PassQueen
a.k.a dumpsqueen.com

Designing and Implementing Cisco Service Provider Cloud Network Infrastructure (SPCNI)

Cisco 300-540

Version Demo

Total Demo Questions: 7

Total Premium Questions: 76

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Architecture	20
Topic 2, Virtualization	21
Topic 3, Provisioning	11
Topic 4, Assurance	7
Topic 5, Security	17
Total	76

QUESTION NO: 1

An OpenStack tenant deploys a virtual firewall on an isolated tenant network. The firewall management interface must be reachable from an external provider network by using a public IPv4 address, while the tenant workload addresses must remain private. Which two Neutron actions must be taken? (Choose two.)

A. Configure a Neutron router with an external gateway.

B. Associate a floating IP address with the firewall management port.

C. Create an additional private tenant subnet.

D. Configure a Kubernetes Ingress resource.

E. Create a security group without an external gateway.

ANSWER: A B

Explanation:

A Neutron router with an external gateway provides Layer 3 connectivity between the tenant network and the external provider network. A floating IP must then be allocated from the external network and associated with the firewall management port to provide the public-to-private address mapping. Creating an additional tenant subnet does not provide external reachability, and a security group controls permitted traffic but does not create the required routing or address association.

QUESTION NO: 2

A customer router in AS 65010 connects to an upstream provider through eBGP neighbor 192.0.2.1. Prefix list CUSTOMER permits only 198.51.100.0/24, and prefix list ACCEPT-UPSTREAM permits the routes that the customer is willing to receive. The customer must advertise only its assigned prefix and must filter received routes from the provider. Which two commands must be configured under the BGP process? (Choose two.)

A. neighbor 192.0.2.1 prefix-list CUSTOMER out

B. neighbor 192.0.2.1 prefix-list CUSTOMER in

C. neighbor 192.0.2.1 prefix-list ACCEPT-UPSTREAM out

D. neighbor 192.0.2.1 prefix-list ACCEPT-UPSTREAM in

E. neighbor 192.0.2.1 next-hop-self

ANSWER: A D

Explanation:

The CUSTOMER prefix list must be applied outbound so that only 198.51.100.0/24 is advertised to the upstream provider. The ACCEPT-UPSTREAM prefix list must be applied inbound to filter routes received from that provider before they are accepted into the BGP table. Reversing either direction applies the policy to the wrong route flow. A route-map is not required when the stated prefix-list policies already express the required filtering logic.

QUESTION NO: 3

Which cloud provider connection permits BGP peering?

A. Azure S2S VPN

B. Azure Bastion

C. AWS Direct Connect

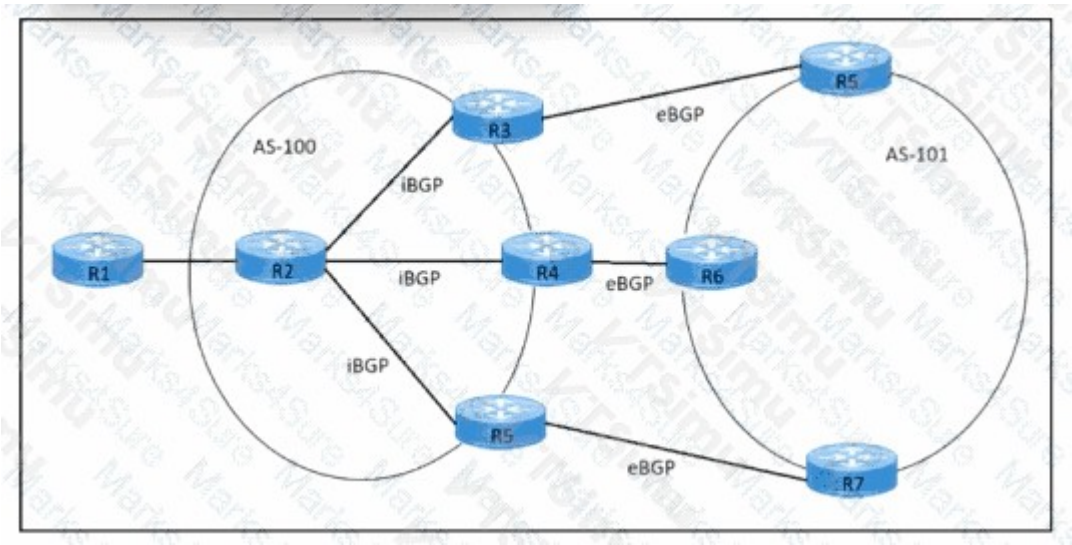
D. AWS-managed VPN

ANSWER: C

Explanation:

AWS Direct Connect is the intended correct answer because it provides dedicated connectivity between a customer or service-provider network and AWS, with routing exchanged through Border Gateway Protocol on a virtual interface. For private connectivity, the customer gateway establishes an external BGP session with AWS across a private virtual interface. This session advertises approved on-premises prefixes to AWS and receives AWS VPC routes, allowing route exchange to adapt dynamically as networks and prefixes change. BGP is fundamental to Direct Connect's routed design: customers specify their autonomous system number, configure the BGP peer addressing and ASN information supplied by AWS, and establish the session from their edge router. This makes Direct Connect a standard choice when a cloud interconnect requires scalable, policy-based Layer 3 route exchange rather than only simple reachability. AWS also documents BGP configuration requirements, including the customer router ASN and BGP peer parameters, for Direct Connect virtual interfaces. See the AWS Direct Connect virtual interface documentation and the AWS Direct Connect User Guide.

