUESTION: 109
An engineer must design a QoS solution for a customer that is connected to an ISP over a 1Gbps link with a
100Mbps CIR. The ISP aggressively drops all traffic received over which is causing numerous TCP retransmissions. The customer is not using any RTP applications but wants to maximize bandwidth usage up to the CIR. Which QoS solution engineer choose?
A. Policing
B. Traffic shaping
C. Policer with markdown
D. Queuing
Answer: B (LEAVE A REPLY)
Traffic shaping is the correct QoS design for a 1-Gbps handoff with a 100-Mbps committed information rate when the provider aggressively drops traffic above the CIR. Cisco distinguishes shaping from policing:
shaping buffers and delays excess packets so the sender transmits at a configured rate, while policing drops or remarks traffic that exceeds the configured rate. Because the customer wants to maximize bandwidth usage up to the 100-Mbps CIR and avoid excessive TCP retransmissions, shaping is the better fit. Policing at the customer edge would simply drop excess traffic locally, creating the same kind of loss problem. A markdown policer is useful when lower-priority traffic should be reclassified, but it does not smooth bursts into the provider's CIR. Queuing helps during congestion, but without shaping the router can still send bursts at the 1- Gbps physical rate into a provider policer. The correct design is to shape outbound traffic to the CIR, then use child-policy queuing if class-based service guarantees are required. Reference topics: traffic shaping, policing, CIR, hierarchical QoS, TCP retransmission avoidance.

NEW QUESTION: 110



freepdfdumps.com

Refer to the exhibit. Due to budget constraints, a customer decided to purchase WAN routers with one LAN and one WAN interface per device. There is a requirement to connect the three sites to ensure high availability without buying additional WAN links. Which design deployment must the customer choose?

A. single-homed full mesh

B. single-homed hub-and-spoke

C. dual-homed hub-and-spoke

D. dual-homed full mesh

Answer: A (LEAVE A REPLY)

A single-homed full-mesh design is the best match when each WAN router has only one WAN interface, no additional WAN links can be purchased, and the customer still wants resilient site-to-site reachability. Single- homed means each site has one physical WAN attachment. Full mesh describes the logical connectivity model between sites, where each site can communicate directly with the others across the WAN service or overlay rather than relying completely on a central hub for intersite traffic. A hub-and-spoke design is simpler, but it creates dependency on the hub for spoke-to-spoke communication and provides poorer resiliency for a three- site design. Dual-homed hub-and-spoke and dual-homed full mesh both require additional WAN connectivity or interfaces, which violates the stated budget and hardware constraints. The best available design is therefore to use the single WAN attachment at each site while building a full-mesh logical topology over the provider or overlay service. Reference topics: WAN topology design, single-homed connectivity, full mesh, hub-and- spoke limitations, high availability under interface constraints.

NEW QUESTION: 111

An enterprise needs to enhance its WAN availability after a recent outage with its only MPLS provider. The proposed solution must have a quick deployment, be affordable, be reliable, and work as a backup for the enterprise ' s primary MPLS connection. Which solution meets these requirements?

A. Contract an internet connection and deploy DMVPN.

B. Deploy BFD echo mode and probe provider PE

C. Deploy an additional WAN router and use a floating static route

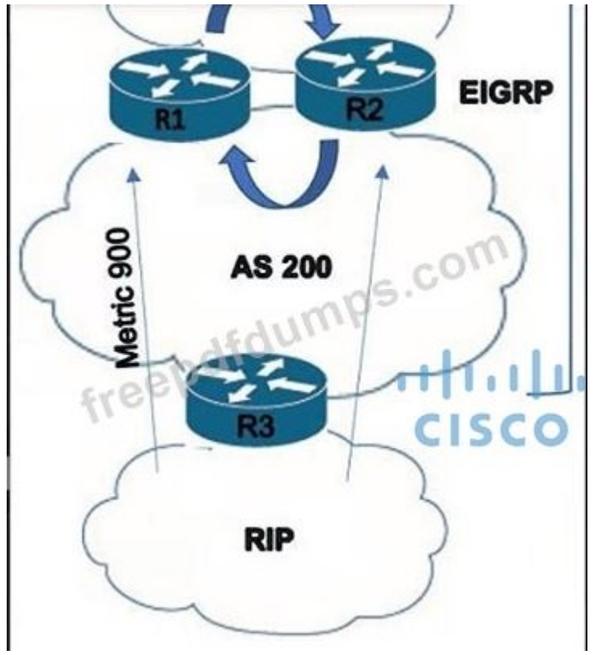
D. Contract another MPLS provider and deploy GET VPN.

Answer: A (LEAVE A REPLY)

Contracting an internet connection and deploying DMVPN is the most practical solution when the enterprise needs a fast, affordable, and reliable backup for a single MPLS provider. DMVPN can use broadband internet transport to build encrypted GRE/IPsec tunnels between sites, providing a backup WAN path without waiting for another private MPLS provider circuit. It supports dynamic routing and

can be deployed incrementally, which makes it suitable after a recent outage when time and cost matter. BFD echo mode toward a provider PE can improve failure detection, but it does not create an alternate transport if the only MPLS provider has an outage. Adding a WAN router and a floating static route also fails unless another path exists. Contracting a second MPLS provider and deploying GETVPN may be robust, but it is slower and more expensive than the requested quick and affordable backup. Reference topics: DMVPN backup WAN, MPLS resiliency, internet VPN, IPsec, branch availability, rapid WAN deployment.

NEW QUESTION: 112



Refer to the exhibit. An architect must design a solution to connect the network behind R3 with the EIGRP network. Which mechanism should be included to avoid routing loops?

A. split-horizon

B. summarization

C. down bit

D. route tags

Answer: (SHOW ANSWER)

Route tags are the correct loop-prevention mechanism when connecting another routing domain or redistributed network behind R3 into an EIGRP network. Redistribution can create route feedback: a route learned from one domain is redistributed into another, then later redistributed back into the original domain as if it were new information. That can produce routing loops or suboptimal path selection, especially in multipoint redistribution designs. Cisco route-map based redistribution commonly uses route tags to mark the origin of redistributed routes. Other redistribution points can then match those tags and prevent the same routes from being reintroduced into the source domain. Split horizon prevents advertisements back out the interface on which a route was learned, but it is not sufficient for multipoint redistribution loop prevention.

Summarization reduces route detail and query scope but does not identify route origin. The OSPF down bit is used in MPLS VPN/OSPF PE-CE scenarios, not general EIGRP redistribution. Therefore, the design should use route tags. Reference topics: EIGRP redistribution, route tagging, routing-loop prevention, route maps, redistribution policy.

NEW QUESTION: 113

An infrastructure team is concerned about the shared memory utilization of a device, and for this reason, they need to monitor the device state. Which solution limits impact on the device and provides the required data?

A. IPFIX

B. static telemetry
C. on-change subscription
D. periodic subscription
Answer: (SHOW ANSWER)
An on-change telemetry subscription is the best solution when the team needs device state data while limiting impact on the device. With periodic telemetry, the device sends subscribed data at a configured interval regardless of whether the value has changed. For state such as shared memory utilization, aggressive periodic polling or streaming can create unnecessary CPU, memory, and bandwidth overhead, especially at scale. On- change telemetry sends an update only when the subscribed data changes, or when a relevant event occurs according to the model and platform behavior. This reduces needless data transmission while still providing timely visibility into meaningful state changes. IPFIX is designed for flow information and traffic accounting rather than detailed device state telemetry. Static telemetry is not the best characterization of the efficient event-driven behavior required here. A periodic subscription can provide the data, but it continuously streams updates and therefore has greater device and collector impact. Cisco model-driven telemetry design favors on- change subscriptions when the operational requirement is state awareness with minimal repeated polling.

NEW QUESTION: 114
Which two statements about VRRP object tracking are true? (Choose two)
A. The priority of a VRRP device can change in accordance with the up or down status of a VRRP object
B. The VRRP interface priority must be manually configured by the administrator
C. A VRRP group can track only one object at a time
D. VRRP can track the status of interfaces and routes
E. VRRP supports only interface tracking
Answer: (SHOW ANSWER)
VRRP object tracking allows the effective priority of a VRRP router to change when a tracked object changes state. Cisco documentation states that VRRP object tracking ensures the best VRRP router is master by altering VRRP priorities according to tracked objects such as interface or IP route states. This directly supports answer A. It also supports answer D, because VRRP can track interfaces and routes through the tracking process. The priority does not have to remain fixed; object tracking is specifically designed so a router can lower its priority when an uplink, route, or critical dependency fails. A VRRP group is not limited to one tracked object in modern Cisco implementations, so option C is too restrictive. VRRP does not support only interface tracking; route tracking is also a common use case, especially when the LAN interface is still up but upstream reachability has failed. In a resilient campus or branch gateway design, object tracking prevents a router from remaining master when it no longer has a valid upstream path, improving deterministic failover.

NEW QUESTION: 115
An engineer is designing a BGP network for a large customer. To permit efficient scaling, the BGP domain is split into clusters. Which peering solution should be used between the route reflectors in different clusters for the BGP routes to be propagated appropriately?
A. The route reflectors should be made dents of each other.
B. The route reflectors should be nonclients with regards to each other.
C. The route reflectors should not have any kind of BGP peering.
D. The route reflectors should have peering through another nonclient router.
Answer: (SHOW ANSWER)
Route reflectors in different clusters should peer with one another as nonclients. BGP route reflection is used to reduce the full-mesh requirement inside an AS. Within a cluster, route-reflector clients peer with their route reflector and do not need to peer with every other client. However, route reflectors still need a way to exchange routes learned from their respective clusters. Cisco BGP route-reflection rules allow a route learned from a client to be reflected to other clients and to nonclient peers. A route learned from a nonclient peer is reflected only to clients. Therefore, route reflectors in different clusters should maintain iBGP peerings as nonclients so that routes can propagate between clusters and then be reflected down to each cluster's clients.
Making route reflectors clients of each other can create confusing policy and cluster behavior. Having no peering would isolate the clusters. Peering only through another nonclient router adds unnecessary indirection and does not improve scale. The correct scalable design is route-reflector-to-route-reflector nonclient peering.
Reference topics: BGP route reflectors, clusters, client and nonclient peers, iBGP scaling, route reflection rules.

NEW QUESTION: 116
Refer to the exhibit.



A network engineer must design a multicast solution based on:
* Many-to-many communications between the users and sources
* Support of up to 50 multicast sources
* Users that must register for steams
Which multicast solution must the engineer select?
A. Any Source Multicast
B. Bidirectional PIM
C. Source-Specific Multicast
D. Multicast VPN
Answer: (SHOW ANSWER)
Bidirectional PIM is the best fit for this multicast design because the application is many-to-many and has a manageable but meaningful number of multicast sources. Cisco describes BIDIR-PIM as a multicast mode designed for many-to-many applications where sources and receivers can be widely distributed. Unlike classic PIM sparse mode, BIDIR-PIM does not build source-specific shortest-path trees for each active source.
Instead, traffic uses a shared bidirectional tree rooted at the rendezvous point, which reduces multicast state in the network and improves scale when many sources may send to the same group. Source-Specific Multicast is optimized for one-to-many distribution where receivers know the source and join an (S,G) channel; it is not the natural many-to-many choice. Any Source Multicast can support many-to-many traffic but usually creates more state because source trees may be built. Multicast VPN is a transport/service architecture, not the specific multicast mode requested. Reference topics: BIDIR-PIM, many-to-many multicast, shared trees, RP- based forwarding, multicast state scaling.

NEW QUESTION: 117
A customer is undergoing a WAN re-architecture and wants to design QoS policies for remote sites that have low bandwidth. What must be considered to have CBWFQ configured in the parent and child policies in an end-to-end QoS design?
A. CBWFQ is only supported in the child policy.
B. CBWFQ is only supported in the parent policy.
C. Traffic shaping is required in the parent policy.

D. Traffic policing is required in the child policy.

Answer: (SHOW ANSWER)

Traffic shaping is required in the parent policy when hierarchical QoS uses CBWFQ in the child policy for low-bandwidth remote sites. In Cisco hierarchical QoS, the parent policy commonly shapes outbound traffic to the real available rate of the WAN service, especially when the physical interface speed is higher than the provider committed rate. The child policy then applies class-based queuing, such as CBWFQ and LLQ, inside the shaped rate. This structure creates an artificial congestion point on the router so the child queues can schedule packets predictably instead of allowing the provider to drop traffic after it exceeds the service rate.

Policing in the child policy would drop excess packets and is not the normal prerequisite for nested CBWFQ.

CBWFQ is supported in child policies within hierarchical designs, and parent shaping is the mechanism that makes the hierarchy effective. This is especially important on low-bandwidth WAN links where voice, video, and business-critical applications require deterministic treatment. Reference topics: hierarchical QoS, CBWFQ, parent shaping, child queuing, WAN QoS design.

NEW QUESTION: 118

An architect is designing a network that will utilize the spanning tree protocol to ensure a loop-free topology.

The network will support an engineering environment where it is necessary for end users to connect their own network switches for testing purposes. Which feature should the architect include in the design to ensure the spanning tree topology is not affected by these rogue switches?

A. BPDU Skew Detection

B. BPDU guard

C. loop guard

D. root guard

Answer: (SHOW ANSWER)

Root guard is the correct feature when the design must allow switches to be attached for testing but prevent those switches from influencing the spanning-tree root placement. Cisco root guard protects the intended STP root by placing a port into a root-inconsistent state if superior BPDUs are received on a port where the root should never appear. This directly addresses the risk created by rogue or lab switches connected by end users:

if a user switch advertises a better bridge priority, the production access topology is not allowed to reconverge around that unauthorized device. BPDU guard is usually used on PortFast edge ports to err-disable a port that receives any BPDU, which is excellent for strict host-only ports but less aligned with the wording that users may connect their own switches for testing. Loop guard protects root or alternate ports from losing BPDUs and accidentally forwarding. BPDU skew detection is a diagnostic feature, not the correct design control.

Therefore, root guard is the appropriate protection against rogue switches affecting STP topology. Reference topics: STP root placement, root guard, superior BPDUs, campus Layer 2 protection.

NEW QUESTION: 119

An architect is designing a network solution for a customer. The customer wants a design with redundancy at the distribution layer and the fastest convergence time possible during a failover. In addition, the company has a small IT support team, so the deployment process must be simple and quick. Which solution must the architect select?

