DUMPSARENA

Juniper Networks Certified Design Specialist - Data Center (JNCDS-DC)

Juniper JN0-1301

Version Demo

Total Demo Questions: 8

Total Premium Questions: 85

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Data Center Design Principles	10
Topic 2, Data Center Network Architectures	42
Topic 3, Data Center Network Components	11
Topic 4, Data Center Network Security	13
Topic 5, Data Center Network Management and Automation	3
Topic 6, Data Center Network Operations	6
Total	85

QUESTION NO: 1

Which two security components are inherent in a Contrail vRouter? (Choose two.)

A. basic tenant separation directional control for Layer 3 and Layer 4 traffic

B. Layer 7 traffic filters

C. security groups

D. encapsulation around each VM residing on the compute node

ANSWER: A C

Explanation:

Contrail vRouter inherently provides basic tenant separation and directional policy enforcement for Layer 3 and Layer 4 traffic. Each virtual network is associated with a routing instance, which separates tenant routing and forwarding domains on compute nodes. The vRouter then applies access-control-list policy in the forwarding path, enabling controls based on IP protocol, source and destination addressing, and transport-layer ports. This provides distributed enforcement close to the workload rather than requiring all traffic to traverse a centralized security appliance.

Security groups are also an inherent Contrail vRouter security component. A security group is a stateful collection of rules that can be associated with virtual-machine interfaces. The vRouter agent programs the applicable rules locally and enforces them for traffic entering or leaving those interfaces. Security-group policy can therefore follow a workload as it is instantiated or moved, while remaining independent of the virtual network's broader policy definition. Juniper describes Contrail security as using network policies and security groups, with enforcement implemented in the vRouter datapath. See Juniper's Contrail security overview and security group overview.

QUESTION NO: 2

Which two statements describe the Contrail vRouter? (Choose two.)

A. It is forwarding plane that runs software on a virtualized server.

B. It extends the network form physical routers and switches into an overlay network hosted in virtual server.

C. It is responsible for providing the management, control, and analytics functions of the virtualized network.

D. It provides the logically centralized control plane and management plane of the system.

ANSWER: A B

Explanation:

Contrail vRouter is the distributed data-plane element deployed on each compute host, so "It is forwarding plane that runs software on a virtualized server." correctly identifies its primary role. It performs packet forwarding for workloads hosted on the server and applies virtual-network connectivity and policy locally, rather than sending every packet to a centralized appliance.

"It extends the network form physical routers and switches into an overlay network hosted in virtual server." also describes the vRouter's function. The vRouter connects virtual-machine or container interfaces to virtual networks and encapsulates traffic for transport across the physical IP fabric. This provides an overlay that allows the logical virtual network to span multiple servers while using the underlay network of physical routers and switches for IP transport. Juniper's Contrail architecture describes the vRouter as the forwarding-plane component installed on compute nodes and explains its role in connecting workloads to overlay virtual networks. See Juniper's Contrail architecture overview and the vRouter overview.

QUESTION NO: 3

You want to use the Intrusion Prevention System features of a vSRX virtual firewall to protect systems from attacks embedded in traffic. In this scenario, which three technologies would be used? (Choose three.)

