

Specialist - Implementation Engineer, PowerEdge MX Modular

EMC DES-4421

Version Demo

Total Demo Questions: 8

Total Premium Questions: 83

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, PowerEdge MX Hardware	18
Topic 2, PowerEdge MX Networking	16
Topic 3, PowerEdge MX Storage	8
Topic 4, PowerEdge MX Management	38
Topic 5, PowerEdge MX Troubleshooting	3
Total	83

QUESTION NO: 1

What is the number of server-facing ports and the bandwidth of each port on the MXG610s?

(Choose two.)

- A. 8x 16 Gbps
- B. 16x 32 Gbps
- C. 16x 16 Gbps
- D. 8x 32 Gbps

ANSWER: B C

Explanation:

The MXG610s Fibre Channel switch module provides 16 server-facing internal FC ports for connecting PowerEdge MX compute sleds through the MX7000 chassis midplane. These ports support 32 Gbps Fibre Channel operation, so **16x 32 Gbps** accurately states both the available server-facing port count and the module's maximum per-port link speed. The same 16 internal ports are also backward compatible with 16 Gbps Fibre Channel operation. Therefore, **16x 16 Gbps** is also valid when the attached host adapter, configuration, or negotiated fabric speed uses 16 Gbps FC. The port count does not change when the negotiated FC speed changes; the distinction is the operating bandwidth of each of the 16 server-facing links. This capability allows MXG610s deployments to support existing 16 Gbps FC environments while providing a migration path to 32 Gbps FC connectivity. Dell's MXG610s specification sheet identifies the 16 internal server-facing FC interfaces and their supported speeds. See the [Dell EMC Networking MXG610s Specification Sheet](#) and the [Dell MXG610s support documentation](#).

QUESTION NO: 2

An administrator is planning Fibre Channel connectivity for MX compute sleds through Fabric B.

Which implementation requirements should be considered?

(Choose three.)

- A. A supported Fibre Channel mezzanine adapter in the compute sled
- B. A supported Fibre Channel I/O module in Fabric B
- C. Appropriate SAN zoning for host and storage identities
- D. An Ethernet VLAN configured on Fabric A
- E. A QuickSync 2 module in the chassis

ANSWER: A B C

Explanation:

Fibre Channel connectivity through Fabric B requires compatible hardware throughout the path. The compute sled must have a supported Fibre Channel mezzanine adapter assigned to Fabric B, and the chassis must contain a supported Fibre Channel I/O module in the corresponding Fabric B slot. The external SAN connection must also be configured with the required Fibre Channel fabric settings, including appropriate zoning for the host WWNs and storage targets. Ethernet switch VLAN configuration alone does not provide Fibre Channel SAN connectivity. Dell documentation should be consulted to validate the specific compute sled, adapter, I/O module, and firmware combinations before deployment. See the [PowerEdge MX7000 documentation](#) and the [Dell Networking MXG610s documentation](#).

QUESTION NO: 3 - (DRAG DROP)

DRAG DROP

A technician is deploying a new MX7000 chassis using the Chassis Deployment Wizard.

What is the correct order of steps followed by the wizard?

Select and Place:

Configure iDRAC and IOM Quick Deploy Settings	STEP 1
Import Profile or join Multi-chassis Group	STEP 2
Configure Activity and Alerts	STEP 3
Define Groups	STEP 4
Update Firmware and configure Proxy	STEP 5

ANSWER:

Configure iDRAC and IOM Quick Deploy Settings	Import Profile or join Multi-chassis Group
Import Profile or join Multi-chassis Group	Configure Activity and Alerts
Configure Activity and Alerts	Configure iDRAC and IOM Quick Deploy Settings
Define Groups	Update Firmware and configure Proxy
Update Firmware and configure Proxy	Define Groups

Explanation:

The Chassis Deployment Wizard follows an order that lets the chassis inherit or establish its overall management context before applying deployment-specific settings. The first step, **Import Profile or join Multi-chassis Group**, is used to bring in an established chassis configuration or associate the new chassis with the appropriate multi-chassis environment. This provides the intended administrative baseline at the start of deployment.

Next, **Configure Activity and Alerts** establishes how the chassis records operational events and sends notifications. With monitoring behavior defined, the wizard moves to **Configure iDRAC and IOM Quick Deploy Settings**. Quick Deploy settings provide the configuration used for server iDRACs and installed I/O modules, helping the chassis consistently apply management-network and related deployment settings as components are discovered or configured.

The workflow then uses **Update Firmware and configure Proxy** to establish firmware readiness and the connectivity method used to obtain update content. Completing firmware and proxy settings before organizing resources supports a more consistent operational state for the deployed infrastructure. Finally, **Define Groups** organizes the deployed resources for ongoing administration, monitoring, and policy-oriented management after the foundational chassis configuration has been completed.

This sequence is consistent with Dell OpenManage Enterprise Modular chassis deployment guidance and the product documentation available through the [Dell OpenManage Enterprise Modular documentation portal](#). The displayed placement correctly follows the wizard sequence from initial profile or group association through final resource grouping.

