ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

What is a benefit of using VXLANs in a cloud-scale environment?

A. extends Layer 2 segments across the underlying Layer 3 infrastructure
B. extends Layer 3 segments across the underlying Layer 2 infrastructure
C. reduces spanning-tree complexity across the Layer 2 infrastructure
D. eliminates the need for a Layer 3 underlay in the service provider infrastructure

Answer: A

Explanation:
In a cloud-scale or data center–scale environment, Virtual Extensible LAN (VXLAN) is used as an overlay technology to transport Layer 2 segments over a Layer 3 underlay network. VXLAN encapsulates Layer 2 Ethernet frames inside UDP/IP packets, allowing broadcast, unknown unicast, and multicast (BUM) traffic and tenant Layer 2 domains to be extended across a routed IP fabric. Key points aligned with Cisco Service Provider Cloud Infrastructure design principles:
VXLAN creates a Layer 2 overlay on top of a Layer 3 underlay.
The VXLAN Network Identifier (VNI) provides a much larger segmentation space than traditional VLANs, enabling multi-tenancy at cloud scale.
Because the underlay is pure Layer 3 (IP routed fabric), VXLAN allows you to interconnect Layer 2 segments between leaf switches or data centers over an IP/MPLS backbone without relying on large Layer 2 domains in the physical network.
Why the options evaluate as follows:
Option A: extends Layer 2 segments across the underlying Layer 3 infrastructure Q
This is the core benefit of VXLAN in cloud-scale designs. VXLAN encapsulates Layer 2 frames into IP/UDP headers, allowing isolated Layer 2 segments (per VNI) to be stretched across a routed IP network. This enables:
Multi-tenant Layer 2 connectivity across a distributed cloud fabric
Mobility of virtual machines or containers while keeping same IP/MAC addressing
Use of an IP-based leaf–spine or service provider underlay for scalability and resiliency
Option B: extends Layer 3 segments across the underlying Layer 2 infrastructure X
This is the opposite of what VXLAN does. VXLAN is explicitly L2-over-L3, not L3-over-L2. Extending pure Layer 3 segments over Layer 2 is not the VXLAN use case.
Option C: reduces spanning-tree complexity across the Layer 2 infrastructure ! (Partially related but not the primary or direct benefit)
In modern designs, the underlay is Layer 3 routed, and VXLAN overlays provide logical Layer 2 segments. This design avoids dependence on spanning tree in the fabric, which indirectly reduces
STP complexity. However, the fundamental, exam-relevant benefit is L2 extension over L3, so C is not the best or most accurate answer compared to A.
Option D: eliminates the need for a Layer 3 underlay in the service provider infrastructure X
VXLAN absolutely requires an IP (Layer 3) underlay for transport. VXLAN tunnels are built over a routed infrastructure (leaf–spine, MPLS/IP core, etc.). It does not remove the need for Layer 3; it depends on it.

Question: 2

An engineer must configure NTP servers in Cisco Enterprise NFVIS. The primary NTP server has an IP address of 192.168.1.1 and the backup NTP server has an IP address of 192.168.2.1. Which two commands must be run to complete the configuration? (Choose two.)

A. system time ntp preferred_server 192.168.1.1

B. utils ntp server add 192.168.2.1 backup
C. system set-manual-time 192.168.1.1 192.168.2.1
D. utils ntp server add 192.168.1.1 primary
E. system time ntp backup_server 192.168.2.1

Answer: A, E

Explanation:

In Cisco Enterprise NFVIS, time synchronization is configured using the system time ntp command structure. NFVIS requires a primary and optionally a backup NTP server to maintain accurate system time for the hypervisor and guest VMs.
Correct NFVIS command syntax for NTP configuration:
system time ntp preferred_server <IP>
This command configures the preferred (primary) NTP server used for system clock synchronization. system time ntp backup_server <IP>
This command configures the backup NTP server, which the system uses if the primary becomes unreachable.
These two commands match Cisco NFVIS time-configuration behavior described in NFV infrastructure design and implementation guidelines.

Why the Correct Answers Are A and E
Option A: system time ntp preferred_server 192.168.1.1
This properly configures the primary NTP server in NFVIS. The preferred server is always the first choice for time synchronization.
Option E: system time ntp backup_server 192.168.2.1
This correctly configures the backup NTP server. If the preferred server fails, NFVIS automatically falls back to the backup server.
Both commands directly match NFVIS's NTP command hierarchy and are the only ones that correctly apply to NFVIS.

Why the Other Options Are Not Correct
Option B uses utils ntp, which is not an NFVIS command.
Option C sets manual time and does not configure NTP servers.
Option D also uses the utils ntp syntax, which applies to other Cisco platforms but not NFVIS.

Question: 3

Refer to the exhibit.

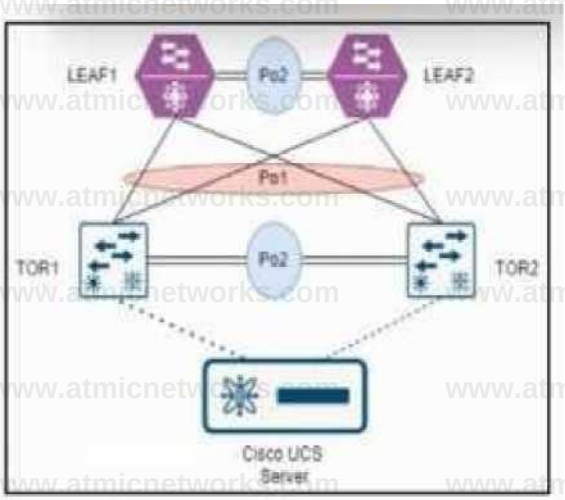

An engineer must connect switch TOR1 and switch TOR2 to switch LEAF1 and switch LEAF2 by using double-sided vPCs. LEAF1 and LEAF2 are already configured as vPC peers. Which action must be taken next to complete the configuration?

A. Add all the switches to the fabric.
B. Configure peering between LEAF1 and LEAF2 and TOR1 and TOR2.
C. Configure MSTP between TOR1 and TOR2.
D. Configure a vPC between TOR1 and TOR2.

Answer: D

Explanation:

In Cisco data center and cloud-scale designs, a double-sided vPC (also called vPC-to-vPC) is used when both ends of a Layer 2 port channel are formed by a pair of switches that operate as vPC peers. In this model:

On the aggregation or leaf side, two switches (in this case, LEAF1 and LEAF2) form a vPC domain with a vPC peer-link and keepalive.

On the access or ToR side, two switches (in this case, TOR1 and TOR2) must also form their own vPC domain with a peer-link and vPC keepalive.

The port-channel that interconnects the two vPC domains is then configured as a vPC on both sides, creating a vPC-to-vPC topology.

The problem statement specifies that LEAF1 and LEAF2 are already configured as vPC peers. For a double-sided vPC to work, the other side (TOR1 and TOR2) must also behave as a single logical entity for the downstream Cisco UCS server and for the upstream vPC connection towards LEAF1 and LEAF2. This is only achieved when TOR1 and TOR2 are configured as vPC peers with:

A vPC domain ID
A vPC peer-link between TOR1 and TOR2
vPC member port-channels towards LEAF1 and LEAF2 and towards the Cisco UCS server

Therefore, the next required step is to configure a vPC between TOR1 and TOR2.

Evaluation of the options:

Option A, “Add all the switches to the fabric,” is generic and not specific to vPC configuration. It does not address the technical requirement to form a vPC domain on the ToR side.

Option B, “Configure peering between LEAF1 and LEAF2 and TOR1 and TOR2,” is incorrect because vPC peering is only configured between the two switches that form each vPC domain (LEAF1–LEAF2 and TOR1–TOR2), not across all four switches together.

Option C, “Configure MSTP between TOR1 and TOR2,” is not required for establishing a double-sided vPC. vPC designs rely on the vPC control plane and the peer-link, not on spanning-tree between the vPC peers for normal operation.

Option D, “Configure a vPC between TOR1 and TOR2,” correctly describes configuring TOR1 and TOR2 as a vPC pair (vPC domain

with peer-link), which is the mandatory step to create a double-sided vPC topology with LEAF1 and LEAF2.

Question: 4

What is a valid connection method between carrier-neutral facilities within the same metro area?

A. OSPF backbone area adjacency
B. private wireless connection
C. DWDM ring
D. CAT6e connection

Answer: C

Explanation:

Comprehensive and Detailed Explanation Based on Designing and Implementing Cisco Service Provider Cloud Network Infrastructure Knowledge

When connecting carrier-neutral facilities (CNFs) or data centers within the same metropolitan area, service providers typically use high-bandwidth, low-latency optical transport methods. The most appropriate and commonly deployed interconnection technology is:

DWDM (Dense Wavelength Division Multiplexing) ring, which provides:

High capacity (10G, 40G, 100G, 400G)

Low latency

Redundancy through ring or mesh topologies

Multi-wavelength multiplexing for cost efficiency

Carrier-grade reliability for metro interconnect services

This aligns with cloud interconnect and metro transport design used in service provider environments.

Evaluation of the Options

A . OSPF backbone area adjacency

This is a routing protocol adjacency, not a physical connection method. It requires a transport link underneath but does not represent the physical interconnect itself.

B . Private wireless connection

Not suitable for CNF or metro DC interconnect because it lacks the bandwidth, reliability, and deterministic performance required for large-scale carrier-grade interconnects.

C . DWDM ring

This is the correct method. DWDM-based metro fiber rings are the standard for connecting carrier neutral facilities in the same metro region.

E. . CAT6e connection

This is limited to short-distance copper Ethernet (tens of meters). It is not used for metro-scale interconnects or between CNFs.

Question: 5

Refer to the exhibit.

```
<output omitted>
!
interface overlay 1
 otv join-interface g1/0
 otv control-group 224.1.1.1
 otv data-group 232.1.1.0/24
 no shutdown
!
<output omitted>
```

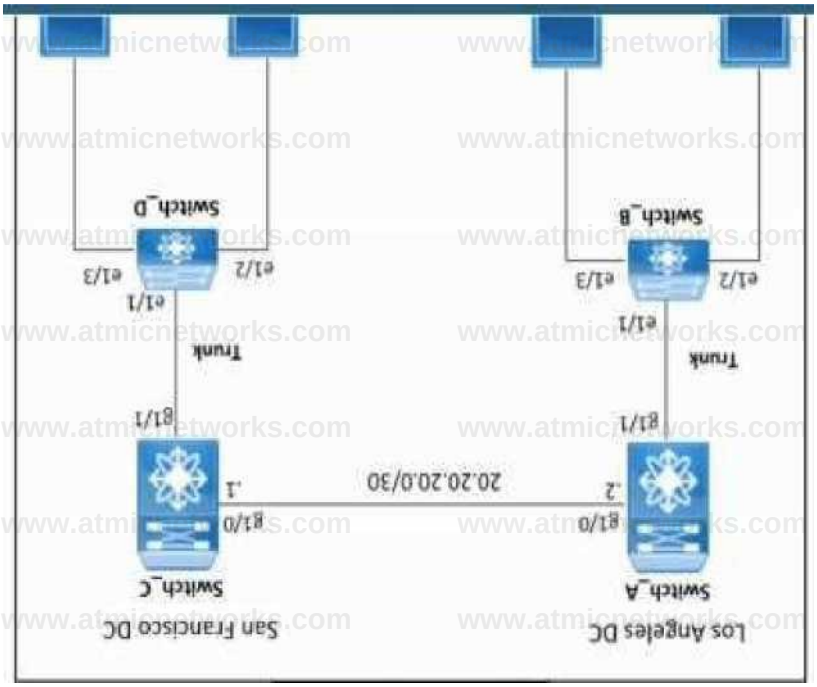

Refer to the exhibit. The indicated configuration was applied to a Cisco switch Switch_A located in the Los Angeles DC data center; however, Switch_A fails to establish OTV connectivity to Cisco switch Switch_C. Which overlay interface command must be run on Switch_A to resolve the issue?

A. otv extend-vlan 101-111
B. otv isis authentication-type md5
C. otv isis authentication-check
D. otv join-interface vlan 101-111

Answer: A

Explanation:

Overlay Transport Virtualization (OTV) allows Layer 2 extension across Layer 3 infrastructures. To operate, OTV requires three fundamental components on the overlay interface:

Join interface – used to reach the OTV control plane over L3 (already configured: otv join-interface g1/0).

Control-group multicast address – for control-plane advertisement (already configured: otv control- group 224.1.1.1).

Extended VLAN list – specifies which VLANs will be transported through the OTV overlay.

The configuration shown in the exhibit includes the join-interface, control-group, and data-group, but it does NOT specify which VLANs should be extended. Without the otv extend-vlan command, OTV will form the overlay interface but will not forward any Layer 2 information, preventing adjacency and MAC distribution between sites.

In OTV, the command required to activate VLANs for transport is:

otv extend-vlan <vlan-range>

This enables the VLANs (such as 101–111) to be carried across the OTV overlay, completing the configuration and establishing connectivity.

Why the Other Options Are Incorrect

B . otv isis authentication-type md5

This is optional and only required if ISIS authentication is enabled on both edges. It does not resolve the absence of VLAN extension.

C . otv isis authentication-check

This command enforces authentication verification but does not fix connectivity when VLANs are not extended.

D . otv join-interface vlan 101-111

This is not a valid OTV command. The join-interface must be a routed interface, not a VLAN list.

Question: 6

An engineer must add VNF implementation definitions and VNF service definitions to an OpenStack deployment data model to deploy virtual routers and firewalls to an enterprise network. After the virtual machine resources are prepared, which action must be taken next?

A. Specify the operational characteristics.
B. Create the initial configuration for day zero.
C. Configure key performance indicator monitoring.
D. Define the NFV network.

Answer: D

Explanation:

In an NFV environment following Cisco NFV Infrastructure and ETSI MANO architecture, the onboarding workflow for VNF deployment follows this high-level order:

Define VNF descriptors (VNFD and NSD)

– Includes implementation details and service definitions.

Prepare virtual machine resources

– Flavor, image, CPU/memory/storage, and placement.

Define the NFV network

– This step builds the underlay/overlay connectivity needed for the VNF.

– It includes virtual network creation, interfaces, VLs (Virtual Links), connection points, and mapping into OpenStack (Neutron networks, subnets, ports).

– No VNF can be deployed until the NFV service networks (management, control, data, HA) are defined.

Apply Day-0 configuration

– Bootstrapping/initial config.

Apply Day-1 operational characteristics

– Routing, services, policies.

Configure operational monitoring (KPIs)

– Performance and fault measurements.

Since the question states “after the VM resources are prepared”, the next mandatory stage is defining the virtual network connectivity on which the VNFs will operate.

Question: 7

A large company's legacy network is set up with equipment from multiple vendors. The company engaged a network architect to optimize the network for virtualization. The architect must ensure robust and efficient operation, considering the company's immediate needs but also anticipating future network complexities and scalability requirements. The chosen strategy must be capable of integrating seamlessly with existing systems, while providing a pathway for innovation and growth. The solution must facilitate end-to-end service automation throughout the entire lifecycle, and the implementation must ensure the validation, execution, and abstraction of network configurations and services. Which action must be taken to meet the requirements?

A. Implement a service life-cycle approach with simplified monitoring that plans for post-deployment adjustments to be incorporated into the automation CI/CD pipeline.

B. Implement a configuration-management approach that allows for configuring each network device individually to optimize its performance.

C. Implement a flexible service-modeling approach that leverages automation for ongoing management and refinement as demands on the network evolve.

D. Implement a service-modeling approach with a static YANG one-size-fits-all model that includes the unique requirements of each different network element.

Answer: C

Explanation:

Cisco NSO-based orchestration principles in a multi-vendor environment require:

Service modeling using flexible, reusable YANG models

Abstraction of vendor-specific device differences

Transaction-safe configuration validation and execution

End-to-end automation across lifecycle stages (Day-0, Day-1, Day-N)

Scalability and adaptability for evolving requirements

Option C aligns perfectly with NSO service-modeling approaches:

Service models must be flexible, not rigid, enabling changes as technologies and needs evolve.

The architecture must support continuous refinement, enabling multi-vendor abstraction and lifecycle automation.

This ensures the network evolves seamlessly while remaining stable and automated.