A. GLBP

B. VRRP

C. HSRP

D. VSS

Answer: D (LEAVE A REPLY)

VSS is the best choice because it provides distribution-layer redundancy with a simplified operational model and fast failover. Cisco Virtual Switching System combines two physical Catalyst switches into one logical switch from the perspective of neighboring devices. This allows the access layer to use Multichassis EtherChannel toward the distribution pair, creating a loop-free Layer 2 topology without relying on STP blocking for normal forwarding. Cisco also describes VSS as using Stateful Switchover and Nonstop Forwarding so the standby chassis can preserve forwarding during supervisor or chassis events. That meets the requirement for very fast convergence while simplifying deployment for a small IT team. GLBP, VRRP, and HSRP are first-hop redundancy protocols. They provide default-gateway resiliency, but they do not merge the distribution layer into a single logical control plane or simplify dual-uplink forwarding the way VSS with MEC does. Reference topics: VSS, Multichassis EtherChannel, distribution-layer redundancy, SSO, NSF, campus high availability.

NEW QUESTION: 120
When a network is designed using IS-IS protocol, which two circuit types are supported? (Choose two.)
A. nonbroadcast multiaccess
B. multiaccess
C. point-to-multipoint
D. nonbroadcast
E. point-to-point
Answer: B,E (LEAVE A REPLY)
IS-IS supports two fundamental circuit types: broadcast multiaccess and point-to-point. A broadcast or multiaccess circuit is used on media such as Ethernet where multiple routers can share the same Layer 2 segment. In that case, IS-IS elects a designated intermediate system to simplify database synchronization on the shared network. A point-to-point circuit is used where exactly two IS-IS routers form an adjacency across a direct link, which is common in modern routed campus, data center, and WAN designs. IS-IS does not use the same nonbroadcast multiaccess and point-to-multipoint network-type model that is familiar from OSPF.
When IS-IS must run across technologies that do not naturally behave like Ethernet broadcast media, the design is typically adapted with point-to-point subinterfaces or appropriate Layer 2 treatment rather than selecting an IS-IS NBMA circuit type. Therefore, from the listed choices, the supported IS-IS circuit types are multiaccess and point-to-point.

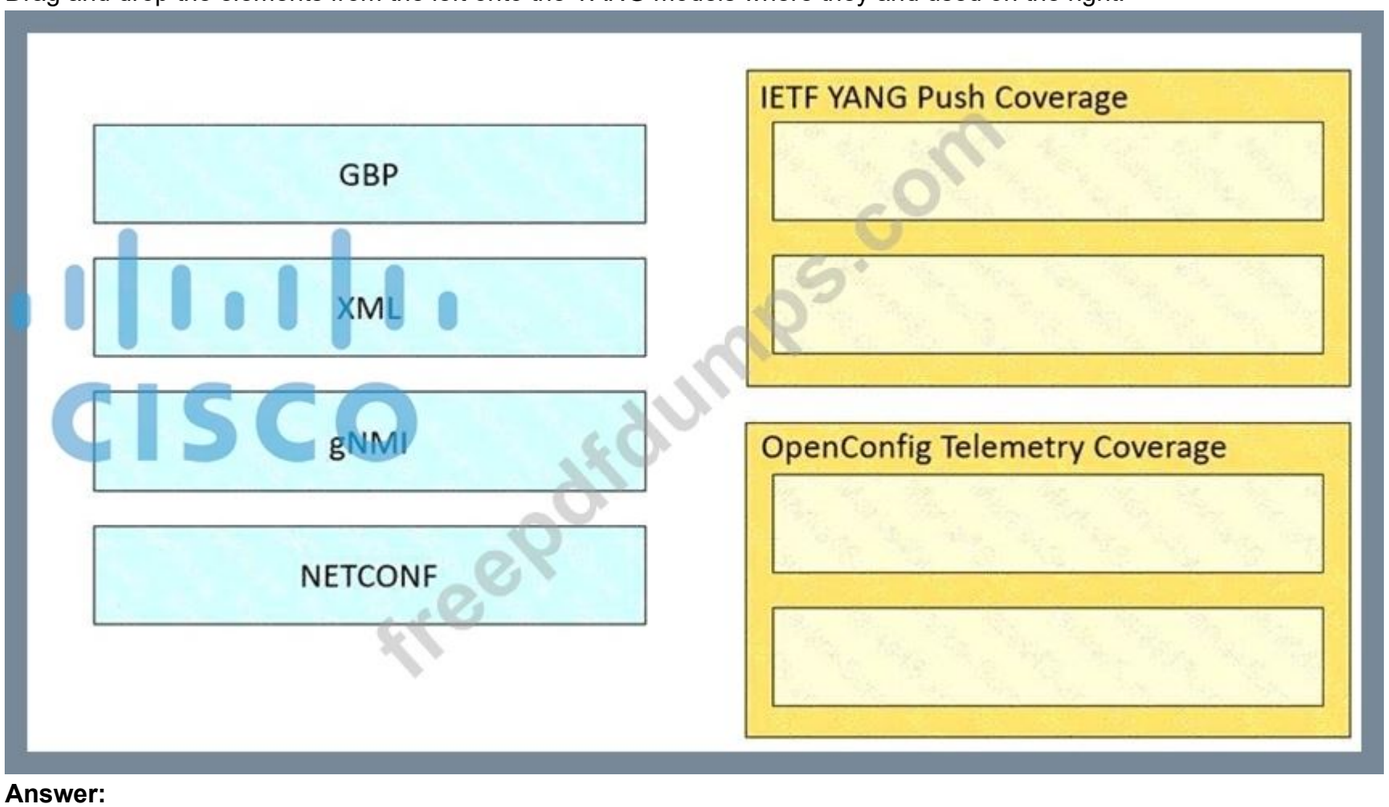
NEW QUESTION: 121
How are wireless endpoints registered in the HTDB in a Cisco SD-Access architecture?
A. Fabric edge nodes update the HTDB based on CAPPWAP messaging from the AP
B. Fabric WLCs update the HTDB as new clients connect to the wireless network
C. Border nodes first register endpoints and then update the HTDB
D. Fabric APs update the HTDB with the clients ' EID and RLOC
Answer: B (LEAVE A REPLY)
Wireless endpoints in Cisco SD-Access are registered through the fabric wireless integration process, and the fabric WLC is the component that updates endpoint information as clients join the wireless network. The Host Tracking Database stores endpoint identifier and location information used by the SD-Access control plane.
For wired endpoints, the fabric edge detects the endpoint and registers the binding. For wireless endpoints, the fabric WLC is integrated with the fabric and has the client session context, so it provides the endpoint information needed for HTDB updates. Border nodes do not first register ordinary wireless endpoints; their role is to connect the fabric to external networks and other domains. Fabric APs also do not own the complete EID-to-RLOC registration function in the way the controller does. The design implication is that fabric wireless requires correct WLC integration with Cisco Catalyst Center, fabric site configuration, control-plane reachability, and policy integration through Cisco ISE. Client registration, mobility, and policy enforcement depend on this controller-to-fabric relationship being stable. Reference topics: SD-Access wireless, fabric WLC, Host Tracking Database, endpoint registration, LISP control plane.

Valid 300-420 Dumps shared by Actual4test.com for Helping Passing 300-420 Exam! Actual4test.com now offer the **newest 300-420 exam dumps**, the Actual4test.com 300-420 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-420 dumps with Test Engine here: https://www.actual4test.com/300-420_examcollection.html (**457** Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

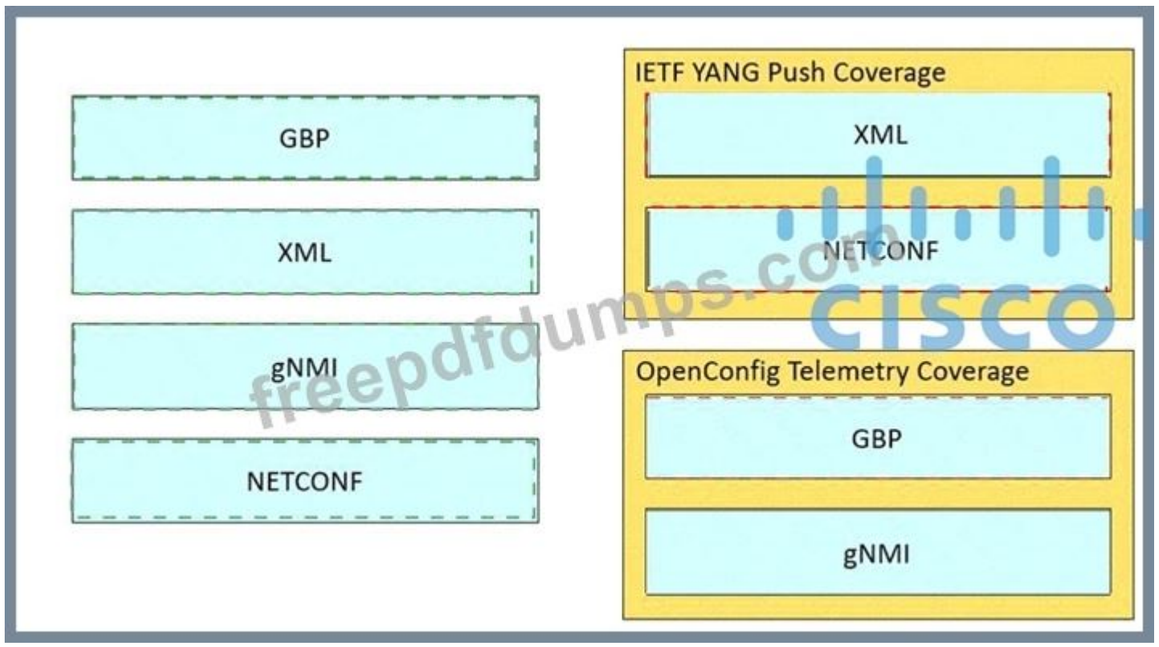
NEW QUESTION: 122
A customer plans to deploy WoL in the enterprise with these high-level design requirements:
DHCP services must be available.
Clients BIOS settings must be set for WoL.
Clients get IP addresses once online.

Spanning-tree PortFast is enabled on the Layer 2 switches.
Which two solutions must the customer select to have a successful deployment? (Choose two.)
A. IP directed broadcast and forward-protocol must be enabled on all the SVI or routed interlaces where the client subnets reside.
B. IP helper-addresses for the client ranges must be enabled on the SVI or routed interface where the WoL server subnet resides.
C. IP helper-addresses for the client ranges must be disabled on the SVI or routed interface where the WoL server subnet resides
D. IP helper-addresses for the WoL server must be enabled on the SVI or routed interface where the client subnets reside
E. IP directed broadcast and forward-protocol must be disabled on all the SVI or routed interfaces where the client subnets reside.
Answer: (SHOW ANSWER)
Wake-on-LAN across routed boundaries requires directed broadcast handling toward the client subnet and helper behavior to forward the magic packet correctly. A sleeping host normally does not maintain a normal IP session, so the WoL magic packet is commonly sent as a directed broadcast to the destination subnet. Cisco designs therefore enable IP directed broadcast on the routed interface or SVI for the client subnet where the sleeping endpoints reside, while controlling the feature carefully to avoid abuse. The ip forward-protocol behavior is also relevant because routers must forward the required UDP broadcast traffic rather than dropping it at the Layer 3 boundary. In addition, an IP helper address toward the WoL server or relay path may be needed on client-facing interfaces so the magic packet can be delivered from the management or server subnet to the client subnet. PortFast helps endpoints become reachable quickly at Layer 2, but it does not replace routed broadcast forwarding. Reference topics: Wake-on-LAN, directed broadcast, ip helper- address, UDP forwarding, PortFast, campus services.

NEW QUESTION: 123
Drag and drop the elements from the left onto the YANG models where they and used on the right.



Answer:



Explanation:

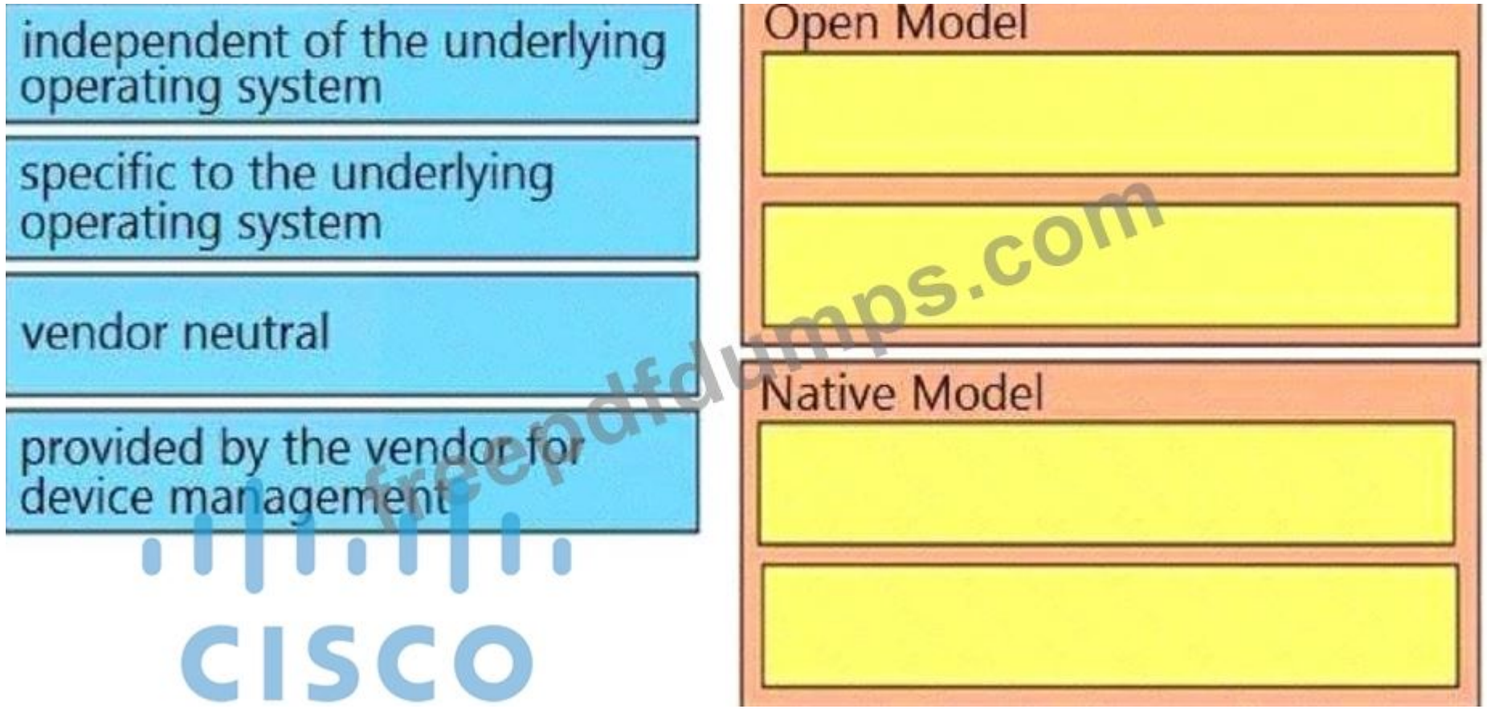
The drag-and-drop answer follows the practical distinction between the YANG model families. IETF models are standards-based models intended to provide a common structure for widely implemented protocol and interface functions. OpenConfig models are vendor-neutral operational and configuration models developed for cross-vendor automation, with a strong focus on consistent telemetry and network state representation.

