

DumpsActual

Input your exam code ...

Our exam study dumps are compiled through several times of checking and verification, which are with the high pass rate up to 98%. All the actual questions & answers are selected and confirmed according to strict standards with assurance for 100% pass.

All Products | Contact now

QUALITY AND VALUE

DumpsActual Practice Exams are written to the highest standards of technical accuracy, using only certified subject matter experts and published authors for development - no all dumps.



TESTED AND APPROVED

We are committed to the process of vendor and third party approvals. We believe professionals and executives alike deserve the confidence of quality coverage these authorizations provide.

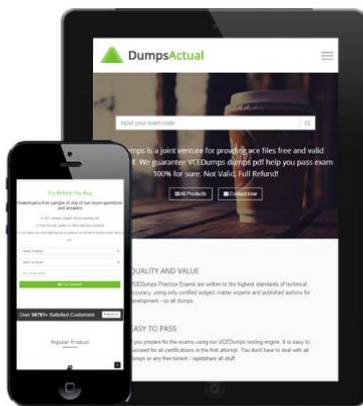
EASY TO PASS

If you prepare for the exams using our DumpsActual testing engine. It is easy to succeed for all certifications in the first attempt. You don't have to deal with all dumps or any free torrent / rapidshare all stuff.



TRY BEFORE BUY

DumpsActual offers free demo of each product. You can check out the interface, question quality and usability of our practice exams before you decide to buy.



Try Before You Buy

Download a free sample of any of our exam questions and answers

- ✓ 24/7 customer support, Secure shopping site
- ✓ Free One year updates to match real exam scenarios
- ✓ If you failed your exam after buying our products we will refund the full amount back to you.

Select a vendor... | Select an exam... | Your email address | Free Download

<http://www.dumpsactual.com>

Achieve 100% pass with the valid & actual exam practice dumps

Exam : **300-420J**

Title : **Designing Cisco Enterprise Networks (300-420日本語版)**

Vendor : **Cisco**

Version : **DEMO**

QUESTION NO: 1

ワイヤレスエンドポイントは、Cisco SD-AccessアーキテクチャのHTDBにどのように登録されますか？

- A. ファブリックエッジノードは、APからのCAPPWAPメッセージングに基づいてHTDBを更新します
- B. 新しいクライアントがワイヤレスネットワークに接続すると、ファブリックWLCがHTDBを更新します
- C. 境界ノードは最初にエンドポイントを登録し、次にHTDBを更新します
- D. ファブリックAPは、クライアントのEIDおよびRLOCでHTDBを更新します

Answer: B

Explanation:

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.

QUESTION NO: 2

エンジニアは、IPv6を既存のIPv4IS-ISネットワークで実行できるようにする設計を作成しています。

IPv4とIPv6のトポロジは完全に一致し、エンジニアはインターフェイスごとのプロトコルごとに同じルーターレベルを使用することを計画しています。どのIS-IS設計が必要ですか？

- A. 遷移機能を有効にしない単一トポロジ
- B. 遷移機能が有効になっている単一トポロジ
- C. 遷移機能が有効になっているマルチトポロジ
- D. 遷移機能を有効にしないマルチトポロジ

Answer: B

Explanation:

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.

QUESTION NO: 3

モバイルサービスプロバイダー「A」は、ISP「B」のIPネットワークバックボーンを基盤トランスポートとして、5Gサポートを開始する予定です。会話型トラフィックは優先転送クラス、ストリーミングサービスは保証転送2クラス、Webブラウジングは保証転送3クラスでマークされます。ISP_Bバックボーンネットワーク上でエンドツーエンドでソリューションを実装する場合、これらの要件を満たすQoSモデルはどれですか？

- A. IntServモデルを用いた6クラスQoS戦略
- B. DiffServモデルを用いた8クラスQoS戦略
- C. IntServモデルを用いた12クラスQoS戦略
- D. DiffServモデルを用いた4クラスQoS戦略

Answer: D

Explanation:

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.

QUESTION NO: 4

gRPC はどのようなセキュリティ機能を提供しますか？

- A. RSA 20*8 暗号暗号化による安全なサーバー・クライアント間トンネルの実装
- B. AES および RSA プロトコルを使用した保存データの強制暗号化
- C. CRCチェック付きのRC6データレベル暗号化を有効にする
- D. TLSを使用したネットワークデバイスと制御システム間の安全な通信をサポート

Answer: D

Explanation:

gRPC provides secure communication between network devices and control systems by supporting TLS. In Cisco programmability and model-driven telemetry designs, gRPC is

commonly used as a high-performance remote procedure call framework that carries structured data encoded with Protocol Buffers. Security is provided through TLS so that the client and server can authenticate and protect the session in transit. The question asks for security functionality, so TLS support is the correct selection. The other options describe incorrect or misleading cryptographic behavior. gRPC does not implement generic RSA tunnel encryption as described in option A, and it does not provide mandatory encryption of data at rest. RC6 with CRC is not the security model used for Cisco gRPC telemetry or programmability. In a secure management design, gRPC should be deployed with certificates, TLS verification, controlled access lists, and appropriate AAA integration so only authorized collectors or controllers can communicate with infrastructure devices. Reference topics: gRPC, TLS, model-driven telemetry, Protocol Buffers, secure automation transport.

QUESTION NO: 5

SD-

Accessアーキテクチャ環境のアンダーレイネットワークを計画する際に、エンジニアはどの要件を考慮すべきでしょうか？

- A. アンダーレイネットワークでは CEF を無効にする必要があります。
- B. エンドポイントはアンダーレイネットワークに含める必要があります。
- C. スパニングツリーのループを回避するには、ネットワーク内に2つ以上のスイッチが必要です。
- D. ネットワークデバイス間でIP接続を確立するには、ルーティングプロトコルを実装する必要があります。

Answer: D

Explanation:

The SD-Access underlay must run a routing protocol to establish IP connectivity between infrastructure devices. Cisco SD-Access uses a routed underlay to provide basic reachability among fabric edge nodes, border nodes, control-plane nodes, and intermediate nodes. Cisco design documentation states that all network elements in the fabric underlay must have IP connectivity through routing. Catalyst Center LAN Automation commonly deploys IS-IS for this purpose, although the underlay can be built manually with another supported routing design. The user endpoints and client subnets are not part of the underlay; they belong to the overlay, where fabric services, segmentation, and policy are applied. CEF should not be disabled, because efficient IP forwarding is necessary for the routed underlay. Two or more switches are not required merely to avoid spanning-tree loops; the SD-Access underlay is normally Layer 3 routed and avoids large Layer 2 loop domains. Reference topics: Cisco SD-Access underlay, routed access, LAN Automation, IS-IS, fabric node reachability.

QUESTION NO: 6

ネットワーク エンジニアは、マルチキャスト

ストリームのスプーフィングを防止し、効率的な帯域幅の利用を保証するマルチキャストソリューションを設計する必要があります。将来、ネットワークは別のマルチキャストドメインとマージされますが、マージには最小限の労力しか必要ありません。お客様の要件を満たす2つのソリューションはどれですか？（2つ選んでください。）

- A. PIM-SSM
- B. IGMPv3
- C. IGMPv2
- D. PIM-SM
- E. MSDP

Answer: A B

Explanation:

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.

QUESTION NO: 7

組織は、2つの異なる自律システムにマルチキャストを展開することを計画しています。彼らのソリューションでは、RPが次のことを行えるようにする必要があります。

- *ドメイン外のアクティブなソースを発見する
 - *他のRPとの接続に基礎となるルーティング情報を使用する
 - *グループに参加しているソースを発表する
- これらの要件をサポートするソリューションはどれですか？

- A. MSDP
- B. SSM
- C. PIM-SM
- D. PIM-DM

Answer: A


Explanation:

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.

QUESTION NO: 8



展示品を参照 従業員ID 4449:30の通信会社に勤務するエンジニア
959 スイッチの STP スケーラビリティを計算して、その数が STP
論理ポートの最大サポート値を下回っていることを確認していますか? スイッチ A
でアクティブな論理インターフェイスの数はいくつですか?