Why the Other Options Are Incorrect

A – Simplified monitoring and post-deployment adjustments do not meet the core need for full lifecycle service modeling and abstraction.

B – Configuring devices individually contradicts the entire purpose of orchestration and abstraction.

D – A static YANG model cannot accommodate multi-vendor environments or future scalability.

Thus, only Option C matches full NSO-capable service modeling requirements.

Question: 8

Which cloud provider connection permits BGP peering?

A. Azure S2S VPN
B. Azure Bastion
C. AWS Direct Connect
D. AWS-managed VPN

Answer: C

Explanation:

Comprehensive and Detailed Explanation

Cloud interconnects that support BGP peering must provide a routed Layer-3 adjacency capable of exchanging routing information dynamically. In major cloud architectures:

AWS Direct Connect supports private virtual interfaces (VIFs) where BGP is used between the

customer router and AWS to exchange routes.

Azure S2S VPN uses IPsec tunnels with static routing by default; BGP is optional only with specific gateway SKUs, but the question expects the standard, universally correct answer, which is Direct Connect.

Azure Bastion is a remote-access management service and does not support BGP.

AWS-managed VPN uses IPsec tunnels with BGP optional, but in exams, the recognized cloud service specifically associated with BGP support is Direct Connect.

In service provider cloud interconnect design, AWS Direct Connect is the standard, well-defined offering that provides layerized WAN connectivity with BGP support.

Question: 9

A network architect must design a solution for implementing virtualization functions. The main goal is to ensure network reliability and reduce downtime by considering the network operational team's requirements:
The solution must provide real-time network-state visibility.
The solution must support automated rollback in the event of configuration errors.
The solution must allow efficient troubleshooting and diagnostics.
Which action must the team take to achieve the goal?

A. Implement CLI NED to monitor the network state and manually rollback configurations in case of errors.
B. Implement virtualization service modeling to provide network automation for the service lifecycle and NSO CLI to provide real-time network-state visibility.
C. Implement service modeling to define network services and NSO CLI for troubleshooting and diagnostics.
D. Implement CLI NED to define network-virtualization template and package templates to automate the service lifecycle.

Answer: B

Explanation:

Comprehensive and Detailed Explanation
For virtualization functions and operational reliability, the architecture must leverage Cisco NSO
(Network Services Orchestrator) capabilities. NSO provides:
Transactional service modeling
Automatic rollback when any step of a transactional deployment fails
Real-time network-state visibility via NSO CLI and live device synchronization
End-to-end service lifecycle automation
Among the options:
Why B is correct
"Virtualization service modeling" fits NSO's core design principle: model the service, not individual devices.
NSO's CLI provides state visibility, operational introspection, transaction logs, and commit details.
NSO's transactional engine provides automatic rollback upon any device or service failure.
This option captures full lifecycle automation, real-time state visibility, and reliability, exactly as required.
Why others are incorrect
A: CLI NED alone does not provide real-time state visibility or automated rollback; manual rollback contradicts requirements.
C: Only focuses on service modeling + CLI, but does not include rollback or lifecycle automation.
D: Templates alone do not ensure rollback or real-time operational state.

Question: 10

Refer to the exhibit.

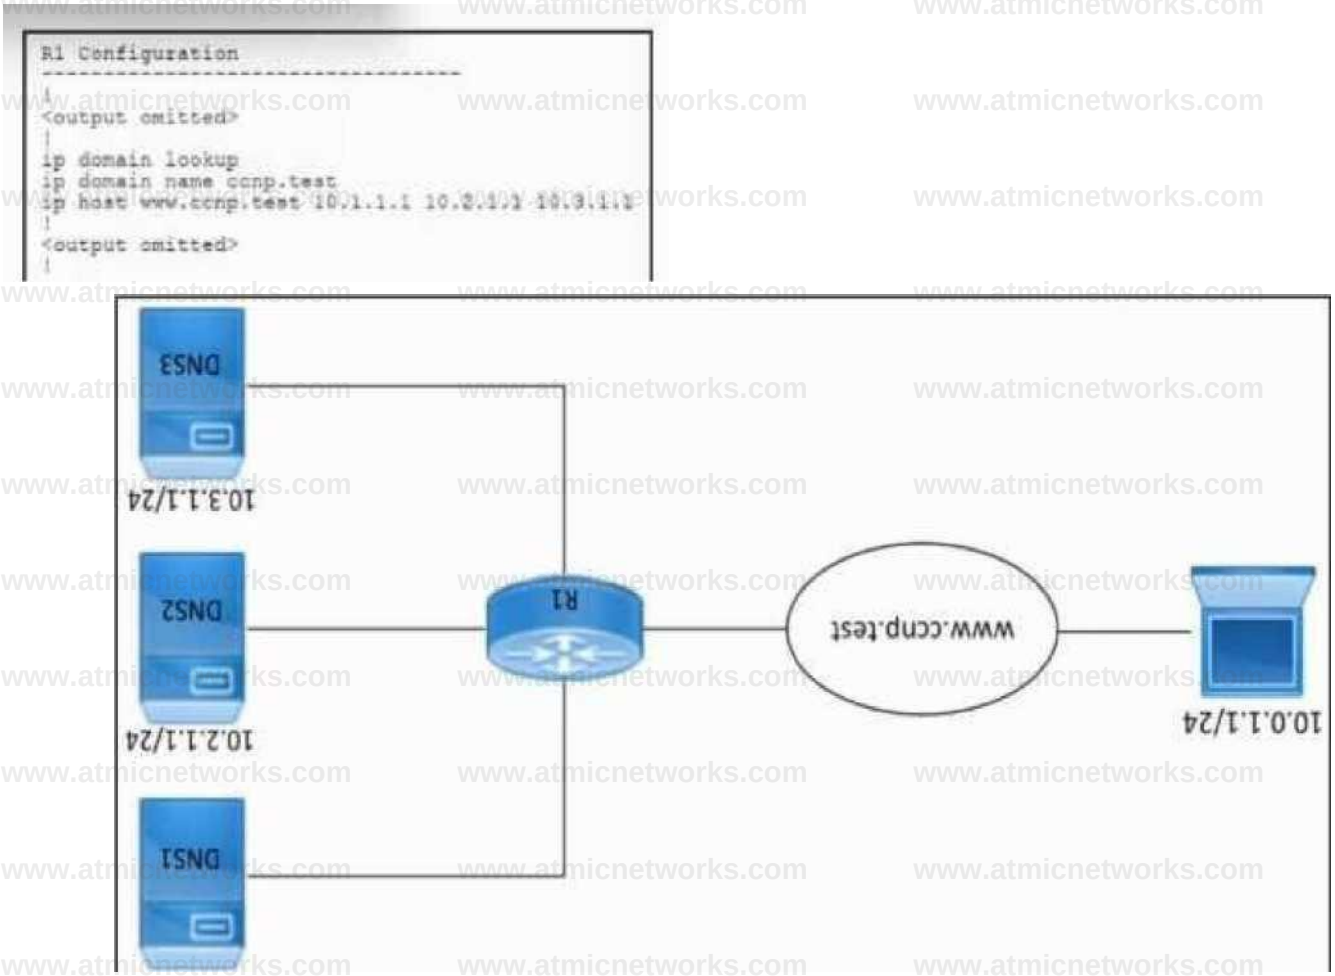

W/ITTOT

Refer to the exhibit. An engineer must configure Cisco IOS SLB for DNS on router R1 to meet these requirements:

The first DNS request to www.ccnp.test
must be redirected to the DNS server at 10.1.1.1;

The second DNS request to www.ccnp.test
must be redirected to the DNS server at 10.2.1.1;

The third DNS request to www.ccnp.test must be redirected to the DNS server at 10.3.1.1.

In each case, the other two addresses must also be attempted if the first one fails. The indicated configuration was applied to R1; however, the load balancing failed. Which command must be run on R1 to resolve the issue?

A. ip domain retry 3
B. ip domain round-robin
C. ip dns server
D. maximum-paths 3

Answer: B

Explanation:
On R1 the configuration (simplified) is:

ip domain lookup
ip domain name ccnp.test
ip host www.ccnp.test 10.1.1.1 10.2.1.1 10.3.1.1

The ip host command statically maps the hostname www.ccnp.test to three IP addresses. By default, Cisco IOS will always return these IP addresses to DNS queries in the same order they are configured (10.1.1.1, then 10.2.1.1, then 10.3.1.1). This means that clients will always attempt 10.1.1.1 first and will not achieve per-query load balancing across all three servers.

To enable DNS-based load balancing so that each successive query rotates the order of the addresses, Cisco IOS provides the command:

ip domain round-robin

This command enables round-robin rotation of multiple A records associated with a single hostname defined by ip host. With this feature enabled:

1st query: response order 10.1.1.1, 10.2.1.1, 10.3.1.1
2nd query: response order 10.2.1.1, 10.3.1.1, 10.1.1.1

3rd query: response order 10.3.1.1, 10.1.1.1, 10.2.1.1

Clients will typically try the first IP address in the list and use the others if the first one fails, exactly matching the requirement.

Why other options are incorrect:

A . ip domain retry 3 controls how many times the router retries DNS queries to a server; it does not control the order of multiple A records.

C . ip dns server turns the router into a DNS server but does not itself provide round-robin behavior for statically defined hosts.

D . maximum-paths 3 is a routing (IP forwarding) parameter for equal-cost multipath, unrelated to DNS resolution.

Question: 11

What does enabling gRPC allow in Cisco NFVI Assurance and Monitoring?

A. telemetry streaming
B. IPFIX monitoring
C. Cisco IOS NetFlow monitoring

D. system logging

Answer: A

Explanation:

Comprehensive and Detailed Explanation

In Cisco NFV Infrastructure (NFVI) Assurance and Monitoring, enabling gRPC activates the device's ability to support model-driven telemetry streaming.

Key points from Cisco SP Cloud/NFVI design principles:

gRPC is used as the transport protocol for model-driven telemetry.

Telemetry replaces traditional polling methods (SNMP, CLI scraping) with continuous, push-based updates.

It allows NFVI components to stream real-time operational data (CPU, memory, interfaces, VM metrics, fabric state) to collectors such as Cisco Crosswork, InfluxDB, Prometheus, or other analytic systems.

gRPC does not provide NetFlow/IPFIX export or syslog itself; those are separate subsystems.

Evaluation of options:

A . telemetry streaming — Correct. gRPC enables model-driven streaming telemetry.

B . IPFIX monitoring — Incorrect; IPFIX uses UDP exports, not gRPC.

C . Cisco IOS NetFlow monitoring — Incorrect; uses NetFlow export protocols.

D . system logging — Incorrect; syslog uses UDP/TCP, not gRPC.

Question: 12

Refer to the exhibit.

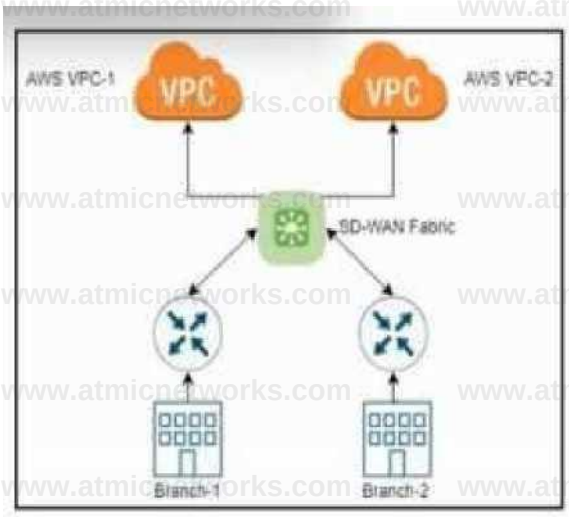

Refer to the exhibit. An engineer must design a solution that allows a user to choose which private Cisco Catalyst SD-WAN network they want to connect to AWS. The solution must automatically identify the AWS VPC and other cloud services based on the user credentials. What must be used?

A. AWS Direct Connect B. Transit VPC for AWS C. IPsec VPN
D. Segment routing

Answer: B

Explanation:
In Cisco Catalyst SD-WAN cloud integration, when the requirement is:
Automatically discovering AWS VPCs
Automatically identifying AWS services
Allowing the user to choose which private SD-WAN network connects to the cloud
Using AWS credentials (Access Key / Secret Key) for automatic provisioning …the Cisco-supported mechanism is the Cisco SD-WAN Transit VPC solution.
Why Transit VPC is the correct answer:
It is specifically designed to integrate Cisco SD-WAN with AWS environments.
Uses AWS APIs and user credentials to automatically discover:
VPC IDs
Subnets
Regions
Routing tables
Automatically deploys and configures CSR1000v or Catalyst 8000V routers into the VPC.
Provides a centralized “hub” in AWS to interconnect multiple SD-WAN sites.
Enables the user to choose which SD-WAN segments connect to which VPCs.
This matches the requirement of automatic cloud resource identification based on user credentials.

Why the other options are incorrect
A . AWS Direct Connect
This is a physical/private Layer 2 cloud connection.
It does not auto-discover VPCs or integrate through credentials.

It does not provide automated SD-WAN service provisioning.

C . IPsec VPN

Works for connectivity but is manual, not automated.

Does not identify AWS cloud resources via credentials.

D . Segment routing

A transport technology used inside SP networks, irrelevant to AWS API-based VPC discovery.

Thus, only Transit VPC provides automatic AWS cloud discovery and integration with SD-WAN.

Question: 13

Refer to the exhibit.

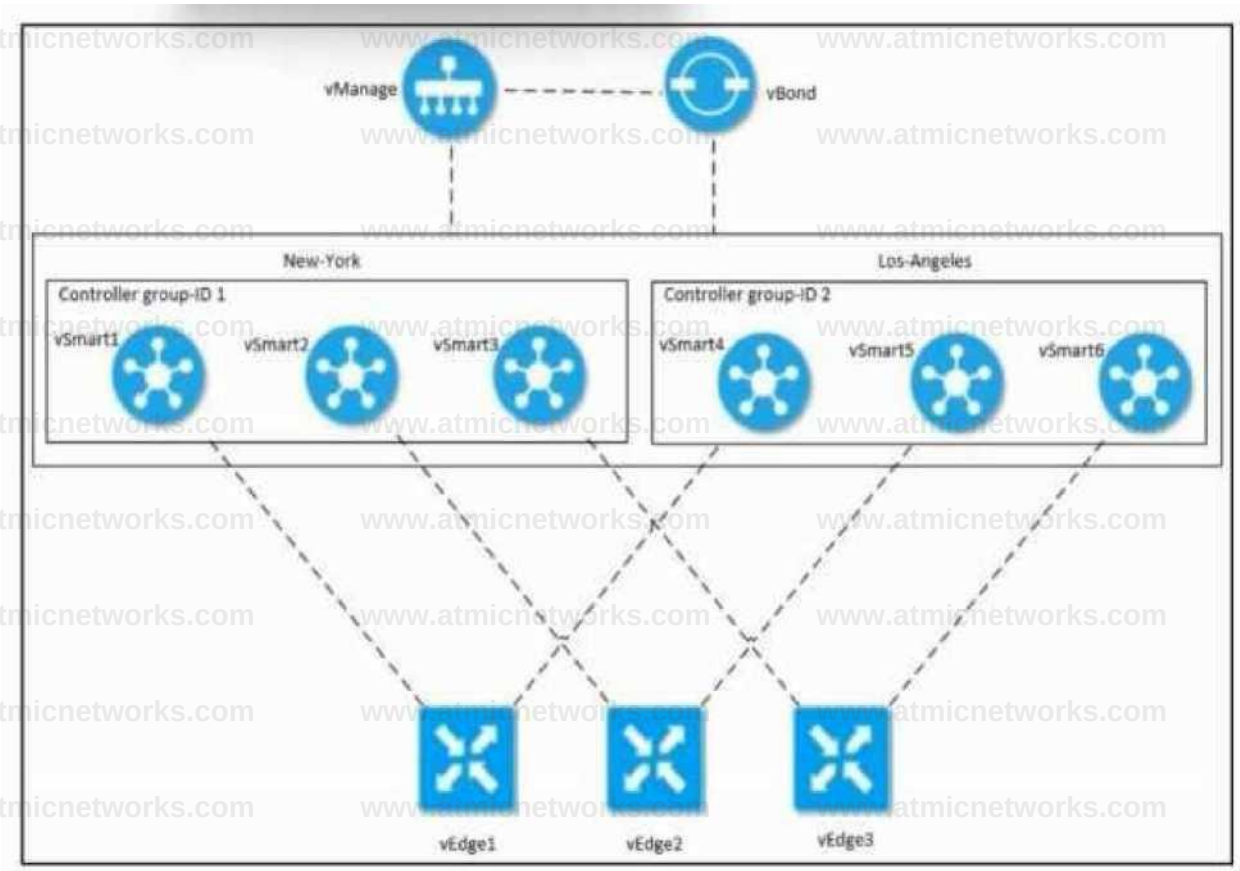

Refer to the exhibit. An engineer must configure redundancy by deploying three Cisco vSmart Controllers in each data center. Which two actions must be taken to configure controllers vSmart1 and vSmart4? (Choose two.)