Cisco native models expose Cisco platform-specific capabilities and usually provide the deepest coverage for IOS XE, IOS XR, or NX-OS features that are not fully represented in a common model. The selected mapping therefore assigns model elements according to origin, portability, and feature depth. In a production automation design, engineers should start with open models when interoperability and long-term portability matter, but they should choose Cisco native models when a required Cisco feature is not available or is incomplete in IETF or OpenConfig. The choice is not cosmetic; it affects payload structure, namespace, validation behavior, and operational consistency across device families. Reference topics: YANG model families, IETF YANG, OpenConfig, Cisco native YANG, model-driven configuration.

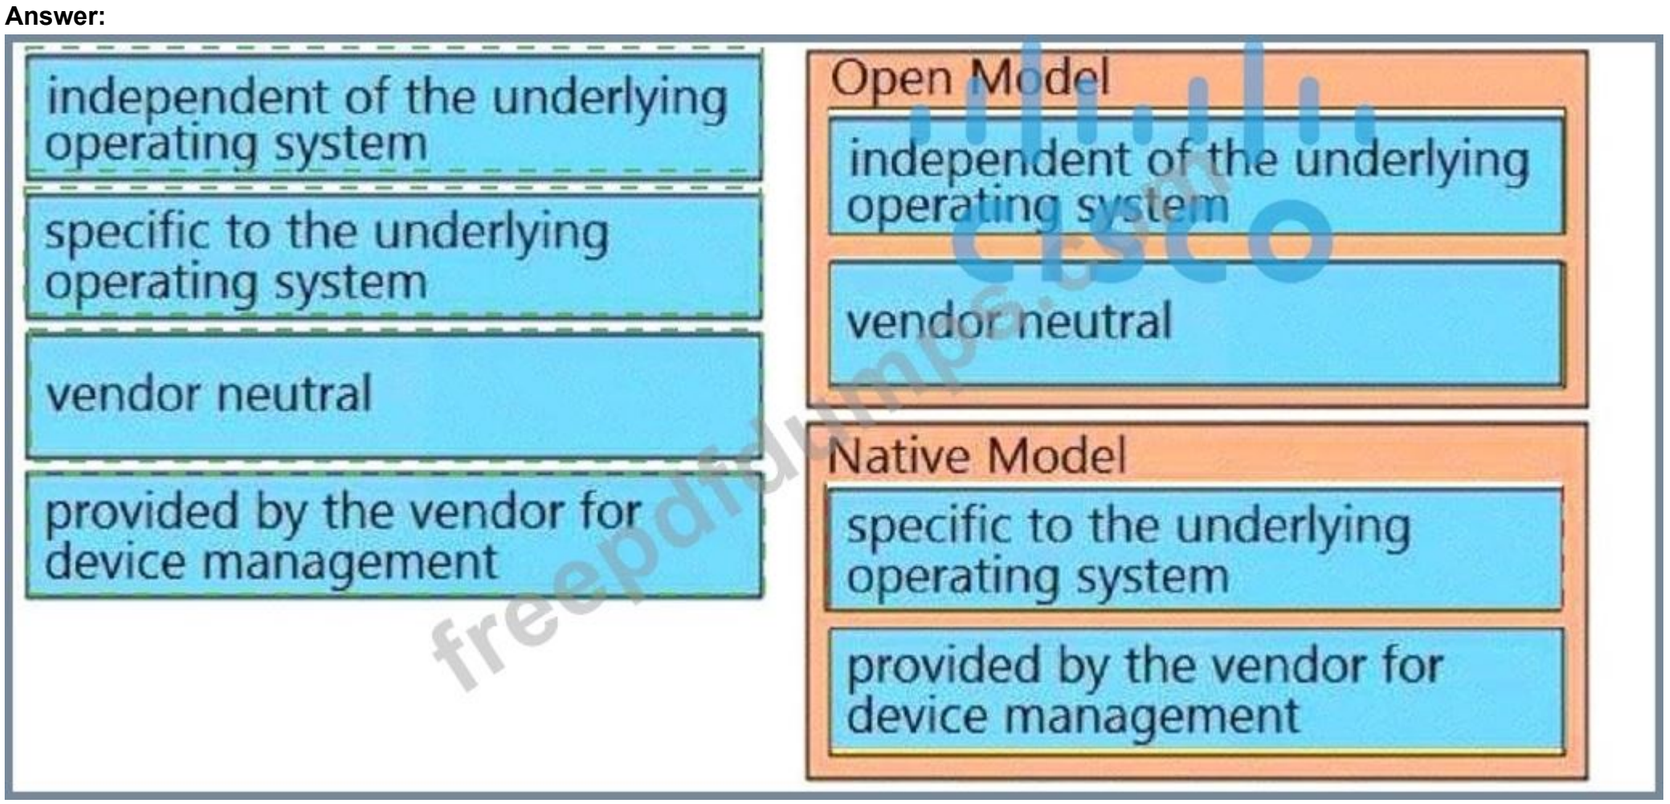
NEW QUESTION: 124

Drag and drop the characteristics from the left onto the Yang model they describe on the right.

Select and Place:



Answer:

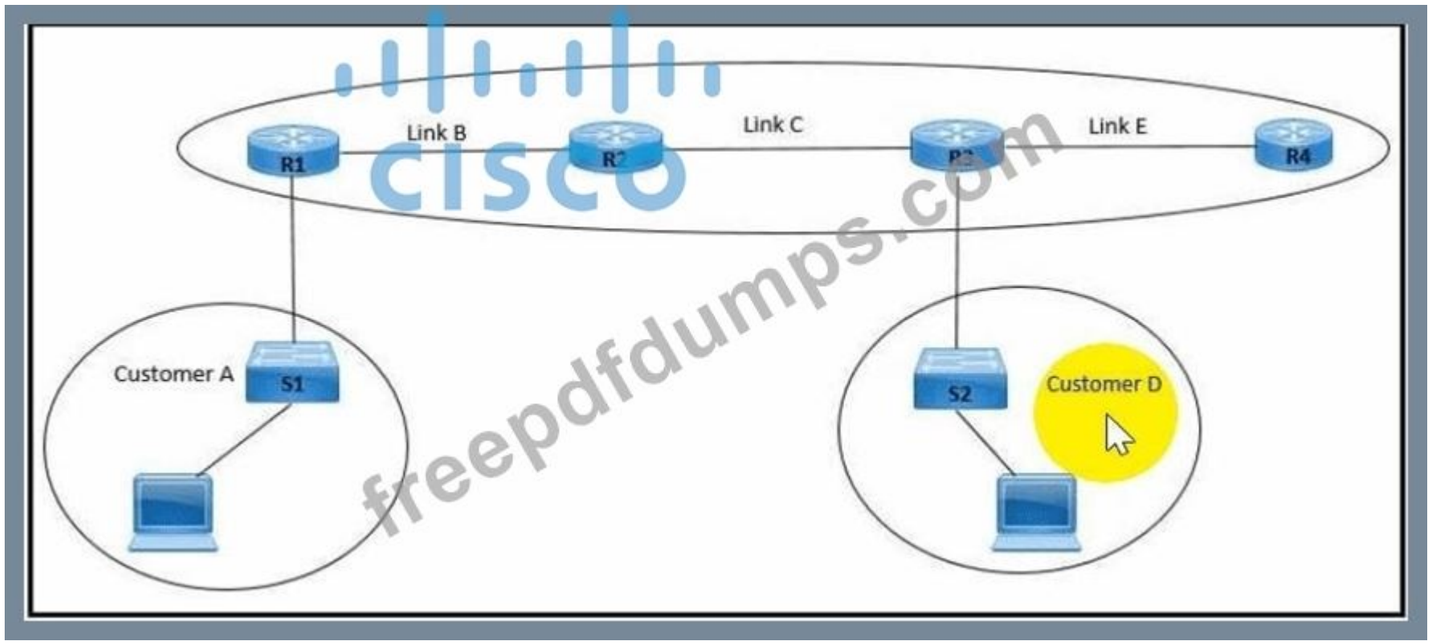


Explanation:

The drag-and-drop mapping is based on how Cisco differentiates YANG model families used for model- driven programmability. IETF models are standards-based and are intended for broad interoperability

across vendors. They are the usual choice when the design priority is a consistent baseline for multivendor configuration or operational state. OpenConfig models are also vendor-neutral, but they are developed by the OpenConfig community and often provide an operationally consistent structure for telemetry and configuration across platforms. Cisco native models are vendor-specific models that expose detailed Cisco IOS XE, IOS XR, or NX-OS features, including capabilities that may not be represented in IETF or OpenConfig models. The design choice is therefore not simply "one model is always best." It depends on whether the requirement prioritizes portability, abstraction, or complete Cisco feature coverage. For a standards-driven multivendor design, use IETF or OpenConfig where sufficient. For Cisco-specific device features, use Cisco native models. Reference topics: YANG models, IETF models, OpenConfig, Cisco native models, NETCONF/RESTCONF programmability.

NEW QUESTION: 125



Refer to the exhibit An architect is designing an IPv4 plan using the 172 20 0.0/16 network The design must maximize the number of subnets and minimize the number of wasted IP addresses In addition, the plan must allocate a subnet to these customers and links

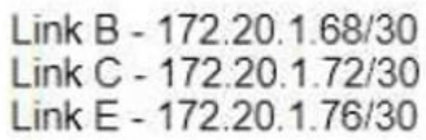
* Customer A, which supports 125 hosts
* Customer D, which supports 62 hosts
* Links B C. and E

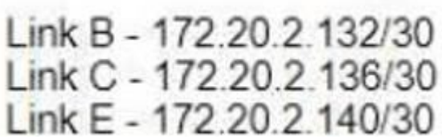
Which two configuration sets meet these requirements ' ? (Choose two)

A.
Customer A - 172.20.0.128/25
Customer D - 172.20.1.0/26

B.
Link B - 172.20.1.70/30
Link C - 172.20.1.74/30
Link E - 172.20.1.78/30

C.
Customer A - 172.20.1.0/24
Customer D - 172.20.2.64/26

D.


E.


Answer: A,C (LEAVE A REPLY)

The selected configuration sets are correct because the addressing plan must use VLSM to maximize subnets and minimize wasted addresses. Customer A requires 125 hosts, so it needs a /25 subnet, which provides 126 usable host addresses. Customer D requires 62 hosts, so it needs a /26 subnet, which provides 62 usable addresses exactly. Point-to-point or small WAN links B, C, and E should use the smallest appropriate link subnets shown in the exhibit, typically /30 for IPv4 point-to-point links when two usable addresses are required. A correct plan must allocate the largest blocks first, align every subnet on the proper boundary, avoid overlap, and avoid assigning a larger-than-needed subnet to any customer or link. Options A and C are the two valid combinations that satisfy those rules in the exhibit. Any option that places a /25 or /26 on the wrong boundary, overlaps with another block, or wastes larger subnets on links would fail the design objective. Reference topics: IPv4 VLSM, subnet boundary alignment, host capacity planning, point-to-point addressing, route summarization readiness.

NEW QUESTION: 126

An engineer is creating a design to enable IPv6 to run on an existing IPv4 IS-IS network. The IPv4 and IPv6 topologies will match exactly, and the engineer plans to use the same router levels for each protocol per interface. Which IS-IS design is required?

A. single topology without enabling transition feature

B. single topology with transition feature enabled

C. multi topology with transition feature enabled

D. multi topology without enabling transition feature

Answer: B (LEAVE A REPLY)

Single-topology IS-IS with the transition feature is the correct design when IPv6 is being added to an existing IPv4 IS-IS network and both topologies must match exactly. In single-topology mode, IS-IS calculates one topology and uses it for both IPv4 and IPv6, which fits the requirement that the IPv4 and IPv6 paths and router levels remain the same on each interface. Cisco documentation notes that single-topology IS-IS IPv6 requires routers to run the same set of address families for adjacency consistency; during migration, transition support allows the network to operate while IPv6 is introduced gradually. Multi-topology IS-IS is used when IPv4 and IPv6 should be allowed to follow different topologies or when not all routers support the same address families. The question explicitly says the topologies will match exactly, so multi-topology is unnecessary. Single topology without the transition feature is risky during phased migration because routers that do not yet run both address families may fail adjacency checks or create blackholing. Reference topics:
IS-IS for IPv6, single-topology IS-IS, transition mode, address-family consistency, IPv6 migration.

NEW QUESTION: 127

What is a feature of the SaaS subscription model?

A. web connection not required

B. access to industrial-strength storage and computing power

C. autonomy and control over hardware

D. tower initial costs

Answer: D (LEAVE A REPLY)

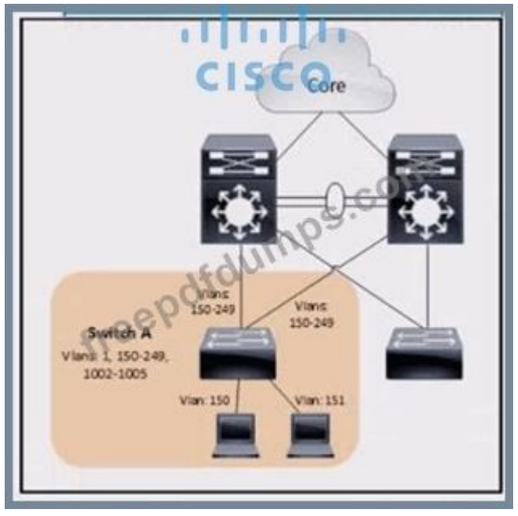
A major feature of the SaaS subscription model is lower initial cost. In Software as a Service, the customer consumes an application provided by a service provider rather than purchasing and operating the full application infrastructure directly. This generally shifts spending from large upfront capital expenditure to recurring operational subscription cost. The provider handles much of the platform, application

maintenance, hosting, scalability, and upgrades, which lowers the customer's initial hardware and software investment. A web connection is normally required for SaaS access, so option A is incorrect. Access to industrial-strength storage and computing power is more characteristic of cloud infrastructure benefits, but it is not the defining SaaS subscription feature in this answer set. Autonomy and control over hardware is reduced, not increased, in SaaS because the provider owns and operates the underlying infrastructure. The option text contains a typo, but "lower initial costs" is the intended and correct feature. Reference topics: SaaS model, cloud subscription economics, operational expenditure, application hosting, cloud service models.

