[Feb-2025 D-PSC-MN-01 Dumps PDF - D-PSC-MN-01 Real Exam Questions Answers [Q85-Q99

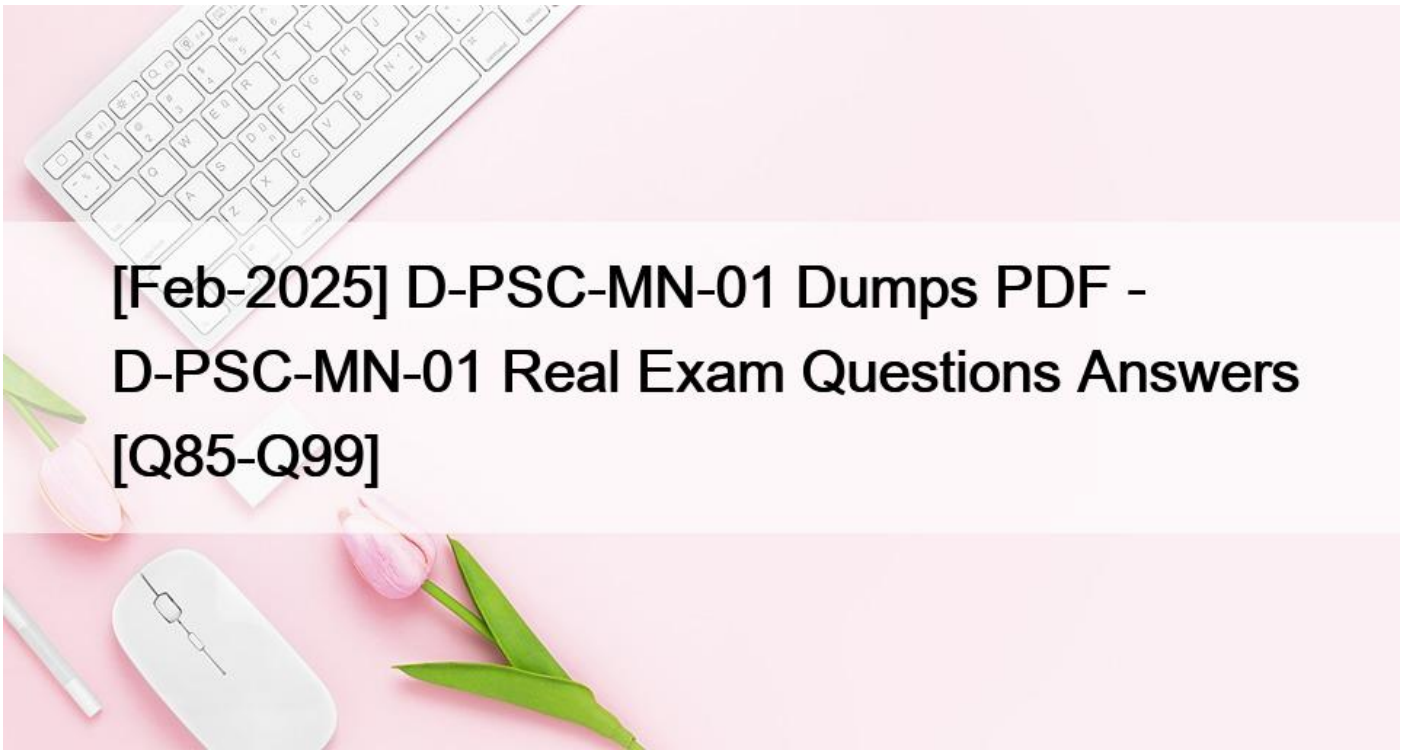

[Feb-2025] D-PSC-MN-01 Dumps PDF - D-PSC-MN-01 Real Exam Questions Answers
D-PSC-MN-01 Dumps 100% Pass Guarantee With Latest Demo

EMC D-PSC-MN-01 Exam Syllabus Topics:

TopicDetailsTopic 1- PowerScale Implementation: Geared towards cluster administrators, this section covers the practical aspects of implementing PowerScale systems. It focuses on creating new clusters, integrating nodes into existing clusters, and applying alternative configurations or accessories to improve performance.Topic 2- PowerScale Hardware Installation: This section evaluates the installation capabilities of hardware engineers to test candidates on their ability to plan and execute the installation of PowerScale hardware, prepare the maintenance site, and adhere to best practices when installing racks, nodes, and related components.Topic 3- PowerScale Hardware Concepts: This part of the exam assesses the expertise of IT infrastructure architects, and emphasizes the ability to differentiate between the architectural elements of PowerScale platforms, understand their specific use cases, and compare design specifications across PowerScale hardware models.Topic 4 - PowerScale Hardware Maintenance: Targeted at system administrators, this section focuses on key tasks involved in maintaining PowerScale hardware. It examines candidates' proficiency in completing initial maintenance steps, preparing nodes for operation, and processing Field Replaceable Unit (FRU) and Customer Replaceable Unit (CRU) requests.

NEW QUESTION 85

How do external networking requirements for a PowerScale cluster typically differ from internal requirements?

* Higher bandwidth and lower latency

* Increased security and access controls
* Use of public IP addresses
* Simpler routing and switching protocols

NEW QUESTION 86

During a PowerScale cluster upgrade, what is the importance of managing versioning effectively?
* To minimize the storage footprint of the upgrade files
* To ensure that all nodes operate with compatible software versions
* To facilitate rollback in case the upgrade fails
* To streamline the upgrade process for remote nodes

NEW QUESTION 87

What is a critical consideration when handling upgrades and patching in a PowerScale environment?
* Scheduling downtime for each node individually
* Ensuring all nodes are upgraded simultaneously to prevent version mismatch
* Testing the upgrades in a staging environment before deployment
* Upgrading the node with the least amount of data first

NEW QUESTION 88

A platform engineer is tasked with adding F600 nodes to an existing Dell EMC PowerScale cluster. After racking and stacking the F600 nodes, they determine that the cluster contains X210 and H400 nodes.

What should the platform engineer consider?
* H400 nodes cannot co-exist with the F600 nodes in the same cluster
* The F600 must use InfiniBand for internal network communication
* F600 nodes cannot co-exist with X210 nodes in the same cluster
* The cluster must be upgraded to Ethernet back-end topology

Adding F600 nodes to an existing cluster requires compatibility in back-end networking.

Key Considerations:

* Back-End Networking:

* F600 Nodes:Use Ethernet for internal communication.

* Existing Nodes (X210 and H400):May be using InfiniBand.

* Action Required:

* Upgrade the cluster’s back-end to Ethernet topology to accommodate F600 nodes.

Why Other Options Are Incorrect:

* Option A:H400 nodes can coexist with F600 nodes if back-end networking is compatible.

* Option B:F600 nodes do not use InfiniBand.

* Option C:F600 nodes can coexist with X210 nodes with the appropriate back-end network.

Dell PowerScale References:

* Dell EMC PowerScale Networking Guide:

* Back-End Network Compatibility:

* Discusses requirements for mixing node types.

* Upgrading Back-End Network:

* Provides steps for transitioning from InfiniBand to Ethernet.

* Best Practices:

* Plan the network upgrade carefully to minimize downtime.

* Consult with Dell EMC support for guidance.

NEW QUESTION 89

What is done with the components when a Gen6 single node is replaced?
* Transferred to the new node
* Returned to Dell
* Sent to WWFA

When a Dell PowerScale Gen6 single node is replaced, the standard procedure is to return the entire faulty node, including all its components, back to Dell. This ensures proper handling, compliance with warranty agreements, and allows Dell to perform failure analysis if necessary.

* Node Replacement Process:

* Faulty Node Identification:A node exhibiting issues is identified for replacement.

* Data Protection:Before replacement, data is protected via OneFS, which ensures data is redistributed across the cluster to prevent data loss.

* Replacement Node Shipment:Dell ships a replacement node to the customer.

* Handling of Components:

* No Component Swapping:Components such as drives, memory modules, and CPUs are not transferred from the old node to the new one.

* Return Procedure:The entire faulty node, with all its components intact, is returned to Dell.

* Purpose of Return:Returning the node allows Dell to:

* Perform Diagnostics:Analyze the faulty components for failures.

* Manage Inventory:Ensure proper accounting of hardware.

* Environmental Compliance:Dispose of or recycle components according to regulations.

* Dell’s Return Material Authorization (RMA) Policy:

* RMA Process:Dell issues an RMA for the faulty node, and the customer is responsible for returning it.

* Shipping Instructions:Detailed instructions are provided to safely package and ship the node back to Dell.