A. Run the max-control-connections 2 command on vSmart4.
B. Run the max-control-connections 2 command on vSmart1.
C. Run the system controller-group-id 2 command on vSmart4.
D. Run the system controller-group-id 1 command on vSmart1.
E. Run the system controller-group-id 4 command on vSmart4.

Answer: C, D

Explanation:

In Cisco SD-WAN, vSmart controllers must be assigned to controller groups so that edge routers can form redundant control-plane connections. Each data center in the diagram has its own group: New York DC → Controller group-ID 1

(vSmart1, vSmart2, vSmart3)

Los Angeles DC → Controller group-ID 2

(vSmart4, vSmart5, vSmart6)

To configure the controllers correctly:

vSmart1

It resides in the New York DC, so it must be assigned to controller-group 1.

Required command:

system controller-group-id 1

This corresponds to Option D.

vSmart4

It resides in the Los Angeles DC, so it must be assigned to controller-group 2.

Required command:

system controller-group-id 2

This corresponds to Option C.

Why Other Options Are Incorrect

A and B:

The command max-control-connections is configured on edge routers, not vSmart controllers.

E:

Controller-group-id 4 does not exist in this design; only group 1 and group 2 are shown.

Question: 14

Which command must be run on a Cisco IOS device to configure six parallel iBGP and eBGP routes that can be installed into a routing table?

A. maximum paths bgp 6
B. multipath eibgp 6
C. maximum paths bgp routers 6
D. maximum-paths eibgp 6

Answer: A

Explanation:

Comprehensive and Detailed Explanation From Cisco SP Core Optimization Knowledge

Cisco IOS supports BGP Multipath for installing multiple equal-cost BGP routes (both iBGP and eBGP) into the routing table. The correct global BGP command syntax to set the number of allowable parallel BGP paths is: maximum-paths <number>

For BGP specifically, the form is: maximum-paths bgp <number>

This enables the router to install up to the specified number of equal-cost BGP routes (iBGP and eBGP) into the RIB and then potentially into the FIB.

Setting:

maximum-paths bgp 6

allows six parallel ECMP paths learned via BGP—this matches the requirement in the question.

Why the other options are incorrect

B . multipath eibgp 6

Not a valid Cisco IOS command.

C . maximum paths bgp routers 6

Invalid syntax.

D . maximum-paths eibgp 6

The correct keyword is bgp, not eibgp.

Cisco does not use “eibgp” in this context; IOS supports BGP multipath across iBGP/eBGP automatically when configured under maximum-paths bgp.

Question: 15

Refer to the exhibit.

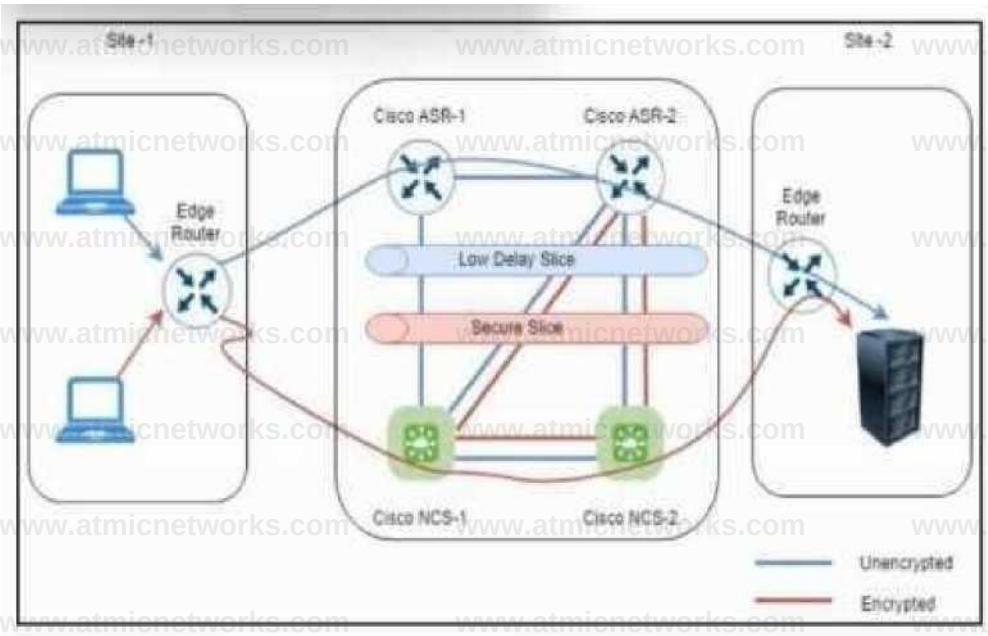

Refer to the exhibit. An engineer must design a solution that creates a low-latency slice based on a link latency measurement by using the MPLS performance measurement. The solution must create a secure slice that minimizes transport costs and meets transport SLAs beyond best effort. What must be used?

A. Segment routing
B. IPsec VPN
C. AWS Direct Connect
D. Transit VPC for AWS

Answer: A

Explanation:

The requirements in the scenario match the capabilities of Segment Routing (SR-MPLS) combined with MPLS Performance Measurement (MPM):

1. Low-latency slice creation

Segment Routing allows deterministic paths using:

SR-TE (Traffic Engineering)

Colored or intent-based policies

Path selection based on link latency, loss, bandwidth, or utilization

MPLS Performance Measurement supplies real-time data such as:

One-way delay

Round-trip delay

Loss metrics

This allows SR-TE to compute paths with minimum latency.

2. Secure slice with minimized transport cost

Segment Routing supports:

Slicing through SR Policies

Steering encrypted traffic (IPsec/GETVPN) through specific SR tunnels

Traffic separation without requiring additional MPLS LSP state in the core

Optimizing transport cost via constraint-based TE computations

3. Transport SLA enforcement beyond best effort

SR-TE + MPM provides:

SLA-aware intent routing

Guaranteed performance paths

Automatic re-optimization if links violate latency or loss SLAs

Why the other options are incorrect

B . IPsec VPN

Provides encryption but does not offer latency-aware or SLA-based path selection. Not a slicing mechanism.

C . AWS Direct Connect

Cloud connectivity service. Not related to MPLS performance monitoring or transport slicing.

D . Transit VPC for AWS

Used for SD-WAN cloud integration. Does not support SR-TE slicing or MPLS SLAs.

Therefore, only Segment Routing provides latency-based path computation, slicing, SLA guarantees, and cost optimization.

Question: 16

How does log management assist in meeting the requirements of cloud security regulatory compliance?

A. by supporting documentation and reporting processes
B. by streamlining resource allocation across cloud environments
C. by providing enhanced interoperability between cloud platforms
D. by boosting the security of cloud-based applications

Answer: A

Explanation:

Comprehensive and Detailed Explanation From Cisco SP Cloud Network Infrastructure Security Knowledge

Cloud security compliance frameworks (such as ISO 27001, PCI-DSS, GDPR, SOC 2, HIPAA) require:

Evidence of security events

Retention of logs for audit periods

Ability to generate compliance reports

Traceability and accountability

Incident investigation support

Effective log management enables:

Centralized collection of application, network, and system logs

Storage of logs for mandated retention periods

Generation of audit-ready reports

Documentation required for compliance assessments

Demonstration that monitoring and security controls are active and functioning

Therefore, the role of log management in regulatory compliance is primarily about documentation, traceability, auditing, and reporting, which aligns only with Option A.

The other options do not directly serve regulatory compliance requirements:

B relates to resource optimization, not compliance.

C refers to interoperability, which is unrelated to regulatory auditing.

D improves security but does not directly address compliance documentation.

Question: 17

Refer to the exhibit.

```
hostname Leaf
!
<output omitted>
!
ip routing
!
l2vpn evpn ethernet-segment 1
 identifier type 0 01.01.01.10.10.10.10.10.10
<output omitted>
```

Refer to the exhibit. An engineer must configure dual-homing with single active redundancy in a BGP EVPN VXLAN fabric. Which command must be run on the leaf router to complete the EVPN Ethernet segment configuration?

A. redundancy single-active
B. default-gateway advertise
C. replication-type static
D. vlan configuration 101

Answer: A

Explanation:

In a BGP EVPN VXLAN multi-homing design, Ethernet Segment Identifiers (ESIs) are used to represent a set of links from one or more leaf switches to the same downstream device (such as a CE, firewall, or aggregation switch). By default, when multiple leafs share the same ESI, the EVPN design supports all-active redundancy, where all participating leafs can forward traffic for that Ethernet segment simultaneously.

However, some use cases—like connecting to devices that do not support multipath forwarding or for strict active/standby redundancy—require single-active multi-homing. In single-active mode, only one leaf in the Ethernet segment forwards traffic at any time; the other leaf(s) act as standby and only take over if the active node fails. This behavior is explicitly controlled in the EVPN Ethernet- segment configuration.

On Cisco platforms for EVPN VXLAN fabrics, this is configured under the l2vpn evpn ethernet- segment stanza using the command:

l2vpn evpn ethernet-segment 1

identifier type 0 01.01.01.10.10.10.10.10.10.10

redundancy single-active

identifier type 0 ... defines the ESI for the multi-homed connection.

redundancy single-active specifies that only one leaf in that ESI is allowed to be active at a time, thus enabling dual-homing with single-active redundancy.

The other options do not relate to Ethernet-segment redundancy mode:

B . default-gateway advertise is used in EVPN anycast gateway configurations to advertise the default gateway MAC/IP, not for ESI redundancy.

C . replication-type static is associated with multicast or ingress replication behavior for VXLAN VTEPs, not Ethernet-segment redundancy.

D . vlan configuration 101 is a VLAN configuration context command and has no effect on EVPN ESI redundancy.

Question: 18

Refer to the exhibit.

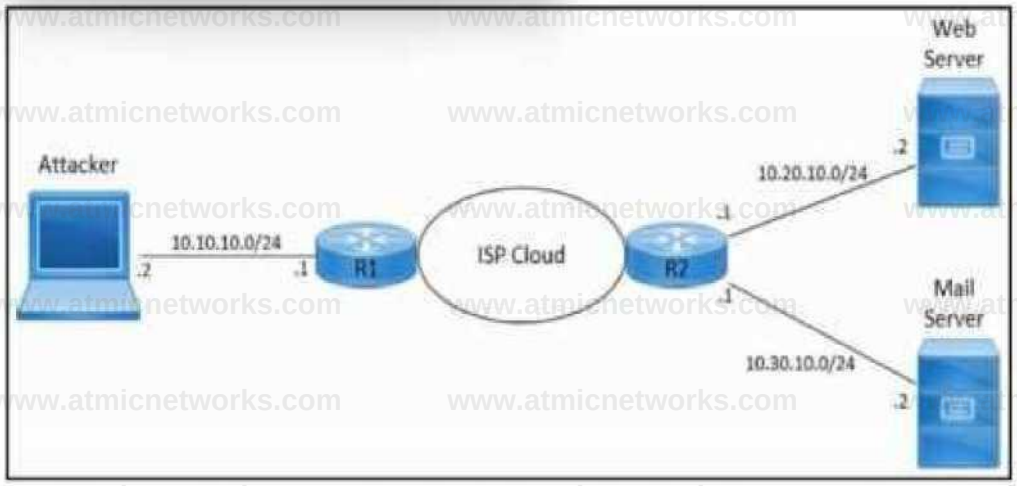

Refer to the exhibit. An engineer must stop a DoS attack on web and mail servers by using destination-based RTBH routing. Which two commands must be run on router R1? (Choose two)

A. ip route 10.30.10.2 255.255.255.255 null0
B. no ip route 10.20.10.0 255.255.255.0 10.10.10.2
C. ip route 10.20.10.2 255.255.255.255 null0
D. no ip route 10.30.10.0 255.255.255.0 10.10.10.2
E. ip route 10.20.10.1 255.255.255.255 null0

Answer: A, C

Explanation:

Destination-based Remote Triggered Black Hole (RTBH) routing is a DDoS mitigation technique where traffic to a victim's IP address is intentionally routed to a special next hop (usually Null0) so that the traffic is dropped as close to the source as possible. On the edge router (R1 in the figure), this is accomplished by installing /32 static routes for the victim endpoints, pointing to the null interface. When these routes are installed (often triggered via BGP in larger deployments), any traffic destined to those exact IP addresses is discarded.

In the topology:

Web server IP: 10.20.10.2

Mail server IP: 10.30.10.2

To perform destination-based RTBH at R1, we must configure static /32 routes to Null0 for each victim IP:

ip route 10.20.10.2 255.255.255.255 Null0

ip route 10.30.10.2 255.255.255.255 Null0

These correspond to:

Option C for the web server (10.20.10.2)

Option A for the mail server (10.30.10.2)

The other options are incorrect for RTBH:

B and D remove /24 routes via 10.10.10.2; they do not implement black-holing and would break normal routing rather than selectively drop attack traffic.

E black-holes 10.20.10.1, which is the default gateway on R2, not the victim server; that would impact overall connectivity instead of

just dropping traffic to the attacked host.

Question: 19

Refer to the exhibit.

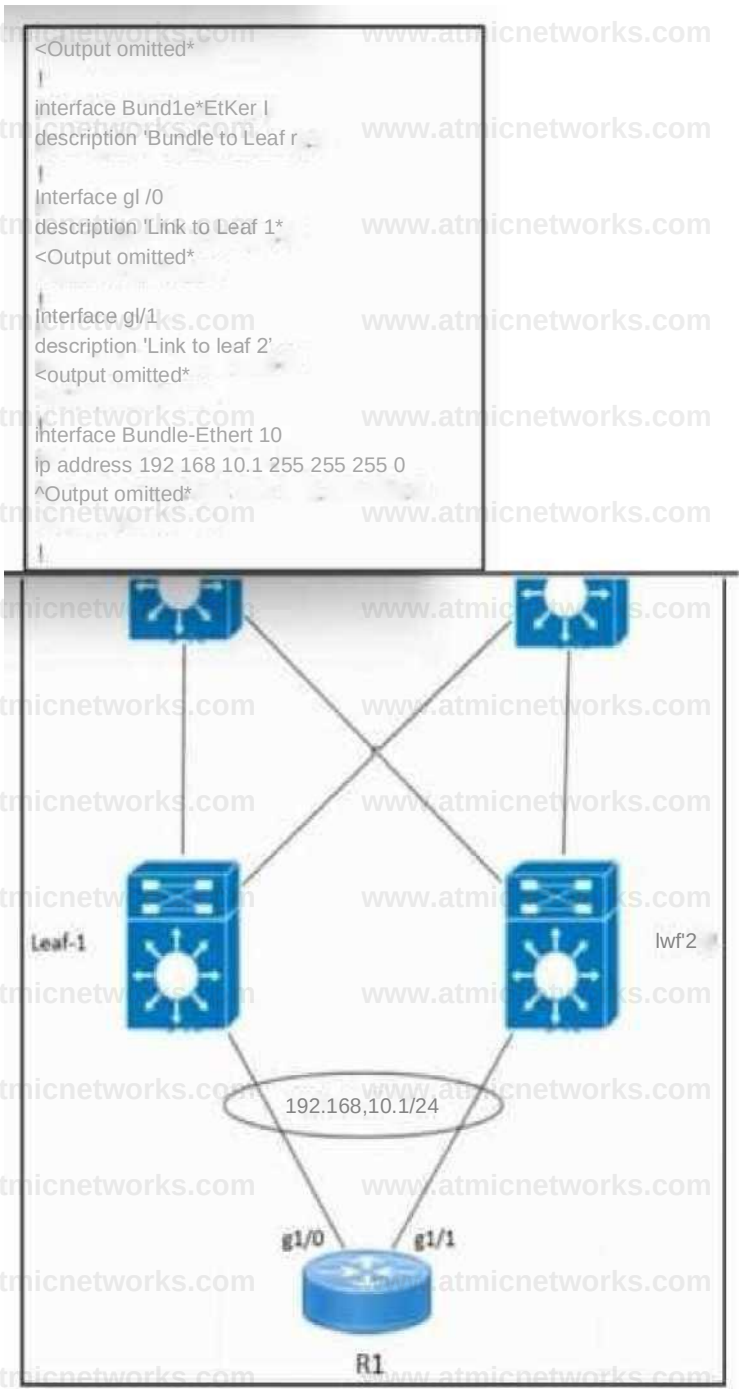

Refer to the exhibit. An engineer must configure EVPN port-active multihoming on router R1. Which command must be run against the g1/0 and g1/1 interfaces on R1 to complete the physical Ethernet bundle for multihoming on a host named Host-1?