NEW QUESTION: 128
Which OSPF area blocks LSA Type 3, 4 and 5, but allows a default summary route?
A. normal
B. stub
C. NSSA
D. totally stubby
Answer: D (LEAVE A REPLY)
A totally stubby OSPF area blocks Type 3, Type 4, and Type 5 LSAs while allowing a default summary route into the area. In OSPF, Type 3 LSAs are interarea summary LSAs, Type 4 LSAs describe reachability to ASBRs, and Type 5 LSAs carry external routes. A normal area accepts all of these LSA types. A standard stub area blocks Type 5 external LSAs and normally Type 4 ASBR summaries, but it still allows Type 3 interarea summaries. An NSSA blocks normal Type 5 external LSAs but allows Type 7 LSAs for locally redistributed external routes. A totally stubby area is more restrictive: the ABR suppresses interarea summaries and external routing detail, then injects a default route so routers inside the area still have reachability outside the area. This design reduces LSDB size, memory consumption, and SPF processing on low-end routers in remote areas. Therefore, the OSPF area type that blocks Type 3, 4, and 5 LSAs but allows a default summary route is a totally stubby area.

NEW QUESTION: 129
Which two considerations must be made regarding the overlay network for a Cisco SD-Access architecture?
(Choose two.)
A. Virtual networks should be used for microsegmentation
B. SGTs should be used for data plane isolation and microsegmentation
C. Virtual networks should be used for data plane isolation only
D. Overlapping IP addresses across different overlay networks should be used to conserve IP addresses
E. Overlapping IP addresses across different overlay networks should be avoided for operational simplicity
Answer: C,E (LEAVE A REPLY)
The correct overlay considerations are that virtual networks provide data-plane isolation and that overlapping IP addresses should be avoided for operational simplicity. In Cisco SD-Access, virtual networks are used for macro segmentation and are commonly mapped to VRFs. This separates routing and forwarding tables between major groups such as corporate users, guests, IoT, and shared services. Security Group Tags provide finer policy control and microsegmentation, but they are not the primary mechanism for data-plane isolation between virtual networks. The phrase "virtual networks should be used for microsegmentation" is therefore inaccurate; microsegmentation is SGT and SGACL driven. Although overlapping IP addresses can technically exist in separate VRF-like contexts, Cisco design practice discourages unnecessary overlap because it complicates troubleshooting, service insertion, external connectivity, and inter-VN communication. A clean overlay address plan improves operations and reduces policy ambiguity. Reference topics: SD-Access overlay design, virtual networks, macro segmentation, SGTs, microsegmentation, overlapping IP avoidance. This also simplifies day-two operations.

NEW QUESTION: 130

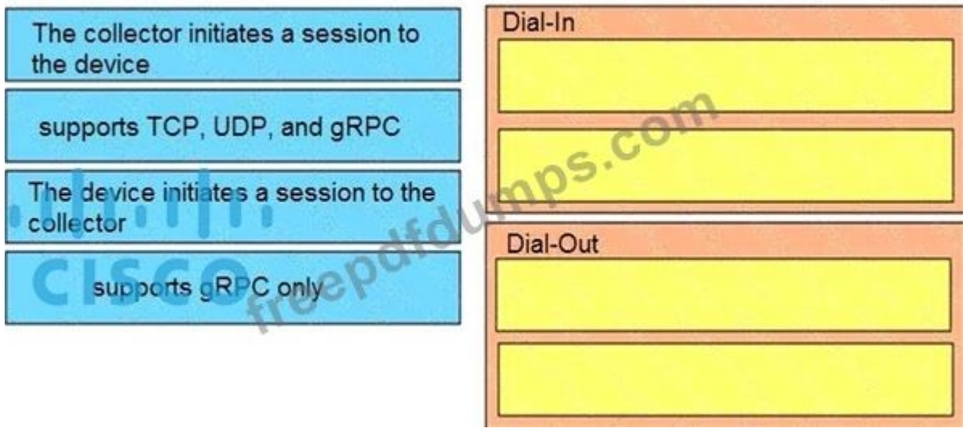


Refer to the exhibit An engineer working for a telecommunication company with an employee ID 4449:30
959 Is calculating STP scalability for switches to ensure that the numbers are below the maximum supported value for STP logical ports How many logical interfaces are active for switch A?

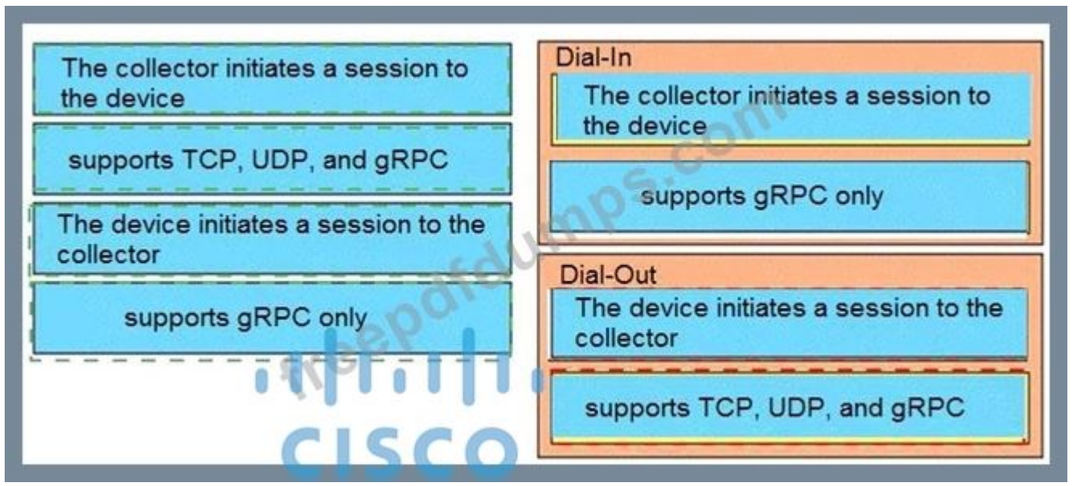
A. 4
B. 307
C. 202
D. 100
Answer: C (LEAVE A REPLY)

The active logical interface count for Switch A is 202 based on the spanning-tree scalability method used for logical ports. In STP design, a logical port is counted per active VLAN instance on each active trunk or access interface. A trunk carrying multiple VLANs therefore consumes multiple logical STP ports, while an access link normally consumes one logical STP port for its VLAN. Cisco campus design guidance uses logical-port counts to ensure that the selected STP mode and platform stay within supported scalability limits. The calculation in the exhibit produces 202 active logical interfaces for Switch A, not simply the number of physical links. Answer A counts only physical interfaces and ignores VLAN instances. Answer D is too low because it does not account for all active VLANs across the relevant links. Answer B overstates the count by including VLANs or links that are not active for Switch A. Reference topics: STP scalability, logical ports, VLAN instances, trunk design, Layer 2 campus sizing.

NEW QUESTION: 131
Drag and drop the characteristics from the left onto the correct telemetry mode on the right.



Answer:



Explanation:

Dial-In: collector initiates a session to the device; supports gRPC only. Dial-Out: device initiates a session to the collector; supports TCP, UDP, and gRPC.

Model-driven telemetry can operate in dial-in or dial-out mode, and the distinction is based on which side initiates the session and how the subscription is established. In dial-in telemetry, the collector initiates a session to the network device and dynamically subscribes to the required sensor paths. This is useful when the collector wants direct control over the subscription lifecycle. Cisco telemetry documentation commonly associates dial-in with gRPC-based dynamic subscriptions on supported platforms. In dial-out telemetry, the device initiates the session toward the collector based on a configured subscription. The subscription exists in the device configuration, and the device streams telemetry data to the configured destination without requiring the collector to start the connection. This design is useful when devices are behind NAT, firewalls, or operational models where the network device should push telemetry to a known collector. Therefore, the mapping is: dial-in equals collector-initiated and gRPC-only in this context; dial-out equals device-initiated and supports the listed transport options. This distinction is central to Cisco model-driven telemetry design.

NEW QUESTION: 132

Which function are fabric intermediate nodes responsible for in an SD-Access Architecture?

A. mapping EIDs to RLOCs

B. encapsulating user traffic in a VXLAN header including the SGT

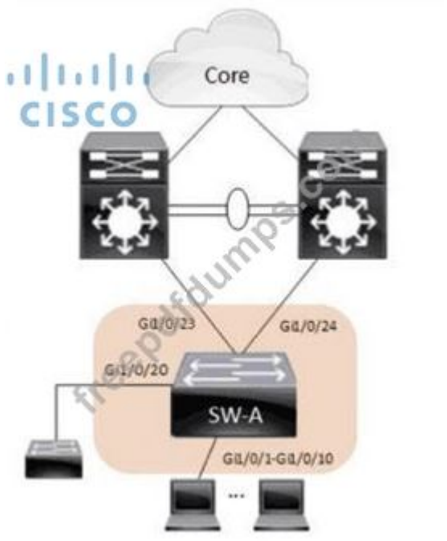
C. registering new endpoints in the HTDB

D. transporting IP packets between edge nodes and border nodes

Answer: D (LEAVE A REPLY)

In Cisco SD-Access, an intermediate node is an underlay transit device. Its responsibility is to provide IP reachability between fabric edge nodes, border nodes, and other fabric infrastructure roles. Intermediate nodes forward routed underlay traffic and do not perform endpoint registration, endpoint lookup, or VXLAN encapsulation for user traffic. Endpoint-to-location mapping is handled through the fabric control plane, which is based on LISP map-server and map-resolver functions. Fabric edge nodes identify connected endpoints, apply policy, act as anycast gateways, and encapsulate user traffic into VXLAN. Border nodes connect the fabric to external networks and handle traffic entering or leaving the fabric. The intermediate node is therefore comparable to a routed distribution or core transport device in a traditional hierarchical campus, except it is participating in the SD-Access underlay. It must provide stable Layer 3 reachability and appropriate MTU for encapsulated traffic but does not maintain the host tracking database or add security group information to VXLAN headers. The correct function is transporting IP packets between edge and border nodes.

NEW QUESTION: 133



Refer to the exhibit. An architect reviews the low-level design of a company ' s enterprise network and advises optimizing the STP convergence time. Which functionality must be to Gi1/0/1-10 to follow the architect ' s recommendation?

A. PortFast
B. root guard
C. UplinkFast
D. BPDU guard

Answer: (SHOW ANSWER)

PortFast is the correct functionality for the Gi1/0/1-10 interfaces when those ports are endpoint-facing access ports. Cisco spanning-tree design uses PortFast to allow edge ports to transition directly to the forwarding state, avoiding the normal listening and learning delay associated with traditional STP convergence. That improves user and endpoint connectivity after link-up events and is especially important for workstations, phones, printers, and access devices that should not participate in the Layer 2 topology. Root guard and BPDU guard are protection features, not the primary convergence optimization requested by the architect.

UplinkFast was used in older STP designs to accelerate recovery for access switches with redundant uplinks, but it is not the edge-port functionality for interfaces Gi1/0/1-10. In a modern design, PortFast should be enabled only on true host-facing ports, normally together with BPDU guard to protect against accidental switch attachment. If the exhibit shows these interfaces as end-user access ports, PortFast is the feature that follows the recommendation to optimize STP convergence time. Reference topics: STP PortFast, edge-port convergence, BPDU guard, Layer 2 campus design.

NEW QUESTION: 134

In a Cisco SD-Access fabric, switch node Is equivalent to an access layer switch In a traditional three-tier campus network design?

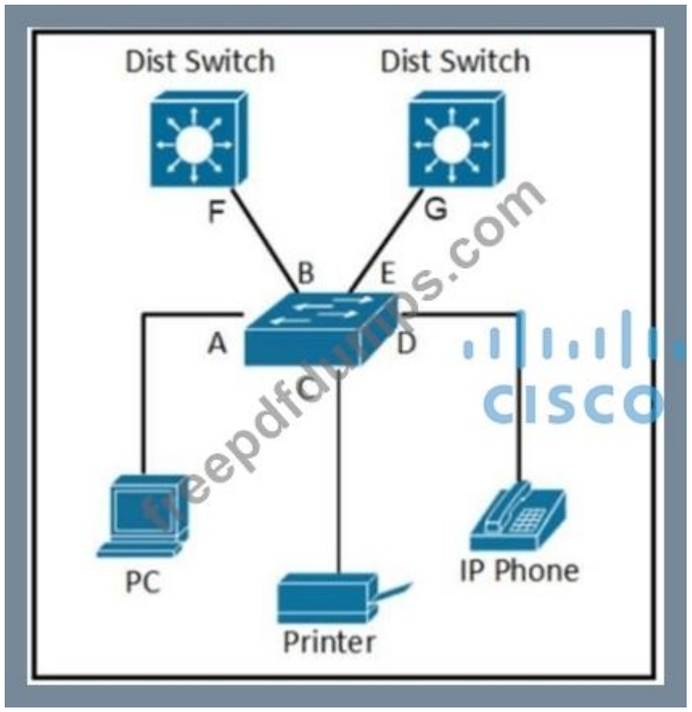
A. edge node
B. border node
C. intermediate node
D. control plane node

Answer: (SHOW ANSWER)

The Cisco SD-Access edge node is equivalent to the access-layer switch in a traditional three-tier campus design. Edge nodes are the fabric devices to which wired endpoints, access points, and other user-facing connections attach. They provide the anycast Layer 3 gateway for connected subnets, classify endpoints into virtual networks, apply policy information, and register endpoint reachability with the SD-Access control- plane node. In traditional campus terminology, this is the role normally performed by access switches at the edge of the network. Border nodes connect the SD-Access fabric to external networks, such as the WAN, data center, internet edge, or shared services. Intermediate nodes provide underlay transport between fabric nodes but do not attach endpoints to the fabric overlay. Control-plane nodes maintain mapping information rather than serving as access-layer attachment points. Understanding the edge-node role is fundamental because SD- Access policy and host mobility begin

where the endpoint enters the fabric. Reference topics: SD-Access edge node, access layer, anycast gateway, endpoint registration, fabric roles.

NEW QUESTION: 135



Refer to the exhibit. Which two points in the network must an engineer configure the ports for explicit trust when using a DiffServ model?

