PassQueen
a.k.a dumpsqueen.com

VPLEX Specialist Exam for Implementation Engineers

EMC E20-260

Version Demo

Total Demo Questions: 25

Total Premium Questions: 254

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, VPLEX Concepts	37
Topic 2, VPLEX Hardware	32
Topic 3, VPLEX Installation	36
Topic 4, VPLEX Configuration	66
Topic 5, VPLEX Management	60
Topic 6, VPLEX Troubleshooting	23
Total	254

QUESTION NO: 1

Which is the best way to identify the root cause of long read latencies using monitors?

A. Check if the %CPU Busy is above 50%

B. Check to see if all volumes are in at least 1 Consistency Group

C. Check the ratio of cache.rhit to cache.miss

D. Check the difference between fe-director.read-lat and storage-volume.read-latency

ANSWER: D

Explanation:

Checking the difference between `fe-director.read-lat` and `storage-volume.read-latency` is the most useful monitor-based method because it separates the host-facing latency observed at the VPLEX front end from the latency incurred while VPLEX reads from the underlying storage volume. `fe-director.read-lat` represents the time associated with servicing host read I/O through a front-end director, while `storage-volume.read-latency` reflects the back-end storage response portion of the operation. Comparing them provides a practical latency decomposition. If the storage-volume value accounts for most of the front-end value, the observed delay is principally associated with the storage path or array response. If front-end read latency is materially higher than storage-volume read latency, the additional time points to processing, queueing, or contention within the VPLEX I/O path. This comparison is therefore actionable: it identifies whether investigation should begin with the back-end storage environment or with VPLEX director and path behavior. Dell provides VPLEX documentation and monitoring references through its VPLEX Series documentation portal; the broader Dell VPLEX support documentation also provides platform context for supported VPLEX environments.

QUESTION NO: 2

You are planning to perform a topology upgrade of VPLEX Local to VPLEX Metro.

What is the process you should follow to complete the upgrade successfully?

A. 1. Upgrade the existing cluster
2. Run the I/O module upgrade script
3. Cable the WAN COM connections
4. Register the cluster
5. Configure the new cluster

B. 1. Run the I/O module upgrade script
2. Cable the WAN COM connections
3. Upgrade the existing cluster
4. Configure the new cluster
5. Register the cluster

C. 1. Upgrade the existing cluster
2. Register the cluster
3. Configure the new cluster
4. Run the I/O module upgrade script
5. Cable the WAN COM connections

D. 1. Cable the WAN COM connections
2. Run the I/O module upgrade script
3. Configure the new cluster
4. Register the cluster
5. Upgrade the cluster

ANSWER: A

Explanation:

The correct process is: “1. Upgrade the existing cluster 2. Run the I/O module upgrade script 3. Cable the WAN COM connections 4. Register the cluster 5. Configure the new cluster.” A Local-to-Metro topology upgrade must first bring the existing VPLEX cluster to the required hardware and software state for Metro operation. The I/O module upgrade script is then run as part of the prescribed conversion work, ensuring that the existing cluster has the appropriate module configuration before the inter-cluster environment is connected. After this preparation, the WAN COM connections can be cabled to establish the communication path between the two clusters. The clusters are then registered so that VPLEX can recognize and manage them as a Metro configuration. Finally, the new cluster is configured, including the required storage, host, and distributed-device configuration activities. This order protects configuration consistency and ensures that each prerequisite is complete before the next stage depends on it. The implementation procedure should always be performed using the release-specific VPLEX installation or upgrade documentation and any generated procedure from Dell’s support tools. See the Dell VPLEX documentation portal and the Dell SolVe procedure generator.

QUESTION NO: 3

You are preparing to migrate the topology of an installed VPLEX Local system. What is a supported upgrade path?