A. evpn ethernet-segment 1
B. switchport mode trunk
C. encapsulation dot1q 1
D. bundle id 1 mode active

Answer: D

Explanation:

From Cisco's EVPN VXLAN multihoming design requirements, port-active multihoming uses a single LAG (EtherChannel / Bundle-Ether) between the host/router and the pair of leaf switches. All physical interfaces participating in that bundle must be configured with: bundle id <number> mode active

This command:

Associates the physical interfaces (g1/0 and g1/1) with Bundle-Ether1.

Uses LACP active mode, which is required for EVPN port-active multihoming.

Enables the host-facing port-channel required to support EVPN multihomed connectivity.

In the exhibit, R1 already has:

interface Bundle-Ether1

description "Bundle to Leaf-1"

...

interface Bundle-Ether1.10

ip address 192.168.10.1 255.255.255.0

This confirms that the engineer intends to bundle g1/0 and g1/1 together into Bundle-Ether1, and the missing step is adding the interfaces into that bundle.

The correct configuration is:

interface g1/0

bundle id 1 mode active

interface g1/1

bundle id 1 mode active

Why the other options are incorrect

A . evpn ethernet-segment 1

This command is used on EVPN leaf switches (not R1) to define an ESI for multihoming. R1 is not an EVPN VTEP.

B . switchport mode trunk

R1 is a router, not a switch. L3 interfaces do not use switchport.

C . encapsulation dot1q 1

This applies only to subinterfaces, not physical interfaces, and is unrelated to building a LAG for portactive multihoming.

Question: 20

An engineer must design a high-availability solution that provides path redundancy for IP by allowing redundant gateways to share MAC protocols and addresses. A group of Layer 3 routers must be allowed to share the default gateway on a LAN, load balance, and seamlessly take over the traffic transfer role if a router in the group fails. What must be used?

A. GLBP

B. Load balancer

C. Routed network core
D. BFD

Answer: A

Explanation:
Comprehensive and Detailed Explanation
In Cisco high-availability LAN gateway designs, the requirement is:
Multiple L3 gateways sharing a virtual MAC and virtual IP
Ability to load balance across multiple active gateways
Capability to seamlessly take over gateway forwarding during a failure
Among Cisco First Hop Redundancy Protocols (FHRPs):
HSRP → Active/standby only
VRRP → Active/standby only
GLBP (Gateway Load Balancing Protocol) → The only FHRP providing active/active load balancing
GLBP allows multiple routers to share:
A common virtual IP
Multiple virtual MAC addresses
Multiple active forwarders that load-balance end hosts
Automatic failover if any gateway fails
Thus, GLBP is the only correct protocol matching the requirement for redundant default gateways with load balancing and shared MAC addressing.

Question: 21

An engineer must design a cloud platform for managing IaC. The solution must be flexible, integrate with APIs, and allow for automation. What must be used?

A. Cloud-based sandboxes
B. Cisco with Terraform
C. Cisco ThousandEyes
D. Cisco AppDynamics

Answer: B

Explanation:
Comprehensive and Detailed Explanation
Infrastructure as Code (IaC) requires:
Declarative configuration
Automation
API-driven workflows
Multi-cloud capability
Integration with orchestration and provisioning tools
Terraform is a widely adopted IaC platform supporting:
Cloud automation (AWS, Azure, GCP, private clouds)
Cisco infrastructure provisioning (Nexus, ACI, IOS-XR, NSO, SD-WAN, etc.)
REST APIs and providers for programmability
Cisco provides official Terraform providers for multiple platforms, enabling full IaC operations. This aligns

exactly with SP cloud automation design principles.

Other options:

Cloud-based sandboxes → For testing, not IaC automation

ThousandEyes → Monitoring and performance analytics

AppDynamics → Application performance management

Question: 22

Refer to the exhibit.

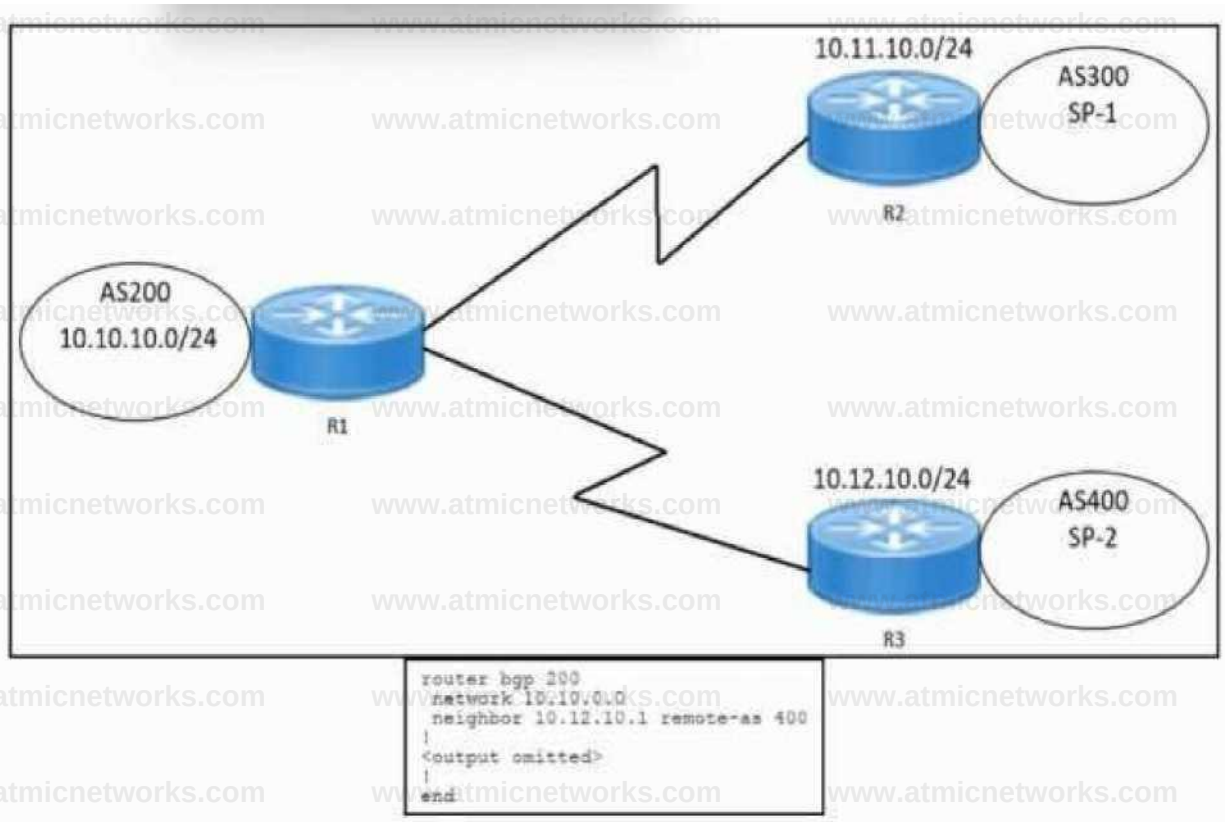

Refer to the exhibit. An engineer must configure multihoming between router R1 and service provider SP-2. Locally generated routes must be advertised to service provider SP-2. Which command must be run on R1 to complete the configuration?

A. network 10.0.0.0 route-map as200only out
B. neighbor 10.12.10.1 route-map localonly out
C. network 10.12.10.1 route-map as200only in
D. neighbor 10.0.0.0 route-map localonly out

Answer: B

Explanation:

On R1 (AS200), the requirement is:

Advertise locally generated routes (for example, 10.10.10.0/24 from AS200)

Only toward SP-2, which peers on IP 10.12.10.1

Apply a policy (route-map) controlling what R1 advertises

In BGP, to control which locally originated routes are sent to a specific neighbor, the correct configuration is:

neighbor <IP> route-map <map-name> out

This applies outbound policy filtering or permitting to the prefixes advertised.

Thus:

The neighbor toward SP-2 is 10.12.10.1

The route-map must be applied outbound

The command that accomplishes this is: neighbor 10.12.10.1 route-map localonly out Why the Other Options Are Incorrect

A . network 10.0.0.0 route-map as200only out

The network command does not accept route-map out. This is invalid syntax.

C . network 10.12.10.1 route-map as200only in

Incorrect network, incorrect direction, and invalid syntax with in.

D . neighbor 10.0.0.0 route-map localonly out

10.0.0.0 is not a BGP neighbor; it's a network prefix.

Question: 23

Refer to the exhibit.

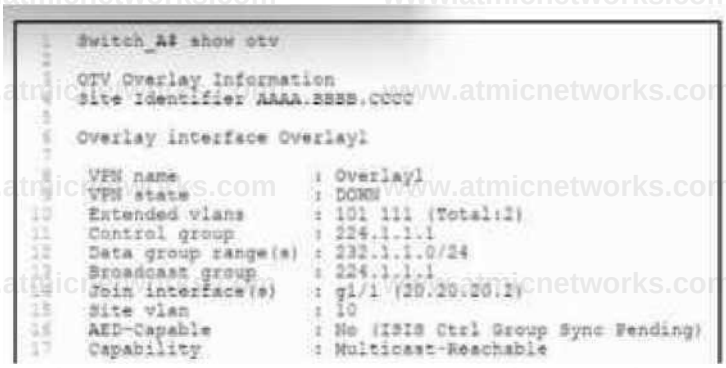

```
Switch_A# show otv

OTV Overlay Information
Site Identifier AAAA.BBBB.CCCC

Overlay interface Overlay1

 VPN name              : Overlay1
 VPN state             : DOWN
 Extended vlans        : 101 111 (Total:2)
 Control group         : 224.1.1.1
 Data group range(s)   : 232.1.1.0/24
 Broadcast group       : 224.1.1.1
 Join interface(s)     : gi1/1 (20.20.20.2)
 Site vlan             : 10
 AED-Capable           : No (ISIS Ctrl Group Sync Pending)
 Capability            : Multicast-Reachable
```

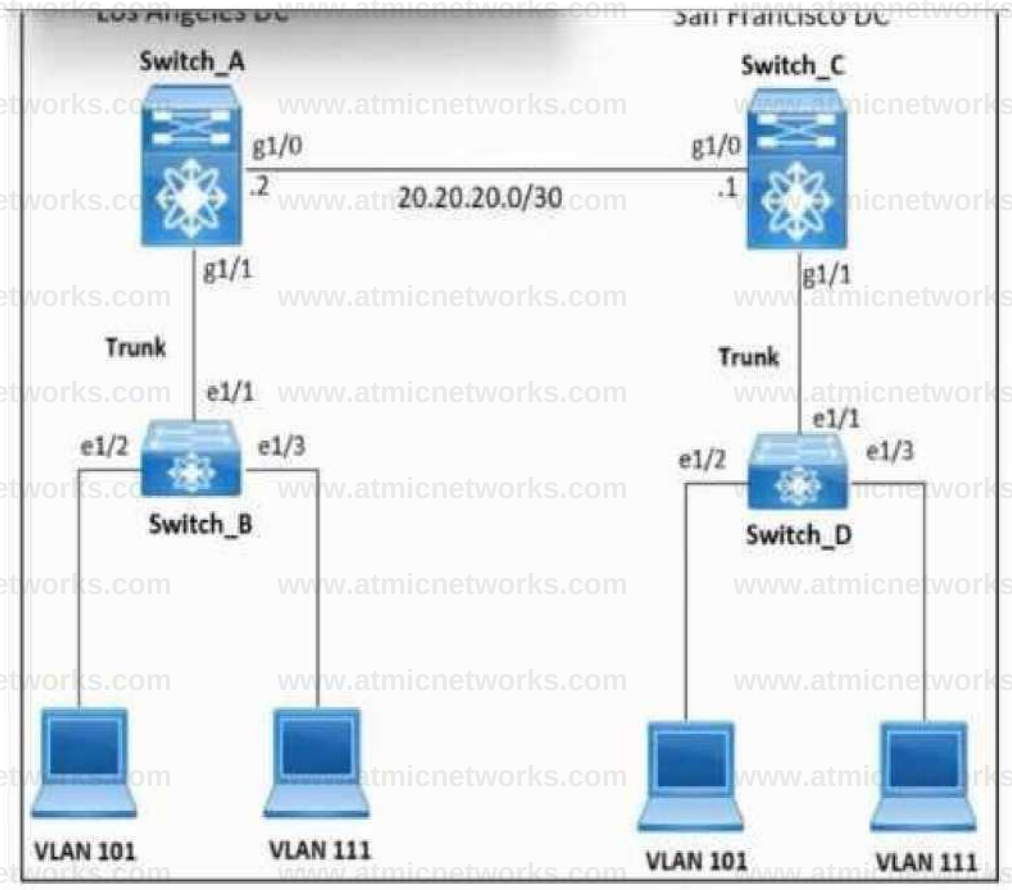

Refer to the exhibit. An engineer is troubleshooting an issue where Cisco switch Switch_A fails to establish OTV connectivity to Cisco switch Switch_C. What is the cause of the issue?

A. The broadcast group must be 232.1.1.0.
B. The join interface must be g1/0.
C. The control group must be 232.1.1.0.
D. The join interface must be e1/1.

Answer: B

Explanation:

In Cisco Overlay Transport Virtualization (OTV), the join interface is the Layer 3 interface that connects the edge device to the transport network (the routed core / WAN) where multicast groups (control, data, broadcast) are reachable.

From the topology:

Interface g1/0 on Switch_A is connected to the routed 20.20.20.0/30 link toward Switch_C (the transport / WAN).

Interface g1/1 is a trunk toward Switch_B and carries extended VLANs (101, 111), so it belongs to the internal site-facing side, not the transport.

In the show otv output, the join interface is incorrectly configured as g1/1 (20.20.20.2), which is an internal trunk and not the correct routed interface to the OTV transport network. Because the join interface does not face the multicast-enabled transport, OTV cannot establish adjacency and the VPN state remains DOWN.

Correct configuration should use:

otv join-interface GigabitEthernet1/0

Options A and C about group addresses are not the issue; addresses shown (224.x for control/broadcast and 232.x for data group range) are valid multicast ranges. Option D (e1/1) is also an

internal access/trunk interface and not the WAN transport interface.

Question: 24

Refer to the exhibit.