- A. 4
- B. 307
- C. 202
- D. 100

Answer: C

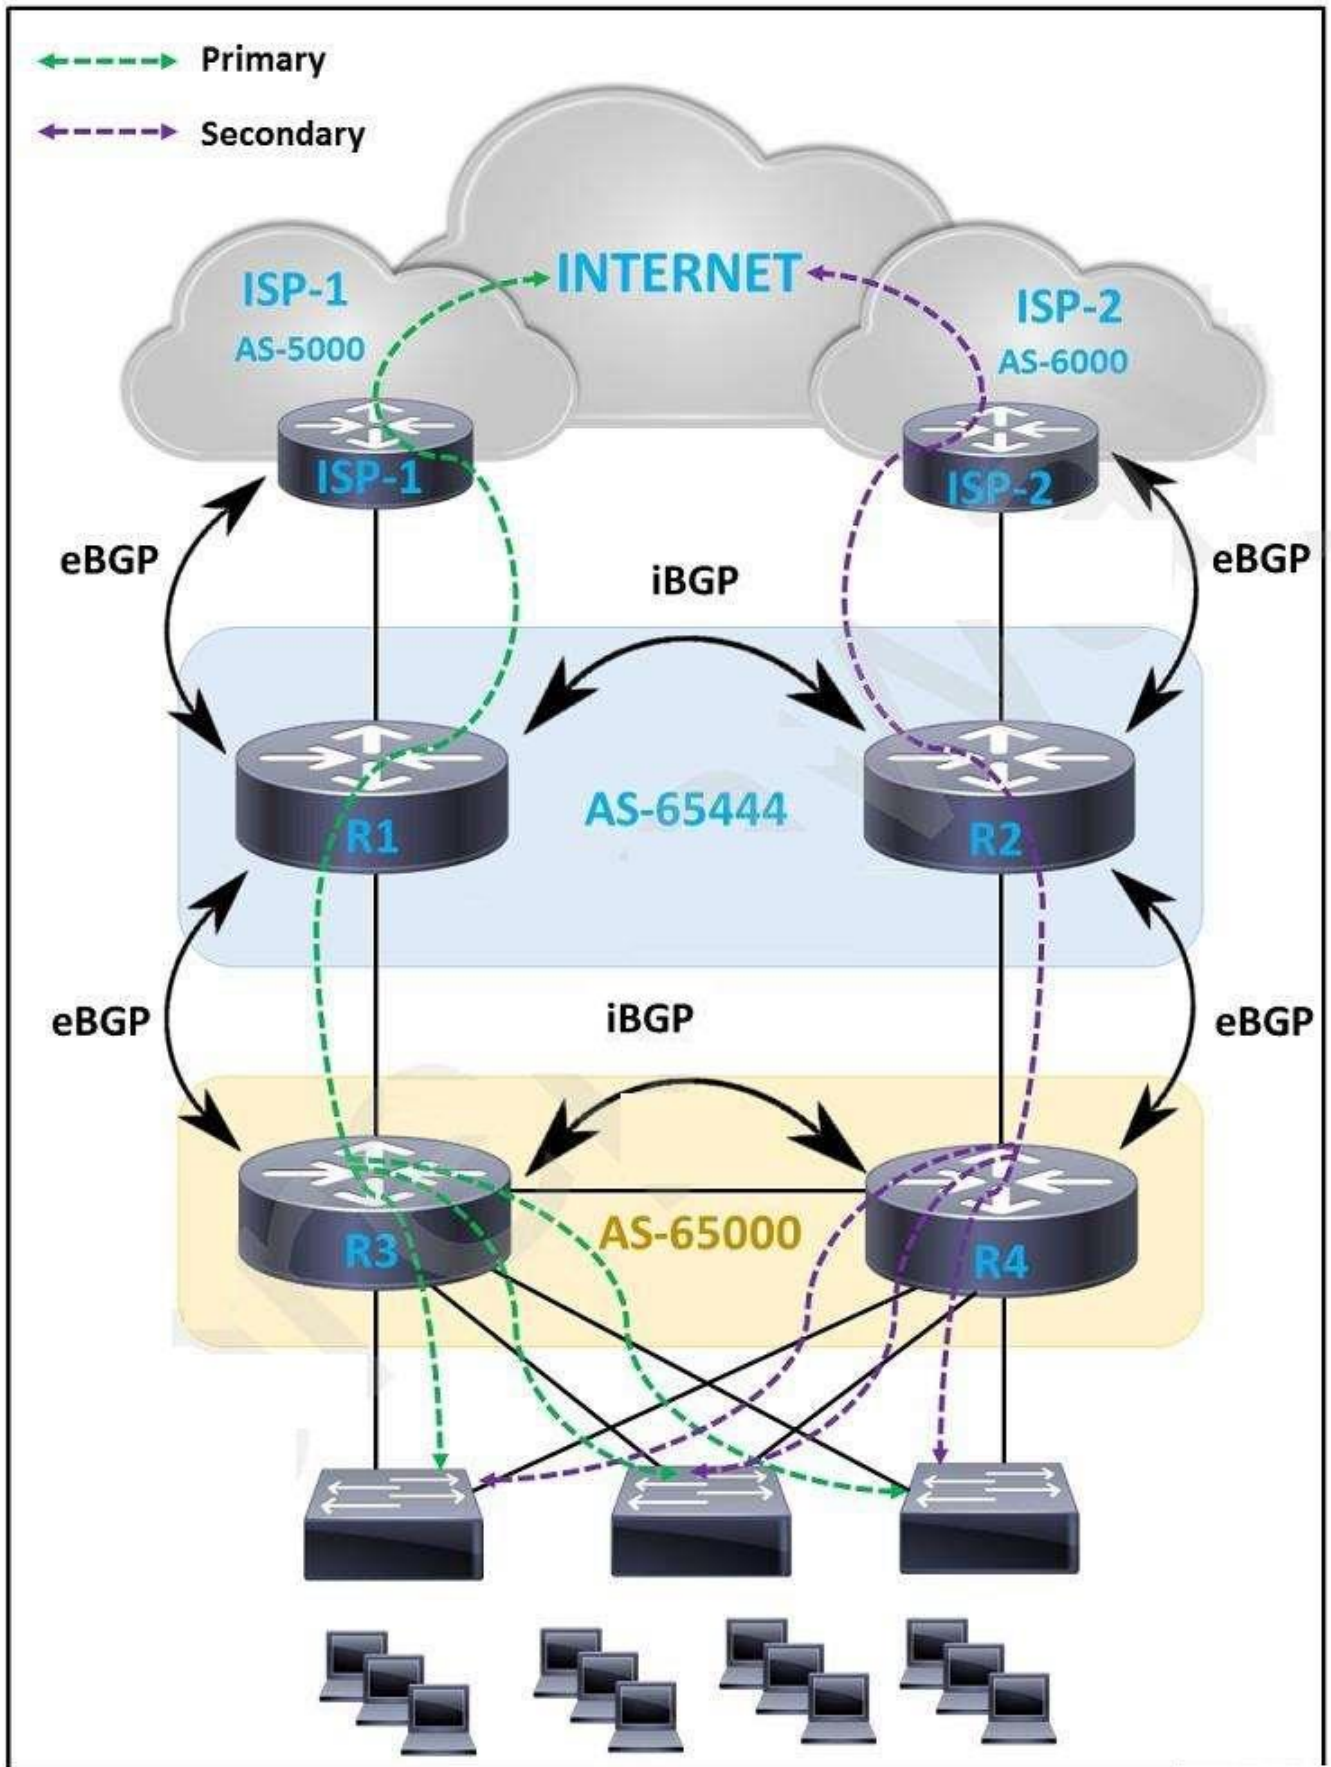
Explanation:

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.

