

VPLEX Specialist Exam for Storage Administrators

EMC E20-562

Version Demo

Total Demo Questions: 8

Total Premium Questions: 87

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, VPLEX Architecture	31
Topic 2, VPLEX Configuration	31
Topic 3, VPLEX Administration	21
Topic 4, VPLEX Troubleshooting	4
Total	87

QUESTION NO: 1 - (DRAG DROP)

What is the correct order of steps to create a virtual volume?

Answer Area

Create device	STEP 1
Claim storage volume	STEP 2
Create virtual volume	STEP 3
Create extent	STEP 4

ANSWER:

Answer Area

Create device	Claim storage volume
Claim storage volume	Create extent
Create virtual volume	Create device
Create extent	Create virtual volume

Explanation:

Creating a VPLEX virtual volume follows the object dependency chain used by VPLEX to virtualize back-end storage. First, the administrator claims a storage volume. Claiming makes a presented back-end LUN available to VPLEX for use and establishes it as a managed storage resource. The claimed storage volume is then used to create an extent. An extent represents the portion of the claimed storage volume that VPLEX will consume; for a typical configuration, it uses the whole claimed volume.

Next, VPLEX creates a device from the extent. The device is the VPLEX storage object that provides the basis for virtualization and can represent a simple device or, in more advanced configurations, participate in more complex device structures. Once the device exists, the administrator creates the virtual volume from that device. The virtual volume is the object ultimately exported through a storage view to hosts, allowing them to access the storage by its VPLEX-provided identity.

Therefore, the required order is: Claim storage volume, Create extent, Create device, and Create virtual volume. This ordering ensures that every object is built only after its required underlying resource is available. Dell's VPLEX product documentation and support resources describe VPLEX administration and provisioning concepts at [Dell VPLEX documentation](#). Additional VPLEX product information is available from [Dell Technologies VPLEX Family](#).

QUESTION NO: 2

Which VPLEX port type provides host access to virtual volumes?

- A. Front-end port
- B. Back-end port
- C. WAN COM port
- D. Management port


ANSWER: A

Explanation:

Front-end ports provide host access to VPLEX virtual volumes. Host initiators connect through the SAN fabric to VPLEX front-end ports, and the initiators, ports, and virtual volumes are associated through a storage view. This enables VPLEX to control masking and presentation of virtual volumes to the appropriate hosts.

Back-end ports connect VPLEX to array storage, while WAN COM ports provide connectivity between clusters in VPLEX Metro configurations. The [Dell VPLEX documentation portal](#) provides additional host-connectivity and storage-view configuration information.

QUESTION NO: 3



Which protocol is presented by "B" in the exhibit?

- A. FC
- B. iSCSI
- C. FCoE
- D. IPFC

ANSWER: A

Explanation:

FC is the protocol presented at “B.” In the VPLEX connectivity model used for this exam, Fibre Channel provides the block-storage transport for the relevant SAN-facing connection. VPLEX is deployed transparently between hosts and back-end storage arrays, virtualizing storage volumes while maintaining SAN connectivity through Fibre Channel ports and fabrics. These FC connections carry SCSI block I/O and are configured through the VPLEX front-end and back-end connectivity design, including zoning and redundant fabric paths. The diagram’s indicated connection therefore represents an FC transport rather than an Ethernet-based storage protocol. This aligns with the VPLEX platform documentation and connectivity guidance available from Dell Technologies, including the [Dell VPLEX documentation page](#) and Dell’s [VPLEX support matrix resources](#), which describe supported VPLEX host and storage connectivity requirements. FC is a lossless SAN protocol designed for reliable, high-performance block storage access and is the protocol identified by the labeled path in the exhibit.

QUESTION NO: 4

Which method of virtual volume expansion requires the volume to be expanded on the array first?

- A. Extent expansion
- B. Storage volume
- C. RAID-C expansion
- D. Concatenation

ANSWER: B

Explanation:

Storage volume is correct because this expansion method begins at the back-end storage array. The administrator first increases the capacity of the array LUN that is presented to VPLEX. After the array has completed the expansion and reports the new capacity, VPLEX can detect the increased size and expand the corresponding storage volume. The additional capacity can then be incorporated into the VPLEX virtual-volume configuration through the applicable VPLEX management workflow. This order is necessary because a VPLEX storage volume represents a physical LUN exported by the array; VPLEX cannot consume capacity that has not yet been provisioned and made available by that underlying LUN. Proper operational practice is to confirm the array-side expansion is complete and that the new size is visible to VPLEX before performing the VPLEX-side expansion steps. Dell’s VPLEX documentation is available through the [Dell VPLEX documentation portal](#), and the platform’s role in virtualizing external storage is summarized on the [Dell VPLEX product page](#).