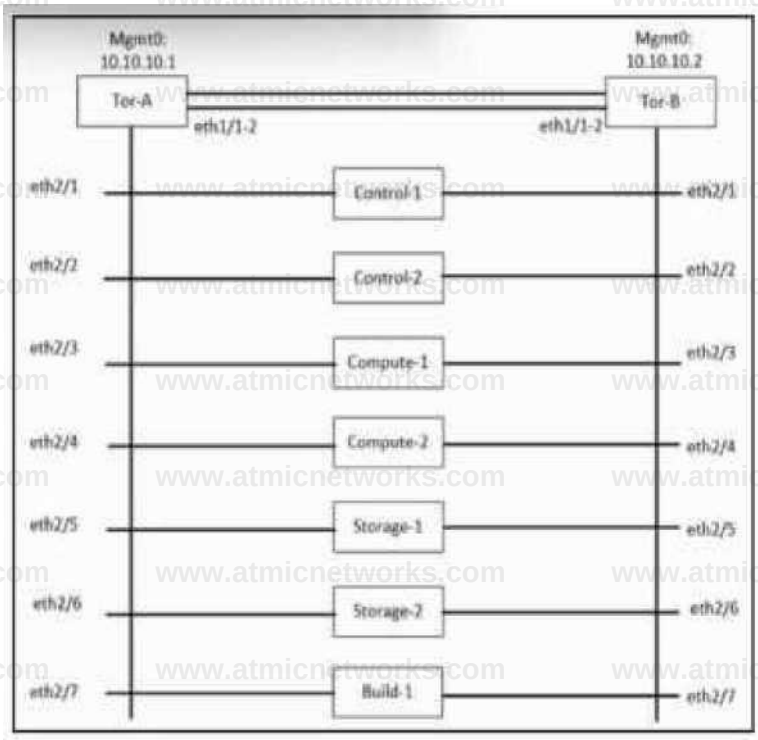

Refer to the exhibit. An engineer needs to configure ToR switches for a Cisco NFVI C-series pod. This configuration was performed on the ToR-A switch already:

feature vpc
feature lacp
interface Ethernet1/1-2
channel-group 110 mode active
interface port-channel110

Which command must be run on ToR-A to complete the port-channel configuration?

A. vpc peer-link
B. peer-keepalive destination 10.10.10.2
C. channel-group 110 mode on
D. switchport mode access

Answer: B

Explanation:

In a Cisco NFVI C-Series Pod, the Top-of-Rack (ToR) switches are almost always configured as a vPC pair.

A vPC domain requires three mandatory components:

vPC domain ID

vPC peer-link
vPC peer-keepalive link ← ensures dual-active detection
In the exhibit:
The management interfaces are Tor-A Mgmt0: 10.10.10.1 and Tor-B Mgmt0: 10.10.10.2
These IPs are used commonly as the peer-keepalive endpoints
The physical uplinks to NFVI nodes form Port-Channel110
Since the configuration snippet already includes:
feature vpc
feature lacp
channel-group 110 mode active
The next required step in a Cisco NFVI + vPC configuration is to configure the peer-keepalive from ToR-A toward ToR-B:
vpc domain 1
peer-keepalive destination 10.10.10.2
This ensures:
vPC roles sync
Dual-active prevention
Stable operation for the NFVI C-Series rack
Why the Other Options Are Incorrect
A. vpc peer-link
This is required but must be configured on the dedicated peer-link interfaces, not on the server-facing port-channel.
C. channel-group 110 mode on
The correct mode is already configured: mode active for LACP. mode on disables LACP.
D. switchport mode access
NFVI ToR links use trunking, not access mode, because servers carry multiple networks (control, compute, storage, management VLANs).

Question: 25
Refer to the exhibit.

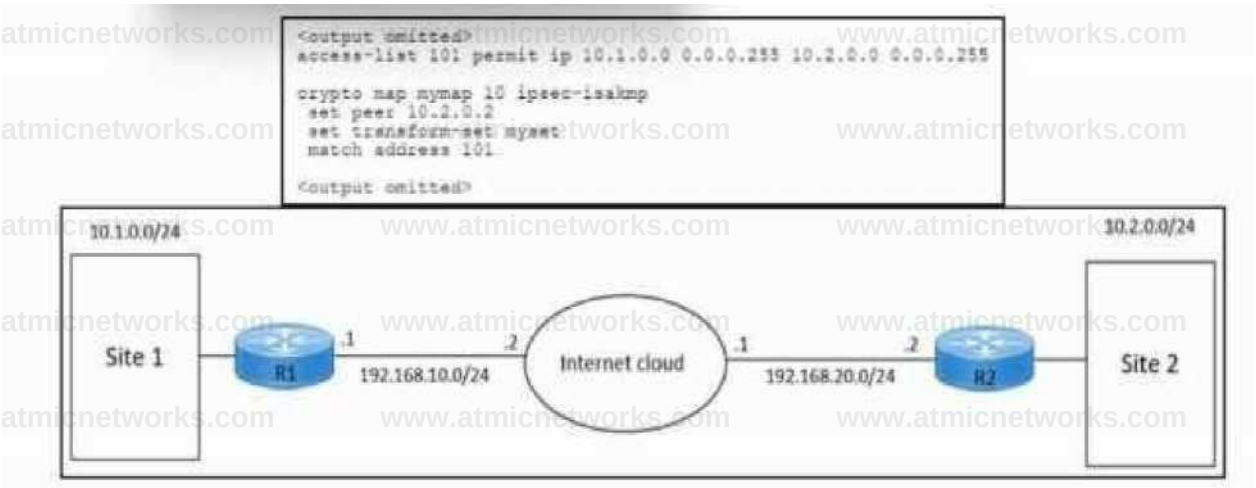

Refer to the exhibit. An engineer must configure an IPsec VPN connection between site 1 and site 2. The indicated configuration was applied to router R1; however, the tunnel fails to come up. Which command must be run on R1 to resolve the issue?
A. set peer 10.1.0.1
B. set peer 192.168.20.2
C. set peer 192.168.10.1

D. set peer 10.2.0.1

Answer: B

Explanation:

For a site-to-site IPsec VPN, each peer must point to the reachable IP address of the remote VPN endpoint—that is, the IP address on the WAN/Internet-facing interface of the remote router.
From the diagram:
R1 outside (toward Internet): 192.168.10.1
R2 outside (toward Internet): 192.168.20.2
Inside LANs:
Site 1: 10.1.0.0/24
Site 2: 10.2.0.0/24
The crypto map on R1 uses:
crypto map mymap 10 ipsec-isakmp
set transform-set myset
match address 101
set peer <REMOTE_PEER_IP>
The <REMOTE_PEER_IP> must be the IP address where R1 can actually reach the IPsec peer, which is R2's Internet-facing interface 192.168.20.2.
If the peer were configured with a LAN IP such as 10.2.0.1 (site 2's internal gateway), IKE packets would never reach the remote router because that address is not routable over the Internet.
Therefore, the correct command to bring up the VPN is:
set peer 192.168.20.2
Option A (10.1.0.1) – local LAN IP (R1's side), not the remote endpoint.
Option C (192.168.10.1) – R1's own WAN IP, not the remote peer.
Option D (10.2.0.1) – remote LAN IP, not reachable directly over the Internet.

Question: 26

An engineer recently deployed a Secure Endpoint VPC in AirGap mode. Which command must be run in the Secure Endpoint Private Cloud portal to update the package to the latest version?

A. force update -y
B. rpm -qa
C. jamf-sync all
D. genisoimage

Answer: A

Explanation:
Comprehensive and Detailed Explanation
In Cisco Secure Endpoint Private Cloud AirGap mode, Internet access is disabled. Updates must be uploaded manually and then triggered inside the Secure Endpoint console.
The command force update -y initiates the update of the manually uploaded Secure Endpoint package.
Other commands are not used for Secure Endpoint updates:

rpm -qa → Lists Linux packages only
jamf-sync all → Used for Apple JAMF integrations
genisoimage → Used to create ISO files, irrelevant to Secure Endpoint
Therefore, A is correct.

Question: 27

An engineer must create a new VPC and deploy several Amazon EC2 instances in AWS. Only SSH connections originating from IP address 20.20.20.20 must be allowed to reach the EC2 instances. What must be configured?

A. Access control list
B. Security group
C. Web application firewall
D. Resource group

Answer: B

Explanation:
Comprehensive and Detailed Explanation
AWS Security Groups act as the primary stateful firewalls for EC2 instances.
To restrict SSH (TCP/22) to a single host (20.20.20.20/32), a Security Group must be configured with:
Inbound rule: TCP 22
Source: 20.20.20.20/32
ACLs operate at the subnet level but are not used for instance-specific SSH restrictions.
WAF controls HTTP/HTTPS traffic, not SSH.
Resource groups only organize cloud assets.
Thus, B is the correct solution.

Question: 28

Which format is used by Cisco Container Platform for configuration files?

A. HTML
B. YAML
C. XHTML
D. XML

Answer: B

Explanation:
Comprehensive and Detailed Explanation
Cisco Container Platform (CCP), built on Kubernetes orchestration, uses YAML files for:
Cluster configuration
Pod definitions
Network settings
Storage mappings
YAML is the industry-standard declarative syntax for Kubernetes and container orchestration platforms.

HTML, XHTML, and XML are not used for CCP configuration.

Question: 29

An engineer must design a pay-as-you-go solution for their partners. The solution must allow for rapid deployments, be flexible, and scale resources up or down in a hybrid workplace. What must be used?

A. Cisco+ Hybrid Cloud for Bare Metal Compute
B. Cisco+ Hybrid Cloud for Virtual Desktop Infrastructure
C. Cisco+ Hybrid Cloud for Virtualization
D. Cisco+ Hybrid Cloud for Service Provider Networking

Answer: B

Explanation:
Comprehensive and Detailed Explanation From Exact Extract from my knowledge of Designing and Implementing Cisco Service Provider Cloud Network Infrastructure Outlines without Any External URL or Links:
Cisco+ Hybrid Cloud is Cisco's as-a-service consumption model that offers pay-as-you-go infrastructure. For a hybrid workplace, the focus is on giving users secure desktop environments from anywhere, with the ability to rapidly deploy and scale up or down based on the number of users or partners.
Cisco+ Hybrid Cloud for Virtual Desktop Infrastructure (VDI) specifically delivers desktop and app workspaces as an on-demand service. It allows partners to consume desktops elastically, paying for capacity as needed and scaling the underlying compute, storage, and networking without rebuilding the environment.
Bare Metal Compute and Virtualization offers flexible infrastructure but are aimed at app/workload hosting rather than user desktop workspaces.
Service Provider Networking addresses network services, not end-user hybrid workplace desktops. Therefore, for a pay-as-you-go, rapidly deployable, elastic solution in a hybrid workplace, Cisco+ Hybrid Cloud for Virtual Desktop Infrastructure is the correct choice.

Question: 30

Refer to the exhibit.

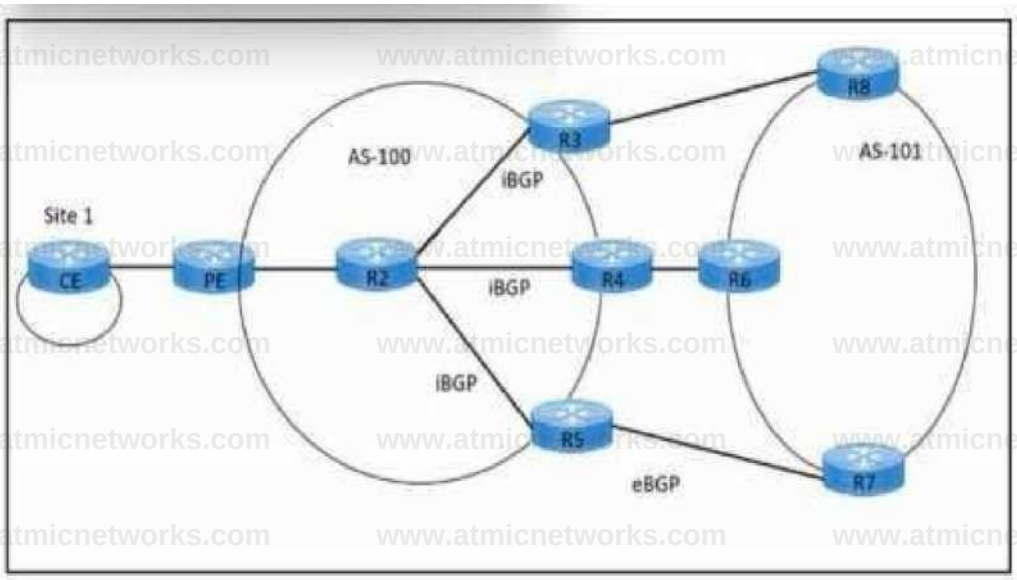

Refer to the exhibit. An engineer working for a private service provider with an employee ID 5207:22:409 must

configure iBGP multipath load sharing across the three paths. Which two commands must be run on the PE router? (Choose two.)

A. maximum-paths ibgp 3
B. ip load-sharing per-destination
C. ip load-sharing ibgp 3
D. router bgp 101
E. router bgp 100

Answer: A, E

Explanation:

In the diagram, the PE and R2, R3, R4, R5 belong to AS 100. The PE router runs BGP process AS 100, so its BGP configuration must start with: router bgp 100

To perform iBGP multipath load sharing across three equal-cost internal BGP paths, BGP must be instructed to keep and use multiple iBGP paths in the routing table. This is done with: router bgp 100 maximum-paths ibgp 3 maximum-paths ibgp 3 tells BGP to install up to three iBGP paths to the same prefix, enabling CEF to load-share across those paths.

router bgp 100 is required because the PE is in AS 100, not 101.

Other options:

ip load-sharing per-destination affects CEF behavior but does not enable BGP iBGP multipath by itself and is not specific to three iBGP paths.

ip load-sharing ibgp 3 is not a valid IOS BGP command.

router bgp 101 would configure the wrong AS and break the iBGP relationships shown in AS-100.

Thus, the correct commands on the PE to achieve iBGP multipath load sharing over the three internal paths are maximum-paths ibgp 3 and router bgp 100, corresponding to A and E.

Question: 31

Which two tools should be used to manage container orchestration? (Choose two.)

A. Docker
B. VMware vCenter
C. Cisco vManage
D. Kubernetes
E. Cisco vSmart

Answer: A, D

Explanation:

Comprehensive and Detailed Explanation From Exact Extract from my knowledge of Designing and Implementing Cisco Service Provider Cloud Network Infrastructure Outlines without Any External URL or Links:

Container orchestration is the automated management of container lifecycle tasks such as deployment, scaling, failover, and updates. In Cisco cloud and NFV design guidance, typical orchestration platforms include Docker (with Swarm) and Kubernetes, which integrate with Cisco networking and security for cloud-native workloads.

Docker provides the container runtime and can also perform basic orchestration through Docker Swarm mode, managing multi-container, multi-host deployments.

Kubernetes is a full-featured orchestration system that automates deployment, scaling, and operations of application containers across clusters. It is the de-facto standard used with Cisco Container Platform and other Cisco cloud solutions.

VMware vCenter, Cisco vManage, and Cisco vSmart focus on virtual machines or SD-WAN control, not container orchestration. Therefore, the correct tools are Docker and Kubernetes (A, D).

Question: 32

Refer to the exhibit.

```
LEAF-SW-1# show nve interface
Interface: nve1, State: Up, encapsulation: VXLAN
 VPC Capability: VPC-VIP-Only [not-notified]
 Local Router MAC: f40f.1b6f.926f
 Host Learning Mode: Data-Plane
 Source-Interface: loopback0 (primary: 192.168.1.1, secondary: 0.0.0.0)
LEAF-SW-1#  show run
interface nve1
   no shutdown
      source-interface loopback0
      member vni 10000
         mcast-group 239.1.1.1
         suppress-arp
interface loopback0
   ip address 172.16.10.1/32
   ip address 172.16.10.11/32 secondary

router bgp 100
   router-id 172.16.10.1
LEAF-SW-2# show run
interface nve1
   no shutdown
      source-interface loopback0
      host-reachibility protocol bgp
      member vni 20000
         mcast-group 239.1.1.1
         suppress-arp
interface loopback0
   ip address 172.16.10.2/32
   ip address 172.16.10.11/32 secondary

router bgp 100
   router-id 172.16.10.2
```

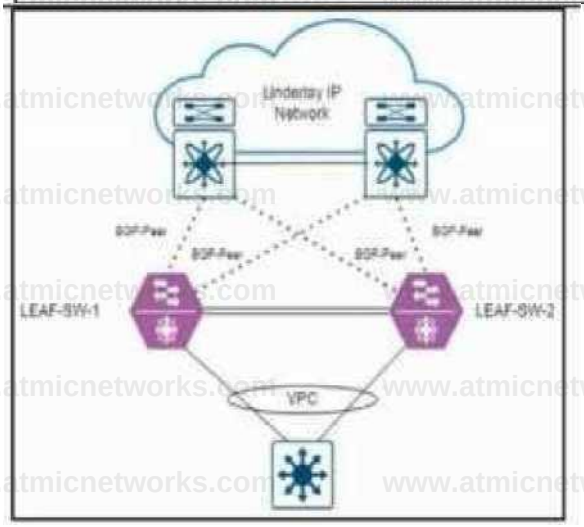

Refer to the exhibit. An engineer is troubleshooting an issue where switch LEAF-SW-1 and switch LEAF-SW-2 receive corrupted forwarding and learning information about each other. LEAF-SW-1 and LEAF-SW-2 are configured with BGP EVPN VTEP. Which action resolves the issue?