A. Local to Geo to Global

B. Local to Metro

C. Local to Global

D. Local to Metro to Geo

ANSWER: B

Explanation:

Local to Metro is a supported VPLEX topology upgrade path. A VPLEX Local configuration provides storage virtualization from a single VPLEX cluster at one site. Upgrading it to VPLEX Metro expands the solution to two clusters at separate sites, enabling distributed volumes and synchronous data access across the Metro configuration. The migration requires the appropriate additional VPLEX hardware, connectivity between the clusters, supported software versions, and adherence to the VPLEX installation and upgrade procedures. Planning must also account for back-end array connectivity, host multipathing, WAN/FC connectivity requirements, and the configuration of consistency groups and distributed devices after the second cluster is introduced. This path is specifically designed to let an existing Local deployment grow into a dual-site Metro deployment while retaining the VPLEX virtualization architecture and managed storage environment. Dell’s VPLEX documentation portal provides the applicable product guides, release notes, and installation documentation for validating the exact hardware and software prerequisites for the deployed VPLEX code level. See the Dell VPLEX Series documentation page and Dell’s VPLEX product documentation knowledge-base article.

QUESTION NO: 4

Refer to the exhibit.

```
VPlexcli:/engines/engine-1-2/directors/director-1-2-A/hardware/ports/A1-FC01> ll
Name                Value
------------------  ------------------
address             0x500014xxxxxxxx11
current-speed       8Gbits/s
enabled             true
max-speed           8Gbits/s
node-wwn            0x500014xxxxxxxxd9
operational-status  ok
port-status         up
port-wwn            0x500014xxxxxxxx11
protocols           [fc]
role                back-end
target-port         -
topology            loop
```

You are implementing a VPLEX Metro configuration, and some BE ports are not logged into the fabric.

Which VPLEX command should be performed?

A. set topology p2p

B. set enabled=true

C. set topology loop

D. set port speed

ANSWER: A

Explanation:

set topology p2p is the required command when VPLEX back-end Fibre Channel ports connect through a switched SAN fabric but are not logging in. A VPLEX port's topology setting must match the connection type used by that port. Point-to-point topology is the appropriate setting for a fabric-attached port, allowing the port to initialize its fabric connection and perform the Fibre Channel fabric login process. Once the topology is set correctly, the switch can recognize the port and establish the required connectivity to the storage-array target ports. This is an essential implementation check before storage volumes can be discovered and used as VPLEX back-end storage. The setting should be applied to the affected back-end ports in the relevant VPLEX cluster and then verified through the VPLEX management interface and the SAN-switch name server or login status. Dell's VPLEX documentation and product-support resources provide the configuration and connectivity guidance used during implementation: Dell VPLEX Series documentation and Dell VPLEX Family support information.

QUESTION NO: 5

What is the total number of engines for a small VPLEX Metro configuration?

A. 1

B. 8

C. 4

D. 2

ANSWER: D

Explanation:

2 is correct. A VPLEX Metro deployment consists of two VPLEX clusters, normally located at separate sites, that provide distributed volume access across the metropolitan-distance configuration. A small VPLEX Metro configuration uses one engine in each cluster. Consequently, the complete Metro system contains two engines: one local engine at each site. Each engine contains the directors that provide the front-end, back-end, and WAN connectivity needed for VPLEX processing and inter-cluster communication. The "small" designation describes the engine count per cluster, not a single-engine Metro system. This distinction is important during implementation planning because the two clusters form the Metro configuration and must be accounted for when validating hardware layout, management, connectivity, and high-availability design. Dell's VPLEX documentation portal provides the platform's product documentation, including installation and configuration material that describes VPLEX cluster and engine architecture: Dell Support: VPLEX Series documentation. Dell's VPLEX family page also identifies VPLEX Metro as a configuration spanning two sites: Dell VPLEX Family.

QUESTION NO: 6

Refer to the exhibit.

The customer deleted ext_2 on a 20 GB LUN in order to make an extent of 10 GB. Why is the customer unable to create a 10 GB extent?