A. B and E

B. F and G

C. A and D

D. C and D

Answer: (SHOW ANSWER)

Explicit QoS trust should be configured at the network boundary points where markings first enter the managed DiffServ domain and are known to be reliable. In a DiffServ design, traffic classification and marking are typically performed at the access edge or at a controlled ingress boundary, and downstream devices then honor DSCP or CoS markings based on the trust boundary. Cisco QoS design stresses that the trust boundary must be intentionally placed; endpoints or unmanaged devices should not be trusted by default because they can mark traffic incorrectly and steal priority treatment. The exhibit answer identifies B and E as the correct trust points, meaning those are the controlled ingress locations where markings should be accepted. Options that trust uncontrolled endpoint-facing points or internal transit locations would either create a security problem or add no practical value. Once traffic is trusted and marked at the correct boundary, queuing, policing, and shaping policies can use the DSCP values consistently across the campus or WAN.

Reference topics: DiffServ, QoS trust boundary, DSCP marking, classification, enterprise QoS design.

NEW QUESTION: 136



Refer to the exhibit A customer requires maximum uptime for the data plane between R1 and R3 running OSPF Which solution must the design include for high availability if the routing process on R2 requires maintenance?

A. BFD on all routers

B. nonstop forwarding on R1 and R3

C. nonstop forwarding on R3 only

D. graceful restart on all routers

Answer: D (LEAVE A REPLY)

Graceful restart is the correct high-availability design when the routing process on an intermediate OSPF router must undergo maintenance while the data plane remains available. Cisco describes graceful restart as a mechanism that allows a restarting router to continue forwarding traffic while neighbors temporarily preserve adjacency and route information. In OSPF, helper-capable neighbors assist the restarting router by maintaining forwarding state during the restart interval, preventing unnecessary route withdrawal and reconvergence. BFD is a failure-detection mechanism, so it would actually accelerate detection of a failure rather than preserve the forwarding path during planned process maintenance. Nonstop forwarding must be enabled on the device performing the control-plane restart and supported by neighbors, but the broad answer choice that captures the required routing-protocol behavior across the design is graceful restart on the participating routers. Configuring it only on R1 and R3 or only on R3 would not support the R2 process maintenance scenario. Reference topics: OSPF graceful restart, NSF, helper mode, control-plane maintenance, data-plane availability.

Valid 300-420 Dumps shared by Actual4test.com for Helping Passing 300-420 Exam! Actual4test.com now offer the **newest 300-420 exam dumps**, the Actual4test.com 300-420 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-420 dumps with Test Engine here: https://www.actual4test.com/300-420_examcollection.html (**457** Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 137

What is the purpose of Cisco vBond as a Session Traversal Utilities for NAT server?

A. allow Cisco Catalyst SD-WAN routers to locate their own mapped IP addresses

B. integrate Cisco SD-Access Wireless into the fabric

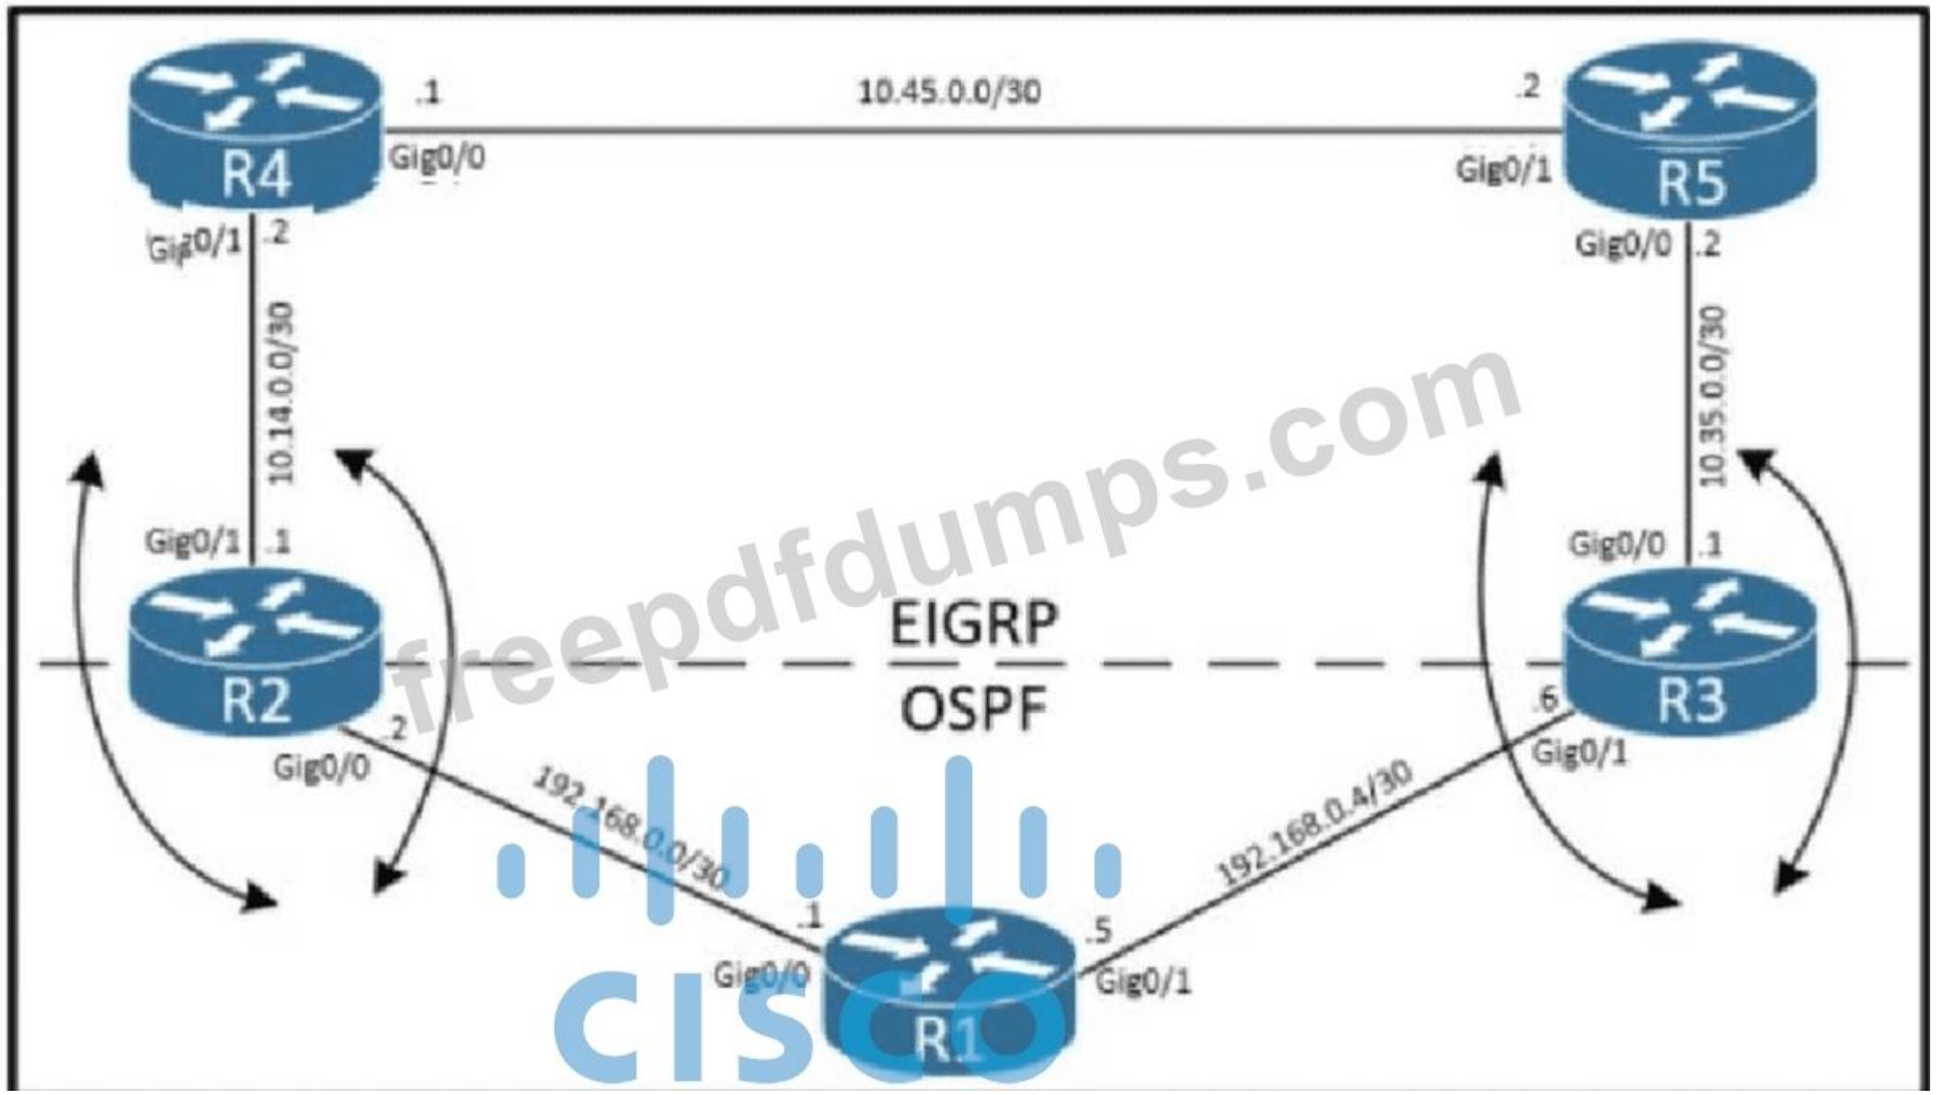
C. secure data traffic between Cisco Catalyst SD-WAN edge routers that use IPsec

D. provide Zero-Touch Provisioning to Cisco Catalyst SD-WAN vEdge devices

Answer: A (LEAVE A REPLY)

As a STUN server, Cisco vBond allows Cisco Catalyst SD-WAN routers to discover their public mapped IP addresses and ports when they sit behind NAT. This function is essential during overlay bring-up because many WAN Edge devices are deployed behind internet circuits, firewalls, or provider NAT devices. vBond helps devices understand how they are seen from the public transport side and facilitates the formation of secure control connections to the SD-WAN controllers. vBond also participates in orchestration and Zero- Touch Provisioning workflows, but the question specifically asks about its purpose as a Session Traversal Utilities for NAT server. That wording points directly to NAT discovery, not general ZTP. vBond does not secure user data traffic between edge routers; data-plane encryption is established between WAN Edge devices. It also has no role in integrating SD-Access wireless into a fabric. Reference topics: Cisco SD-WAN vBond, STUN, NAT traversal, orchestration, WAN Edge onboarding. This is fundamental for internet transports.

NEW QUESTION: 138



Refer to the exhibit. As part of a design review of redistribution, a client requested that R2 be preferred over R3 for traffic passing toward the EIGRP domain. Which method meets this design requirement?

A. Redistribute EIGRP into OSPF with metric-type E1 on R2 and metric-type E2 on R3.

B. Remove the mutual redistribution on R3.

C. Redistribute OSPF into EIGRP with metric 10000 100 255 1 1500 on R2 and metric 10 1000 255 1 1500 on R3.

D. Redistribute EIGRP into OSPF with metric-type E2 on R2 and metric-type E1 on R3.

Answer: (SHOW ANSWER)

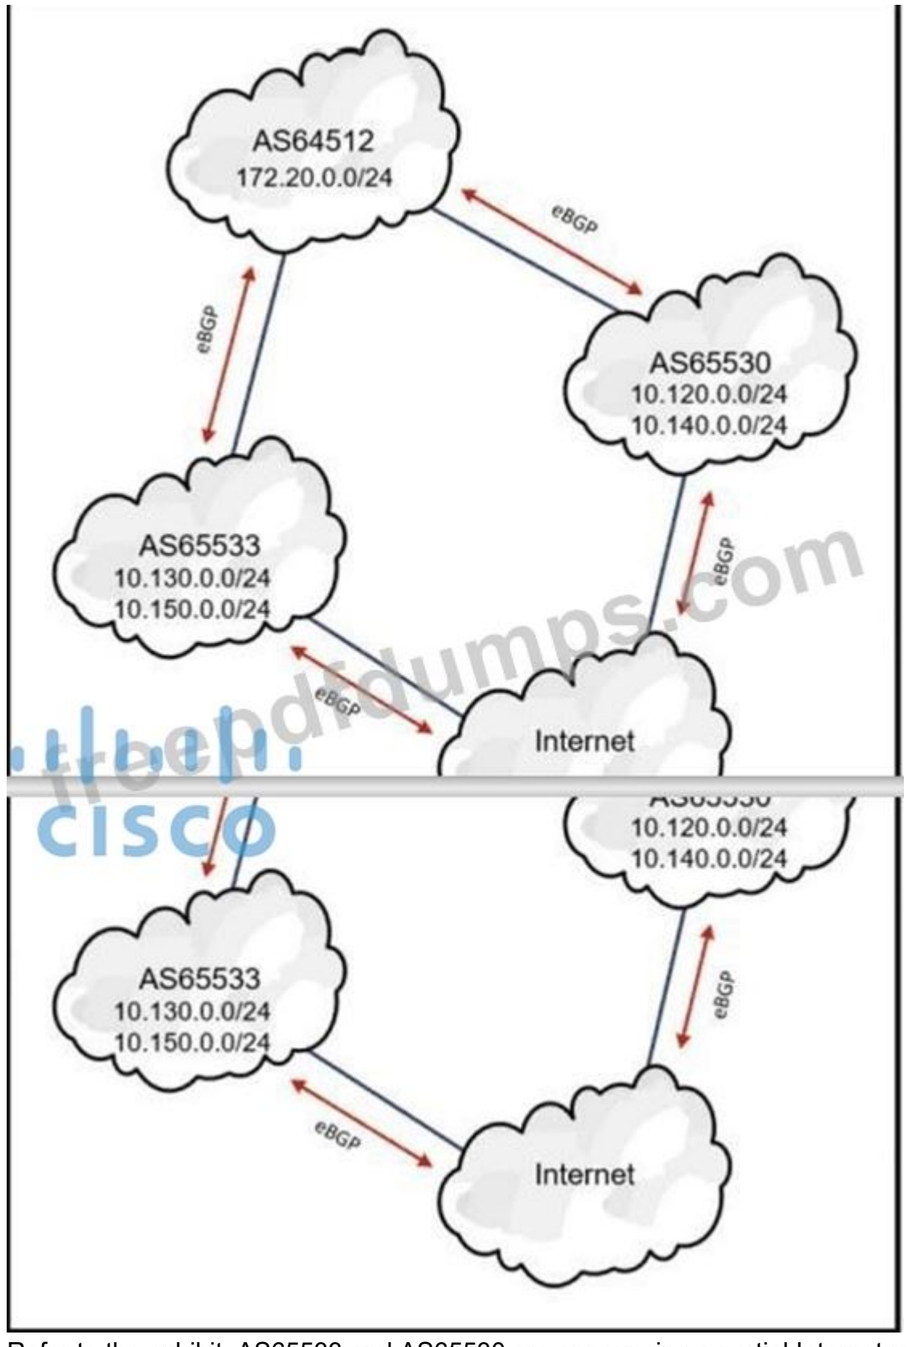
The design must influence how EIGRP-learned routes are injected into the OSPF domain, so R2 should redistribute EIGRP into OSPF as metric-type E1 and R3 as metric-type E2. Cisco OSPF documentation distinguishes external Type 1 and Type 2 routes. An E1 route includes the external cost plus the internal OSPF cost to reach the ASBR, while an E2 route uses only the external metric. Cisco path selection prefers O E1 over O E2 for the same external destination. Therefore, making R2 advertise the EIGRP routes as E1 causes OSPF routers to prefer R2 as the exit toward the EIGRP domain. Option C adjusts OSPF-to-EIGRP redistribution, which is the wrong direction for traffic passing toward EIGRP from OSPF. Removing redistribution on R3 would reduce redundancy. Option D reverses the preference and would favor R3.

freepdfdumps.com

Reference topics: OSPF external route types, EIGRP-to-OSPF redistribution, ASBR preference, E1 versus E2 metrics.

