

DUMPSBOSS.

**Specialist - Implementation Engineer
PowerEdge Version 2.0**

EMC DES-4122

Version Demo

Total Demo Questions: 8

Total Premium Questions: 82

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

**support@dumpsboss.co
dumpsboss.co**

Topic Break Down

Topic	No. of Questions
Topic 1, PowerEdge Server Hardware	22
Topic 2, PowerEdge Server Configuration	23
Topic 3, PowerEdge Server Deployment	1
Topic 4, PowerEdge Server Management	27
Topic 5, PowerEdge Server Troubleshooting	9
Total	82

QUESTION NO: 1 - (DRAG DROP)

Match each RAID level to its characteristic.

Level	Characteristic
RAID 10	Highest disk overhead
RAID 5	Failure of one drive affects overall performance
RAID 6	Requires minimum of four drives
RAID 1	Used for high performance databases

ANSWER:

Level	Characteristic
RAID 10	RAID 1
RAID 5	RAID 5
RAID 6	RAID 6
RAID 1	RAID 10

Explanation:

RAID 1 maps to **Highest disk overhead** because it mirrors data: each block written to one drive is also written to its mirror partner. This duplication provides redundancy and can improve read availability, but it consumes approximately half of the installed raw capacity for protected storage. For example, two equal-sized drives in a RAID 1 mirror provide the usable capacity of only one drive. Dell describes RAID 1 as mirrored disks, with the second disk maintaining a copy of the first disk's data: [Dell RAID overview](#).

RAID 5 maps to **Failure of one drive affects overall performance**. RAID 5 distributes parity information across member drives. Following a single-drive failure, reads involving data from the failed member require reconstruction from the surviving data and parity blocks. That reconstruction work increases I/O activity and can reduce application performance until the replacement drive is rebuilt.

RAID 6 maps to **Requires minimum of four drives**. RAID 6 uses two independent parity values distributed across the drives, allowing the array to tolerate failure of up to two drives. Four drives are the minimum needed to provide the required data and dual-parity layout. Dell's RAID reference identifies RAID 6 as dual-parity RAID and specifies the minimum drive requirement: [Dell PERC RAID levels documentation](#).

RAID 10 maps to **Used for high performance databases**. It stripes data across mirrored drive pairs, combining parallel I/O with the resilience of mirroring. This design avoids parity-calculation overhead on writes and provides strong random read and write performance, characteristics that suit database workloads with frequent transactional activity.

QUESTION NO: 2

A Dell EMC PowerEdge server has been updated with an LCC and iDRAC DUP to the latest firmware when performing a firmware rollback?

- A. Both the iDRAC and the LCC can be rolled back
- B. The LCC can be rolled back; the iDRAC cannot be rolled back
- C. The iDRAC can be rolled back; the LCC cannot be rolled back
- D. The iDRAC cannot be rolled back and the LCC cannot be rolled back

ANSWER: C

Explanation:

The iDRAC can be rolled back; the LCC cannot be rolled back. On PowerEdge servers, iDRAC firmware maintains a rollback-capable firmware image mechanism that permits returning iDRAC to its previously installed firmware version when a rollback is supported for that release. This capability is intended to restore the iDRAC management controller software after an update, typically through the iDRAC web interface or a supported update workflow. Lifecycle Controller firmware is delivered as part of the iDRAC with Lifecycle Controller update package, but it does not have an independently supported rollback mechanism. Therefore, an administrator planning firmware maintenance should treat a Lifecycle Controller update as non-reversible and verify the target DUP, release notes, prerequisites, and supported downgrade path before applying it. Dell's iDRAC documentation describes firmware update and rollback behavior for the controller, while the Lifecycle Controller documentation and update guidance identify the relationship between iDRAC and Lifecycle Controller firmware. See Dell's [iDRAC with Lifecycle Controller documentation](#) and the Dell knowledge-base article on [updating PowerEdge firmware](#).

QUESTION NO: 3

What is a requirement to operate GPUs in a Dell EMC PowerEdge R740 server?

- A. All GPU cards must be installed at the same time
- B. All GPU cards must be of the same type but can be a different model
- C. All GPU cards must be of the same type and model
- D. All GPU cards must be installed on the same riser card.

ANSWER: C

Explanation:

All GPU cards must be of the same type and model is the applicable PowerEdge R740 configuration requirement. Dell validates GPU configurations as a homogeneous set, meaning that every installed accelerator must be the same supported GPU product and model. This consistency is important because the server's GPU enablement design—including supported riser and slot layouts, power cabling, auxiliary GPU power allocation, thermal profile, firmware validation, and driver support—is qualified around matched accelerator configurations. A uniform GPU set also provides predictable capabilities and resources to the operating system and any virtualization, AI, VDI, or compute software using the accelerators.

Before installation, the selected GPU model must also be supported for the specific R740 configuration, and the system must include the appropriate GPU-capable risers, power supplies, cables, and cooling components specified by Dell. These installation dependencies are documented in the PowerEdge R740 owner's manual and the server technical guide. Consult Dell's [PowerEdge R740 Owner's Manual](#) for GPU installation and configuration instructions, and the [PowerEdge R740 Technical Guide](#) for supported hardware configuration details.

QUESTION NO: 4

A company that has very specific firmware requirements has acquired Dell EMC PowerEdge servers. Any replaced parts of the type must match the original parts identically. The procedure must be as simple as possible with minimal risk.

What action should a technician take to help accomplish this goal?