QUESTION NO: 9

展示品を参照してください。



エンジニアは、ISP-1 が常に ISP-2 よりも優先されるように WAN ソリューションを設計する必要があります。ISP-2

経路のパスはバックアップと見なされ、ISP-1へのパスがダウンしている場合にのみ使用する必要があります。エンジニアはどのソリューションを選択する必要がありますか？

A. R1:

- ISP-1に広告されるルート: 0x ASパスの先頭に追加
- ISP-1から受信したルート: HIGHローカル優先度
- R2にアドバタイズされたルート: アクションなし
- R2から受信したルート: コミュニティNO-EXPORT

R2:

- ISP-2:5x ASパスの先頭に広告されるルート
- ISP-2から受信したルート: LOWローカル優先度
- R1に広告されるルート: コミュニティ NO-ADVERTISE
- R1から受信したルート: アクションなし

B. R1:

- ISP-1に広告されるルート: 0x ASパスの先頭に追加
- ISP-1から受信したルート: HIGHローカル優先度
- R2にアドバタイズされるルート: コミュニティNO-EXPORT
- R2から受信したルート: アクションなし

R2:

- ISP-2に広告されるルート: 5x ASパスの先頭追加
- ISP-2から受信したルート: LOWローカル優先度
- R1にアドバタイズされたルート: アクションなし
- R1から受信したルート: アクションなし

C. R1:

- ISP-1に広告されるルート: 0x ASパスの先頭に追加
- ISP-1から受信したルート: LOWローカル優先度
- R2にアドバタイズされるルート: コミュニティ NO-ADVERTISE
- R2から受信したルート: アクションなし

R2:

- ISP-2に広告されるルート: 5x ASパスの先頭追加
- ISP-2から受信したルート: HIGHローカル優先度
- R1にアドバタイズされたルート: アクションなし
- R1から受信したルート: コミュニティ NO-ADVERTISE

D. R1:

- ISP-1に広告されるルート: 5x ASパスの先頭追加
- ISP-1から受信したルート: LOWローカル優先度
- R2にアドバタイズされるルート: コミュニティ NO-ADVERTISE
- R2から受信したルート: アクションなし

R2:

- ISP-2に広告されるルート: 0x ASパスの先頭に追加
- ISP-2から受信したルート: HIGHローカル優先度
- R1にアドバタイズされるルート: コミュニティNO-EXPORT
- R1から受信したルート: アクションなし

Answer: B

Explanation:

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.

QUESTION NO: 10

エンジニアは、サービスプロバイダーとのデュアルBGPピアリングソリューションの設計を担当しています。設計は次の条件を満たす必要があります。

*ルーターは、サブネットマスクが/24より大きいプレフィックスを学習しません。

*ルーターは、マスクの長さのみに基づいて、ルーティングテーブルに含めるルートを決めます。

*ルーターは、サービスプロバイダーの構成に関係なく、この選択を行います。

エンジニアはどのソリューションを設計に含める必要がありますか？

A. ルートマップとアクセスリストを使用して目的のネットワークをブロックし、ルートマップをインバウンドのBGPネイバーに適用します。

B. ルートマップとプレフィックスリストを使用して目的のネットワークをブロックし、ルートマップをBGPネイバーのアウトバウンドに適用します。

C. IPプレフィックスリストを使用して目的のネットワークをブロックし、IPプレフィックスリストをBGPネイバーのアウトバウンドに適用します。

D. IPプレフィックスリストを使用して目的のネットワークをブロックし、IPプレフィックスリストをインバウンドのBGPネイバーに適用します。

Answer: D

Explanation:

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.

QUESTION NO: 11

エンジニアは、数百のスイッチとVLANを含む大規模なレイヤー2ドメインを設計する必要があります。エンジニアの主な目標は次のとおりです。

*すべてのリンクの帯域幅を効率的に利用する

*レイヤー2ループを回避する

*スイッチのCPUとメモリへの影響を最小限に抑えます

エンジニアはどのテクノロジーを設計に含める必要がありますか？

A. PVST+

B. Rapid PVST+

C. MST

D. RSTP

Answer: C

Explanation:

Multiple Spanning Tree is the correct technology for a large Layer 2 domain with hundreds of switches and VLANs when the goals are efficient link use, loop prevention, and low CPU and memory impact. MST maps multiple VLANs into a smaller number of spanning-tree instances, allowing the designer to load-balance groups of VLANs across alternate paths without running a separate STP instance for every VLAN. Cisco Rapid PVST+ improves convergence compared with classic STP, but it maintains a separate instance per VLAN, which becomes costly at scale. PVST+ has similar per-VLAN scaling concerns and slower convergence behavior than rapid variants. Plain RSTP improves convergence but does not provide the same VLAN-to-instance mapping flexibility for hundreds of VLANs. MST is standards-based and particularly useful in multivendor environments where VLAN grouping and predictable root placement are required. The design should define MST regions consistently, map VLANs deliberately to instances, place roots near the Layer 3 gateway, and keep the number of instances small enough to protect supervisor CPU and memory. Reference topics: MST, RSTP, Rapid PVST+, VLAN-to-instance mapping, Layer 2 scalability

QUESTION NO: 12

vEdgeルーターの冗長性が設計されている場合、どのFHRPがサポートされますか？

A. HSRP

B. OMP

C. GLBP

D. VRRP

Answer: D

Explanation:

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.

QUESTION NO: 13

SD-WAN エッジ ルーターではどのキューイング構造が使用されますか？

- A. FIFO
- B. LLQ+WFQ
- C. 1P-4Q-2T
- D. 優先度

Answer: B

Explanation:

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.

QUESTION NO: 14

Cisco SD-Access アーキテクチャのオーバーレイ ネットワークに関して、考慮すべき 2 つの点は何ですか？

(2つ選択してください。)

- A. マイクロセグメンテーションには仮想ネットワークを使用する必要があります
- B. SGT はデータプレーンの分離とマイクロセグメンテーションに使用する必要があります
- C. 仮想ネットワークはデータプレーンの分離にのみ使用する必要があります
- D.

IPアドレスを節約するために、異なるオーバーレイネットワーク間で重複するIPアドレスを使用する必要があります。

E. 運用の簡素化のため、異なるオーバーレイネットワーク間での IP アドレスの重複は避けるべきです。

Answer: C E

Explanation:

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.

QUESTION NO: 15

エンジニアは、顧客をインターネットに接続するためのソリューションを設計する必要があります。このソリューションには、サービスプロバイダーからのCIRが50Mbpsのレイヤー3回線が含まれます。プロバイダーのスイッチから顧客のルーターへのハンドオフは1Gbpsです。音声トラフィックの途切れに関する潜在的な問題を防ぐために、エンジニアはどのソリューションを含める必要がありますか？

- A. ルーターへの接続の帯域幅を減らします。
- B. 親ポリシングポリシーを使用して階層型QoSを実装します。
- C. 親シェーピングポリシーを使用して階層型QoSを実装します。
- D. ルーターインターフェイスに帯域幅ステートメントを追加します。

Answer: C

Explanation:

Hierarchical QoS with a parent shaping policy is the correct solution for a 1 Gbps physical handoff where the contracted provider rate is only 50 Mbps. The customer router can transmit at the physical interface speed, but the provider policer or service edge expects traffic to remain within the committed information rate. Without shaping, bursts leave the customer router at 1 Gbps and can be dropped by the provider, which is especially damaging to voice traffic because drops and jitter create choppy audio. Cisco QoS design commonly uses a parent shaper set to the provider CIR and then applies child queuing policies under that shaped rate. This allows voice and other critical classes to be prioritized before packets are transmitted into the constrained service. A parent policing policy would drop or mark excess traffic locally rather than smoothing bursts.

Reducing physical bandwidth is not normally possible on a provider Ethernet handoff, and the interface bandwidth statement changes routing/QoS reference values but does not enforce the CIR. Therefore, parent shaping is required. Reference topics: hierarchical QoS, traffic shaping, CIR, LLQ, WAN edge QoS.

QUESTION NO: 16

ISISを実行しているルーターは、高いCPUと帯域幅の使用率を示しています。エンジニアは、ルーターがL1 /

L2として構成されていて、L1とL2のネイバーがあることを発見しました。問題に対処するために設計を最適化するステップはどれですか。

- A.このルーターを各インターフェースのDISにします
- B.L1 / L2ルーターでデフォルトルートを実バタイズするデフォルトの動作を無効にします
- C.ルーターをL1またはL2に構成します
- D.各インターフェースをL1またはL2回線タイプとして構成します