A. protocol decoders

B. zero-day protection

C. traffic normalization

D. application firewall

E. application tracking

ANSWER: A B C

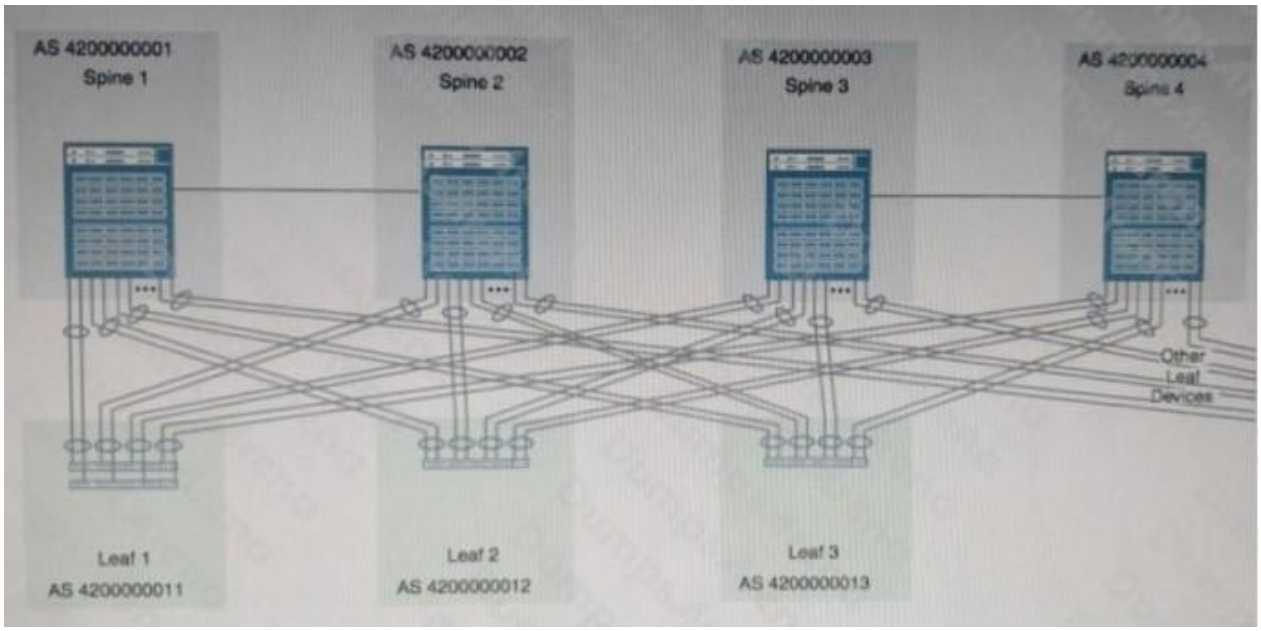
Explanation:

IPS inspection on a vSRX uses protocol decoders, zero-day protection, and traffic normalization to identify and mitigate attacks carried within permitted traffic flows. Protocol decoders interpret the structure and fields of supported network and application protocols, allowing IPS inspection to recognize malformed or suspicious protocol behavior rather than treating traffic as opaque data. Traffic normalization removes ambiguities in packet and stream construction, such as fragmentation and TCP segmentation techniques, so that IPS inspection evaluates traffic consistently and attackers cannot evade detection through alternate encodings of the same session.

Zero-day protection is provided through behavioral and anomaly-oriented inspection capabilities that complement signature-based detection. This enables the IPS service to identify suspicious protocol behavior and exploit characteristics even when a specific attack signature is not yet available. Together, these technologies provide protocol-aware inspection, evasion resistance, and protection against previously unknown threats. Juniper describes IPS as a security service that detects and prevents attacks using protocol and traffic inspection mechanisms; see the Juniper IPS overview and the Juniper SRX Series security-services overview.

QUESTION NO: 4

Click the Exhibit button. You are creating a resilient IP fabric for a data center based on a Clos architecture. Referring to the exhibit, which statement is true?



A. A physical link must be added between spine 2 and spine 3 and between spine 1 and spine 4.

B. Additional physical links between spine 1 and spine 2 and between spine 3 and spine 4 must be added for resiliency.

C. No physical links need to be added or removed.

D. The physical links between spine 1 and spine 2 and between spine 3 and spine 4 must be removed.

ANSWER: D

Explanation:

The physical links between spine 1 and spine 2 and between spine 3 and spine 4 must be removed. A Clos-based IP fabric uses distinct device roles and interconnects adjacent tiers: leaf devices connect upward to spine devices, while spine devices provide the transit layer between leaves. Spine devices are not interconnected with one another. Resiliency and equal-cost multipath forwarding are achieved through multiple independent leaf-to-spine links, so a leaf can use multiple spine paths and the fabric can tolerate a link or spine failure without requiring links inside the spine tier.

