

DUMPSBOSS.

Dell VPLEX Deploy Achievement

EMC D-VPX-DY-A-24

Version Demo

Total Demo Questions: 8

Total Premium Questions: 80

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

When are the front-end ports enabled during a VPLEX installation?

- A. Before launching the VPLEX EZ-Setup wizard
- B. Before creating the metadata volumes and backup
- C. After exposing the storage to the hosts
- D. After creating the metadata volumes and backup

ANSWER: D

Explanation:

After creating the metadata volumes and backup is correct because VPLEX deployment protects its configuration state before enabling host-facing access. Metadata volumes hold the distributed system's critical configuration information, including details needed to maintain and recover the VPLEX environment. Once those metadata volumes are created, a metadata backup establishes a recoverable baseline for the newly configured system.

Enabling the front-end ports at this point is an intentional sequencing control. It ensures that host connectivity is introduced only after the VPLEX configuration and its metadata are established and safeguarded. The front-end ports provide the host-facing paths through which initiators access VPLEX virtual volumes, so activating them follows the foundational configuration and backup tasks rather than preceding them. This order supports a controlled transition from initial deployment to host integration and storage presentation.

Dell documentation and support resources for VPLEX configuration, deployment procedures, and product documentation are available through the [Dell VPLEX documentation portal](#) and the [Dell Support portal](#). These resources should be consulted for the exact procedure applicable to the installed VPLEX release and topology.

QUESTION NO: 2

A VPLEX Metro cluster is being installed for a company that is planning to create distributed volumes with 200 TB of storage. Based on this requirement, and consistent with

EMC best practices, what should be the minimum size for logging volumes at each cluster?

- A. 10 GB
- B. 16.5 GB
- C. 20 GB

ANSWER: C

Explanation:

20 GB is the minimum logging-volume size required at each cluster for 200 TB of distributed-volume capacity. VPLEX Metro uses logging volumes to maintain the distributed-device write log, which protects data consistency when the inter-cluster communication path is interrupted or when recovery processing is required. EMC sizing guidance allocates logging-volume capacity according to the amount of distributed storage configured: 10 GB is required for each 100 TB of distributed-volume capacity. Therefore, a planned 200 TB distributed configuration requires 20 GB of logging capacity at each VPLEX cluster. The logging volumes must be available locally at both clusters, because each cluster must be able to record the information needed to preserve and recover distributed-device consistency independently. This capacity should be provisioned before creating the distributed devices, rather than treating the basic minimum for a small configuration as sufficient for the full planned capacity. Dell Technologies' VPLEX documentation describes distributed-device configuration and the role of logging volumes in maintaining consistency across Metro clusters; see the [Dell VPLEX product documentation](#) and the [VPLEX documentation portal](#).


QUESTION NO: 3 - (SIMULATION)

What is the correct order of steps to migrate from an old array to a new one without disruption using VPLEX?

Steps

- Connect new storage array to VPLEX
- VPLEX admin monitors progress of the migration; I/O to host continues
- VPLEX admin initiates mobility job(s)
- Once volumes on the new array are fully synchronized, the admin commits the migration
- VPLEX discovers the new array and the admin creates migration target devices

Answer area



ANSWER: Check the explanation for the answer

Explanation:

Place the migration steps in this order:

This is nondisruptive because VPLEX maintains the host-facing virtual volume while the mobility job copies and synchronizes data from the old-array device to the new migration target. Commit is performed only after synchronization completes.

QUESTION NO: 4

What are the requirements to upgrade a VPLEX from VS2 to VS6?

- A. GeoSynchrony 6.0 minimumBoth VS2 and VS6 at same code levelSame number of enginesWWN zoning
- B. GeoSynchrony 5.0 minimumBoth VS2 and VS6 at same code levelSame number of enginesWWN zoning
- C. GeoSynchrony 6.0 minimumBoth VS2 and VS6 at same code levelSame number of enginesWWN zoning and temporary FC Local COM I/O modules in VS2
- D. GeoSynchrony 6.0 minimumBoth VS2 and VS6 at same code levelSame number of enginesWWN zoning and temporary FC Local COM I/O modules in VS6

ANSWER: C

Explanation:

GeoSynchrony 6.0 minimum
Both VS2 and VS6 at same code level
Same number of engines
WWN zoning and temporary FC Local COM I/O modules in VS2

is correct because a nondisruptive VS2-to-VS6 hardware upgrade is performed as a controlled mixed-hardware transition. GeoSynchrony 6.0 or later provides the required VS6 platform support, and the participating hardware must be at the same supported GeoSynchrony code level before the hardware conversion can proceed. Maintaining the same engine count preserves the cluster's supported topology during the replacement process and is a prerequisite of the defined upgrade procedure.