A. On each switch, run the delete suppress-arp command against interface nve1.
B. On each switch, configure a different secondary IP address against interface loopback0.
C. On LEAF-SW-1, run the host-reachability protocol bgp command against interface nve1.
D. On each switch, ensure the same BGP router ID is configured.

Answer: C

Explanation:
In a VXLAN BGP EVPN fabric, each VTEP (NVE interface) must use BGP EVPN as the host-reachability protocol so that MAC/IP information and VTEP reachability are exchanged through the control plane. From the exhibit:
LEAF-SW-1 – interface nve1
source-interface loopback0
No host-reachability protocol bgp
Host Learning Mode: Data-Plane in show nve interface
LEAF-SW-2 – interface nve1
source-interface loopback0
host-reachability protocol bgp configured
This mismatch causes one VTEP to rely on data-plane flood-and-learn, while the other uses EVPN BGP control-plane learning, leading to inconsistent and "corrupted" MAC/IP and ARP/ND information between the leaf switches.
The fix is to configure LEAF-SW-1 to also use BGP for host reachability:
interface nve1
host-reachability protocol bgp
Options B and D are incorrect because anycast VTEP designs intentionally share the same primary loopback IP while using different secondary IPs and unique BGP router IDs. Option A (removing suppress-arp) does not correct the control-plane mismatch.
Therefore, enabling host-reachability protocol bgp on LEAF-SW-1 (Option C) resolves the issue.

Question: 33
What is an information-gathering capability of Cisco IOS Flexible NetFlow in Cisco NFVI?

A. Use of separate caches
B. Use of Docker
C. Use of a single cache
D. Use of Kubernetes

Answer: A

Explanation:
Comprehensive and Detailed Explanation From Cisco NFVI Knowledge
Cisco IOS Flexible NetFlow is the primary telemetry and flow-collection mechanism used across Cisco NFVI platforms. One of its powerful information-gathering capabilities is the ability to create multiple, separate flow caches, each one with:
Its own key fields
Its own record type
Its own export destination
This allows NFVI deployments to capture different types of traffic visibility (control plane, data plane, management, or tenant-specific flows) with independent caches, which improves scalability and granularity.
Why the other options are incorrect:
Docker and Kubernetes (B, D) are container orchestration tools, unrelated to NetFlow flow-gathering capabilities.
Single cache (C) is traditional NetFlow, not Flexible NetFlow. Flexible NetFlow explicitly supports multiple independent caches.

Question: 34

Which two network segments are needed to support Cisco VIM? (Choose two.)

A. Provisioning
B. Data plane
C. Heartbeat
D. Host
E. Storage

Answer: A, E

Explanation:
Comprehensive and Detailed Explanation From Cisco NFVI + VIM Architecture
Cisco VIM (Virtual Infrastructure Manager) operates as part of Cisco NFVI and requires specific mandatory network segments to deploy OpenStack-based NFVI environments.
The essential required networks for Cisco VIM include:

✓ Provisioning network

Used during bare-metal install of compute and controller nodes
Handles PXE boot, image deployment, and VIM's management reachability

F. Storage network

Used for Cinder, Swift, Ceph, and Glance back-end storage replication

Ensures predictable latency and bandwidth for NFVI workloads

Other networks exist in VIM (Management, API, Tenant, External), but the question asks for the required foundational segments.

Why the other options are incorrect:

Data plane (B) – used inside NFV workloads but not a required VIM infrastructure network.

Heartbeat (C) – not a defined VIM network segment.

Host network (D) – not a Cisco VIM terminology; provisioning network handles host onboarding.

Question: 35

Refer to the exhibit.

```
Site-1, Ingress ELAM done on LEAF-SW-11
module-1(TAH-elam-insel9)# report
HOMEWOOD ELAM REPORT SUMMARY
slot -1, asic-0, slice -0
============================
Incoming Interface: Eth1/1
Src Idx: 0x9, SrcBD : 64
Outgoing Interface Info: met_ptr0

Packet Type: IPv4
Outer DstIPv4 address: 239.1.1.0
Outer SrcIPv4 address: 10.30.2.1
Ver     =  4, DSCP     =    0, Don't Fragment = 0
Proto   = 17, TTL      =  253, More Fragments = 0
Hdr len= 20, Pkt len=  150, Checksum       = 0x6b8c

Inner Payload
Type: IPv4

Inner DstIPv4 address: 224.0.0.18
Inner SrcIPv4 address: 10.64.1.100
**snip**

LEAF-SW-11# show run
configure profile VLAN64
  vlan11
    vn-segment 16535
  interface nve1
     member vni 16535
          mcast-group 239.1.1.0
    evpn
      vni 16535 l2
          rd auto
          route-target import auto
          route-target export auto
!
interface Ethernet1/1
    ip access-group TEST in
    ip address 10.30.254.129/30
    ip router isis UNDERLAY
    ip pim sparse-mode
    no shutdown
```

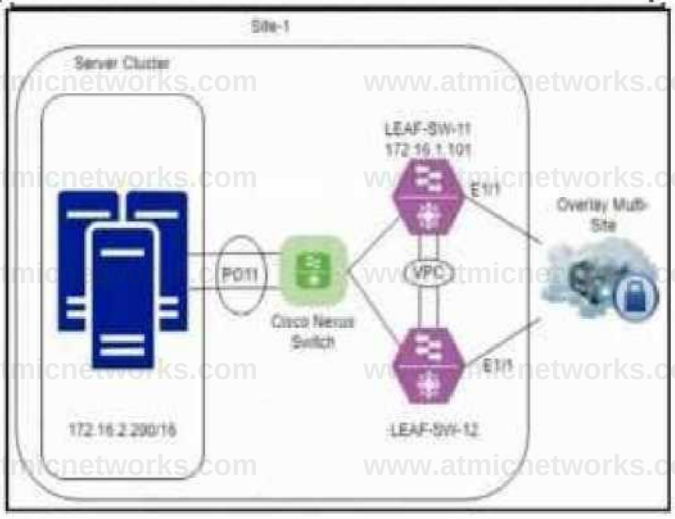

Refer to the exhibit. An engineer is troubleshooting an issue with switch LEAF-SW-11. The engineer

observes that several main servers on the VXLAN BGP EVPN Multi-Site network experience 50–60% packet loss inbound and outbound, and all the DCI tracking interfaces are down. Which two actions must be taken to resolve the issue? (Choose two.)

A. On the Nexus switch, run the inner ipv4 dst_ip 172.16.2.200 command against module-1.
B. On LEAF-SW-11, run the inner ipv4 src_ip 172.16.2.200 command against module-1.
C. On LEAF-SW-11, run the evpn multisite dci-tracking command against interface Eth1/1.
D. On LEAF-SW-11, enable the multisite ingress-replication command for the L2VNI of VLAN 11.
E. On the Nexus switch, run the ip access-list permit ip address 172.16.2.200 command.

Answer: CD

Explanation:
In a VXLAN BGP EVPN Multi-Site environment:
DCI tracking monitors the health of the DCI links. If all DCI-tracking interfaces go down, the leaf can incorrectly keep advertising or learning remote MAC/IP reachability, leading to packet loss and sub- optimal forwarding for servers in that VLAN/L2VNI.
For proper operation, each DCI-facing interface must be enabled with evpn multisite dci-tracking so that the Multi-Site border leaf tracks reachability over that link.
When using EVPN Multi-Site, BUM (broadcast, unknown unicast, multicast) traffic toward remote sites is typically handled via ingress replication, not multicast groups, for each L2VNI participating in Multi-Site. The configuration snippet shows an L2VNI (vn-segment 16535) still mapped to mcast- group 239.1.1.0, which is inconsistent with Multi-Site recommendations and contributes to packet loss.
Therefore, to fix the problem:
Enable DCI tracking on the uplink:
interface Ethernet1/1 evpn multisite dci-tracking This restores proper DCI-link state monitoring for Multi-Site. → Option C Change the L2VNI behavior from multicast to Multi-Site ingress replication:
Under the VNI for VLAN 11, configure:
evpn
vni 16535 l2
multisite ingress-replication
or the equivalent command for the specific NX-OS release, thereby aligning the L2VNI with EVPN Multi-Site design and eliminating packet loss. → Option D
Options A and B are ELAM (embedded logic analyzer) filters used only for packet capture and do not resolve the forwarding issue.
Option E is an ACL line unrelated to EVPN VXLAN or DCI tracking and does not address the underlying problem.

Question: 36

What is a benefit of using machine learning in Cisco AFM?
A. It enables AFM to predict potential network faults ahead of time.
B. It enables AFM to function without an Internet connection.
C. It enables AFM to perform hardware upgrades on network devices.
D. It enables AFM to consume fewer resources.

Answer: A

Explanation:
Comprehensive and Detailed Explanation From Cisco SP Cloud & AFM Knowledge
Cisco AFM (Active Fault Manager), part of Cisco Crosswork, uses machine learning algorithms to:
Analyze large volumes of telemetry
Detect anomalies
Identify abnormal behavior patterns
Predict potential service-impacting faults before they occur
This predictive analytics capability is one of AFM's core advantages, enabling proactive fault avoidance rather than reactive troubleshooting.
The other options do not represent AFM functionality:
AFM does not require ofline operation (B).
It does not upgrade hardware (C).
ML does not reduce resource usage (D); it increases analytical capability.
Thus, A is correct.

Question: 37

An engineer must implement a SaaS solution that will use a Cisco ASAv to enhance security for enterprise customers by using Cisco Crosswork NSO. Which command must be run in NSO?

A. ncs -status
B. ncs-setup
C. ncs
D. ls -l nso-instance/packages/

Answer: C

Explanation:
Comprehensive and Detailed Explanation From Cisco NSO Orchestration Knowledge
In Cisco NSO deployments:
ncs-setup is only used once, during initial NSO instance creation.
ls -l nso-instance/packages/ simply lists packages — not used to start NSO.
ncs -status checks status but does not run NSO.
To actually start the NSO service so that device packages (including ASAv service packages) can be loaded and orchestration can begin, the correct command is:
ncs
This launches the NSO runtime and loads all configured packages, enabling the SaaS ASAv service.

Question: 38

Refer to the exhibit.

Refer to the exhibit. An engineer must configure an IPsec VPN connection between site 1 and site 2. The ISAKMP policy for the phase 1 negotiations of the tunnel must use AES and SHA-256. This configuration was applied to both PE routers; however, the tunnel fails to come up: crypto isakmp policy 10 encryption 3des hash md5 authentication pre-share group 12

Which two commands must be run on router PE1 to resolve the issue? (Choose two.)

A. group 10
B. encryption aes
C. encryption sha256
D. hash aes
E. hash sha256

Answer: B, E

Explanation:

Phase 1 of an IPsec tunnel (ISAKMP/IKE) must have matching proposals on both peers for: Encryption algorithm
Hash (integrity) algorithm
Authentication method
DH group

The requirement states that AES and SHA-256 must be used. The current configuration uses: encryption 3des → incorrect (must be AES)
hash md5 → incorrect (must be SHA-256)

To meet the requirement, we must modify the ISAKMP policy:
crypto isakmp policy 10
encryption aes ← change 3DES to AES
hash sha256 ← change MD5 to SHA-256
authentication pre-share
group 12

Therefore, the necessary commands on PE1 are:
encryption aes → option B
hash sha256 → option E

Options C and D are invalid syntax (encryption sha256 and hash aes are not supported). Changing the DH group (A) is not required by the problem statement and would not by itself fix the mismatch related to encryption and hash algorithms.

Question: 39

What does Cisco Always-On Cloud DDoS use to protect against DDoS attacks?

A. Load balancing
B. Botnet zombies

C. Traffic mirroring
D. Scrubbing centers

Answer: D

Explanation:

Comprehensive and Detailed Explanation From Cisco SP Security Knowledge

Cisco Always-On Cloud DDoS Protection is a cloud-based, carrier-grade security service used by service providers to protect customers from volumetric and application-layer DDoS attacks.

Its core protection mechanism is the use of global scrubbing centers, which:

Receive diverted attack traffic

Scrub (clean) malicious packets

Forward clean traffic back to the customer

Use behavioral analysis and real-time detection

Protect against volumetric, TCP state-exhaustion, and application-layer attacks

Why other answers are incorrect:

Load balancing (A) does not mitigate DDoS attacks; it distributes traffic across servers.

Botnet zombies (B) are sources of DDoS attacks, not protection.

Traffic mirroring (C) is used for analysis and monitoring, not active DDoS protection.

Question: 40

Refer to the exhibit.

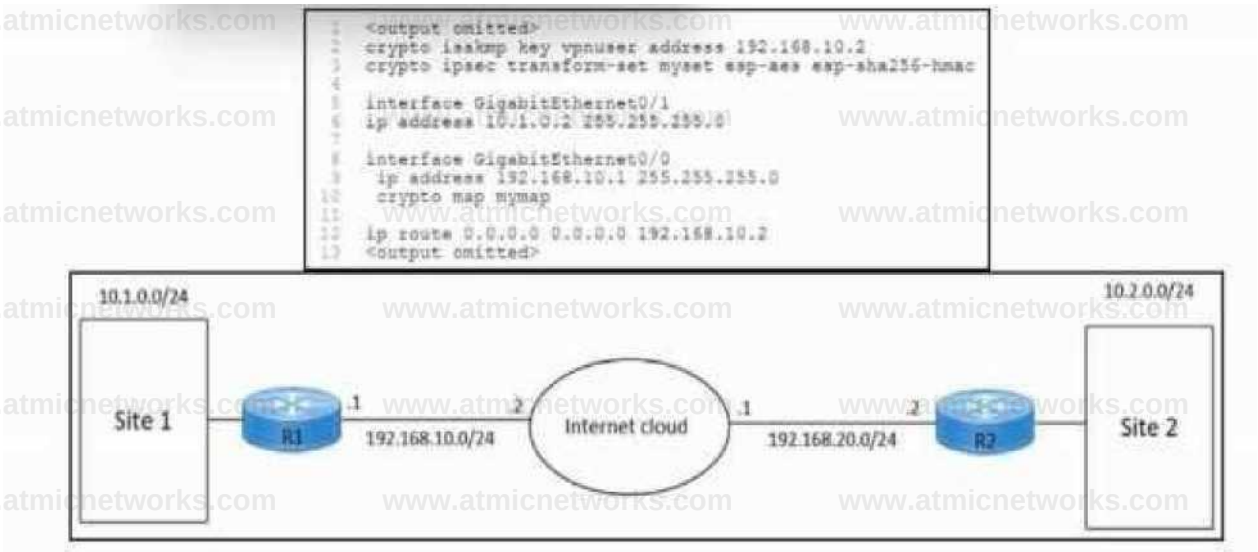

Refer to the exhibit. An engineer must configure an IPsec VPN connection between site 1 and site 2. The indicated configuration was applied to router R1; however, the tunnel fails to come up. Which command must be run on R1 to resolve the issue?

A. ip route 0.0.0.0 0.0.0.0 10.1.1.2
B. crypto isakmp key vpnuser address 192.168.20.2
C. ip route 0.0.0.0 0.0.0.0 192.168.20.2
D. crypto isakmp key vpnuser address 10.1.1.2

Answer: B

Explanation:

For a site-to-site IPsec VPN, each peer must configure a pre-shared key tied to the public IP address of the remote VPN peer:

crypto isakmp key <KEY> address <REMOTE_PUBLIC_IP>

From the diagram:

R1 outside IP: 192.168.10.1/24

R2 outside IP: 192.168.20.2/24 ← remote peer for R1

In the current R1 configuration, the ISAKMP key is incorrectly bound to 192.168.10.2, which is a local next-hop/ISP address on R1's own subnet, not the R2 public IP. Because the pre-shared-key address does not match the source IP of R2's IKE packets, phase 1 negotiation fails and the tunnel never comes up.