Answer: D

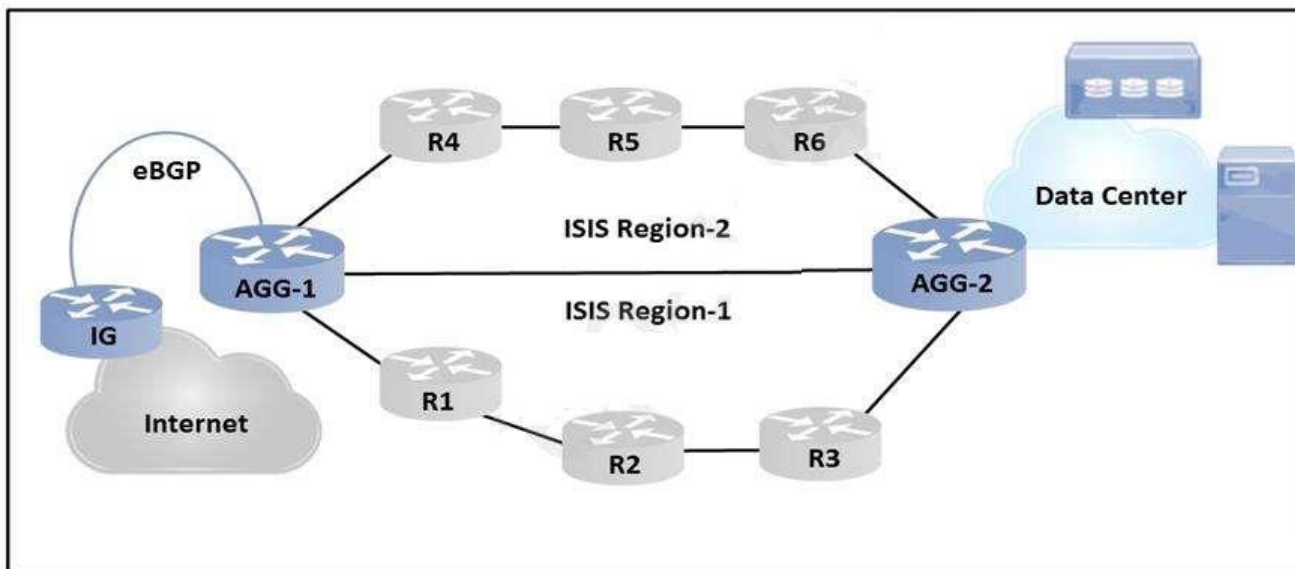
Explanation:

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.

QUESTION NO: 17

展示品を参照してください。



アーキテクトは、エンタープライズ顧客向けの IGP ソリューションを設計する必要があります。設計では、次の内容をサポートする必要があります。

物理リンクフラップの影響は最小限に抑えられます。

アクセス ルータは、リンク障害発生後すぐに収束する必要があります。

建築家が設計に含めるべき ISIS ソリューションはどれですか? (2 つ選択してください。)

- A. BGP から IS-IS への再配布を使用して、レベル 1 エリア内のすべてのインターネット ルートをアドバタイズします。

B. IS-IS インターフェイスとループバック IP

アドレスをインターネットとデータセンターに向けてアドバタイズします。

C. SPF および PRC 間隔を短縮して、収束時間を改善します。

D. ネットワーク全体でレベル 1/レベル 2 隣接関係を確立するように、すべてのアクセスルータと集約ルータを設定します。

E. アクセスルータを設定してレベル 1 隣接関係を確立し、集約ルータを設定してレベル 1/レベル 2 隣接関係を確立します。

Answer: C E

Explanation:

The correct IS-IS design combines tuned SPF/PRC behavior with a proper Level 1 and Level 2 hierarchy.

Reducing SPF and PRC intervals, within platform and stability limits, helps routers react more quickly to topology changes after a link failure. However, fast timers alone are not enough; the hierarchy must contain instability so physical link flaps in the access layer have limited impact on the rest of the network. Access routers should form Level 1 adjacencies inside their local area, while aggregation routers should operate as Level 1/Level 2 routers to connect the access area to the Level 2 backbone. This design gives access routers a simple local view and uses aggregation as the interarea boundary. Running all access and aggregation routers as L1/L2 across the network increases the size of the backbone and exposes more routers to wider flooding, which is the opposite of the requirement.

Redistributing Internet routes into Level 1 is also poor design because it bloats the access area. Therefore, the right choices are C and E. Reference topics: IS-IS hierarchy, L1/L2 design, SPF and PRC tuning, access-layer convergence, fault containment.

QUESTION NO: 18

ある企業は、既存のネットワーク インフラストラクチャ内に IPv6

を導入したいと考えています。現在のインフラストラクチャ機器はすべて IPv6

をサポートしており、企業は追加機器の購入を必要としない移行戦略を望んでいます。計画では、運用管理コストを低く抑える必要があります。IPv6

マルチキャストをサポートし、DNS

を使用してアプリケーションを移行できるようにする必要があります。企業はどの戦略を選択する必要がありますか？

A. ハイブリッド ISATAP トンネル モデル

B. ハイブリッド手動トンネルモデル

C. サービスブロックモデル

D. デュアルスタックモデル

Answer: D

Explanation:

Dual stack is the correct IPv6 migration strategy because the existing infrastructure already supports IPv6, the company wants low operational overhead, and applications should migrate naturally using DNS. Cisco describes dual stack as running IPv4 and IPv6 simultaneously on the same infrastructure, allowing hosts and applications to use either protocol during the migration period. It does not require additional tunneling equipment, avoids encapsulation overhead, supports native IPv6 multicast, and allows application

preference to follow DNS responses. ISATAP and manual tunnels introduce overlay complexity and are primarily transition mechanisms for connecting IPv6 islands over IPv4 infrastructure. A service block model may be useful for specific service insertion, but it does not provide a whole-network migration strategy. Since every current infrastructure device supports IPv6, the cleanest and most operationally straightforward model is to enable IPv6 alongside IPv4 and migrate services gradually. Reference topics: IPv6 migration, dual-stack deployment, DNS-based protocol selection, native IPv6 multicast, transition strategy design.

QUESTION NO: 19

エンジニアは、レイヤ3ルータが1つしかない小さなブランチサイト用にEIGRPネットワークを設計しています。エンジニアは、ルータがローカルLAN上で不要なマルチキャストメッセージを送信することなく、リモートEIGRPネイバーにローカルLANネットワークをアドバタイズすることを望んでいます。エンジニアはどのような行動を取るべきですか？

- A. EIGRPの代わりに、このサイトに静的なデフォルトルートを使用します
- B. networkコマンドとパッシブインターフェイス機能を使用してローカルLANをアドバタイズします
- C. redistribute connectedコマンドを使用してローカルLANネットワークを再配布します
- D. ローカルLANサブネットをスタブネットワークとしてアドバタイズします

Answer: B

Explanation:

The branch router should advertise the local LAN with the EIGRP network command and make the LAN-facing interface passive. A passive interface allows the connected network to be included in EIGRP advertisements while suppressing EIGRP hello packets and neighbor formation on that interface. This is exactly what the engineer needs: remote EIGRP neighbors learn the LAN subnet, but unnecessary multicast EIGRP messages are not sent onto the user LAN where no EIGRP neighbors should exist. Replacing EIGRP with a static default route would not meet the requirement to advertise the local LAN through EIGRP. Redistributing connected routes can work, but it is less clean for a simple LAN interface and can introduce external-route behavior or policy complexity. Advertising the subnet as a stub network does not suppress multicast hellos on the LAN interface by itself. Cisco design best practice is to use passive interfaces on LAN-facing interfaces that should be advertised but should not form routing adjacencies. Reference topics: EIGRP passive interface, network statements, branch routing, unnecessary neighbor prevention, multicast control traffic.

QUESTION NO: 20

同じ2つのBGP機能のうち、同じAS番号を共有するeBGPネイバー間のルート交換が成功するのはどれですか。（2つ選択してください。）

- A. advertise-best-external
- B. bestpath as-path ignore
- C. client-to-client reflection
- D. as-override
- E. allow-as-in

Answer: D E

Explanation:

When eBGP neighbors share the same autonomous system number, normal BGP loop-prevention behavior can block route acceptance because the receiving router sees its own AS in the AS_PATH. Two BGP features address this design condition. The allow-as-in feature permits a BGP router to accept routes even when its own AS appears in the AS_PATH a configured number of times. This is often used in dual-homed or migration designs where the same customer AS appears at multiple sites. The as-override feature is typically used by a provider edge router to replace the customer AS in the AS_PATH with the provider AS before advertising the route to another customer site that uses the same AS. Together, these features allow successful route exchange in same-AS eBGP scenarios while preserving loop-prevention intent through controlled policy. Advertise-best-external influences advertisement of best external paths and does not solve same-AS rejection. Bestpath as-path ignore affects path selection, not route acceptance. Client-to-client reflection is an iBGP route-reflector behavior, not an eBGP same-AS solution.