Dell PowerScale References:

* Dell EMC PowerScale Field Replacement Unit (FRU) Installation and Replacement Guide:

* Outlines the procedures for replacing nodes and the requirement to return faulty units to Dell.

* Dell EMC PowerScale OneFS Administration Guide:

* Provides information on cluster maintenance and node management.

* Warranty and Support Agreements:

* Specify the obligations for returning faulty hardware under service contracts.

NEW QUESTION 90

Which two backend switches support 100 GbE?
* DCS-7308
* S5232-ON
* 29264-ON
* D4040
The two backend switches that support 100 GbE are:

* B. S5232-ON

* C. Z9264-ON

* Dell EMC Networking S5232-ON:

* Features:

* A 1U high-density switch with 32 ports of 100 GbE QSFP28.

* Supports 10/25/40/50/100 GbE speeds through breakout cables.

* Use Cases:

* Ideal for high-performance backend networks in PowerScale clusters.

* Dell EMC Networking Z9264-ON:

* Features:

* A 2U switch offering 64 ports of 100 GbE QSFP28.

* Provides extensive scalability for large network deployments.

* Use Cases:

* Suitable for large-scale PowerScale clusters requiring extensive bandwidth.

* Why Options A and D Are Incorrect:

* A. DCS-7308:

* The DCS-7300 series is modular and may support 40/100 GbE, but it’s not commonly associated with PowerScale backend networks.

* Not a standard recommendation for PowerScale deployments.

* D. D4040:

* The Dell Networking D4040 is a 40 GbE switch, not supporting 100 GbE natively.

* Does not meet the requirement for 100 GbE support.

Dell PowerScale References:

* Dell EMC PowerScale Network Design Considerations:

* Discusses recommended switches for backend networking, including 100 GbE options.

* Switch Specification Sheets:

* S5232-ON Data Sheet:

* Details port configurations and capabilities.

* Z9264-ON Data Sheet:

* Provides specifications and deployment scenarios.

* Dell EMC Best Practices Guide:

* Recommends network hardware compatible with PowerScale clusters.

NEW QUESTION 91

An engineer runs ini_reformat_node command.

What are they attempting to do?
* Reformat the mirrored FEC data.
* Reformat a node quickly to repurpose a node.
* Reformat the mirrored journals.
* Reformat a node to securely erase all data.

The isi_reformat_node command is a utility used on Dell PowerScale (Isilon) clusters toreformat a node and securely erase all dataon it. This command initializes the node’s storage media, effectively wiping all user data, metadata, and system configurations from the node’s drives.

Purpose of isi_reformat_node:

* Secure Data Erasure:It ensures that all data is securely erased, which is essential when decommissioning a node or repurposing it for a different use.

* Node Recovery or Repurposing:It prepares the node for re-integration into the cluster or for use in a different cluster by resetting it to a factory-like state.

Usage Scenarios:

* Decommissioning a Node:When permanently removing a node from a cluster and ensuring no residual data remains.

* Repurposing Hardware:When reassigning the node to a different cluster or role and needing to eliminate all previous configurations and data.

* Recovering from Corruption:In cases where the node’s data is irreparably corrupted, reformatting allows for a clean start.

Key Points:

* Data Loss Warning:Running isi_reformat_node will result in complete data loss on that node. It’s crucial to ensure that the data is backed up or that the node’s data is no longer needed.

* Cluster Impact:Before reformatting, the node should be appropriately prepared, and the cluster should be informed to avoid any data protection issues.

* Secure Erasure Standards:The command follows secure erasure standards to prevent data recovery through forensic methods.

References:

* Dell PowerScale OneFS CLI Administration Guide- Details on using isi_reformat_node and its implications.

* Dell PowerScale OneFS Administration Guide- Procedures for safely removing and reformatting nodes.

* Dell Knowledge Base Article- Best practices for decommissioning and reformatting nodes in a PowerScale cluster.

NEW QUESTION 92

An engineer wants to create a 4-node cluster after rack and stack.

What port must they use to start the installation?
* COM

* FE
* BE
* iDRAC

When an engineer wants to create a 4-node cluster after rack and stack, they must use theiDRAC (Integrated Dell Remote Access Controller)port to start the installation.

* Understanding iDRAC:

* Remote Management:

* iDRAC provides out-of-band management capabilities, allowing administrators to manage and monitor nodes remotely.