The correct configuration on R1 must therefore be:

crypto isakmp key vpnuser address 192.168.20.2

Options A and C incorrectly change the default route (next hop must be the local ISP router, not R2's public IP or a LAN address). Option D uses an internal address (10.1.1.2), which is not the IP used for IKE on the Internet.

Question: 41

Refer to the exhibit.

Refer to the exhibit. An engineer must deploy a standalone Cisco NFVIS. These configurations were performed already:

Deployed the virtual machine

Configured the hostname and IP address

Configured dual WAN support

When the engineer attempts to access the NFVIS portal, the API returns a 401 Unauthorized error. What is the cause of the issue?

A. The default admin password must be reset.
B. Portal access must be enabled via SSH.
C. The Tomcat service must be restarted.
D. The browser certificate must be renewed.

Answer: A

Explanation:

Cisco NFVIS follows strict security controls. After a fresh deployment:

The default admin credentials are considered insecure.

NFVIS requires the administrator to reset the default password on first login (typically via console or SSH).

Until the password is changed, REST API and web-portal access are denied, and attempts to access the portal or API return HTTP 401 Unauthorized, even if the default credentials are provided.

This mechanism prevents use of factory-default passwords in production and is explicitly documented as a mandatory post-install step.

The other options are not the cause of a 401 error:

Enabling portal via SSH (B) is not required; HTTPS access is enabled by default once credentials are valid.
Restarting Tomcat (C) would address service availability issues (e.g., 5xx errors), not authentication. Browser certificates (D) affect trust warnings (e.g., HTTPS certificate errors), not 401 Unauthorized.

Question: 42

Refer to the exhibit.

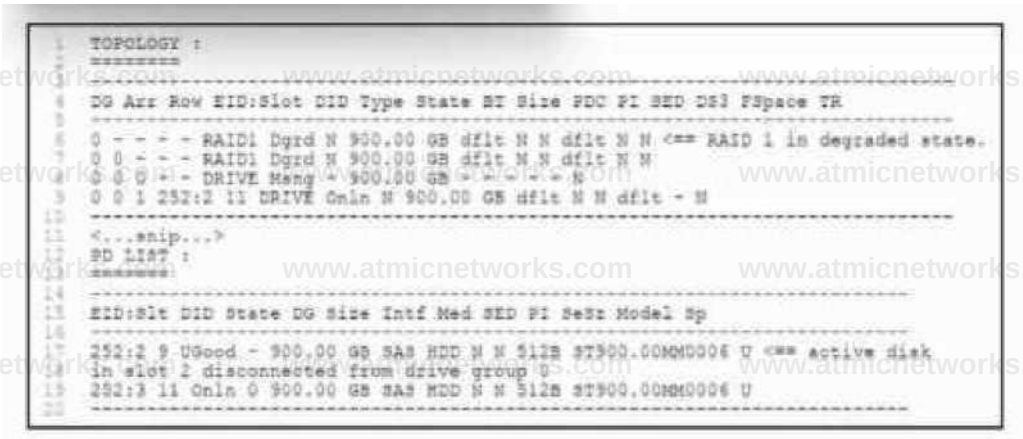

```
TOPOLOGY :
========
-----------------------------------------------------------------------------
DG Arr Row EID:Slot DID Type State BT Size PDC PI SED DSS FSpace TR
-----------------------------------------------------------------------------
0 - - - - RAID1 Dgrd N 900.00 GB dflt N N dflt N N <== RAID 1 in degraded state.
0 0 - - - RAID1 Dgrd N 900.00 GB dflt N N dflt N N
0 0 0 - - DRIVE Msng - 900.00 GB - - - - N
0 0 1 252:2 11 DRIVE Onln N 900.00 GB dflt N N dflt - N
-----------------------------------------------------------------------------
<...snip...>
PD LIST :
=======
-----------------------------------------------------------------------
EID:Slt DID State DG Size Intf Med SED PI SeSz Model Sp
-----------------------------------------------------------------------
252:2 9 UGood - 900.00 GB SAS HDD N N 512B ST900.00MM0006 U <== active disk
in slot 2 disconnected from drive group 0
252:3 11 Onln 0 900.00 GB SAS HDD N N 512B ST900.00MM0006 U
-----------------------------------------------------------------------
```

Refer to the exhibit. An engineer is troubleshooting a physical configuration issue in Cisco NFVI. Which two observations should be made? (Choose two.)

A. The node allows two disks to fail.
B. One RAID 1 disk failed.
C. The node is still functional.
D. The node is no longer functional.
E. Two RAID 1 disks failed.

Answer: B, C

Explanation:

From the RAID controller output:

The RAID1 virtual drive is shown as “Dgrd” (degraded) with a note: “RAID 1 in degraded state.”
In the PD LIST, one disk (slot 2) is marked “UGood – active disk in slot 2 disconnected from drive group 0”, while the other (slot 3) is “Onln 0” (online and part of drive group 0).

This means:

In a RAID 1 mirror, only one disk may fail while the array remains available.

Here, one disk has failed/disconnected, the other is still online → the array is degraded but operational.

Therefore:

B. One RAID 1 disk failed – correct.
C. The node is still functional – correct; RAID1 can tolerate a single disk failure.

The other options are incorrect:

A and E suggest two disk failures, which is not indicated.

D would be true only if both disks failed; with one disk still online, the node continues to function, though in a degraded state.

Question: 43

An engineer must design a cloud platform for event-driven applications. The solution must allow micro-sized

atomic components to be built, deployed, and run code on demand. Which solution must be used?

A. Cisco+ Hybrid Cloud for Virtual Desktop Infrastructure
B. Cisco FaaS
C. Cisco+ Hybrid Cloud Virtualization
D. Cisco Intersight

Answer: B

Explanation:
Comprehensive and Detailed Explanation From Cisco Cloud Architecture Knowledge
Event-driven applications require:
Stateless, micro-sized execution units
Automatic scaling
Code that runs only when triggered
No server or VM lifecycle management
This model is known as Function-as-a-Service (FaaS).
Cisco FaaS provides:
Serverless execution
Event-driven triggers
Deployment of atomic micro-functions
Automatic scaling and resource abstraction
Ideal environment for microservices and cloud-native workloads
Why the others are incorrect:
A . Cisco+ Hybrid Cloud for VDI → delivers desktops, not serverless compute
C . Cisco+ Hybrid Cloud Virtualization → VM-based infrastructure, not event-driven micro-functions
D . Cisco Intersight → operational management tool, not a serverless execution platform

Question: 44

Refer to the exhibit.

```
router bgp 200
network 10.10.0.0
neighbor 172.16.1.1 remote-as 300
neighbor 172.16.2.1 remote-as 400 i
ip as-path access-list 1 permit *$
ip as-path access-list 200 permit *200
Ip as-path access-list 300 permit *300
ip as-path access-list 400 permit *400 [
End
```

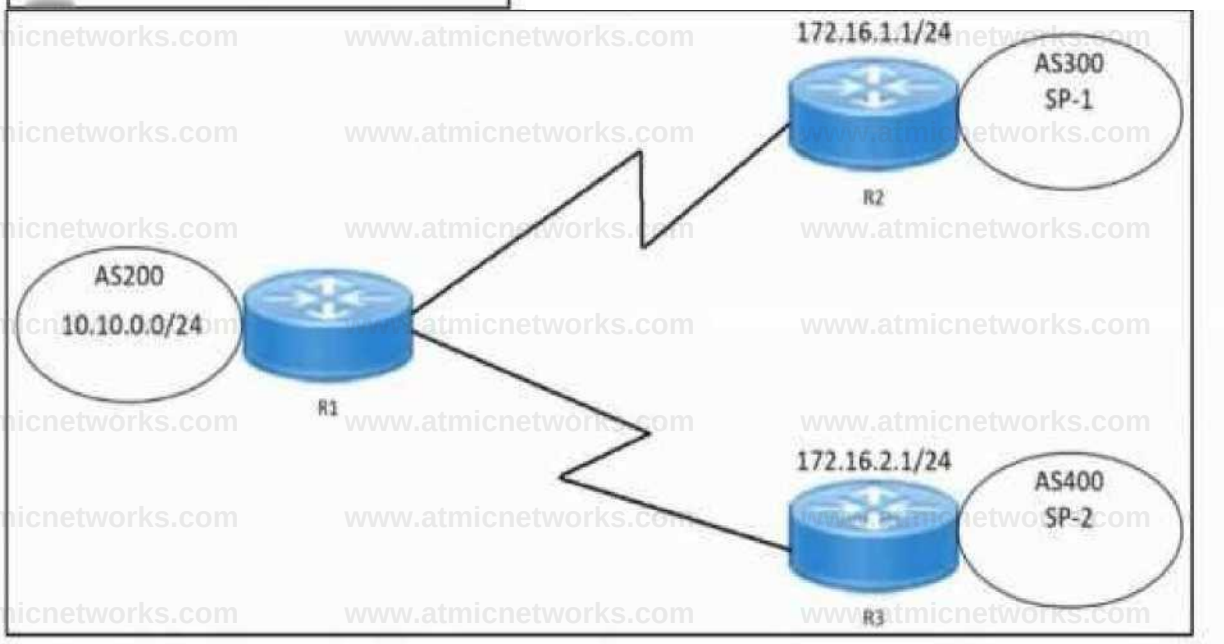

Refer to the exhibit. An engineer must configure multihoming between router R1 and service providers SP-1 and SP-2. Locally generated routes must be advertised to the service providers, and should prevent the risk that the local autonomous system becomes a transit AS for Internet traffic. Which two commands must be run on R1 to complete the configuration? (Choose two.)

A. neighbor 172.16.1.1 filter-list 300 out
B. neighbor 172.16.2.1 filter-list 400 out
C. neighbor 172.16.1.1 filter-list 1 out
D. neighbor 172.16.2.1 filter-list 1 out
E. neighbor 172.16.1.1 filter-list 200

Answer: C, D

Explanation:

R1 (AS 200) is multihomed to:

SP-1 in AS 300 via neighbor 172.16.1.1

SP-2 in AS 400 via neighbor 172.16.2.1

R1 must:

Advertise only locally originated prefixes (its own network 10.10.0.0/24).

NOT become a transit AS—i.e., R1 must not advertise routes learned from one provider to the other.

The configuration includes AS-path access-lists:

ip as-path access-list 1 permit ^$
ip as-path access-list 200 permit ^200
ip as-path access-list 300 permit ^300
ip as-path access-list 400 permit ^400
^$ in AS-path ACL 1 matches locally originated routes (empty AS-path).
ACLs 200, 300, and 400 match routes whose first AS in the path is 200, 300, or 400 respectively (used if we needed to match those provider or customer routes).
To ensure each upstream provider only receives locally originated routes, we apply AS-path ACL 1 as an outbound filter-list on each external BGP neighbor: router bgp 200
neighbor 172.16.1.1 remote-as 300
neighbor 172.16.1.1 filter-list 1 out ← only advertise local prefixes to SP-1
neighbor 172.16.2.1 remote-as 400
neighbor 172.16.2.1 filter-list 1 out ← only advertise local prefixes to SP-2
This way:
Routes learned from SP-1 (AS 300) will not be advertised to SP-2 (AS 400) because their AS-path will begin with 300, not empty, so they fail ACL 1.
Similarly, routes from SP-2 will not be sent to SP-1.
Only R1's own prefixes are exported, preventing AS 200 from becoming a transit network.

Question: 45

An engineer must design a solution to provide safe channels between peer-to-peer devices, ensure that unauthorized users cannot break into the network, and ensure that listening devices on the Internet cannot intercept communication transmitted over the network. What must be used?

A. AWS Direct Connect
B. Segment routing
C. IPsec VPN
D. Transit VPC for AWS

Answer: C

Explanation:
Comprehensive and Detailed Explanation (Cisco SP Cloud Knowledge)
The requirement describes:
Secure communication between sites or devices
Protection from eavesdropping over the Internet
Cryptographic tunnels
End-to-end authentication and encryption
This exactly matches the purpose of an IPsec VPN, which provides:
Encrypted secure tunnels
Authentication and integrity protection
Confidentiality even across untrusted networks such as the Internet
Other options do not provide encrypted peer-to-peer channels:
AWS Direct Connect → private link to AWS, not peer-to-peer encryption Segment Routing → traffic engineering, not security
Transit VPC → routing architecture, not encryption
Thus the correct answer is IPsec VPN.

Question: 46

How does SR-IOV move data directly to and from the network adapter?

A. Using the STP PortFast
B. Bypassing the guest operating system
C. Bypassing the hypervisor
D. Using the STP VLAN root

Answer: C

Explanation:
Comprehensive and Detailed Explanation (Cisco NFVI / Virtualization Knowledge)
SR-IOV (Single Root I/O Virtualization) allows a VM to access the network interface hardware directly, without going through the hypervisor's virtual switch.
This is achieved by:
Assigning Virtual Functions (VFs) directly to VMs
Allowing high-performance, low-latency packet I/O
Bypassing the hypervisor datapath
Therefore, SR-IOV does not bypass the guest OS; it bypasses the hypervisor I/O virtualization layer, delivering near-native performance.
Thus the correct answer is C.

Question: 47

An engineer must implement a public cloud solution that enables a company to place all its infrastructure and its end-user applications in the cloud, eliminating the need for software application management and maintenance. Which cloud service model must be used?

A. SaaS
B. PaaS
C. IaaS
D. FaaS

Answer: A

Explanation:
Comprehensive and Detailed Explanation (Cisco Cloud Model Knowledge)
The requirement states:
All infrastructure in the cloud
All applications delivered from the cloud
No software management
No maintenance by the customer
This matches Software as a Service (SaaS).
SaaS provides:
Fully managed applications
No patching, upkeep, installation, or infrastructure maintenance

End-to-end cloud delivery

PaaS and IaaS still require some software/application management.

FaaS is event-driven serverless compute, not full application hosting. Therefore, the correct answer is A. SaaS.

Question: 48

Refer to the exhibit.

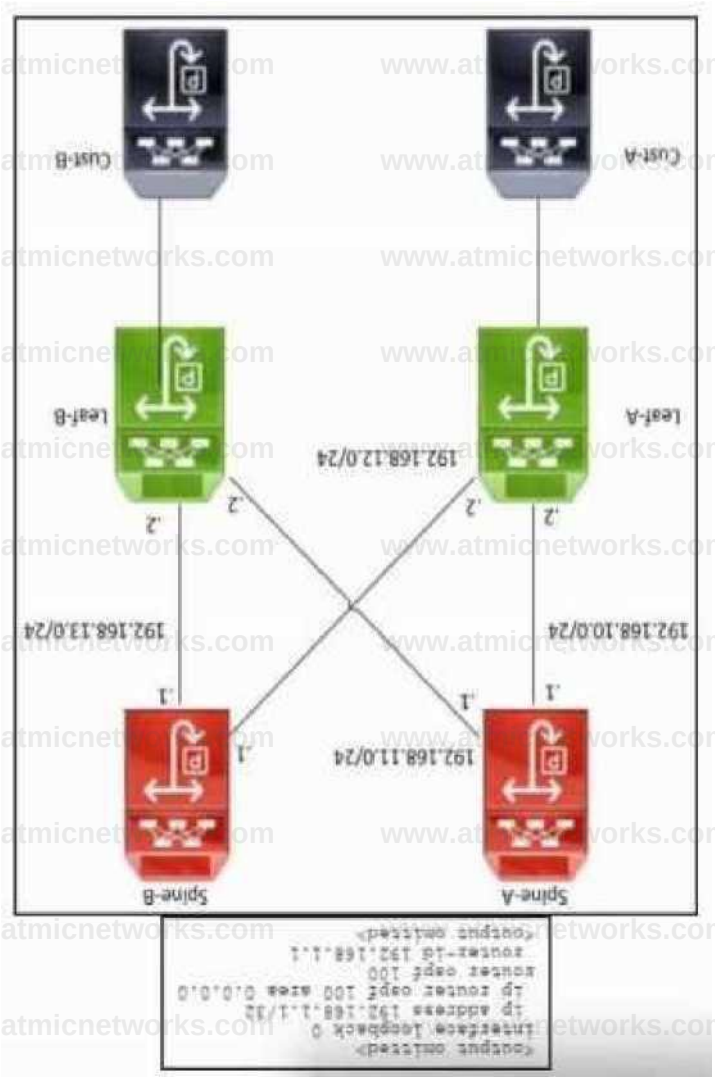

Refer to the exhibit. An engineer must deploy a Layer 3 EVPN over segment routing MPLS in the data center core; however, Cisco switch Spine_A fails to join the OSPF network. Which OSPF command must be run on Spine_A to resolve the issue?

