

DUMPS ARENA

Dell PowerScale Maintenance Exam

EMC D-PSC-MN-01

Version Demo

Total Demo Questions: 11

Total Premium Questions: 110

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, PowerScale Hardware	80
Topic 2, PowerScale Software	22
Topic 3, PowerScale Troubleshooting	8
Total	110

QUESTION NO: 1

What is a function of the IDSDM in Dell EMC PowerScale F600 nodes?

- A. Primary backup for PSI information
- B. With two micro SD cards, it is a redundant backup for PSI information and temporary storage for FW upgrade binaries.
- C. With one micro SD card, it is a read-only backup for PSI information. With two micro SD cards, it is temporary storage for FW upgrade binaries.
- D. Secondary backup for PSI information.

ANSWER: D

Explanation:

Secondary backup for PSI information is correct. The Internal Dual SD Module (IDSDM) in a PowerScale F600 node provides protected, removable storage for a backup copy of Platform-Specific Information (PSI). PSI contains the platform identity and configuration information required for the node to be correctly recognized and operated within the PowerScale environment. Maintaining this copy separately from the primary platform storage provides a recovery path when the primary PSI information is unavailable or requires restoration during hardware service.

The IDSDM design incorporates microSD media to provide reliable persistence for this backup function. During maintenance, this component should be handled as a system-configuration device: follow ESD procedures, retain the existing media when directed by Dell service procedures, and avoid replacing or removing it casually because its contents are associated with the node platform identity. This role supports serviceability and recovery of the node rather than general-purpose operating-system, user-data, or firmware-binary storage. Dell's PowerScale documentation portal provides the applicable hardware installation and service documentation for the F600 platform, while Dell's product page identifies the F600 as part of the PowerScale all-flash node family: [Dell PowerScale F600 documentation](#) and [Dell PowerScale F600 product information](#).

QUESTION NO: 2

A customer plans a OneFS software upgrade during a maintenance window. The customer requires client file service to remain available throughout the activity.

Which two actions should be completed before starting the upgrade? Select two.

- A. Verify that the cluster is healthy
- B. SmartFail all healthy nodes before the upgrade
- C. Run the applicable pre-upgrade checks
- D. Reformat each node before applying the upgrade
- E. Disable SmartConnect for the duration of the upgrade

ANSWER: A C

Explanation:

Verify that the cluster is healthy and **run the applicable pre-upgrade checks** are correct. A rolling or parallel upgrade relies on sufficient healthy cluster resources to continue serving clients while nodes transition through the upgrade. Pre-upgrade checks identify conditions that can prevent a supported upgrade or cause the upgrade to stop.

SmartFailing healthy nodes reduces resiliency and available resources. Reformatting nodes destroys their local configuration and is not an upgrade preparation activity. Disabling SmartConnect is also unsuitable because clients require the supported connection mechanism to reconnect as nodes are upgraded.

QUESTION NO: 3

A field engineer is preparing to replace a failed PowerScale hardware component and needs Dell-approved guidance before beginning work.

Which two activities can Dell SolVe Online provide for this maintenance task? Select two.

- A. Provide step-by-step service procedures for the applicable platform.
- B. Provide validated solution diagrams and topology information.
- C. Automatically install the latest OneFS release on the cluster.
- D. Replace the requirement to verify cluster health after service.
- E. Collect and submit all cluster logs without administrative action.

ANSWER: A B

Explanation:

SolVe Online provides current, guided, model-specific service procedures, including the preparation, replacement, and verification activities required for supported hardware work. It also provides solution diagrams and topology information that can help the engineer validate physical connectivity. SolVe is not a replacement for OneFS software media management or cluster log collection tools.

QUESTION NO: 4

Which two backend switches support 100 GbE?

- A. DCS-7308
- B. S5232-ON
- C. Z9264F-ON
- D. D4040

ANSWER: B C

Explanation:

S5232-ON and Z9264F-ON support 100 GbE and are appropriate high-bandwidth switch platforms for PowerScale backend networking. The S5232-ON platform provides thirty-two QSFP28 interfaces capable of operating at 100 GbE, with flexibility for lower-speed connectivity through breakout configurations where required. This makes it suitable for backend designs that need substantial node-to-node bandwidth while retaining a compact switch footprint.