NEW QUESTION: 139
A company wants to switch from static routing to a dynamic routing protocol to ease the administrative and operational overhead. The network topology is hub and spoke, and the branches use DMVPN back to the hub using two 100 Mbps internet connections. Both links must be used due to spikes in traffic, and routing must take traffic utilization of the links into account. Also, the branch routers have limited memory and CPU resources. Which routing protocol and design solution must the company choose?
A. iBGP with the hub routers set up as route reflectors and branches set up as clients
B. OSPF deployed in area 0 with branch routers connecting from area 1
C. ISIS with the hub and spoke routers configured in two different areas
D. EIGRP with branch routers as stub routers using ECMP
Answer: D (LEAVE A REPLY)
EIGRP with branch routers configured as stubs and equal-cost multipath is the best fit for this hub-and-spoke DMVPN design. The branch routers have limited memory and CPU resources, so they should not be forced to maintain unnecessary topology information or answer broad queries from the rest of the network. EIGRP stub routing limits the routes a branch advertises and prevents it from being used as a transit path, which reduces query scope and improves stability. Because the two internet links are both 100 Mbps, equal-cost multipath can use both links for active forwarding during traffic spikes. OSPF can operate over DMVPN, but a single area or hub-and-spoke design may create more overhead and less efficient query containment for small branch routers. IS-IS is less common for this WAN overlay use case, and iBGP with route reflectors can scale but does not natively consider link utilization or provide the same simple branch stub behavior. Reference topics:
EIGRP stub, DMVPN, ECMP, branch routing, query containment, WAN convergence.

NEW QUESTION: 140



Refer to the exhibit. AS65533 and AS65530 are announcing a partial Internet routing table as well as their IP subnets. An architect must create a design that ensures AS64512 become a transit AS. Which filtering solution must the architect choose?

A. Maximum-prefix

B. No-advertise

C. Next-hop

D. No Export

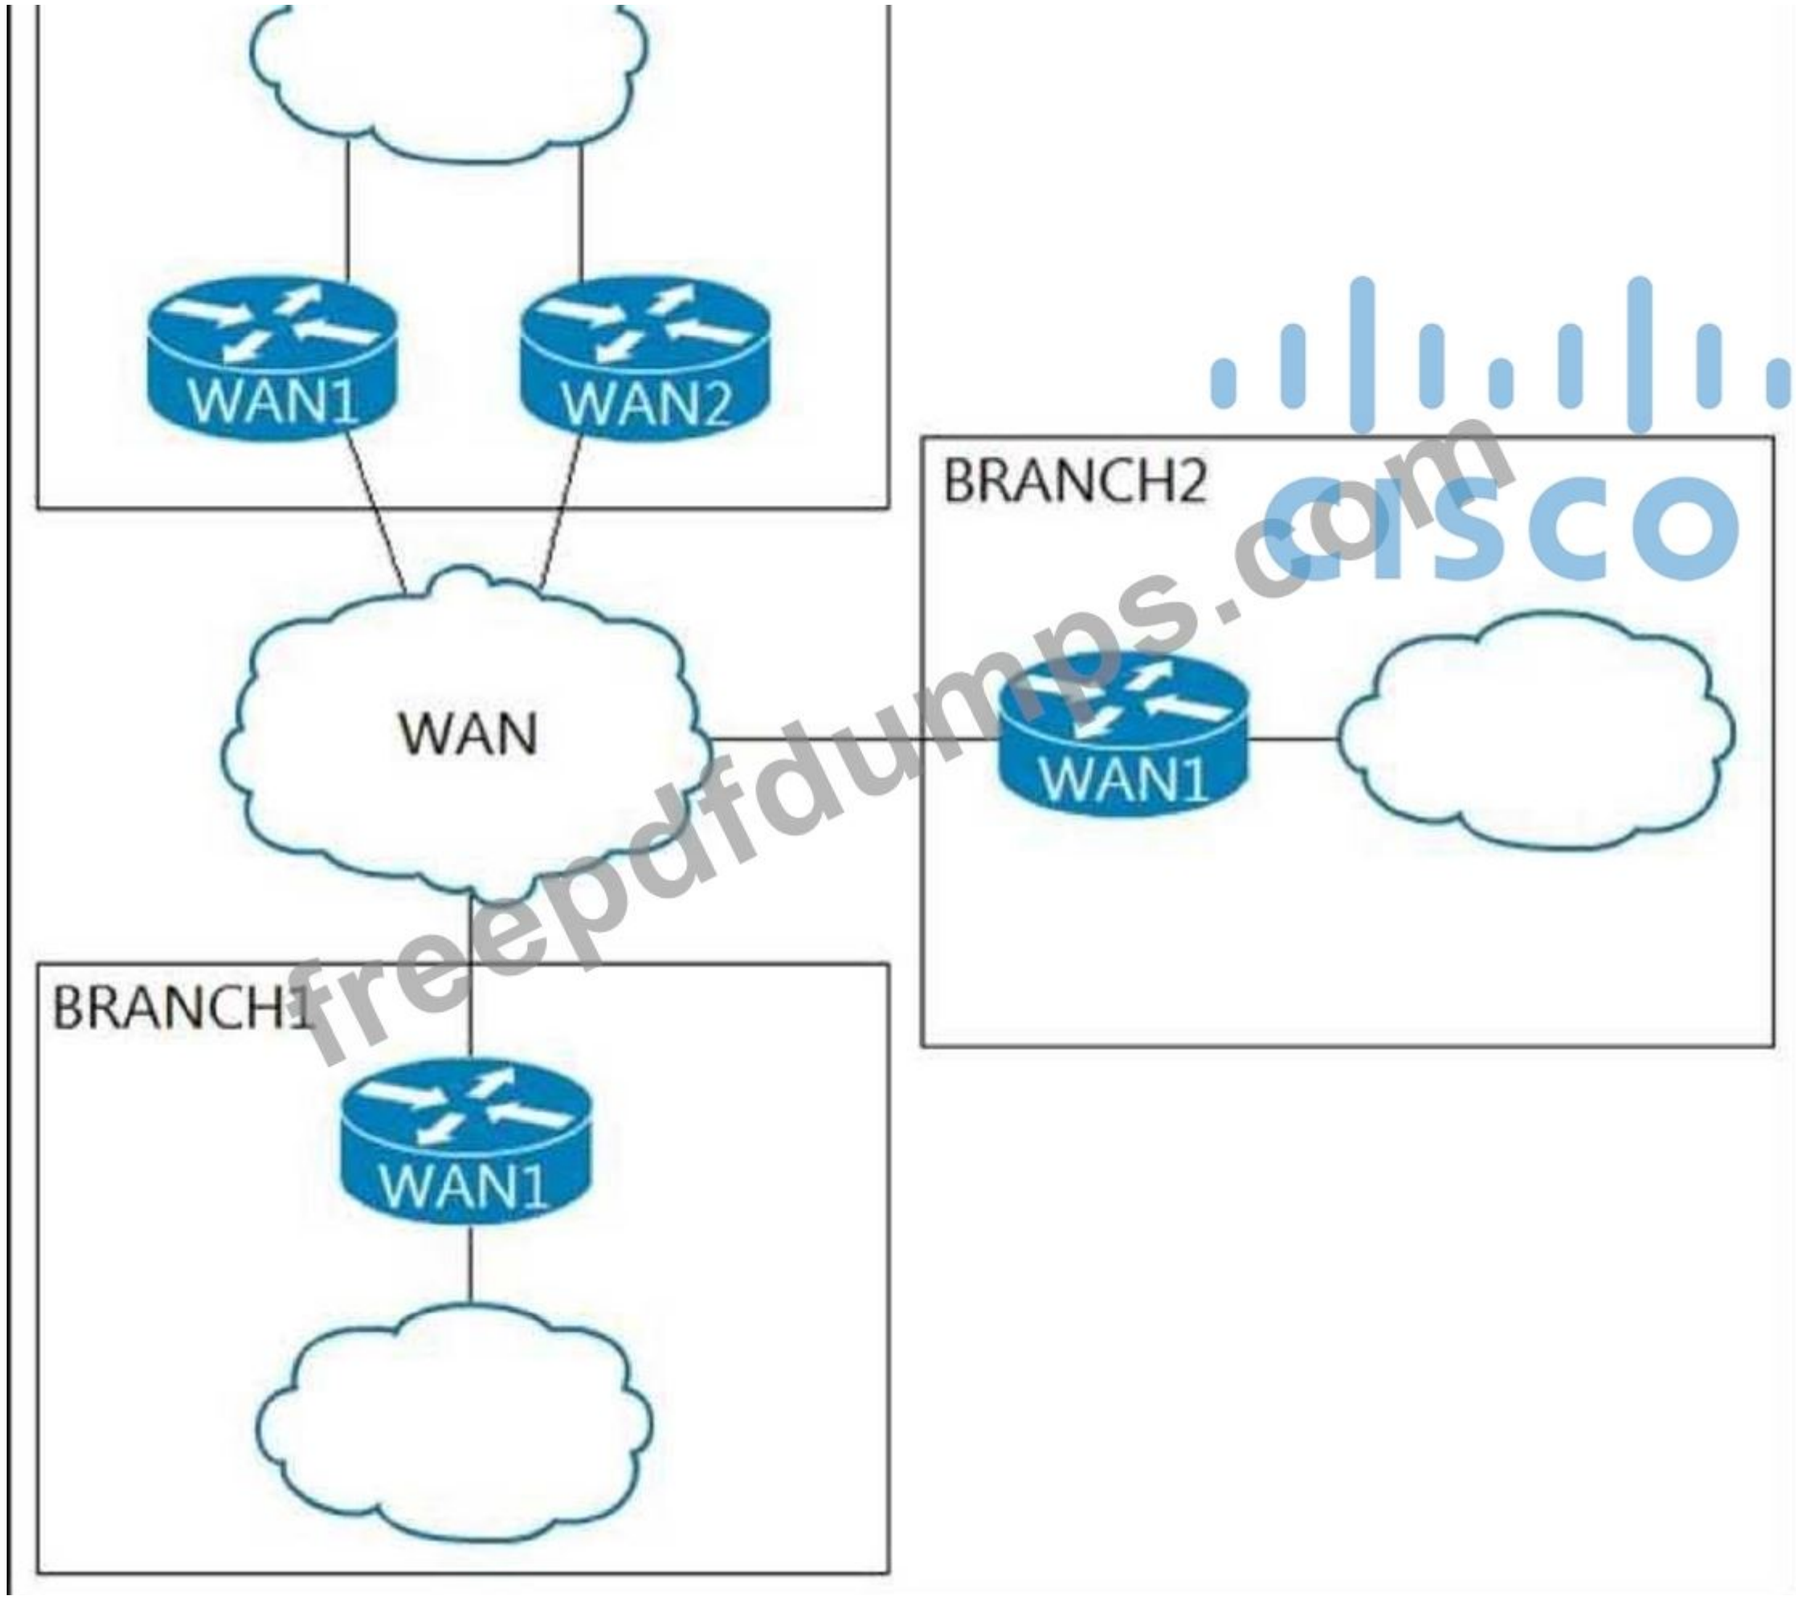
Answer: D (LEAVE A REPLY)

The No-Export community is the appropriate filtering mechanism when a design must prevent an autonomous system from becoming an unintended transit path. In a multihomed BGP design, a customer AS can accidentally advertise routes learned from one provider to another provider. If that happens, external traffic may cross the customer network to reach destinations beyond it, making the customer a transit AS. Applying No-Export to the appropriate learned or advertised routes instructs the receiving AS not to advertise those routes outside its own AS boundary. This helps contain route propagation and prevents third parties from learning paths through AS64512. Maximum-prefix protects a router from receiving too many routes; it is not a transit-prevention policy. No-advertise is more restrictive and prevents advertisement to all peers, which may break required internal or provider reachability. Next-hop is a path attribute used for reachability and path selection, not a route-leak prevention mechanism. In production, No-Export should be combined with explicit prefix filters so policy does not depend on community handling alone. Reference topics: BGP No- Export, transit AS prevention, route-leak control, BGP communities, multihoming.

NEW QUESTION: 141
Which control-plane technology allows the same subnet to exist across multiple network locations?
A. LISP
B. VXLAN
C. FabricPath
D. ISE mobility services
Answer: A (LEAVE A REPLY)
Locator/ID Separation Protocol is the control-plane technology that allows endpoint identity to be separated from endpoint location. In Cisco SD-Access, LISP is used by the fabric control plane to map endpoint identifiers to routing locators, so an endpoint subnet can remain logically consistent even when endpoint reachability is provided from different network locations. This is why LISP is associated with host mobility and with designs where the same subnet or endpoint identity can exist across multiple places without forcing the physical topology to be a single stretched Layer 2 domain. VXLAN is the data-plane encapsulation used to carry overlay traffic, but VXLAN itself is not the mapping control plane. FabricPath is a Cisco Layer 2 multipath technology and does not provide LISP-style endpoint-location mapping. ISE mobility services are identity and policy functions, not a routing locator mapping system. Therefore, the correct control-plane technology is LISP. Cisco SD-Access design material describes the control plane as the mapping system that maintains endpoint-to-device relationships.