Keeping links between spine switches introduces nonstandard intra-tier connectivity that is unnecessary in a Clos fabric and can complicate the intended predictable, scalable topology. The design should instead preserve a consistent leaf-to-spine structure, with each leaf connected to the appropriate spine switches. This arrangement supports the scale-out behavior and multipath characteristics expected of a data center IP fabric. Juniper describes Clos fabrics as a leaf-and-spine architecture designed for scalable, highly available data center connectivity; see the Juniper Clos fabric overview and Juniper's data center networking solutions.

QUESTION NO: 5

You want to provide Layer 2 connectivity to a backup data center or to a disaster recovery site. What would you use to accomplish this task?

A. Data Center Interconnect

B. network redundancy

C. SDN

D. SAN protocols

ANSWER: A

Explanation:

Data Center Interconnect is the appropriate solution because it connects geographically separate data centers and can extend Layer 2 services between them. This capability is particularly useful when a backup or disaster-recovery location must support workload mobility, application recovery, or consistent network segmentation across sites. A DCI design transports Ethernet connectivity over the selected intersite infrastructure, allowing VLAN-based services and associated Layer 2 reachability to be delivered between the primary and recovery data centers.

In a Juniper data center design, DCI is a use case rather than a single mandatory protocol. The implementation can use technologies appropriate to the distance, resiliency, scale, and operational model, such as EVPN-VXLAN overlays with an IP fabric and intersite connectivity. EVPN provides a standards-based control plane for Ethernet virtual private networks, while VXLAN supplies the overlay encapsulation used to extend Layer 2 segments across a routed underlay. The essential requirement described is therefore the interconnection of data centers, not merely internal redundancy or storage access. Juniper's EVPN-VXLAN overview describes this architecture, and the EVPN-VXLAN DCI overview covers its data-center-interconnect use.

QUESTION NO: 6

You plan to deploy a set of SRX Series devices in an active/passive chassis cluster with an active/passive network infrastructure. Which two statements are true in this scenario? (Choose two.)

A. End-host traffic is serviced through both SRX nodes during non-failure conditions.

B. End-host traffic is serviced through one SRX node during non-failure conditions.

C. In non-failure conditions, the end-host traffic across the fabric link is present.

D. In non-failure conditions, the end-host traffic across the fabric link is eliminated.

ANSWER: B D

Explanation:

In an active/passive SRX chassis cluster, one node is the primary node for redundancy group 0 and, under normal operating conditions, is also active for the data-plane redundancy groups that own the traffic interfaces. Therefore, end-host traffic is serviced through one SRX node during non-failure conditions. The secondary node remains synchronized and prepared to assume traffic processing if a monitored failure causes redundancy-group failover.

With an active/passive network infrastructure, the connected switching or routing topology directs normal traffic to the interfaces of the active SRX node. This design avoids a normal-state forwarding path in which traffic arrives at one cluster member and must be sent across the chassis-cluster fabric link to the other member for processing or egress. Consequently, end-host traffic across the fabric link is eliminated during non-failure conditions. The fabric link remains essential for cluster control, state synchronization, and any traffic handling required during failover or exceptional forwarding scenarios, but it is not part of the intended steady-state end-host data path in this deployment model.

Juniper describes SRX chassis-cluster node roles, redundancy groups, and fabric-link functions in the SRX chassis cluster overview and its redundancy groups documentation.

QUESTION NO: 7

What would you use to balance traffic loads between servers in a Clos Layer 3 data center fabric?