The Z9264F-ON platform extends this capability with sixty-four QSFP28 interfaces supporting 100 GbE. Its higher port density is useful where a larger PowerScale cluster or a more highly consolidated backend fabric requires many 100 GbE connections. In both cases, the QSFP28 100 GbE interface capability is the key characteristic that satisfies the question's backend-switch requirement. Dell's product-support documentation for these platforms provides the applicable installation, hardware, and technical-reference materials: [Dell PowerSwitch S5232F-ON documentation](#) and [Dell PowerSwitch Z9264F-ON documentation](#).

QUESTION NO: 5

A customer plans to upgrade OneFS during an approved maintenance window. Before starting the upgrade, the engineer finds that the cluster is still restriping after a previous hardware event and that available disk space does not meet the upgrade requirement.

Which two conditions must be remediated before the upgrade is started? Select two.

- A. The active restriping activity
- B. The insufficient upgrade disk space
- C. The use of SmartConnect for client access
- D. The presence of a separate management network
- E. The node bezel installation

ANSWER: A B

Explanation:

The cluster should not be upgraded while it is degraded or actively restriping because protection and normal cluster operation have not yet returned to a stable state. The required upgrade disk space must also be available before the OneFS image is installed. Resolving these conditions first reduces upgrade risk and supports the selected rolling or parallel upgrade method.

QUESTION NO: 6

What is the isi diagnostics gather command used for?

- A. Diagnose the CPU.
- B. Diagnose the TPM
- C. Gather performance data.
- D. Gather cluster logs

ANSWER: D

Explanation:

Gather cluster logs is correct because `isi diagnostics gather` creates a diagnostic gather containing troubleshooting information from the PowerScale cluster. The command is intended to collect the logs, configuration details, and other system diagnostic artifacts that administrators or Dell Support need to investigate a cluster condition. A gather is especially useful because the relevant evidence for an issue can be distributed across nodes; collecting it through the OneFS diagnostic utility provides a consolidated support bundle rather than requiring individual logs to be manually retrieved from each node.

Administrators commonly run this command when directed by Dell Support or while preparing information for a service request. The resulting gathered data can be reviewed locally or uploaded through the appropriate support workflow, allowing support personnel to correlate cluster events, services, configuration state, and node-level records during diagnosis. In this context, "cluster logs" accurately represents the command's purpose: it is a broad diagnostic log collection operation, not a targeted hardware diagnostic for a single component. Dell's PowerScale documentation resources are available through the [Dell PowerScale OneFS documentation portal](#), and Dell's [Knowledge Base](#) provides support procedures and troubleshooting guidance for PowerScale environments.

QUESTION NO: 7

When a platform engineer disconnects and removes the power supply from a Dell EMC PowerScale Gen 6 node, they see an LED on the inside of the compute module.

What does the LED indicate?

- A. Power supply was not removed correctly and the node is shutting down.
- B. Platform engineer has exactly two minutes to swap the power supply.

C. Visible LED is from the front panel of the node.

D. Pair node is supplying redundant power.

ANSWER: D

Explanation:

Pair node is supplying redundant power. In PowerScale Gen 6 modular platforms, the compute modules are installed as a paired configuration within the chassis, and the power design provides redundancy across that pair. After one power supply is disconnected and removed, the illuminated internal LED is a service indicator that voltage remains present because the companion compute module's power supply is still feeding the shared power path. This indication is important during maintenance: removal of one supply does not necessarily make the compute module electrically unpowered. The engineer should therefore continue to follow the documented field-replacement procedure and treat the module appropriately while redundant power remains active. The LED confirms that the platform's redundant-power design is sustaining operation during the power-supply replacement, rather than signaling a required immediate shutdown. Dell's PowerScale documentation library provides the platform-specific installation and replacement manuals, including hardware service guidance and safety procedures. See the [Dell PowerScale OneFS documentation page](#) and Dell's [PowerScale Hardware Information Guide](#) for the applicable Gen 6 hardware documentation and component-handling instructions.