QUESTION NO: 21

エンジニアは、定義済みのビジネス要件に基づいてトラフィックをクラスに分離できるスケラブルな QoS

アーキテクチャを設計する必要があります。また、設計では、差別化サービスコードポイントを QoS 優先度記述子値として使用し、少なくとも 10 レベルの分類をサポートする必要があります。エンジニアはどの QoS テクノロジーを設計に含める必要がありますか？

- A. 返信
- B. ディフサーブ
- C. ベストエフォート
- D. インターサーブ

Answer: B

Explanation:

DiffServ is the correct scalable QoS architecture for classifying traffic according to business requirements and using DSCP values as the priority descriptor. Cisco QoS design distinguishes DiffServ from IntServ. IntServ uses per-flow signaling with RSVP and can reserve resources, but that model is not operationally scalable across large enterprise networks. DiffServ is class-based. Traffic is classified and marked at the network edge, then each device applies per-hop behavior to those classes using local QoS policy. DSCP provides a six-bit field in the IP header, which supports more than ten classification values and allows standardized treatment classes such as EF, AF, class selector, and default. Best effort provides no differentiated treatment, and RSVP is a signaling mechanism associated with IntServ rather than the requested scalable class-based architecture.

"Interserv" is not the correct Cisco QoS model term. Therefore, the engineer should design around DiffServ, with a clear trust boundary, marking policy, queuing model, and consistent treatment across WAN and campus boundaries. Reference topics: DiffServ, DSCP, per-hop behavior, class-based QoS, scalable QoS architecture.

QUESTION NO: 22

NETCONFでは、RESTCONFと比較して、どのエンコーディング言語がサポートされていますか？

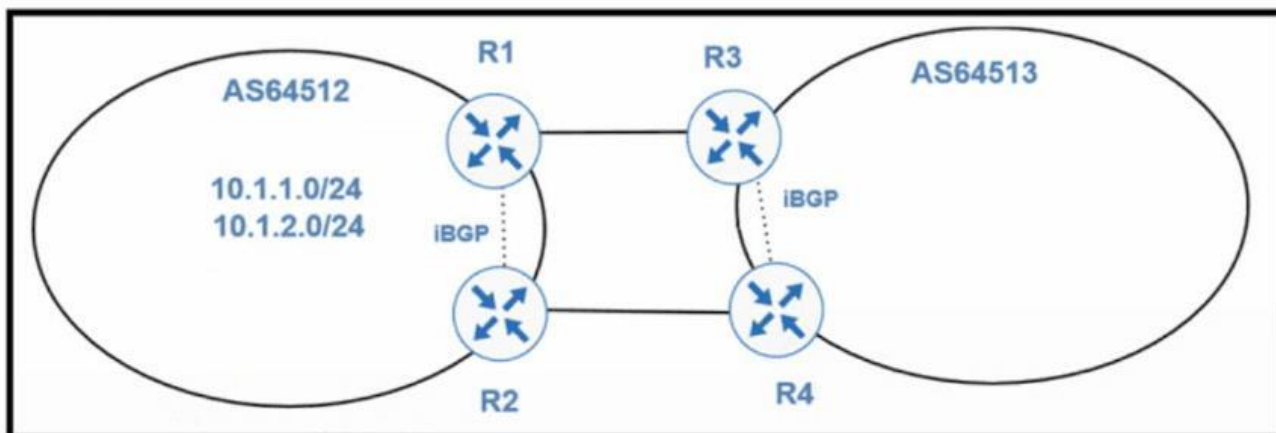
- A. NETCONFはXMLとJSONをサポートし、RESTCONFはXMLをサポートします。
- B. NETCONFはXMLをサポートし、RESTCONFはJSONをサポートします。
- C. NETCONFはJSONをサポートし、RESTCONFはXMLをサポートします。
- D. NETCONFはXMLをサポートし、RESTCONFはXMLとJSONをサポートします。

Answer: D

Explanation:

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.

QUESTION NO: 23



図を参照してください。アーキテクトは、上流のサービスプロバイダーに接続するリンクの負荷分散を必要とする顧客向けに BGP ポリシーを設計します。顧客には次の要件があります。* ネットワーク 10.1.1.0/24 宛ての着信トラフィックは R3-R1 リnkを通過する必要があり、リンクに障害が発生した場合は、すべての着信トラフィックが R4-R2 リnkを通過する必要があります。

* ネットワーク 10.1.2.0/24 宛ての着信トラフィックは R4-R2 リnkを通過する必要があり、リンクに障害が発生した場合は、すべての着信トラフィックは R3-R1 リnkを通過する必要があります。

建築家はどのソリューションを選択する必要がありますか？

- A. * R1は、set as-path prepend 64512 64512を使用して、ネイバーにルートマップを適用したプレフィックス10.1.2.0/24をアナウンスする必要があります。

* R2 は、set as-path prepend 64512 64512 を使用して、ネイバーにルートマップを適用したプレフィックス 10.1.1.0/24 をアナウンスする必要があります。

B. * R1 は、コミュニティ属性 64513:300 を持つプレフィックス 10.1.2.0/24 と、コミュニティ属性 64513:200 を持つプレフィックス 10.1.1.0/24 をアナウンスする必要があります。

* R2 は、コミュニティ属性 64513:200 を持つプレフィックス 10.1.2.0/24 と、コミュニティ属性 64513:300 を持つプレフィックス 10.1.1.0/24 をアナウンスする必要があります。

C. * R1 は、set as-path prepend 64512 64512 を使用して、ネイバーに適用されたルートマップとともにプレフィックス 10.1.1.0/24 をアナウンスする必要があります。

* R2 は、set as-path prepend 64512 64512 を使用して、ネイバーにルートマップを適用したプレフィックス 10.1.2.0/24 をアナウンスする必要があります。

D. * R1 は、コミュニティ属性 64513:200 を持つプレフィックス 10.1.2.0/24 と、コミュニティ属性 64513:300 を持つプレフィックス 10.1.1.0/24 をアナウンスする必要があります。

* R2 は、コミュニティ属性 64513:300 を持つプレフィックス 10.1.2.0/24 と、コミュニティ属性 64513:200 を持つプレフィックス 10.1.1.0/24 をアナウンスする必要があります。

Answer: A


Explanation:

The correct policy is to prepend the AS path on the nonpreferred advertisement for each prefix. For inbound BGP traffic engineering, the enterprise normally influences remote autonomous systems by changing attributes that those systems see when selecting a path. AS-path prepending makes a route appear less attractive by adding repeated instances of the local AS number to the AS path. In this design, network 10.1.1.0/24 should enter through R3-R1, so R2 must advertise that prefix with a prepended AS path to make the R4- R2 path less preferred unless the primary path fails. Conversely, network 10.1.2.0/24 should enter through R4- R2, so R1 must advertise that prefix with prepending to make the R3-R1 path less preferred. This creates inbound load sharing while preserving failover because the prepended route remains available as a backup.

Local preference affects outbound path selection inside the local AS, not inbound selection by the provider.

Reference topics: BGP traffic engineering, AS-path prepending, inbound path control, prefix-specific policy, multihomed WAN design.

QUESTION NO: 24



図を参照してください。顧客は、OSPF を実行している R1 と R3 間のデータ

プレーンの最大稼働時間を要求しています。R2 のルーティングプロセスにメンテナンスが必要な場合、高可用性のために設計にどのソリューションを含める必要がありますか。

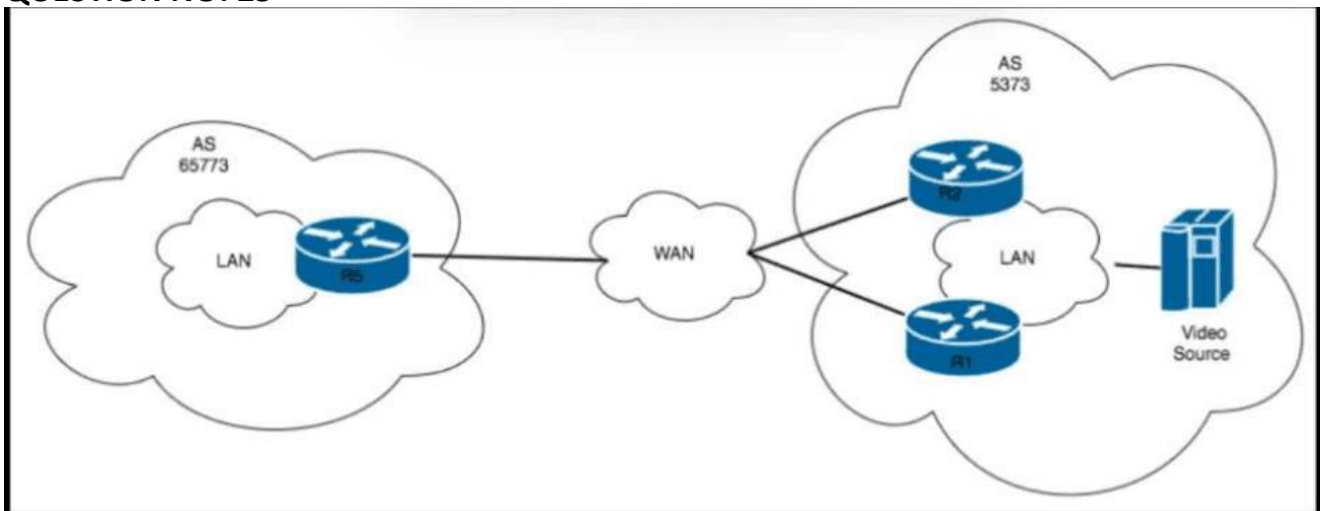
- A. すべてのルータ上のBFD
- B. R1 と R3 でのノンストップ転送
- C. R3のみのノンストップ転送
- D. すべてのルータでグレースフルリスタート

Answer: D

Explanation:

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.