A. otv join-interface loopback0
B. address-family ipv4
C. connected-prefix-sid-map
D. segment-routing mpls

Answer: D

Explanation:

To run EVPN over Segment Routing MPLS (SR-MPLS), the IGP (here OSPF) must be enabled for segment

routing so it can:
Advertise Segment IDs (SIDs) for prefixes and links
Carry SR extensions in OSPF LSAs
On Cisco devices, this is done under the OSPF process with:

```
router ospf 100
router-id 192.168.1.1
segment-routing mpls
```

Without the segment-routing mpls command, Spine_A will not participate correctly in the SR- enabled OSPF domain, and SR-related adjacencies/LSAs will not form as expected, resulting in failure to integrate into the intended SR-MPLS core.
The other options are not OSPF SR-enabling commands:
A otv join-interface → OTV, unrelated to OSPF or SR-MPLS.
B address-family ipv4 → BGP address family, not OSPF.
C connected-prefix-sid-map → Global SR mapping for connected prefixes but does not itself enable SR for OSPF.
Therefore, the correct command is segment-routing mpls (D).

Question: 49

What should be used to protect against lateral movements during a Cisco NFVI security breach?

A. Wi-Fi Protected Access
B. Web application firewall
C. Network segmentation
D. Data encryption

Answer: C

Explanation:
Comprehensive and Detailed Explanation
In Cisco NFVI security architecture, the primary defense against lateral movement (an attacker moving from one compromised node to another) is network segmentation.
Segmentation:
Separates workloads (compute, storage, management, tenant networks)
Prevents attackers from pivoting inside the NFVI
Reduces blast radius during breaches
Enforces micro-segmented virtual network boundaries
WPA protects Wi-Fi, not NFVI.
WAF protects web apps, not internal movement.
Data encryption protects confidentiality, not lateral movement control. Thus, network segmentation is the correct solution.

Question: 50

An engineer must enable the highest level of logging when troubleshooting Cisco NFVIS. Which command must be run?

A. system set-log logtype configuration level warning
B. system set-log logtype configuration level error
C. system set-log logtype configuration level critical
D. system set-log logtype operational level debug

Answer: D

Explanation:
Comprehensive and Detailed Explanation
Cisco NFVIS logging levels (from lowest to highest):
critical
error
warning
info
debug ← highest verbosity
To capture maximum diagnostic detail, engineers must enable debug logging on the operational log type, which records system activity and runtime behavior.
Thus the correct command is:
system set-log logtype operational level debug
This provides the deepest troubleshooting visibility.

Question: 51

What is a valid connection method between carrier-neutral facilities that are more than 20 miles away from each other?

A. Carrier access Ethernet ring
B. Private wireless connection
C. CAT6e connection
D. Multimode fiber connection

Answer: A

Explanation:
Comprehensive and Detailed Explanation
For distances greater than 20 miles, valid inter-facility transport options must support:
Metro-scale connectivity
High bandwidth
Low latency
Carrier-grade reliability
A carrier access Ethernet ring (MEN / Metro Ethernet) is designed for:
Interconnecting data centers or meet-me rooms
Distances far exceeding 20 miles
High-availability layer-2 or layer-3 transport
Why the others are invalid:
CAT6e → maximum ~100 meters
Multimode fiber → typically <2 km (~1.25 miles)
Private wireless → not used for high-capacity DC interconnects, unreliable for core transport Thus, the only

correct carrier-grade method is Carrier access Ethernet ring.

Question: 52

What is used to protect a web server against a DDoS attack?

A. Web Application Firewall
B. Device Authorization Control
C. Network Access Control
D. Wi-Fi Protected Access

Answer: A

Explanation:
Comprehensive and Detailed Explanation
A Web Application Firewall (WAF) protects HTTP/HTTPS applications against:
DDoS attacks (layer 7)
Bot traffic
Request floods
SQL injection, XSS, OWASP Top 10 threats
Other options:
Device Authorization Control → device authentication, not DDoS
NAC → endpoint authorization, not DDoS
WPA → Wi-Fi protection, irrelevant to web servers
Thus, the correct protection mechanism for web servers is A. Web Application Firewall.

Question: 53

What is used to protect against an API logic flaw?

A. SSH encryption at rest
B. Data encryption at rest
C. Remediation of vulnerabilities
D. Data encryption in transit

Answer: C

Explanation:
Comprehensive and Detailed Explanation
An API logic flaw is a weakness in the API's business logic — not in encryption or transport security. It occurs when:
API functions are misused
Business rules are bypassed
Security validation is missing
Workflow logic is incorrect
These issues cannot be solved through encryption (at rest or in transit). They require:
✓ Vulnerability remediation
✓ Fixing API code logic
✓ Updating API validation, flow control, and authentication logic

Thus the correct answer is C. Remediation of vulnerabilities.

Question: 54

An engineer must implement a solution on a Cisco ASR 1000 Series router to protect against DDoS attacks. DDoS traffic must be dropped by transmitting Flowspec attributes to edge routers, instructing them to generate an ACL via class-maps and policy-maps. The engineer already configured BGP neighbors. Which action must be taken next?

A. Configure Flowspec for the BGP address-family
B. Set the BGP routing process
C. Activate the BGP neighbors
D. Configure the route reflector

Answer: A

Explanation:
Comprehensive and Detailed Explanation
BGP Flowspec allows routers to distribute traffic-filtering rules using BGP NLRI.
To enable Flowspec, after neighbors are configured, the essential next step is:
✓ Activate the Flowspec address-family under BGP
Example:
router bgp 65000
address-family ipv4 flowspec
neighbor X.X.X.X activate
exit-address-family
This enables:
FlowSpec NLRI exchange
Distribution of drop rules (rate-limit, redirect, null route, etc.)
Automatic ACL/class-map/policy-map generation on edge routers
Why the other options are incorrect:
B . Set BGP routing process → already done when neighbors were configured
C . Activate neighbors → only makes sense inside an address-family; flowspec AF must be enabled first
D . Configure route reflector → optional and not required for Flowspec to operate
Thus, the correct next step is A. Configure Flowspec for the BGP address-family.

Question: 55

Refer to the exhibit.

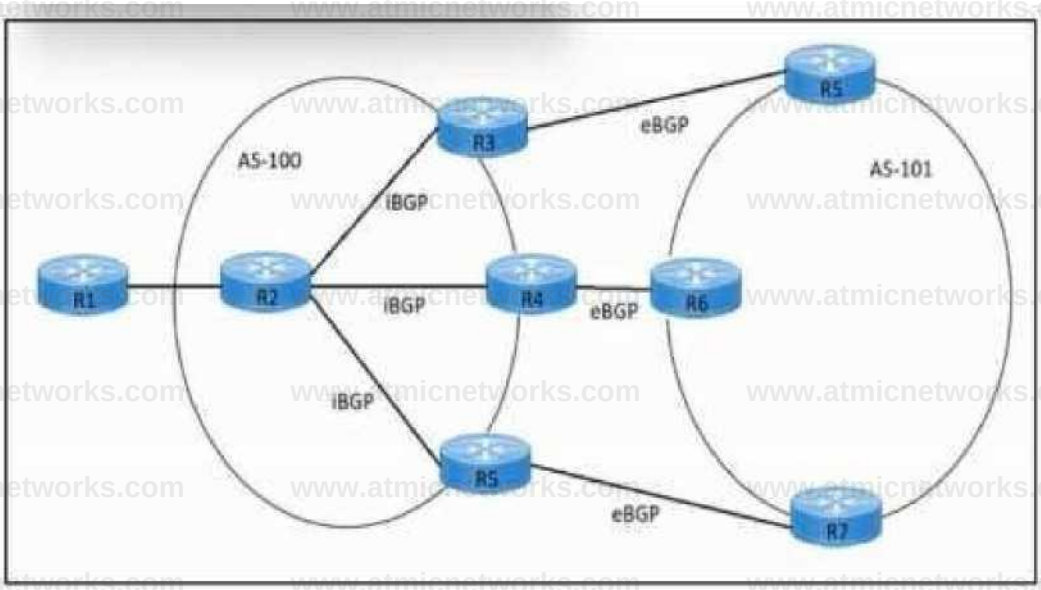

Refer to the exhibit. An engineer must configure iBGP multipath load sharing across three paths.
Which two commands must be run on router R2? (Choose two.)

A. router bgp 100
B. ip load-sharing ibgp 3
C. maximum-paths ibgp 3
D. router bgp 101
E. ip load-sharing per-destination

Answer: AC

Explanation:
Router R2 is inside AS 100 and has three iBGP paths (via R3, R4, R5) toward AS 101. To perform iBGP multipath across these three equal-cost paths, BGP must:
Run the correct BGP process for AS 100
Allow installation of multiple iBGP paths in the routing table
This is done with:
router bgp 100
maximum-paths ibgp 3
router bgp 100 – enters the BGP process for AS 100 (correct AS per diagram).
maximum-paths ibgp 3 – tells BGP to keep up to 3 iBGP paths to the same destination, enabling CEF to load-share across them.
Other options:
ip load-sharing ibgp 3 – not a valid command.
router bgp 101 – wrong AS number.
ip load-sharing per-destination – controls CEF hashing but does not enable BGP to install multiple iBGP paths by itself.

Question: 56

Refer to the exhibit.

#4Kk»i ,tei'»lcr - Ooi'-f Awpllce'-ion Contatf r<
loaded: loaded (/llb/syitend/kysten/docker.service; disabled; vendor preset; enabled) active: inactive (dead)
TriggeredOy: • docker.socket

Does: https//docs.docker.com

Refer to the exhibit. An engineer is troubleshooting a Cisco NFVI issue where the management node fails to start. Which service must be restarted to resolve the issue?

A. docker-kibana
B. docker
C. kube-apiserver
D. docker-cobbler

Answer: B

Explanation:
In Cisco NFVI, the management node relies heavily on Docker containers for:
NFVIS management functions
VIM services
Orchestration components
If the management node fails to start and the system shows:
docker.service: inactive (dead)
…then all Docker-based platform services also fail to start.
The correct recovery action is to restart the Docker engine:
systemctl restart docker
This brings up:
All NFVI-required Docker containers
Management services
REST APIs and cluster components
Why other answers are incorrect:
docker-kibana → Only affects Kibana logging container
docker-cobbler → Used for provisioning, not core NFVI management
kube-apiserver → Part of Kubernetes cluster, but relies on Docker; restarting it won't help until
Docker is running
Thus, the correct answer is B. docker.

Question: 57

What is a capability of a Cisco NFVIS SNMP trap?

A. Monitors the activities of a network host
B. Controls the activities of a network host
C. Sends an unsolicited notification to the SNMP manager
D. Retrieves an SNMP object variable from the MIB

Answer: C

Explanation:
Comprehensive and Detailed Explanation
SNMP traps in Cisco NFVIS (and in SNMP generally):
Are unsolicited notifications
Sent from the NFVIS device to the SNMP manager
Indicate alarms, changes, or significant operational events
Do not require polling
Examples:

Disk failures
VM crashes
Host status changes
Resource alarms
Why the others are wrong:
A describes SNMP monitoring (done by the manager with GET requests)
B SNMP cannot control host activities
D SNMP GET retrieves a variable, but traps send unsolicited notifications
Thus, the correct answer is C.

Question: 58

Refer to the exhibit.

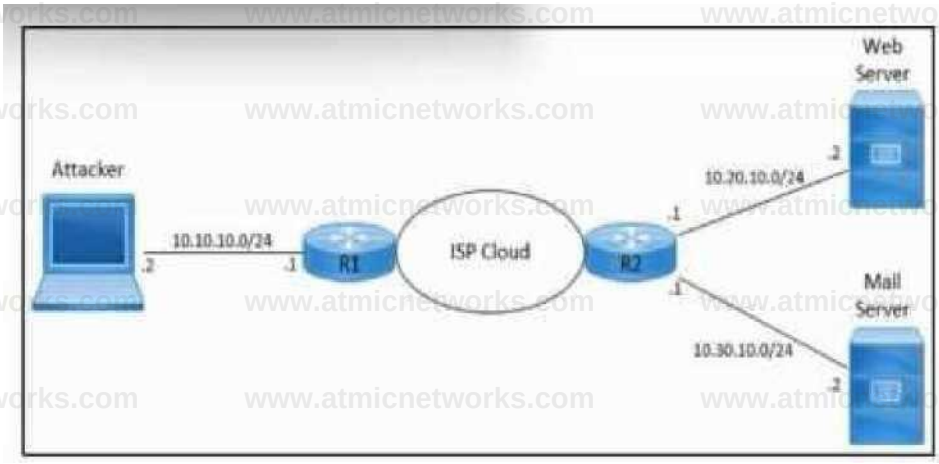

Refer to the exhibit. An engineer must stop DDoS attacks on web and mail servers by using an ACL. Which two commands must be run on router R17? (Choose two.)

A. access-list 101 deny ip 10.10.10.2 255.255.255.255 10.20.10.2 255.255.255.255
B. access-list 101 deny ip 10.0.0.0 0.255.255.255 10.10.0.2 0.0.0.0
C. access-list 101 deny ip 10.10.10.2 255.255.255.255 10.30.10.2 255.255.255.255
D. access-list 101 deny ip 10.10.10.2 0.0.0.0 10.20.10.2 0.0.0.0
E. access-list 101 deny ip 10.10.10.2 0.0.0.0 10.30.10.2 0.0.0.0

Answer: DE

Explanation:
The attacker's IP is:
10.10.10.2
The servers under attack are:
Web Server: 10.20.10.2
Mail Server: 10.30.10.2
We must deny traffic from attacker → servers. Correct ACL format uses host wildcards (0.0.0.0): deny ip 10.10.10.2 0.0.0.0 10.20.10.2 0.0.0.0 deny ip 10.10.10.2 0.0.0.0 10.30.10.2 0.0.0.0 These match D and E.

Question: 59

What is a benefit of a carrier-neutral data center?

A. Reduction in foot traffic

B. Path diversity on multiple carriers
C. Lower electrical costs
D. Avoidance of a price lock-in

Answer: B

Explanation:

Carrier-neutral facilities allow customers to connect to multiple telecom carriers, providing:
Physical path diversity
Redundancy
Competitive options
High availability for interconnects
Thus, path diversity on multiple carriers is the primary benefit.

Question: 60

Which type of cyberattack does Cisco Umbrella DNS-layer security effectively help mitigate?

A. Phishing and malware-based attacks
B. DDoS attacks targeting specific servers
C. Brute force attacks on user accounts
D. Advanced persistent threats and zero-day exploits

Answer: A

Explanation:
Cisco Umbrella DNS-layer security:
Blocks malicious domains used in phishing, malware, C2 communications, and ransomware
Stops threats before connections are made
Uses DNS-based filtering and threat intelligence
It does not mitigate:
DDoS (needs scrubbing centers)
Brute force login attempts
Zero-day exploits directly
Thus, A is correct.

Question: 61

An engineer attempts to kill a NETCONF session. The session ID is equal to the current session. What is the effect of this action?

A. An invalid-value error is returned, and the current user is not logged out.
B. The session is terminated, and all the connected users are logged out.
C. The configuration is saved, and the current user is logged out.
D. The configuration is removed, and the current user is logged out.

Answer: A

Explanation:

NETCONF rules state:
A user cannot terminate their own active session.
Attempting to kill the same session returns invalid-value error.
The session continues running normally.
Thus:
No logout
No termination
No configuration impact
Therefore, A is correct.