* Virtual Console Access:

* Offers a virtual console for accessing the node’s BIOS and performing installations.

* Initial Cluster Setup with iDRAC:

* Accessing the Node:

* Connect to each node’s iDRAC interface using its IP address.

* Starting the Installation:

* Use the iDRAC virtual console to interact with the node as if you were physically present.

* Benefits:

* No need for physical access to the COM port or direct console connections.

* Allows for remote configuration and reduces the time required for setup.

* Why iDRAC is Used Over Other Ports:

* COM Port:

* The COM (serial) port can be used for initial setup, but it requires physical access and serial cables.

* Less convenient compared to iDRAC’s remote capabilities.

* FE (Front-End) and BE (Back-End) Ports:

* FE and BE ports are used for data network connections, not for initial setup or management.

* The initial configuration cannot be performed through these ports.

* Procedure for Using iDRAC:

* Step 1: Connect to iDRAC

* Ensure that iDRAC network interfaces are connected and configured with IP addresses.

* Access iDRAC via a web browser using the IP address.

* Step 2: Launch Virtual Console

* Log in to the iDRAC interface.

* Launch the virtual console to access the node’s system interface.

* Step 3: Perform Initial Configuration

* Use the virtual console to run the OneFS installation wizard.

* Configure cluster settings, networking, and other parameters.

* Dell PowerScale Best Practices:

* Using iDRAC for Installation:

* Recommended for its convenience and efficiency.

* Allows for consistent setup procedures across multiple nodes.

* Network Preparation:

* Ensure iDRAC interfaces are properly connected to the management network.

* Verify network settings to allow access from the engineer’s workstation.

Dell PowerScale References:

* Dell EMC PowerScale OneFS Installation Guide:

* Provides detailed steps on installing OneFS using iDRAC.

* Dell EMC PowerScale Networking Guide:

* Discusses management network configurations, including iDRAC setup.

* Dell EMC iDRAC User’s Guide:

* Offers comprehensive information on using iDRAC features.

* Knowledge Base Articles:

* Article ID 000125678:”Initial Cluster Configuration Using iDRAC”

* Article ID 000125679:”Best Practices for Remote Installation on PowerScale Nodes”

NEW QUESTION 93

Which cluster management tool provides a graphical user interface for PowerScale clusters?
* CLI
* PAPI
* Serial Console
* WebUI

NEW QUESTION 94

When testing connectivity to the Dell EMC PowerScale cluster through the Web Administration interface, what should be typed in the address line of the browser?
* https://<clustername>:8080
* http://<clustername>:80
* http://<clustername>:8082
* https://<clustername>:8081
To access the Dell EMC PowerScale cluster through the Web Administration interface, you need to use a web browser to connect to the cluster’s management port using HTTPS on the appropriate port.

Correct URL Format:

* https://<clustername>:8080

* Protocol:HTTPS (Secure HTTP)

* Cluster Name/IP:Replace <clustername> with the cluster’s hostname or IP address.

* Port Number:8080, which is the default port for the PowerScale Web Administration interface.

Why Option A is Correct:

* HTTPS Protocol:

* Ensures secure communication with the cluster’s web interface.

* Port 8080:

* Default port for accessing the Web Administration interface on PowerScale clusters.

* Combining Both:

* Using https://<clustername>:8080 directs the browser to the correct interface.

Why Other Options Are Incorrect:

* Option B (http://<clustername>:80):

* Uses HTTP on port 80, which is not the default for the administration interface.

* Option C (http://<clustername>:8082):

* Incorrect protocol and port.

* Option D (https://<clustername>:8081):

* Port 8081 is not the default for the web interface.

Dell PowerScale References:

* Dell EMC PowerScale OneFS Web Administration Guide:

* Accessing the Web Interface:

* Specifies that the Web Administration interface is accessible via HTTPS on port 8080.

* Provides the URL format for connecting to the cluster.

* Example:

* If your cluster’s hostname is powerscale-cluster, you would enter:

arduino

Copy code

https://powerscale-cluster:8080

* Security Note:

* Certificate Warnings:

* Browsers may display a security warning due to self-signed certificates.

* You can proceed by adding an exception or install a valid SSL certificate.

* Best Practices:

* DNS Configuration:

* Ensure that the cluster’s hostname is resolvable via DNS.