QUESTION NO: 25



図を参照してください。VoD コンテンツ作成を専門とする会社には、WAN リンクで接続された別のマルチキャスト ドメインに 2 つのオフィスがあります。オフィス間では BGP 通信が確立されています。クライアントは各オフィスの LAN 内にあります。AS5373 では、R2 が RP として選択されています。AS65773 でクライアントに VoD コンテンツを配信するには、ネットワーク アーキテクトは何を設計する必要がありますか。

- A. MSDP

- B. Auto-RP を備えた PIM ASM
- C. PIM SSM
- D. BSR 付き PIM ASM

Answer: A

Explanation:

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.

QUESTION NO: 26

IPv4のみのネットワークトポロジを使用しているお客様が、IPv4トポロジサービスを維持しながらIPv6接続を有効にしたいと考えています。お客様は、IPv4サービスをIPv6トポロジに移行してから、IPv4トポロジを廃止する予定です。これらの要件をサポートするトポロジはどれですか？

- A.デュアルスタック
- B.6VPE
- C.6to4
- D.NAT64

Answer: A

Explanation:

Dual stack is the migration topology that supports IPv6 while preserving existing IPv4 services during the transition. In a dual-stack design, hosts and network devices run IPv4 and IPv6 in parallel, allowing applications and services to move gradually from IPv4 to IPv6 without forcing an immediate cutover. That matches the customer's plan: keep the current IPv4 topology and services, introduce IPv6 connectivity, migrate services, and eventually decommission the IPv4 topology. 6VPE is used to carry IPv6 VPN routes across an MPLS IPv4 provider core, which is a service-provider VPN use case rather than a general enterprise migration topology. 6to4 is an automatic tunneling mechanism and is not appropriate as a stable enterprise migration plan. NAT64 enables IPv6-only clients to reach IPv4 resources through translation, but it does not preserve the complete IPv4 topology while services migrate. Cisco migration guidance treats dual stack as the cleanest operational model when infrastructure supports both protocols, because it avoids translation side effects and allows DNS and application behavior to determine which protocol is used.

QUESTION NO: 27

Cisco SD-Access アーキテクチャにおけるファブリック コントロール プレーンの目的は何ですか？

- A. ファブリック内で G6AC ポリシーを作成、伝播、適用する
- B. BGPルートリフレクタ機能を備えたトランジットノードを作成する
- C. 複数のサブネットを1つのRLOCに拡張する
- D. エンドポイントと場所のマッピングを作成して解決する

Answer: D

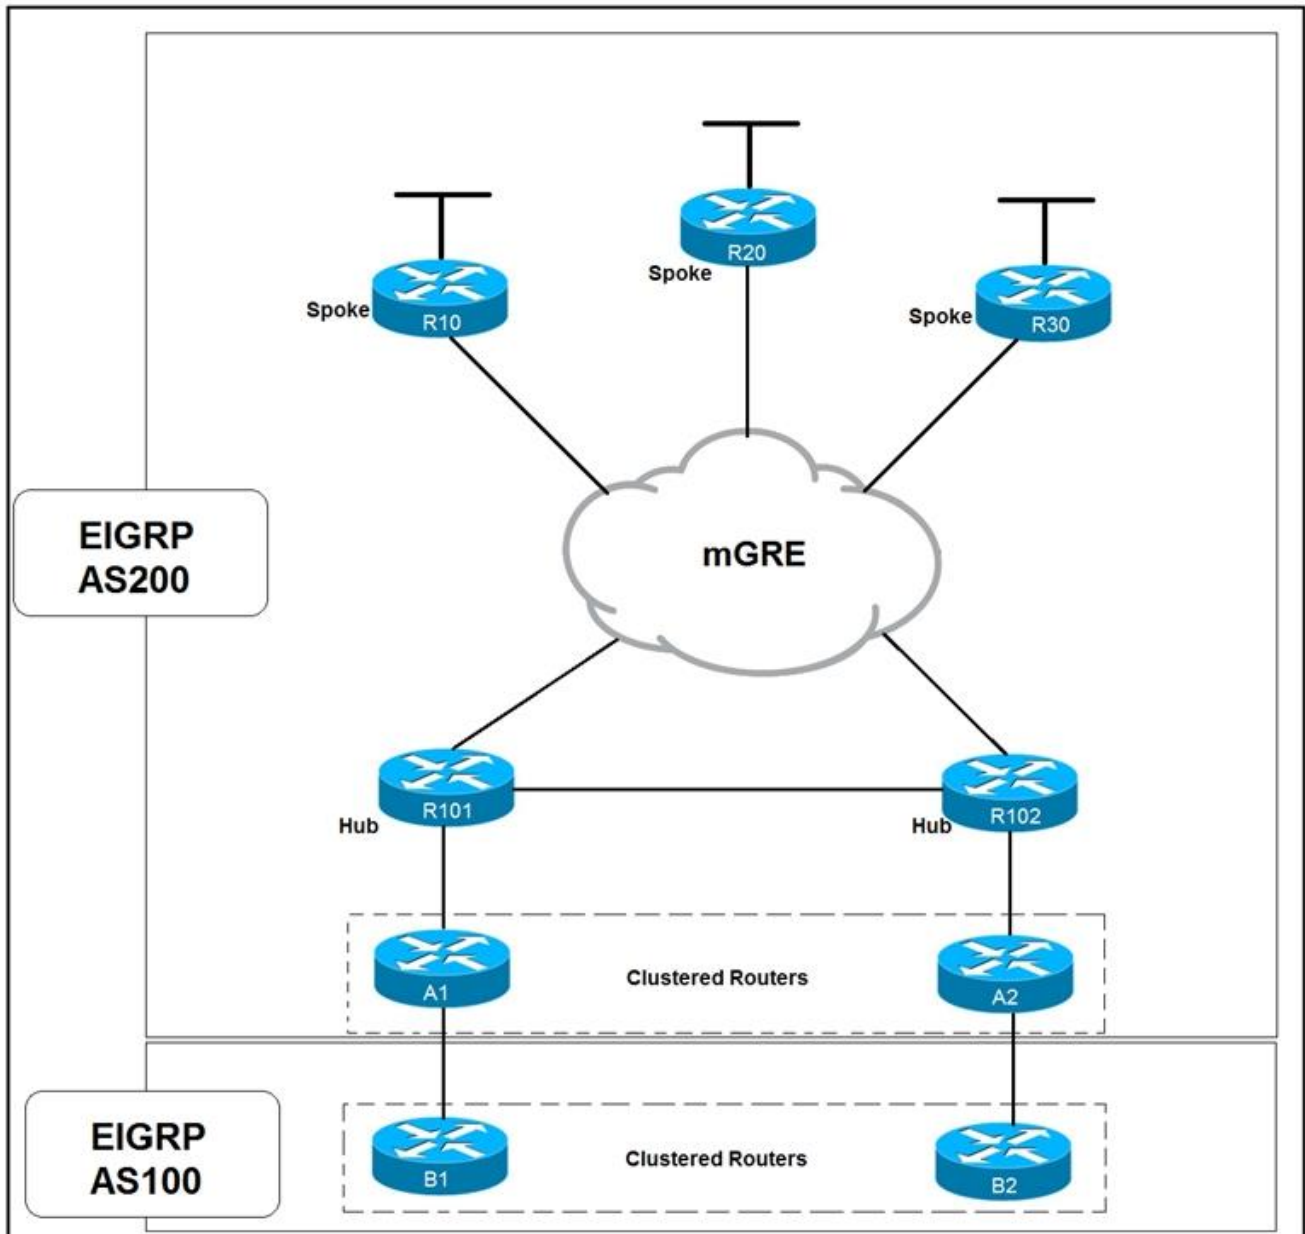
Explanation:

The fabric control plane in Cisco SD-Access creates and resolves endpoint-to-location mappings. SD-Access separates endpoint identity from physical topology by using LISP in the control plane. Endpoint addresses are treated as Endpoint Identifiers, and fabric node locations are represented by Routing Locators. When an endpoint appears on a fabric edge node, the edge registers that endpoint-to-edge binding with the control-plane node. When another edge needs to send traffic to that endpoint, it queries the control plane to resolve the current location and then uses the fabric data plane to encapsulate traffic toward the correct destination.