QUESTION NO: 4

A company has several MX7000-systems, four FX2 solutions, one M1000e, and several PowerEdge rack servers.

Which management interface is recommended for managing this range of systems?

- A. OpenManage Modular, no optional plugins required
- B. OpenManage Essentials
- C. OpenManage Modular with optional plugins
- D. OpenManage Enterprise

ANSWER: D

Explanation:

OpenManage Enterprise is the recommended management interface because it provides centralized, enterprise-scale lifecycle management across the company's mixed PowerEdge infrastructure. It is designed to discover, inventory, monitor, configure, update, and manage PowerEdge servers from one web-based console, avoiding the need to use separate management interfaces for each server form factor. For modular environments, OpenManage Enterprise can integrate with the relevant modular-management capabilities and extensions to provide a unified operational view alongside rack servers. This makes it appropriate for an estate that includes MX7000 modular systems, FX2 converged infrastructure, M1000e blade infrastructure, and standalone PowerEdge rack servers.

OpenManage Enterprise also supports consistent policy-based configuration, firmware and driver compliance reporting, alert handling, and deployment workflows. These functions help administrators apply common operational practices across a heterogeneous server fleet while retaining the platform-specific controls required by modular systems. Dell positions OpenManage Enterprise as its centralized infrastructure-management console for PowerEdge environments. See the [Dell OpenManage Enterprise product page](#) and the [OpenManage Enterprise documentation](#).

QUESTION NO: 5

Which management solution is embedded in the PowerEdge MX7000 chassis and provides chassis-level inventory, health monitoring, power control, and configuration management?

- A. OpenManage Enterprise–Modular
- B. OpenManage Mobile
- C. Lifecycle Controller
- D. SmartFabric OS10

ANSWER: A

Explanation:

OpenManage Enterprise–Modular is the embedded management solution for the PowerEdge MX7000 chassis. It runs on the chassis Management Modules and provides management of chassis infrastructure, including compute and storage sled inventory, I/O modules, power supplies, fans, alerts, firmware workflows, and Multi-Chassis Management operations. Individual compute sled iDRAC controllers continue to provide server-specific out-of-band management. See the [OpenManage Enterprise–Modular documentation](#) for supported chassis-management functions.

QUESTION NO: 6 - (SIMULATION)

Match each deployment logical block with its attributes.

ANSWER: See the explanation for the answer

Explanation:

Match the PowerEdge MX deployment logical blocks as follows:

Compute block — MX compute sleds (for example, PowerEdge MX740c/MX840c) installed in an MX7000 chassis; supplies the CPU and memory resources for workloads.

Storage block — MX5016s storage sleds and their drives; provides shared/direct-attached SAS storage capacity to compute sleds through the storage fabric.

Network block — Ethernet fabric and I/O modules, such as MX9116n Fabric Switching Engines, MX5108n Ethernet switches, or pass-through modules; provides server connectivity and external network uplinks.

Management block — the redundant MX7000 management modules running OpenManage Enterprise–Modular (OME-M), with iDRAC on each sled; provides chassis, sled, firmware, template, and policy management.

In short: Compute = MX compute sleds, Storage = MX5016s/drives, Network = fabric I/O modules and uplinks, and Management = OME-M/management modules and iDRAC.

QUESTION NO: 7

Which server I/O identities can be virtualized by FlexAddress in a PowerEdge MX environment?

(Choose two.)

- A. MAC addresses
- B. World Wide Names
- C. iDRAC IP addresses
- D. Compute sled service tags

ANSWER: A B

Explanation:

FlexAddress virtualizes MAC addresses for Ethernet connectivity and World Wide Names for Fibre Channel connectivity. These identities can be persistently assigned by the chassis so that supported replacement hardware can assume the established network and SAN identities. IP addresses are configured through network services or management interfaces and are not FlexAddress identities. A system service tag is a physical system identifier and is not virtualized by FlexAddress. See the [PowerEdge MX7000 documentation](#) for FlexAddress requirements and supported adapter configurations.

QUESTION NO: 8

Which MX7000 feature allows an administrator to virtualize network WWN and MAC addresses?

- A. FlexAddress
- B. MX Scalable Architecture
- C. Identity Pool
- D. Virtual Top of Rack

ANSWER: A

Explanation:

FlexAddress is the MX7000 feature that virtualizes and persistently assigns server I/O identities, including MAC addresses for Ethernet adapters and World Wide Names (WWNs) for Fibre Channel adapters. Rather than tying these identifiers permanently to the physical network adapters in a compute sled, FlexAddress allows them to be assigned at the chassis

level. This gives an administrator identity persistence when a compute sled is replaced or moved within the modular infrastructure: the replacement hardware can assume the established MAC and WWN identities without requiring corresponding network or SAN zoning changes. This capability simplifies hardware service operations and helps preserve the server's expected connectivity configuration across lifecycle events. Dell describes FlexAddress as providing persistent, chassis-assigned virtual MAC addresses and WWNs for server I/O. See the [Dell PowerEdge MX7000 Technical Guide: FlexAddress](#) and the [PowerEdge MX7000 Owner's Manual](#).