* Alternatively, use the cluster’s IP address in the URL.

* Browser Compatibility:

* Use a modern browser that supports the necessary encryption protocols.

* Firewall Settings:

* Verify that network firewalls allow traffic on port 8080.

NEW QUESTION 95

A company must ensure their PowerScale cluster can handle many active client connections. What must they do when designing their system?
* Include a Leaf-Spine backend network
* Use the P100 node.
* Add maximum RAM in each node.
* Add L3 cache to the nodes.

To ensure a Dell PowerScale cluster can handle many active client connections, the company shouldinclude a Leaf-Spine backend networkin their system design.

* Understanding Network Topologies:

* Leaf-Spine Architecture:

* A high-performance network topology designed to handle large amounts of east-west (node-to-node) traffic.

* Consists of two network layers: leaf switches (access layer) and spine switches (aggregation layer).

* Every leaf switch connects to every spine switch, providing multiple pathways and reducing bottlenecks.

* Benefits for PowerScale Clusters:

* Scalability:

* Supports a large number of nodes and client connections without significant degradation in performance.

* Low Latency:

* Reduces hop count between any two endpoints, minimizing latency.

* High Throughput:

* Provides increased bandwidth to accommodate many active connections.

* Redundancy:

* Multiple pathways between nodes enhance fault tolerance.

* Handling Many Active Client Connections:

* Network Bandwidth:

* A Leaf-Spine network ensures sufficient bandwidth is available for client connections and data movement.

* Load Balancing:

* Distributes client connections evenly across the network to prevent overloading any single path.

* Reduced Contention:

* Minimizes network congestion, leading to improved client experience.

* Why Other Options Are Less Suitable:

* Option B (Use the P100 node):

* P100 nodes are accelerator nodes that enhance performance but do not specifically address handling many client connections.

* Option C (Add maximum RAM in each node):

* While increasing RAM can improve performance, it does not directly impact the cluster’s ability to handle numerous client connections.

* Option D (Add L3 cache to the nodes):

* Adding L3 cache improves data retrieval speeds but does not significantly affect network capacity for client connections.

Dell PowerScale References:

* Dell EMC PowerScale Network Design Considerations:

* Discusses network topologies and their impact on cluster performance.

* Dell EMC PowerScale Network Design Considerations

* Dell EMC PowerScale Best Practices Guide:

* Recommends network architectures for optimal performance.

* Dell EMC PowerScale Best Practices

* Knowledge Base Articles:

* Article ID 000123002:”Implementing Leaf-Spine Architecture for PowerScale Clusters”

* Article ID 000123003:”Scaling Client Connections in Dell PowerScale Environments”

NEW QUESTION 96

A platform engineer cables, powers up, and creates a new cluster of F600 nodes. While performing verification checks, no clients can access the cluster and the engineer cannot ping the nodes.

What is the probable issue?
* Manual reboot was not done following the cluster join.
* Network was connected to the F600 rNDC ports.
* Nodes were connected to the management ports.
* Internal network was cabled improperly.
When setting up a new Dell EMC PowerScale F600 cluster, proper cabling is crucial for network connectivity and cluster

functionality. The F600 nodes have multiple network ports:

* Front-End (Client) Ports: For client network access.

* Back-End (Internal) Ports: For node-to-node communication within the cluster.

* Management Ports: Dedicated for node management and should not be used for client traffic.

Probable Issue:

* Incorrect Cabling to Management Ports: If the nodes are connected to the management ports instead of the front-end network ports, clients will not be able to access the cluster, and the nodes will not respond to pings on the client network.

Key Points:

* Management Ports Usage: Reserved for administrative tasks like cluster management and should not carry client data.

* Impact of Misconfiguration: Connecting client networks to management ports leads to inaccessible services for clients.

Dell PowerScale References:

* Dell EMC PowerScale Networking Configuration Guide:

* Chapter on Network Connectivityemphasizes the correct usage of network ports.

* Illustrations showing proper cabling of front-end and management ports.

* Installation Best Practices:

* Verify all network connections according to the hardware installation guide.

* Use the correct network ports for client and management traffic.

NEW QUESTION 97

What type of privileges are required to perform tenant-specific administration?
* Zone role-based access control
* System root privileges
* System admin privileges
* Role-based access control

Tenant-specific administration involves managing specific zones or tenants within a Dell EMC PowerScale environment. This requires a certain level of privileges that allows administrators to control resources without having full system-wide access.