Group-Based Access Control policy creation and enforcement involve Cisco ISE, security group tags, and policy components rather than the basic mapping purpose of the fabric control plane. BGP route reflection is not the SD-Access control-plane function, and extending multiple subnets to one RLOC is not the main purpose. Reference topics: SD-Access control plane, LISP, EID-to-RLOC mapping, endpoint registration, map resolution.

QUESTION NO: 28

展示品を参照してください。



どのソリューションが EIGRP 収束時間を短縮しますか？

- A. 1秒未満のタイマーを有効にする
- B. ホールドタイムの値を増やす
- C. デッドタイム値を増やす
- D. スポーク上でスタブルーティングを有効にする

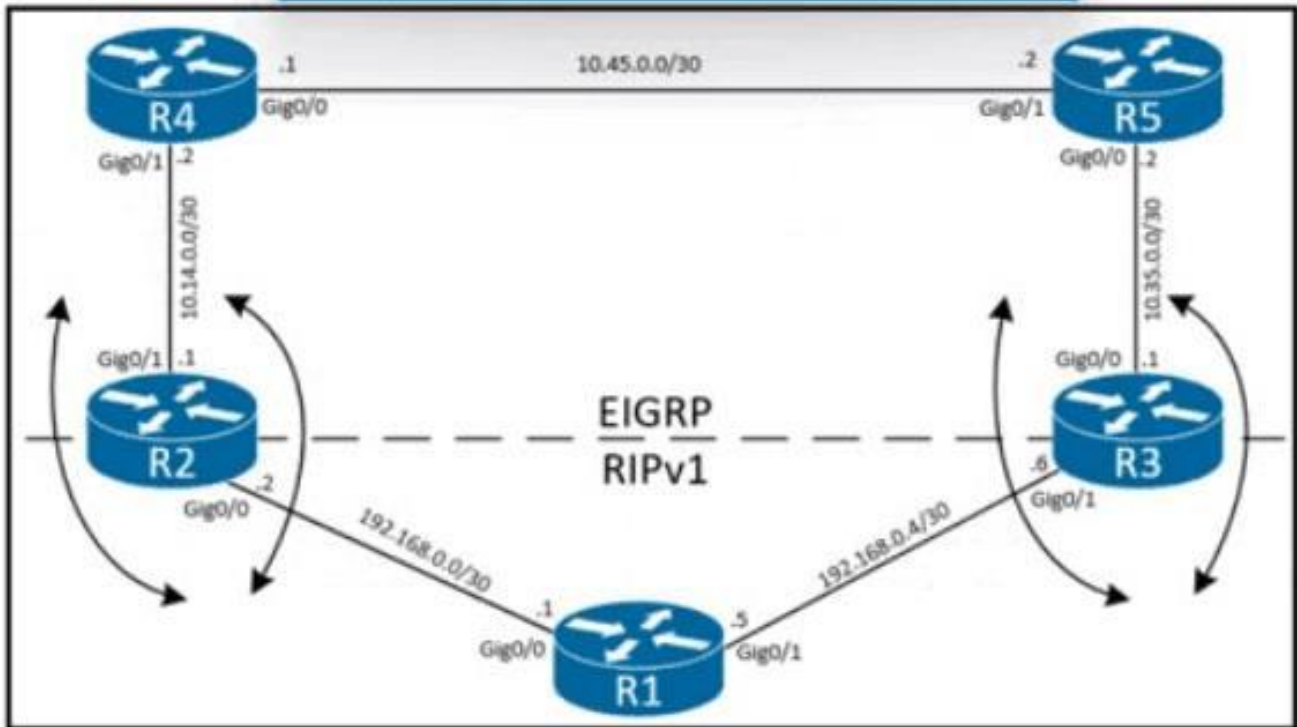
Answer: D

Explanation:

Enabling EIGRP stub routing on the spokes decreases convergence time in a hub-and-spoke design by reducing unnecessary queries and preventing spokes from being treated as transit routers. EIGRP convergence can be delayed when a router loses a route and must query many neighbors to determine whether an alternate path exists. In remote-site or spoke designs, a spoke normally has no useful alternate path for networks beyond itself, so querying it wastes time and can contribute to stuck-in-active conditions. Cisco EIGRP stub routing allows the spoke to advertise only permitted route types and informs upstream routers not to query the spoke for routes outside its scope. This creates a cleaner query boundary

and accelerates convergence after link or route failures. Subsecond timers can detect neighbor loss more quickly, but they do not address the broader EIGRP query behavior shown in the design. Increasing hold or dead timers would slow detection, not improve convergence. Therefore, in the displayed hub-and-spoke EIGRP design, enabling stub routing on the spokes is the correct solution to decrease convergence time.

QUESTION NO: 29



展示を参照してください。エンジニアが顧客向けの再配布ソリューションを設計しています。顧客は最近別の会社を買収し、RIPv1
を実行する新しいネットワークを会社の既存のネットワークと統合することを決定しました。マルチポイント双方向再配布によってルーティング
ループが発生しないようにするには、エンジニアはどの再配布手法を選択する必要がありますか。

- A. EIGRPプロセス下での配布リストの受信はRIPv1で学習したプレフィックスを拒否する
- B. EIGRPプロセスの下でアウトバウンドに配布リストを配布し、RIPv1で学習したプレフィックスを拒否する
- C. RIPv1プロセスの下でアウトバウンドに配布リストを配布し、EIGRP学習プレフィックスを拒否する
- D. RIPv1プロセス下での配布リストの受信はEIGRP学習プレフィックスを拒否する

Answer: C

Explanation:

The correct choice is to use a distribute list outbound under the RIP process to stop EIGRP-learned prefixes from being advertised back into the RIP domain. In multipoint two-way redistribution, the major risk is route feedback: a route learned from one protocol is redistributed into another protocol and then reintroduced back into the original domain

through another redistribution point. RIPv1 is especially limited because it is classful and has no native route tagging, so filtering must be carefully placed when route tags are unavailable. By filtering outbound advertisements under the RIP process, the design prevents routes that originated in the EIGRP domain from being sent into RIP where they could be relearned elsewhere and create loops or suboptimal paths. Inbound filtering on only one process does not reliably stop feedback at all redistribution points. A stronger modern design would use route tags with protocols that support them, but given the answer choices, outbound filtering toward RIP is the appropriate loop-prevention method. Reference topics: route redistribution, distribute lists, route feedback, RIPv1 limitations, EIGRP integration.

QUESTION NO: 30

エンジニアは、ISP にシングルホーム接続している企業向けにルーティングソリューションを設計する必要があります。この企業の目標は、CE デバイスと PE デバイス間で BGP を実行することです。BGP の実行をサポートするために、この企業は ARIN からパブリック AS 番号と IP サブネットを取得しました。エンジニアはどのソリューションを選択する必要がありますか？

- A. * 顧客はパブリックIPサブネットをISPに通知します
* ISP は顧客にデフォルト ルートを通知します。
- B. * 顧客はパブリックIPサブネットをISPに通知します
* ISPは顧客にBGPテーブルを通知する
- C. * ISP は顧客のパブリック IP サブネットをアナウンスします。
* ISP は部分的な BGP テーブルを顧客に通知します。
- D. * 顧客はISPにデフォルトルートを通知する
* ISPは顧客にデフォルトルートを通知します

Answer: A

Explanation:

A single-homed enterprise that runs BGP with one ISP should advertise its assigned public prefix to the ISP and receive a default route from the ISP. Because there is only one upstream provider, the customer does not need the full Internet table to make path-selection decisions; all Internet destinations exit the same provider path. Receiving only a default route conserves memory and CPU on the CE router and simplifies operations.