- A. Configure servers to access the remote firmware repository

- B. Enable the Match Firmware of Replaced Part setting in the system BIOS
- C. Install and configure a repository manager
- D. Enable the Collect System Inventory On Restart (CSIOR) attribute

ANSWER: B


Explanation:

Enable the Match Firmware of Replaced Part setting in the system BIOS. This PowerEdge Lifecycle Controller feature is designed specifically for environments that require replacement components to retain firmware consistency with the component they replace. When the setting is enabled, Lifecycle Controller preserves the firmware version information for supported hardware components and, after a replacement is detected, automatically updates the replacement part to the previously installed matching firmware level. This avoids a technician having to identify, obtain, and manually apply a particular firmware package during a service event, reducing both effort and the possibility of deploying an unintended version.

The feature is particularly appropriate when an organization has validated a specific firmware baseline and wants service replacements to conform to that baseline rather than simply receive the newest available update. Dell documents the setting as *Match Firmware of Replaced Part* in Lifecycle Controller/System Setup and describes its use for automatically matching replacement-part firmware. Ensure the required firmware packages remain available in the Lifecycle Controller local repository or through the organization's configured update source so the automatic remediation can complete. See Dell's [Lifecycle Controller User's Guide: Match Firmware of Replaced Part](#) and the [Dell PowerEdge Lifecycle Controller documentation](#).

QUESTION NO: 5 - (SIMULATION)

You are Investigating a disk error on a Dell EMC PowerEdge server and discover a Hot Spare has not been assigned to the Virtual Disk Group. Use the simulator to assign the Global Hot Spare. When you have finished using the simulator, click Next.



ANSWER: See the explanation for the answer

Explanation:

Assign a disk as the Global Hot Spare:

1. From the iDRAC menu, select **Storage**.
2. Open **Physical Disks**.
3. Identify an available/unused physical disk in **Ready** state that is compatible with the virtual disk group.
4. Select that disk and choose **Assign Global Hot Spare** from the available action menu.
5. Click **Apply** (or **Apply Now**) to submit the configuration job.
6. Confirm the disk status changes to **Global Hot Spare**, then click **Next** in the simulator.

A global hot spare is not tied to only one virtual disk; it can automatically replace a failed compatible member disk in the controller's virtual disks.

QUESTION NO: 6

Which card connects to the PCI bus, is physically smaller than a on the system board of a Dell EMC PowerEdge 14G server?

- A. NDC
- B. OCP
- C. PCIe
- D. FPGA

ANSWER: A

Explanation:

NDC is correct. A Network Daughter Card is a purpose-built network adapter that installs in a dedicated connector on the PowerEdge system board. In 14th-generation PowerEdge systems, the NDC interface is connected through PCIe lanes, so it provides PCIe-based network connectivity without consuming one of the server's standard, full-size PCIe expansion slots. Its compact daughter-card form factor is physically smaller than a conventional PCIe expansion card and is designed to integrate closely with the motherboard and chassis. Depending on the server configuration, an NDC can provide multiple Ethernet ports and may support different network speeds or media types. This design gives PowerEdge systems flexible integrated networking while preserving PCIe slots for devices such as storage controllers, accelerators, or additional adapters. Dell's PowerEdge R740 documentation identifies the NDC as a separately installed system-board component and provides procedures for its removal and installation. Dell also describes network daughter cards as supported network-adapter options in the platform technical documentation. See the [Dell PowerEdge R740 Owner's Manual](#) and the [Dell PowerEdge R740 Technical Specifications](#).

QUESTION NO: 7

Which mobile devices are supported by Quick Sync 2.0 in 14G PowerEdge servers?

- A. Android devices only
- B. iOS, Android, and Windows Phone devices
- C. iOS and Android devices only
- D. iOS devices only

ANSWER: C

Explanation:

iOS and Android devices only is correct. Dell EMC Quick Sync 2.0 on 14th-generation PowerEdge servers is designed for local, at-the-server management using a supported mobile device running the Dell OpenManage Mobile application. The feature uses the Quick Sync 2 bezel/module to establish a wireless connection that lets an administrator view server inventory and health information, access logs, and perform supported management tasks without first connecting through the normal network-management path. Dell's 14G PowerEdge documentation identifies the OpenManage Mobile client platforms

used with Quick Sync 2.0 as Apple iOS and Google Android. This device support is significant during installation, troubleshooting, or datacenter walk-up administration, where a technician can securely interact with the server from a phone or tablet. For product documentation and supported-management information, see Dell's [PowerEdge R740 manuals](#) and the [Dell OpenManage Mobile support information](#).

QUESTION NO: 8

A technician is sent on-site to troubleshoot a server that produced an error message during the most recent POST, but then completed the POST and booted to the operating system. The user cannot remember the specific error and no entry was made in the error log relating to this error.

How can the technician find out what the error message was?

- A. Access the Lifecycle Controller and view the last boot cycle video capture entry
- B. Reload the BIOS firmware
- C. Export the Server profile and analyze for any POST errors
- D. Access the iDRAC and view the last boot cycle video capture entry

ANSWER: D

Explanation:

Access the iDRAC and view the last boot cycle video capture entry is correct because iDRAC provides the Boot Capture capability, which records the server console output during the most recent boot process. This capture preserves visual POST output, including messages displayed before the operating system starts, allowing a technician to review an error that was seen on the physical console but was not retained in the system event log. The technician can connect to the iDRAC web interface, open the server's console or boot-capture area, and replay or view the captured boot sequence to identify the exact POST message. This is particularly useful for transient initialization warnings and BIOS/firmware messages that do not necessarily create a Lifecycle Log or System Event Log entry. Because the server completed POST and booted normally, reviewing the stored capture is non-disruptive: it does not require restarting the server or reproducing the condition. Dell documents Boot Capture as an iDRAC feature for capturing and reviewing video output from server boot activity; see the [Dell iDRAC9 User's Guide: Boot Capture](#) and the [Dell Integrated Dell Remote Access Controller overview](#).