QUESTION NO: 5

What is a characteristic of VPLEX Witness in a VPLEX Metro configuration?

- A. It is not in the host I/O data path
- B. It mirrors distributed-volume data
- C. It presents virtual volumes to hosts
- D. It replaces the WAN COM links between clusters

ANSWER: A

Explanation:

VPLEX Witness is not in the host I/O data path. It acts as an independent arbitration component that assists the VPLEX clusters when they lose communication with one another. Witness helps distinguish a site failure from an inter-cluster partition so VPLEX can protect distributed-volume consistency while allowing the appropriate cluster to continue I/O.

Host read and write I/O continues to be processed by VPLEX directors and the back-end storage arrays. Witness does not present volumes to hosts, process production I/O, or mirror application data. Additional information is available from the [Dell VPLEX documentation portal](#).

QUESTION NO: 6

Refer to the exhibit.



	Name	Capacity	Health	Status	Use
1	VPD83T3:600601603e4023004aa97306d45fe411	79.99 GB	OK	OK	meta-data
2	VPD83T3:600601603e4023004a291b84a806e611	6.00 GB	OK	OK	unclaimed
3	VPD83T3:600601603e40230048291b84a806e611	6.00 GB	OK	OK	claimed
4	VPD83T3:600601603e40230047d92079f916e611	27.00 GB	OK	OK	used

Which displayed storage volume is available to be used for the creation of an extent?

- A. 1
- B. 2
- C. 3
- D. 4

ANSWER: A

Explanation:

Volume 1 is available for creation of an extent because it is displayed as an eligible, unclaimed back-end storage volume. In VPLEX, an extent is the initial virtual-storage object created from a storage volume that has been presented by the storage array. Before VPLEX can use a storage volume in this way, the volume must be visible to the appropriate VPLEX engine, be in a healthy/usable state, and not already be consumed by another VPLEX object or reserved for another purpose. The displayed attributes for volume 1 identify it as meeting those prerequisites, so it can be selected when creating an extent. Creating the extent establishes the VPLEX representation of the selected array LUN; that extent can subsequently be incorporated into a device, which may then be used to create a virtual volume for host presentation. This object progression is fundamental to VPLEX provisioning and ensures that physical storage is first claimed and managed by VPLEX before it is exported to hosts. Dell's VPLEX documentation resources are available through the [Dell VPLEX documentation page](#); product context and supported VPLEX capabilities are also described on the [Dell VPLEX family page](#).

QUESTION NO: 7

What must be done before a claimed VPLEX storage volume can be unclaimed?

- A. Remove all extents created on the storage volume
- B. Remove the storage volume from every storage tier
- C. Place the storage volume in a storage view
- D. Convert the storage volume to a distributed device

ANSWER: A

Explanation:

All extents created on the storage volume must be removed before the storage volume can be unclaimed. An extent represents all or part of a claimed storage volume and can be used by a VPLEX device. VPLEX cannot release the underlying storage volume while an extent continues to reference its capacity.

The administrator must first remove any dependent VPLEX objects in the required order, ensuring that devices and virtual volumes using the extent have been appropriately removed or migrated. After no extents remain on the storage volume, it can be unclaimed and returned to an available state. Refer to the [Dell VPLEX documentation portal](#) for storage-volume administration procedures.

QUESTION NO: 8

What is required before a host can detect the virtual volumes presented by the VPLEX?

- A. Virtual volumes can only be detected after a reboot
- B. RAID configuration must be enabled for Virtual volumes
- C. EZ Provisioning wizard must be run on the host
- D. Host must initiate a bus-scan of the HBAs

ANSWER: D

Explanation:

Host must initiate a bus-scan of the HBAs. After a VPLEX virtual volume has been created, mapped to a storage view, and the appropriate initiators have been registered, the host must discover the newly presented SCSI device paths. A bus scan, also commonly called an HBA rescan or storage rescan, causes the host's Fibre Channel HBAs and operating-system SCSI layer to query the fabric for available targets and LUNs. The host can then identify the VPLEX virtual volume, create the applicable device entries, and make the volume available for subsequent multipathing and operating-system configuration. The exact command or procedure is host-OS-specific—for example, a Linux host may scan its SCSI hosts, while VMware ESXi uses an adapter or storage rescan—but the underlying requirement is the same: initiate discovery of the newly exposed VPLEX volume. This is the normal discovery workflow after changing VPLEX masking or storage-view presentation and does not inherently require a host reboot. Dell's VPLEX documentation describes making virtual volumes visible through storage views and then performing host-side discovery and configuration. See the [Dell VPLEX product documentation](#) and the [Dell guidance for rescanning newly presented LUNs on Linux hosts](#).