PassQueen
a.k.a dumpsqueen.com

Nokia Data Center Fabric Fundamentals

Nokia 4A0-D01

Version Demo

Total Demo Questions: 5

Total Premium Questions: 56

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

When autoboot flag is set in grub.cfg, which of the following is NOT performed as part of the Zero Touch Provisioning (ZTP) autoboot process?

A. Uses the DHCP server to provide the IP address and default gateway

B. Uses the Web or FTP server to provide the RPM packages and any necessary images

C. Downloads and executes provisioning scripts

D. Uses Nokia's Fabric Services System to deploy the digital sandbox automatically

ANSWER: D

Explanation:

Uses Nokia's Fabric Services System to deploy the digital sandbox automatically is the activity that is not part of the SR Linux ZTP autoboot workflow. The autoboot flag directs the node to begin its initial, unattended provisioning sequence at startup. This process is designed to establish the node's basic network reachability and obtain the information and artifacts needed to configure the node, using the ZTP mechanisms available during boot. Its purpose is bootstrap provisioning of the network device itself, rather than deployment of an external fabric-management platform or a lab environment.

Fabric Services System is a separate Nokia platform used for fabric automation and lifecycle operations. A digital sandbox is likewise an environment created and managed through FSS workflows; it is not an intrinsic boot-stage action executed by the SR Linux ZTP client because an autoboot setting is present in `grub.cfg`. Consequently, enabling autoboot does not cause the node to install, instantiate, or automatically deploy FSS or a digital sandbox. Nokia's SR Linux documentation describes ZTP as a device bootstrap mechanism, while Fabric Services System documentation treats sandbox management as an FSS capability with its own deployment and operational workflow. See the Nokia SR Linux Zero Touch Provisioning documentation and the Nokia Fabric Services System documentation.

QUESTION NO: 2

Based on the exhibit output, which of the following is NOT the reason the user runs the â€œbaseline diffâ€□ command?

A. User admin is not able to commit their changes.

B. Another user has committed changes.

C. User admin wants to check if there are differences between the current running config and their baseline datastore.

D. User admin notices that â€œ!â€□ appears in the CLI prompt.

ANSWER: A

Explanation:

User admin is not able to commit their changes. is the correct selection because an inability to commit is not the operational purpose of the `baseline diff` command. In the SR OS MD-CLI configuration workflow, a baseline represents the running-configuration state against which a user's current configuration session is compared. The `baseline diff` command is an inspection command: it displays configuration changes between that baseline and the current running configuration. It therefore helps an administrator understand whether the running configuration has changed since the baseline was established, including changes that can occur while more than one administrator is working on the router.

The command provides visibility before the user decides how to proceed with their own pending configuration work. A commit is a separate configuration action governed by the configuration session, validation, locking, and conflict-handling behavior; simply being unable to commit is not what `baseline diff` is designed to diagnose or resolve. Nokia's MD-CLI documentation describes baseline-based configuration comparison as part of its configuration-management workflow. See the Nokia SR OS MD-CLI configuration management documentation and the Nokia SR OS MD-CLI overview.

QUESTION NO: 3

A workload moves from Leaf1 to Leaf2 while retaining its MAC address. Both leaf routers advertise EVPN MAC/IP Advertisement routes for the same MAC address during convergence. Which attribute allows remote VTEPs to identify the advertisement associated with the more recent location?

A. The MAC Mobility sequence number

B. The route distinguisher

C. The VXLAN network identifier

D. The Ethernet Segment identifier

ANSWER: A

Explanation:

The MAC Mobility extended community carries a sequence number that identifies the relative freshness of MAC reachability advertisements. When the same MAC is advertised from a new location with a higher sequence number, remote VTEPs use the newer advertisement. The route distinguisher makes a route unique, but it does not indicate that a MAC has moved. The VNI identifies the service, and the ESI identifies an Ethernet Segment.

QUESTION NO: 4

Which of the following is NOT how prefixes are typically learned and advertised in a leaf and spine data center?

A. Spine routers advertise learned prefixes to leaf routers directly using eBGP.

B. Spine routers learn prefixes from leaf routers directly using eBGP.

C. Each leaf router advertises the prefixes of its own locally connected servers.

D. Each leaf router learns the prefixes of remote servers from other leaf routers directly using eBGP.

ANSWER: D

Explanation:

Each leaf router learns the prefixes of remote servers from other leaf routers directly using eBGP. is the relationship that is not typical in a leaf-and-spine fabric. The physical and routing topology is deliberately built so that every leaf connects northbound to every spine, while leaf-to-leaf links and BGP adjacencies are normally absent. A leaf originates or redistributes reachability for its locally attached endpoints and advertises those prefixes to its directly connected spine peers. Each spine consequently receives endpoint reachability from the originating leaf and advertises the eligible routes to its other directly connected leaf peers. A destination leaf therefore learns a remote server prefix from one or more spine routers, not from the remote server's leaf router through a direct eBGP session.

This arrangement produces a scalable, repeatable Clos fabric: adding a leaf does not require it to establish routing sessions with every existing leaf, and the spines provide the common routing layer between all leaves. Equal-cost paths through multiple spines can also be installed when the routing policy and BGP path attributes permit them. Nokia's BGP documentation describes the use of BGP peers and route advertisement in the routing control plane, while its EVPN documentation covers the data-center overlay use cases commonly deployed over this leaf-spine underlay. See Nokia SR Linux BGP documentation and Nokia SR Linux EVPN documentation.

QUESTION NO: 5

A leaf-spine underlay has four equal-cost paths between two leaf routers. A single long-lived TCP flow is transmitted continuously while all four paths remain available. Which behavior is expected from ECMP forwarding?

A. Each packet is forwarded through a randomly selected spine router.

B. The flow is replicated through every equal-cost path.

C. The flow normally remains mapped to one eligible next hop.

D. The flow is forwarded only after all equal-cost paths have been tested.

ANSWER: C

Explanation:

ECMP uses a hash derived from packet-header fields to select one eligible next hop for a flow. Packets belonging to the same TCP flow normally produce the same hash result and remain on the selected path, helping preserve packet ordering. ECMP distributes multiple flows across equal-cost paths; it does not normally spray packets of one flow across all paths.