NEW QUESTION: 142
An engineer is working with NETCONF and Cisco NX-OS based devices. The engineer needs a YANG model that supports a specific feature relevant only to Cisco NX-OS. Which model must the engineer choose?
A. Native
B. IEEE
C. OpenConfig
D. IETF
Answer: (SHOW ANSWER)
A Cisco native YANG model is the right choice when the engineer needs to configure a capability that is specific to Cisco NX-OS. YANG models fall broadly into open standards-based models and vendor-native models. IETF and OpenConfig models are useful for multivendor consistency, but they intentionally model common behavior and may not cover every platform-specific feature. Cisco native models are generated to represent Cisco operating system configuration and operational data more closely, including features that are unique to IOS XE, IOS XR, or NX-OS. Cisco developer guidance commonly positions native models as the model family to use when full platform feature coverage is required. IEEE is not the relevant model family for Cisco NX-OS feature-specific configuration in NETCONF. OpenConfig is useful when the design objective is common cross-vendor abstraction, but that is not the question's requirement. Therefore, for an NX-OS-only feature that must be configured through NETCONF/YANG, the engineer should choose the Cisco native model for NX-OS.

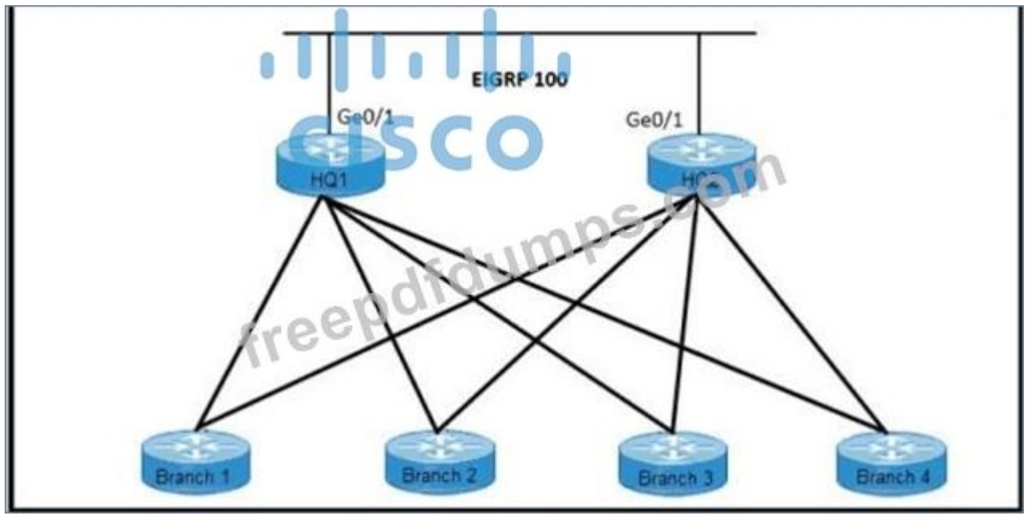
NEW QUESTION: 143

freepdfdumps.com

Refer to the exhibit. An architect must design an IP addressing scheme for a multisite network connected via a WAN transit. The campus site must accommodate 12,000 devices and the branch sites must accommodate
1,000 devices. Which address scheme optimizes network device resources, contains convergence events to the different blocks of the network, and ensures future growth of the network?
A. Campus: 10.0.0.0/18
* Branch1: 10.0.192.0/21
* Branch2: 10.0.200.0/21
B. * Campus: 10.0.0.0/16
* Branchi: 10.255.0.0/20
* Branch2: 10.255.16.0/20
C. * Campus: 10.0.0.0/10
* Branch1: 10.64.0.0/10
* Branch2: 10.128.0.0/10
D. * Campus: 10.0.0.0/20
* Branch1: 10.0.64.0/21
Branch2: 10.0.128.0/21
Answer: A (LEAVE A REPLY)
Option A is the best addressing design because it gives each site enough capacity while avoiding excessive address waste and supporting summarization. A campus requiring 12,000 devices needs at least 14 host bits because 2^13 minus two is only 8190 usable addresses, while a /18 provides 16,382 usable addresses. That leaves room for growth without assigning an unnecessarily huge block. Each branch requiring 1,000 devices needs at least 10 host bits; a /22 provides 1,022 usable addresses, but /21 provides 2,046 usable addresses and additional growth capacity. The selected branch blocks, 10.0.192.0/21 and 10.0.200.0/21, are contiguous on valid /21 boundaries and can be summarized cleanly where appropriate. Option D fails because a /20 supports only 4094 usable addresses and cannot accommodate the 12,000-device campus. Option B overallocates a /16 to the campus and /20s to branches, wasting resources. Option C is grossly oversized and operationally poor.
Reference topics: hierarchical IPv4 addressing, route summarization, VLSM, host-capacity planning, convergence containment.

NEW QUESTION: 144



Refer to the exhibit. An architect must create a stable and scalable EIGRP solution for a customer. The design must:
*conserve bandwidth, memory, and CPU processing
*prevent suboptimal routing
*avoid any unnecessary queries

Which two solutions must the architect select? (Choose two.)
A. route summarization
B. prefix lists
C. distribute lists
D. stub routing
E. static redistribution
Answer: A,D (LEAVE A REPLY)
Route summarization and EIGRP stub routing are the two correct design techniques for a stable and scalable EIGRP deployment. Summarization reduces the number of prefixes advertised upstream, conserves bandwidth and memory, hides route flaps behind the summary boundary, and reduces the size of the EIGRP topology table. Stub routing prevents branch or access devices from being used as transit routers and limits the scope of EIGRP queries. This is critical because uncontrolled EIGRP query propagation can increase convergence time and processing load, especially in larger hierarchical networks. Prefix lists and distribute lists can filter routes, but filtering alone does not provide the same query-boundary and topology-hiding benefits as summarization and stubs. Static redistribution increases complexity and can create loop or policy risks if not controlled with tags and metrics. The question specifically asks to conserve bandwidth, memory, and CPU, prevent suboptimal routing, and avoid unnecessary queries; summarization and stub routing are the direct Cisco design answers. The architect should summarize at hierarchy boundaries and configure access or branch routers as EIGRP stubs where they should not provide transit. Reference topics: EIGRP summarization, EIGRP stub routing, query scope, hierarchical routing.

NEW QUESTION: 145
An engineer is designing a campus network with Cisco Catalyst 95CO switches in the aggression layer. The design requires running nonblocking Layer 2 MEC from the aggregation layer to the access layer. The Catalyst switches are located on different campus floors for availability reasons, and each access switch veil contam a single VLAN. Which technology must the engineer choose for the aggregation switches in the design?
A. VPC
B. VSS
C. StackWise Virtual
D. StackWise-180
Answer: (SHOW ANSWER)
StackWise Virtual is the correct technology for Cisco Catalyst 9500 aggregation switches when the design requires nonblocking Layer 2 multichassis EtherChannel from aggregation to access. The question specifically describes Catalyst 9500 switches located on different floors for availability, with access switches dual-homed through a Layer 2 MEC design. Traditional VSS was used on older Catalyst platforms such as Catalyst 4500 and 6500 families. For Catalyst 9000 platforms, Cisco uses StackWise Virtual to virtualize two physical switches into a single logical switching system. This enables multichassis EtherChannel, removes the need for STP to block one uplink, improves convergence, and provides device-level redundancy. vPC is a Nexus data-center technology, not a Catalyst campus aggregation technology. StackWise-180 refers to physical stack bandwidth behavior on specific access platforms and does not provide the same distributed campus aggregation function. Therefore, the design should use StackWise Virtual on the Catalyst 9500 aggregation pair. Reference topics: Catalyst 9500, StackWise Virtual, multichassis EtherChannel, campus aggregation, Layer 2 resiliency.

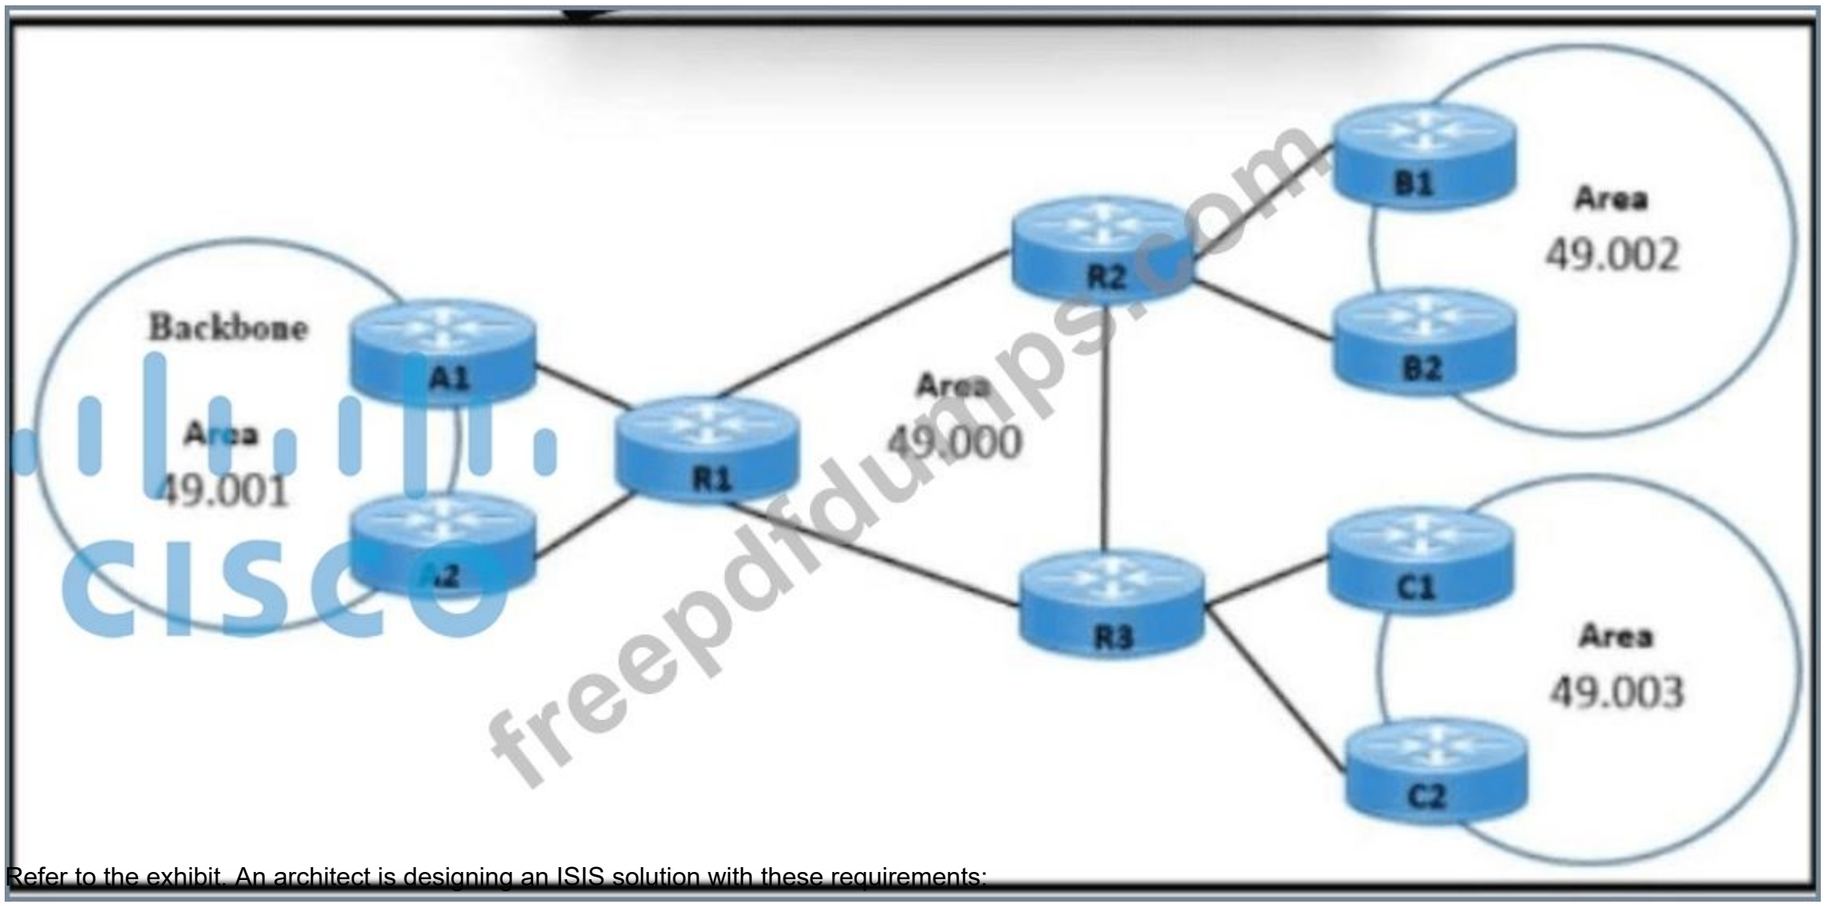
NEW QUESTION: 146
An architect must develop a campus network solution that includes:
logically segmented and isolated networks
ability to communicate between network segments when required
support for overlapping IP addresses
widely available technologies to avoid purchasing specialized equipment Which solution must the architect select?
A. VSS with IGP
B. 802.1Q with HSRP
C. vPC with HSRP
D. VRF-Lite with OSPF

Answer: D (LEAVE A REPLY)
VRF-Lite with OSPF satisfies the requirements because it creates multiple isolated Layer 3 routing tables without requiring a service-provider MPLS core or specialized hardware. Each VRF has its own forwarding table, interfaces, and route exchange process, so overlapping IP address space can exist in different virtual networks. When communication between segments is required, the design can use controlled route leaking, a firewall, or a shared-services VRF. OSPF can run per VRF to exchange routes within each isolated segment using widely supported routing technology. VSS, vPC, and plain 802.1Q with HSRP can improve Layer 2 or device redundancy, but they do not provide independent Layer 3 routing tables with overlapping addressing.
A multihop MPLS design can also support many VPNs, but it is more complex and beyond the stated need for widely available technologies without specialized equipment. VRF-Lite is therefore the correct enterprise campus segmentation choice. Reference topics: VRF-Lite, virtual routing and forwarding, inter-VRF routing, overlapping IP address support, campus segmentation.