The Fibre Channel zoning must be prepared so that the replacement hardware has the necessary connectivity to the storage and host fabrics. In addition, temporary FC Local COM I/O modules are installed in the VS2 hardware during the transition. These modules provide the required local communication connectivity while VS2 and VS6 hardware coexist as the upgrade is staged. This temporary component requirement is specific to the VS2 side of the conversion and is documented through Dell's controlled service procedures. The exact supported versions, cabling, zoning, and sequencing should be validated in the applicable SolVe procedure before performing the upgrade. See [Dell SolVe](#) and the [Dell VPLEX documentation portal](#).

QUESTION NO: 5

An administrator starts an expansion of a VPLEX virtual volume. The operation is rejected because a rebuild is active on the volume.

What is the appropriate next action?

- A. Disable the consistency group and restart the expansion
- B. Force the expansion while the rebuild is active
- C. Wait for the rebuild to complete, then retry the expansion
- D. Create a new metadata backup, then restart the expansion

ANSWER: C

Explanation:

The administrator should allow the rebuild to complete successfully before retrying the expansion. A rebuild is a recovery operation that must finish before VPLEX accepts a capacity change affecting the volume layout.

Forcing an expansion during rebuild activity is not an appropriate method to bypass the protection. Changing consistency-group membership or creating a metadata backup does not remove the rebuild-related restriction.

QUESTION NO: 6

Refer to the exhibit.



A storage administrator is experiencing an unplanned site outage. After recovery of the network, they are still unable to access the management server.

Which port on the back of the management server can the administrator connect to with a hardcoded IP address to access the VPLEX cluster and troubleshoot the issue?

- A. C
- B. A
- C. D
- D. B

ANSWER: B

Explanation:

The port labeled **A** is the appropriate connection for this recovery scenario. It provides the service-management access path used to reach the VPLEX management environment when normal management-network connectivity is unavailable. An administrator can directly attach a workstation to this interface and configure the workstation with the required hardcoded address in the service network range. This establishes an out-of-band path to the management server, allowing access to the VPLEX cluster for fault isolation and recovery activities.

This access method is particularly valuable after an unplanned outage because restoration of physical network connectivity alone does not guarantee that the normal management path, routing, or management services are available. Using the

designated service interface gives the administrator a controlled local connection from which to verify cluster status, collect diagnostic information, and perform the required troubleshooting steps. The port identification must be matched to the rear-panel labeling in the exhibit; the service-access interface shown there is labeled **A**. Dell's VPLEX product documentation and support resources provide the platform-specific service and management procedures: [Dell VPLEX documentation](#) and [Dell VPLEX Product Support page](#).

QUESTION NO: 7

How much cache is available in a VPLEX VS2 dual engine setup?

- A. 288 GB
- B. 144 GB
- C. 128 GB
- D. 72 GB

ANSWER: B

Explanation:

144 GB is the available cache capacity for a VPLEX VS2 dual-engine configuration. A VS2 engine contains two directors, and each VS2 director provides 36 GB of cache. Therefore, one engine has 72 GB of total director cache: 36 GB multiplied by two directors. A dual-engine configuration has two such engines, for a total of four directors. Multiplying 36 GB per director by four directors gives 144 GB of installed cache.

This calculation uses the physical VPLEX VS2 director cache specification and the standard engine composition. Understanding the count of directors is essential: the dual-engine topology doubles the director count from two to four, and therefore doubles the per-engine cache total from 72 GB to 144 GB. This cache supports VPLEX distributed-cache operations and helps maintain coherent access to storage volumes across the VPLEX directors. Dell's VPLEX support documentation is available through the [Dell VPLEX documentation portal](#); Dell also provides VPLEX product and lifecycle information through its [VPLEX family page](#).

QUESTION NO: 8

Connectivity has been restored after a WAN outage. The storage administrator of a VPLEX VS6 now wants to verify management connectivity between MMCS-A on Cluster-1

and MMCS-A on Cluster-2.

Which command can the administrator run to determine if the remote management IP is reachable?

- A. connectivity validate-wan-com
- B. 11 /engines/ ** /ports
- C. sudo /usr/sbin/ipsec statusall
- D. vpn status

ANSWER: D

Explanation:

`vpn status` is the appropriate command because VPLEX Metro management communication between the management consoles is maintained through the VPLEX VPN connection. After a WAN interruption, this command checks the operational state of that VPN relationship and reports the connectivity information needed to determine whether the peer MMCS management IP can again be reached. This makes it a direct validation step for MMCS-to-MMCS management-path recovery, rather than a check of storage-engine ports or a general configuration validation.

In a dual-cluster VPLEX deployment, confirming this management VPN status is important after restoring WAN service because the management consoles use the intercluster management network for remote administrative communication and system coordination. The command should be run from the relevant VPLEX management environment, where its output can be reviewed for the remote endpoint and the established VPN state. Dell's VPLEX documentation and product-support resources provide the applicable operational command references and platform guidance: [Dell VPLEX Series Documentation](#) and [Dell VPLEX VS6 Support Documentation](#).