QUESTION NO: 4



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: A C

Explanation:

`router bgp 100` is required because R2 belongs to autonomous system 100 in the topology. This command enters the BGP routing process in which the iBGP neighbor and address-family settings are configured. The multipath setting must be applied under that BGP process, rather than under an unrelated routing process.

`maximum-paths ibgp 3` permits BGP to install as many as three eligible equal iBGP paths for the same prefix. Normally, BGP selects and installs only one best path. With this command, when the paths learned through the three iBGP peers satisfy BGP equal-path requirements, R2 can retain all three paths for forwarding. CEF then performs the actual traffic

distribution across the installed next hops according to its configured hashing behavior. The value of three specifically matches the required number of parallel iBGP paths in the exhibit.

Multipath eligibility still depends on the candidate routes having the necessary equivalent BGP path-selection attributes and resolvable next hops. Cisco documents BGP multipath behavior and the maximum-paths configuration in its BGP multipath technical documentation and IOS XE BGP configuration guide.

QUESTION NO: 5



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

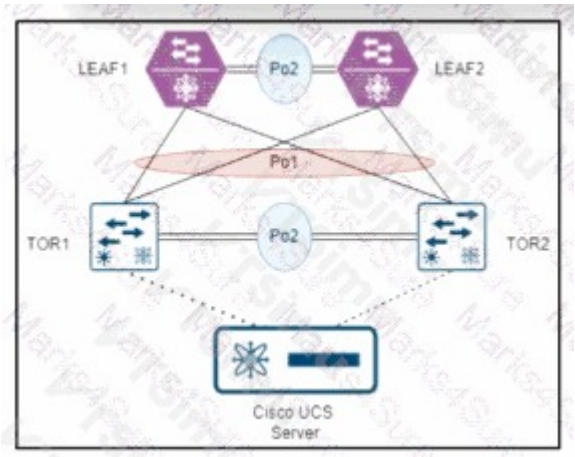
D. docker-cobbler

ANSWER: B

Explanation:

The service to restart is **docker**. Cisco NFVI deployments of this generation use Docker as the container runtime for the management-node platform services. The management node hosts the cloud-management and orchestration functions as a set of containers, so those functions depend on the Docker daemon being active and able to start containers. When the exhibit indicates that `docker.service` is inactive or dead, the appropriate recovery action is to restart the Docker service, for example with `systemctl restart docker`. Restarting the daemon restores the container runtime, allows its configured containers to be brought up, and enables the dependent management services to initialize normally. The service state can then be verified with `systemctl status docker`, and container recovery can be checked with `docker ps -a`. This aligns with Cisco NFVI operational troubleshooting, where management-plane services are containerized and Docker availability is a prerequisite for their operation. Cisco's NFVI documentation describes the management infrastructure and its container-based services; see the Cisco NFVI Installation Guide. For the standard systemd service-management commands used for this recovery, see the systemctl documentation.

QUESTION NO: 6



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:

Configure a vPC between TOR1 and TOR2. A double-sided vPC, also called a vPC-to-vPC topology, requires a vPC domain at each end of the interswitch port-channel. Because LEAF1 and LEAF2 are already operating as vPC peers, they provide one logical port-channel endpoint toward the ToR layer. TOR1 and TOR2 must likewise be configured as vPC peers so that they provide the corresponding logical endpoint for the same Layer 2 port-channel.

Specifically, the ToR pair needs its own vPC domain, a peer-keepalive path, and a vPC peer-link. After that peer relationship is established, the uplinks from TOR1 and TOR2 to the leaf pair can be bundled into matching port channels and assigned a common vPC identifier. This permits active forwarding across links from both ToR switches and both leaf switches while maintaining the loop-avoidance and consistency mechanisms provided by vPC. The peer-link also carries required synchronization and handles traffic for certain failure or orphan-port scenarios. Cisco documents that a vPC domain is formed by two peer switches using a peer-link and peer-keepalive connection, which is the required foundation for this topology. See the Cisco Nexus 9000 Series NX-OS vPC Configuration Guide and Cisco's Virtual PortChannel overview.

QUESTION NO: 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:

Implement a flexible service-modeling approach that leverages automation for ongoing management and refinement as demands on the network evolve. This approach directly supports the required lifecycle-oriented orchestration of services in a heterogeneous network. Cisco Network Services Orchestrator uses service models, commonly defined with YANG, to describe an intended service independently of the device-specific configuration needed to realize it. The orchestrator maps that service intent to the appropriate network-element implementations, providing a consistent abstraction layer across vendors and platforms.

Flexible service modeling is essential because the organization must accommodate both current legacy infrastructure and later changes in scale, topology, technologies, and service requirements. A reusable model can be extended and refined without forcing operators to redesign the automation strategy whenever the underlying network evolves. NSO also provides transactional service deployment behavior, including validation and coordinated configuration changes, helping maintain reliable execution across affected devices. This allows automation to cover initial provisioning, modifications, assurance-related updates, and retirement activities as part of the service lifecycle. Cisco describes NSO's model-driven, multivendor orchestration capabilities at Cisco NSO Developer Hub and provides its product overview at Cisco Network Services Orchestrator.