Zone Role-Based Access Control:

* Definition:

* Zone role-based access control allows administrators to assign permissions to users or groups for specific zones or tenants.

* Functionality:

* Enables delegation of administrative tasks within a particular zone without granting system-wide privileges.

* Benefits:

* Enhances security by limiting access.

* Simplifies management in multi-tenant environments.

Why Other Options Are Incorrect:

* System Root Privileges (Option B):

* Root privileges grant full access to the entire system, which is not necessary for tenant-specific tasks.

* System Admin Privileges (Option C):

* Similar to root, system admin privileges are broader than needed.

* Role-Based Access Control (Option D):

* While general role-based access control is related,zonerole-based access control is specifically designed for tenant administration.

Dell PowerScale References:

* Dell EMC PowerScale OneFS Administration Guide:

* Chapter on Access Control and Authentication:

* Details how to implement zone role-based access control.

* Provides instructions on assigning roles to users within specific zones.

* Best Practices:

* Regularly review and update access permissions.

* Use the principle of least privilege to enhance security.

NEW QUESTION 98

When setting up internal networking for a PowerScale cluster, what is a crucial requirement?
* Public IP addresses for all nodes
* Private interconnects between nodes
* Unrestricted internet access
* Serial connection between nodes

NEW QUESTION 99

Which Dell EMC PowerScale Gen 6 components are considered CRUs?

* SSD, power supply, fan, and bezel
* Back-end switch, power supply, disk drive, and bezel
* SSD, power supply, disk drive, and bezel
* SSD, memory card, disk drive, and bezel

In Dell EMC PowerScale Gen 6 systems, components are categorized as either Customer Replaceable Units (CRUs) or Field Replaceable Units (FRUs). CRUs are parts that customers are authorized to replace themselves without the need for a Dell EMC service technician. Understanding which components are CRUs is essential for efficient maintenance and minimal downtime.

CRUs in PowerScale Gen 6:

* SSD (Solid-State Drive) and Disk Drive:

* Both SSDs and HDDs are considered CRUs.

* They can be hot-swapped without powering down the node.

* Customers can replace failed drives to restore full functionality.

* Power Supply:

* Power supplies are hot-swappable and designated as CRUs.

* Nodes typically have redundant power supplies; replacing one does not require shutting down the node.

* Bezel:

* The front bezel is considered a CRU.

* It can be removed and replaced without tools.

* Serves both aesthetic and airflow management purposes.

Why Option C is Correct:

* SSD:As a storage component, the SSD is a CRU.

* Power Supply:Designated as a CRU for quick replacement.

* Disk Drive:Refers to HDDs, also CRUs.

* Bezel:Easily removable by the customer.

Why Other Options are Incorrect:

* Option A:Includesfan, which is not typically a CRU in Gen 6 nodes; fans are usually considered FRUs due to the complexity of replacement and potential risks.

* Option B:Listsback-end switch, which is a complex component requiring a service technician (FRU).

* Option D:Includesmemory card(DIMM), which is generally not a CRU because replacing memory modules requires technical

expertise and poses risks if not done correctly.

Dell PowerScale References:

* Dell EMC PowerScale Hardware Replacement Guide:

* Chapter on Customer Replaceable Units (CRUs):

* Lists the components that customers can replace.

* SSDs, HDDs, power supplies, and bezels are specified as CRUs.

* Safety Instructions:

* Provides guidelines on how to safely replace CRUs.

* Emphasizes the use of ESD (Electrostatic Discharge) precautions.

* Dell EMC PowerScale OneFS Administration Guide:

* Disk Management Section:

* Explains how to handle drive replacements.

* Details the hot-swapping procedure for drives.

* Best Practices:

* Before Replacing a CRU:

* Verify the component has indeed failed.

* Ensure you have the correct replacement part.

* During Replacement:

* Follow proper shutdown procedures if required (though CRUs are typically hot-swappable).

* Use appropriate ESD protection.

* After Replacement:

* Confirm that the system recognizes the new component.

* Monitor the system for any errors or warnings.

Dumps Real EMC D-PSC-MN-01 Exam Questions [Updated 2025:
https://www.dumpsmaterials.com/D-PSC-MN-01-real-torrent.html]