QUESTION NO: 8 - (SIMULATION)

What Is the correct order of steps to Install a PowerScale Gen6 in a rack?

Steps

- Carefully push the chassis off the lift arms and into the rack.
- Release the lift casters and carefully slide the chassis into the cabinet as far as the lift allows.
- Slide the first few inches of the back of the chassis onto the supporting ledge of the rails.
- Install two mounting screws at the top and bottom of each rail to secure the chassis to the rack.
- Align the chassis with the rails that are attached to the rack.

Correct order

The diagram shows the correct order of steps indicated by arrows. The first arrow points right, corresponding to the first step. The second arrow points left, corresponding to the second step. The third arrow points right, corresponding to the third step. The fourth arrow points right, corresponding to the fourth step. The fifth arrow points right, corresponding to the fifth step.

ANSWER: See the explanation for the answer

Explanation:

Correct installation order:

This sequence keeps the chassis supported by the lift while it is initially seated on the rail ledges, then secures it only after it is fully positioned in the rack.

QUESTION NO: 9

A platform engineer is planning the rack deployment for Dell PowerScale A3000 nodes. The rack must support the high-density platform according to the supported Dell rack design.

Which rack solution should be used?

A. Titan-D

B. Titan-P

C. Titan-HD

ANSWER: C

Explanation:

Titan-HD is the supported rack solution for A3000 nodes. The high-density A3000 platform requires a rack design that accommodates its chassis, rail, weight, power, cooling, and service-clearance requirements. Titan-D and Titan-P are not the specified rack solution for this platform.

QUESTION NO: 10

A platform engineer is preparing the back-end network for a 12-node F200 cluster that uses 100 GbE switch ports and 25 GbE node-facing connections.

Which two configuration details are required for the breakout-cable design? Select two.

- A. Use a 1:2 breakout cable
- B. Use a 1:4 breakout cable
- C. Configure the 100 GbE switch port for breakout operation
- D. Use a 1:8 breakout cable
- E. Connect the back-end interfaces to the management network

ANSWER: B C**Explanation:**

Use a 1:4 breakout cable and **configure the 100 GbE switch port for breakout operation** are correct. One 100 GbE QSFP28 interface can be divided into four 25 GbE lanes, providing the required node-facing connections. The switch interface must be configured to recognize the breakout arrangement.

A 1:2 or 1:8 cable does not provide the supported four 25 GbE lanes from the 100 GbE port. Connecting the back-end interfaces to the management network would prevent the intended private cluster interconnect design.

QUESTION NO: 11

Which slot on the R640 hardware is used for front-end connectivity on a Dell EMC PowerScale F600 node?

- A. PCIe Slot 3
- B. rNDC Slot
- C. PCIe Slot 1
- D. PCIe Slot 2

ANSWER: B**Explanation:**

The **rNDC Slot** is used for front-end connectivity on a PowerScale F600 node. The F600 is built on the Dell PowerEdge R640 platform, whose rack Network Daughter Card architecture provides a dedicated, integrated location for the front-end Ethernet adapter. This design supplies the client-facing network interfaces through which SMB, NFS, S3, and other client traffic reaches the cluster. Using the rNDC location for this purpose also preserves the standard PCIe expansion capacity for the node components and configurations that require those slots. When servicing or verifying an F600 node, identify the installed rNDC and its port type, then confirm that its cabling is connected to the intended client-access network and that the configured front-end interfaces match the cluster network design. The physical slot assignment is important because node hardware layouts and adapter locations vary by PowerScale node model; service personnel should use the model-specific hardware documentation rather than applying a layout from a different platform. Dell's PowerScale documentation resources are available through the [PowerScale Info Hubs](#), and product documentation can be accessed from the [PowerScale F600 support documentation page](#).