The customer must still originate its public prefix toward the ISP so the provider can advertise reachability to the Internet. The ISP should not be the originator of the customer-owned ARIN prefix toward the customer; it may advertise that prefix to the Internet after receiving it from the customer. Requesting the full BGP table or even a partial table adds unnecessary scale for a single-homed site. Advertising a default route to the ISP is also wrong because the ISP already owns upstream Internet reachability. Reference topics: single-homed BGP, CE-PE peering, default route, public prefix advertisement, Internet edge design.

QUESTION NO: 31

ネットワーク エンジニアは、IP アドレス 172.16.15.12/32 のループバック インターフェイスを構成するスクリプトを準備します。会社のセキュリティポリシーに準拠するには、'Content-type' を使用します。スクリプトに「application/yang-data+json」が追加されました。ネットワーク

デバイスへの接続を保護する必要があります。
この要件を満たすために、ネットワーク エンジニアはどのコード
スニペットを使用する必要がありますか？

```
{
  "ietf-interfaces:interface": {
    "name": "Loopback0",
    "type": "iana-if-type:softwareLoopback",
    "enabled": true,
    "ietf-ip:ipv4": {
      "address": [
        {
          "ip": "172.16.15.12",
          "netmask": "255.255.255.255"
        }
      ]
    },
    "ietf-ip:ipv6": {}
  }
}
```

```
{
  "interface": "ietf-loopback" {
    "name": "Loopback0",
    "enabled": true,
    "address": "ipv4"
    [
      {
        "ip": "172.16.15.12",
        "netmask": "255.255.255.255"
      }
    ]
  },
  "address : "ipv6": {}
}
```

```
{
  "ietf-interfaces:interface": {
    "name": "Loopback0",
    "type": "iana-if-type:softwareLoopback",
    "enabled": true,
    "address_ipv4": {
      {
        "ip": "172.16.15.12",
        "netmask": "255.255.255.255"
      }
    }
  },
  "address_ipv6": {null}
}
```

```
{
  "ietf-interfaces:interface": {
    "name": "Loopback0",
    "type": "iana-if-type:softwareLoopback",
    "enabled": true,
    "address_ipv4": {
      {
        "ip": "172.16.15.12",
        "netmask": "0.0.0.0"
      }
    }
  },
  "address_ipv6": {}
}
```

- A. オプションA
- B. オプションB
- C. オプションC
- D. オプションD

Answer: A

Explanation:

The correct code snippet must use a secure HTTPS-based RESTCONF transaction with the YANG JSON media type. The requirement explicitly states that the content type is application/yang-data+json and that the connection to the network device must be secured. In Cisco IOS XE programmability, RESTCONF uses HTTP methods against YANG-modeled resources, and secure deployments use HTTPS rather than clear-text HTTP. A loopback

interface with address 172.16.15.12/32 would be configured by sending a JSON payload that follows the selected YANG model structure, commonly an IETF or Cisco native interface model, with the correct content-type and authentication headers. The key security indicator in the option must therefore be HTTPS/TLS, not an insecure HTTP URL or a non-YANG content type. NETCONF over SSH would also be secure, but the stated content-type points to RESTCONF with JSON encoding. Option A is the valid snippet because it aligns the payload format with a secure transport. Reference topics: RESTCONF, YANG JSON encoding, application/yang-data+json, HTTPS, IOS XE programmability.

QUESTION NO: 32

コストが等しくないロードバランシングを可能にする2つのルーティングプロトコル (2つ選択してください。)

- A.EIGRP
- B.IS-IS
- C.BGP
- D.OSPF
- E.RIPng

Answer: A C

Explanation:

EIGRP and BGP are the two listed routing protocols that can support unequal-cost load balancing in Cisco enterprise designs. EIGRP performs unequal-cost load sharing by using the variance command. Feasible successor routes whose feasible distance falls within the variance multiplier can be installed alongside the successor route, provided they satisfy EIGRP loop-prevention rules. This allows traffic to use links with different composite metrics while still maintaining loop-free forwarding. BGP can also support unequal-cost distribution in specific designs, such as using BGP multipath features with additional bandwidth-based mechanisms or policy to influence traffic sharing across nonidentical paths. OSPF and IS-IS commonly support equal-cost multipath, but they do not natively perform EIGRP-style unequal-cost load balancing based on different path metrics. RIPng is also limited by hop-count behavior and is not used for unequal-cost multipath in this context. The correct design answer is therefore EIGRP and BGP. In production, the engineer must still validate hardware forwarding behavior, policy consistency, and whether the traffic-sharing method is per-flow or per-prefix.

QUESTION NO: 33

アーキテクトは、企業ネットワーク機器を管理するための計画を設計する必要があります。その設計には、以下の点を考慮する必要があります。

- * すべてのネットワーク機器に専用の管理インターフェースがあるわけではありません
- * すべてのデバイスのすべてのIP対応インターフェースに到達可能である必要があります
- * 暗号化は、サポートしているすべてのデバイスで使用する必要があります

建築家はどちらの解決策を選択すべきか？

- A. KVMサーバー
- B. インバンド
- C. 帯域外
- D. ターミナルサーバー

Answer: B

Explanation:

In-band management is the correct design because the management plan must reach all IP-enabled interfaces, and not every device has a dedicated management port. In-band management uses the production IP network for administrative access, allowing management stations to reach loopbacks, SVIs, routed interfaces, or front-panel interfaces through normal routing. Cisco management guidance distinguishes in-band management from out-of-band designs that depend on a separate physical management network or console path. Because the question requires reachability to all IP-enabled interfaces, in-band is the only listed model that naturally supports that scope. Security must still be enforced with encrypted protocols where supported, such as SSH and HTTPS, plus ACLs, AAA, and management-plane protection where appropriate. A terminal server provides console access, not routed IP management to every interface. A KVM server is a server-management tool, not a network-device management architecture. Out-of-band management is preferred for isolation, but it does not fit devices without dedicated management access. Reference topics: in-band management, secure management protocols, SSH, HTTPS, network management design.

QUESTION NO: 34

エンジニアは、データセンター内のルーターにすべてのWANルートがアドバタイズされないように、EIGRPを設定する必要があります。どのような操作を行うべきでしょうか？

- A. スタブルータを受信専用モードに設定します。
- B. デフォルトルートのみを通知します。
- C. ローカルサブネットを要約します。
- D. スタブルータを分散モードで設定します。

Answer: A

Explanation:

The action is to configure the EIGRP stub router in receive-only mode. Cisco EIGRP stub routing limits which routes a router advertises to neighbors, reducing query scope and improving stability in hub-and-spoke or WAN designs. The receive-only keyword is the most restrictive stub option: the router receives routes from neighbors but does not advertise any routes to them. That directly satisfies the requirement that WAN routes must not be advertised to the routers in the data center. Advertising only a default route still advertises routing information, so it does not meet the wording. Summarizing local subnets reduces the number of routes, but it still advertises reachability. The "distributed mode" option is not the correct EIGRP stub behavior for this requirement; redistributed stub mode would allow redistributed routes to be advertised. Receive-only is therefore the precise design control for stopping route advertisement from the stub router. Reference topics: EIGRP stub routing, receive-only mode, query scoping, route advertisement control, WAN-to-data-center design.

QUESTION NO: 35

Cisco SD-Access

では、仮想ネットワークによってセグメンテーションが作成され、ユーザーとリソースを分離できます。

このタイプのセグメンテーションはどのように説明されますか？

- A. マクロ
- B. ベトナム語間
- C. マイクロ
- D. 伸ばした

Answer: A

Explanation:

Virtual networks in Cisco SD-Access provide macro segmentation. Macro segmentation separates major user, device, or service groups into distinct virtual networks, each with its own routing and forwarding context.

This is comparable to VRF-based separation in traditional enterprise designs. For example, corporate users, guests, building systems, and IoT devices can be placed into separate virtual networks so their routing domains remain isolated unless controlled inter-VN communication is explicitly provided. Microsegmentation is different; it is normally enforced inside or across virtual networks using Security Group Tags and Security Group ACLs, allowing policy decisions based on group identity rather than only IP subnets or VLANs. Inter-VN describes communication between virtual networks, not the segmentation type itself.

Stretched segmentation is not the standard SD-Access term. Therefore, the separation created by virtual networks is macro segmentation, while SGT-based policy provides finer-grained microsegmentation. Reference topics:

Cisco SD-Access virtual networks, macro segmentation, VRF, SGT, microsegmentation, policy enforcement.