A. MC-LAG

B. multicast protocols

C. ECMP routing

D. BGP export policies

ANSWER: C

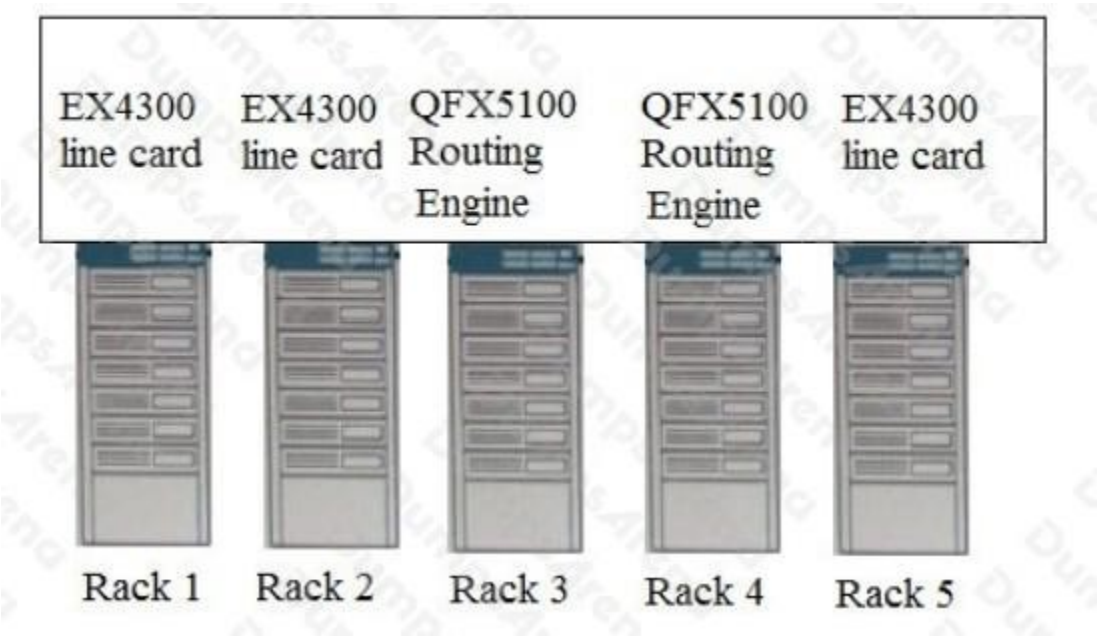
Explanation:

ECMP routing is used to balance traffic loads in a Clos Layer 3 data center fabric. A Clos fabric provides multiple equal-cost paths between leaf switches and destination servers through its spine layer. When the routing protocol installs these equal-cost next hops in the forwarding table, ECMP enables the device to distribute traffic flows across them rather than selecting only one path. This uses the available spine links efficiently, provides high aggregate bandwidth, and allows the fabric to continue forwarding traffic over the remaining equal-cost paths if a link or spine device fails.

In Juniper data center designs, BGP is commonly used as the underlay routing protocol for an IP fabric, and its multipath capabilities can install several paths with the same cost. The forwarding plane then performs hash-based load balancing, typically using packet-header fields so that packets in an individual flow remain ordered while different flows can use different paths. This behavior is fundamental to the scalable, resilient east-west traffic model of a Layer 3 Clos topology. Juniper's IP fabric documentation describes ECMP as a core mechanism for distributing traffic across equal-cost routes: Juniper IP Fabric Overview. Additional ECMP configuration and forwarding details are available in the Junos ECMP Load Balancing documentation.

QUESTION NO: 8

Click the Exhibit button. You are designing a Layer 2 Ethernet fabric. Your customer wants to know which protocol in your design will ensure a loop-free topology. Referring to the exhibit, which protocol meets this requirement?



A. STP

B. VCCP

C. BGP

D. EVPN

ANSWER: B

Explanation:

VCCP is correct because the Virtual Chassis Control Protocol is the control-plane protocol used by Juniper Virtual Chassis Fabric systems to establish and maintain the fabric topology. It discovers the interconnected switches, elects the appropriate control-plane roles, and manages the Virtual Chassis Ports that form the fabric links. When redundant interswitch paths are present, VCCP identifies the topology and prevents Layer 2 forwarding loops by placing the required redundant fabric links into a nonforwarding state. This allows the fabric to retain physical redundancy while presenting and operating as one logical switching system, rather than relying on independent switches to make separate loop-prevention decisions. The protocol also reacts to topology changes, enabling the fabric to preserve a loop-free forwarding topology when links or members change. Juniper's Virtual Chassis Fabric documentation describes VCCP as the protocol responsible for discovering and maintaining the fabric, including its loop-free topology behavior. See the Juniper Virtual Chassis Fabric overview and the Juniper documentation on understanding Virtual Chassis Fabric.