NEW QUESTION: 147
A company has many spoke sites with two data centers. The company wants to exchange the routing information between the data centers and the spoke sites using EIGRP. All locations belong to a single AS.
and auto-summarization Is disabled. Which two actions must the company choose? (Choose two.)
A. Exchange all routes between locations
B. Summarize the routes between the hubs.
C. Make each spoke site router a stub router
D. Summarize the routes from spokes to the hubs.
E. Split the network into two separate ASs
Answer: C,D (LEAVE A REPLY)
The scalable EIGRP hub-and-spoke design should make each spoke a stub router and summarize routes from the spokes toward the hubs. Cisco EIGRP documentation states that stub routing improves network stability, reduces resource utilization, and simplifies branch or edge router configuration. Stub spokes limit the routes they advertise and prevent the hub routers from sending unnecessary queries through remote sites.
Summarizing spoke routes toward the hubs further reduces routing-table size and contains convergence events when a subnet fails behind a spoke. Exchanging all routes between every location would increase memory, CPU, and query scope, which is the opposite of good EIGRP scale design. Summarizing only between the hubs may be useful in a dual-data-center design, but it does not solve the spoke scaling requirement. Splitting the network into two autonomous systems is unnecessary when the requirement states that all locations belong to a single AS. Reference topics: EIGRP hub-and-spoke design, stub routing, route summarization, query scope reduction, branch scale.

NEW QUESTION: 148
Mobile service provider "A" is going to launch 5G support with the ISP "B" IP network backbone as the underlay transport. Conversational traffic type will be marked with expedited forwarding class, streaming services will be marked with assured forwarding 2 class, and web browsing will be marked with assured forwarding 3 class. Which QoS model meets these requirements if the solution is to be implemented end to end on the ISP_B backbone network?
A. 6-class QoS strategy with IntServ model
B. 8-class QoS strategy with DiffServ model
C. 12-class QoS strategy with IntServ model
D. 4-class QoS strategy with DiffServ model
Answer: (SHOW ANSWER)
The correct model is a 4-class QoS strategy with DiffServ. The traffic classes listed in the question are explicitly DiffServ per-hop behaviors: Expedited Forwarding for conversational traffic, Assured Forwarding class 2 for streaming, and Assured Forwarding class 3 for web or interactive traffic. Cisco QoS guidance uses DiffServ to classify and mark traffic with DSCP values, then apply per-hop behavior at each node along the path. This gives scalable end-to-end treatment across a provider backbone without maintaining per-flow reservations. The four service treatments implied are conversational, streaming, interactive/web, and background or best effort, which matches the stated service categories without unnecessary class expansion.
IntServ is not appropriate for an ISP backbone because it uses RSVP-style per-flow signaling and state, which scales poorly compared with DiffServ PHB handling. The 8-class and 12-class options add complexity beyond the stated mappings. Reference topics: DiffServ, DSCP, EF, AF2, AF3, per-hop behavior, provider QoS design.

NEW QUESTION: 149
What are two advantages of the Cisco SD-WAN technology 9 (Choose two)
A. Improved application experience
B. Easier deployment
C. Optimized cloud connectivity
D. Proactive network management
E. Consistent connectivity
Answer: A,C (LEAVE A REPLY)
Improved application experience and optimized cloud connectivity are two principal advantages of Cisco SD- WAN technology. Cisco SD-WAN separates the control plane from the data plane and applies centralized policy to choose the best available transport path for business applications. Features such as application-aware routing, performance measurement, path quality monitoring, forward error correction, packet duplication for selected traffic, and QoS integration help improve user experience for voice, video, SaaS, and other critical applications. SD-WAN also improves cloud connectivity by allowing branches to access trusted SaaS and IaaS resources directly instead of backhauling all traffic through a central data center. Cloud OnRamp functions further simplify connectivity to cloud-hosted workloads and SaaS applications. Easier deployment and proactive management may be operational benefits in some designs, but the strongest technology advantages in this answer set are application experience and optimized cloud access. Consistent connectivity is too generic and does not identify a specific SD-WAN design value. Reference topics: Cisco SD-WAN benefits, application-aware routing, direct internet access, Cloud OnRamp, SaaS optimization.

NEW QUESTION: 150


Refer to the exhibit. An architect is designing an ISIS solution with these requirements:

* The backbone area will grow to 50 routers in the next 12 months.
* Routers A1 and A2 must avoid suboptimal routing.
* Summarization and route-leaking should be allowed in areas 49.002 and 49.003.
Which solution must the architect select?
A. area 49.000 L1, area 49.001 L2, area 49.002 L2, and area 49.003 L2
B. area 49.000 L1, area 49.001 L1, area 49.002 L2, and area 49.003 L2
C. area 49.000 L2. area 49.001 L1, area 49.002 L1, and area 49.003 L1
D. area 49.000 L2. area 49.001 L2, area 49.002 L1, and area 49.003 L1
Answer: C (LEAVE A REPLY)
The correct IS-IS hierarchy places area 49.000 as Level 2 and areas 49.001, 49.002, and 49.003 as Level 1. In IS-IS, Level 2 forms the interarea backbone, while Level 1 areas contain local topology details. This design supports growth in the backbone because Level 2 routers exchange interarea reachability without carrying every local topology change from every access area. Areas 49.002 and 49.003 can perform summarization and route leaking at their Level 1/Level 2 boundaries, allowing routers such as A1 and A2 to avoid suboptimal routing by learning selected more-specific routes from the backbone when necessary. If the nonbackbone areas were made Level 2, the design would flatten the hierarchy and expose more of the network to topology changes. If the backbone were Level 1, interarea routing would be broken or badly constrained. The selected hierarchy preserves scaling, allows route control, and avoids unnecessary propagation of local link flaps.
Reference topics: IS-IS hierarchy, Level 1 areas, Level 2 backbone, route leaking, summarization, scaling.

NEW QUESTION: 151
How does OMP behave in a Cisco Catalyst SD-WAN architecture if no policy is defined?
A. To allow a hub-and-spoke topology for WAN Edge routers to communicate via the central location
B. To allow a point-to-point topology for WAN Edge routers to communicate from the central location to remote locations
C. To allow all WAN Edge routers to communicate using a full mesh topology
D. To block all communication between WAN Edge routers
Answer: (SHOW ANSWER)
Without centralized policy, OMP allows all WAN Edge routers in the same VPN to communicate in a full- mesh overlay. Cisco Catalyst SD-WAN uses OMP as the control-plane protocol between WAN Edge routers and SD-WAN controllers. WAN Edge devices advertise OMP routes, TLOC routes, and service routes to the controllers, and the controllers distribute the reachable routes back to the relevant edge devices. Cisco policy documentation describes the default unpolicied behavior as permissive: routes are accepted and advertised so sites in the same VPN learn reachability to each other. That default route distribution produces a full-mesh data-plane capability, even though control connections remain hubbed through the controllers. Hub-and- spoke, regional isolation, point-to-point restriction, or service-chained topologies require centralized control or data policy to filter routes or steer traffic. Blocking all communication is the opposite of the default behavior. Reference topics: Cisco SD-WAN OMP, default route distribution, full-mesh overlay, centralized control policy, VPN segmentation.

Valid 300-420 Dumps shared by Actual4test.com for Helping Passing 300-420 Exam! Actual4test.com now offer the **newest 300-420 exam dumps**, the Actual4test.com 300-420 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-420 dumps with Test Engine here: https://www.actual4test.com/300-420_examcollection.html (**457** Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 152
Which two border nodes are available in the Cisco SD-Access architecture? (Choose two.)
A. extended border
B. edge border
C. internal border

D. anywhere border
E. intermediate border
Answer: (SHOW ANSWER)
Internal border and anywhere border are valid Cisco SD-Access border node types. A border node connects the SD-Access fabric to networks outside the fabric, but different border roles are used depending on what external reachability is required. An internal border provides connectivity to known internal networks such as shared services, data center prefixes, or enterprise routes that should be reachable from inside the fabric. An anywhere border combines external and internal border behavior and can provide broader exit capabilities depending on the fabric design. The commonly described border roles also include external border, which provides default or Internet-facing exit from the fabric, but that option is not available in the listed choices.
"Extended border," "edge border," and "intermediate border" are not the correct named border node types.
The answer is therefore internal border and anywhere border. Border placement is a major design decision because it influences route advertisement, VN-to-VRF mapping, external connectivity, transit behavior, and failure-domain boundaries for SD-Access fabrics.

NEW QUESTION: 153
A customer's environment includes hosts that support IPv6-only. Several of these hosts must communicate with a public web server that has only IPv4 domain name resolution. Which solution should the customer use in this environment?
A. utilize NAT64 to translate the addresses
B. Implement NAT44 at the edge of the customer network
C. use 6to4 and a tunnel to translate the addresses
D. implement 6PE to resolve hostname resolution
Answer: (SHOW ANSWER)
NAT64 is the correct solution when IPv6-only hosts must communicate with an IPv4-only public web server.
NAT64 translates IPv6 packets from the client side into IPv4 packets toward the server and translates the return traffic back into IPv6. In most practical designs, NAT64 is paired with DNS64 so IPv6-only clients can receive synthesized AAAA records for IPv4-only destinations. NAT44 is only IPv4-to-IPv4 translation and cannot support IPv6-only hosts. 6to4 is an IPv6-over-IPv4 tunneling mechanism, not an IPv6-to-IPv4 translation service for reaching an IPv4-only server. 6PE carries IPv6 reachability across an MPLS IPv4 core and does not solve IPv6-only client access to an IPv4-only domain. The question specifically states that the hosts support IPv6-only and the public web server has only IPv4 domain name resolution. That is the classic NAT64/DNS64 use case. The architect must also plan stateful translation capacity, logging requirements, IPv4 address pool sizing, and DNS behavior. Reference topics: NAT64, DNS64, IPv6 transition, IPv6-only hosts, IPv4-only server access.

NEW QUESTION: 154
Which encoding languages are supported in NETCONF compared to RESTCONF?
A. NETCONF supports XML and JSON, and RESTCONF supports XML.
B. NETCONF supports XML, and RESTCONF supports JSON.
C. NETCONF supports JSON, and RESTCONF supports XML.
D. NETCONF supports XML, and RESTCONF supports XML and JSON.
Answer: D (LEAVE A REPLY)
NETCONF supports XML, while RESTCONF supports both XML and JSON. Cisco programmability documentation describes NETCONF as an XML-based protocol that exchanges RPC requests and replies such as get, get-config, and edit-config using XML encoding. NETCONF is tightly aligned with YANG- modeled data, but the protocol message format itself is XML. RESTCONF exposes YANG-modeled data through an HTTP or HTTPS API and can represent payloads in XML or JSON. Cisco DevNet material commonly summarizes the comparison as NETCONF equals XML, while RESTCONF equals XML or JSON. That is why option D is accurate. Option A is wrong because NETCONF does not use JSON in the same way RESTCONF does. Option B is incomplete because RESTCONF is not limited to JSON. Option C reverses the protocol encodings. This distinction matters in automation design because tool selection, content type, parsing, and API workflows differ by protocol. Reference topics: NETCONF, RESTCONF, YANG, XML encoding, JSON encoding, model-driven programmability.
NEW QUESTION: 155
What is one function of the vSmart controller in an SD-WAN deployment?

A. orchestrates vEdge and cEdge connectivity
B. responsible for the centralized control plane of the SD-WAN network
C. provides centralized network management and a GUI to monitor and operate the SD-WAN overlay
D. provides a data-plane at branch offices to pass traffic through the SD-WAN network
Answer: (SHOW ANSWER)
The vSmart controller provides the centralized control plane in a Cisco SD-WAN deployment. It establishes secure control connections with WAN Edge routers and exchanges routing, topology, security, and policy information using OMP. The vSmart controller does not forward user data traffic; data-plane forwarding occurs directly between WAN Edge routers over IPsec or other supported encapsulations. Its function is to distribute reachability information, TLOC information, service routes, centralized policies, application-aware routing policies, and segmentation information that determines how the overlay operates. vBond orchestrates onboarding, authentication facilitation, and NAT traversal for control connections. vManage provides centralized management, monitoring, configuration templates, and operational GUI/API functions. vEdge or Cisco IOS XE SD-WAN routers provide the branch or site data plane and local routing functions. Therefore, among the choices, the vSmart controller is responsible for the centralized control plane of the SD-WAN network. Correctly separating orchestration, management, control, and data-plane roles is a core Cisco SD- WAN design principle.

NEW QUESTION: 156
An engineer must design a management network for a customer ' s enterprise network. The design must:
* provide the ability to grant and revoke access privileges
* allow only protocols SSH, NTP, FTP, and SNMP
* restrict access to management Interfaces
Which solution must the engineer choose to meet the requirements?
A. in-band
B. enterprise internal private
C. out-of-band
D. mGRE
Answer: C (LEAVE A REPLY)
An out-of-band management network best satisfies the requirement to grant and revoke administrative access, restrict management protocols, and limit exposure of management interfaces. Out-of-band management separates administrative traffic from the production data plane by using dedicated management interfaces, management switches, terminal infrastructure, or separate routing. This isolation improves security because a compromise or congestion event in the production network does not automatically provide access to device management planes. It also allows stronger access controls, protocol restrictions, logging, and administrative segmentation for SSH, NTP, FTP, and SNMP. In-band management uses the production network and is more exposed to outages or data-plane attacks. An enterprise internal private network is too vague and does not guarantee management-plane isolation. mGRE is a tunneling technique, not a complete management access architecture. A professional design should also include AAA, role-based authorization, ACLs, management VRFs where appropriate, logging, and monitoring. Reference topics: out-of-band management, management- plane security, AAA, access control, protocol restriction, network operations design.

NEW QUESTION: 157
Which two overlay network design considerations must be made for a Cisco SD-Access network? (Choose two.)
A. LAN automation for deployment
B. Layer 3 to the access design
C. Reduce subnets and simplify DHCP management
D. Dedicated IGP process for the fabric
E. Avoid overlapping IP subnets
Answer: C,E (LEAVE A REPLY)
The SD-Access overlay should reduce subnet sprawl and avoid overlapping IP subnets unless there is a specific operational requirement. Cisco SD-Access creates overlay virtual networks over an IP underlay and uses anycast gateway, VXLAN encapsulation, LISP control-plane mapping, and policy constructs to simplify endpoint mobility and segmentation. Reducing the number of subnets simplifies DHCP scope management, default-gateway placement, and operational troubleshooting. Avoiding overlapping address space across overlay networks is also a practical design recommendation because

overlapping subnets complicate shared services, fusion routing, security logging, and policy troubleshooting. LAN automation and a dedicated IGP process are underlay deployment considerations, not overlay design choices. Layer 3 to the access design is also associated with routed underlay and campus fabric transport, not the overlay segmentation model. A clean overlay design should define virtual networks, scalable groups, IP pools, shared-services reachability, and external connectivity before deployment. Reference topics: Cisco SD-Access overlay design, anycast gateway, virtual networks, DHCP simplification, overlapping subnet avoidance.

Valid 300-420 Dumps shared by Actual4test.com for Helping Passing 300-420 Exam! Actual4test.com now offer the **newest 300-420 exam dumps**, the Actual4test.com 300-420 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-420 dumps with Test Engine here: https://www.actual4test.com/300-420_examcollection.html (**457** Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)