A. Storage array must be re-discovered.

B. ext_2 cannot be deleted.

C. Chunks of the storage need to be contiguous.

D. The defrag property is set to false.

ANSWER: C

Explanation:

Chunks of the storage need to be contiguous. A VPLEX extent is defined as a continuous range of blocks on one backing storage volume; it cannot be assembled from separate, nonadjacent free regions of that LUN. Deleting `ext_2` releases only the specific block range previously occupied by that extent. If the remaining free capacity on the 20 GB LUN is split by another existing extent or by allocated ranges, VPLEX may show enough aggregate free capacity but still cannot allocate one uninterrupted 10 GB range. The requested extent therefore requires a contiguous 10 GB block range on the storage volume. This is an extent-layout and allocation issue, rather than a capacity-total issue. Reviewing the storage-volume layout in the VPLEX management interface identifies the usable contiguous free ranges before creating or resizing extents. Dell's VPLEX documentation resources are available through the VPLEX Series documentation page; Dell also provides current VPLEX support resources through its VPLEX Series support page.

QUESTION NO: 7

Refer to the exhibit.

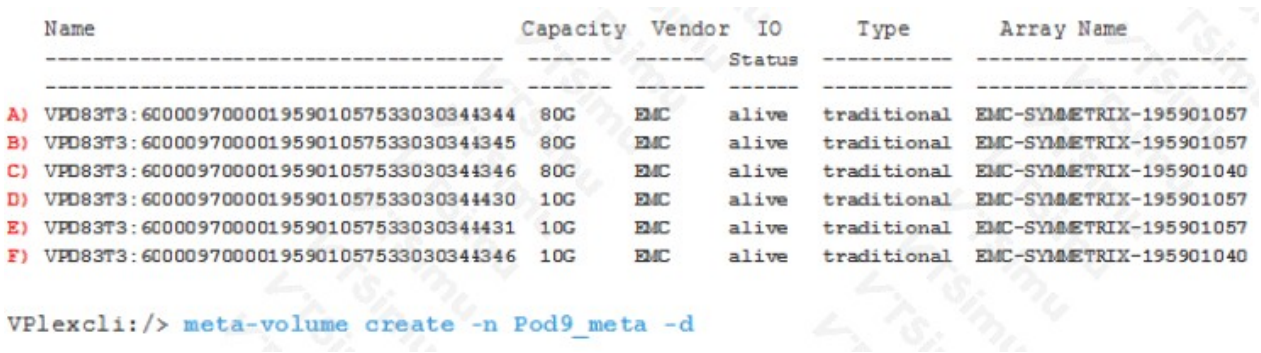

```
   Name                                    Capacity  Vendor  IO      Type        Array Name
   --------------------------------------  --------  ------  Status  ----------  --------------------
   --------------------------------------  --------  ------  ------  ----------  --------------------
A) VPD83T3:6000097000019590105753303034434 4  80G       EMC     alive   traditional  EMC-SYMMETRIX-195901057
B) VPD83T3:6000097000019590105753303034434 5  80G       EMC     alive   traditional  EMC-SYMMETRIX-195901057
C) VPD83T3:6000097000019590105753303034434 6  80G       EMC     alive   traditional  EMC-SYMMETRIX-195901040
D) VPD83T3:6000097000019590105753303034443 0  10G       EMC     alive   traditional  EMC-SYMMETRIX-195901057
E) VPD83T3:6000097000019590105753303034443 1  10G       EMC     alive   traditional  EMC-SYMMETRIX-195901057
F) VPD83T3:6000097000019590105753303034434 6  10G       EMC     alive   traditional  EMC-SYMMETRIX-195901040

VPlexcli:/> meta-volume create -n Pod9_meta -d
```

Which two volumes should be used to complete the displayed command?

A. B and C

B. A and B

C. D and F

D. D and E

ANSWER: A

Explanation:

B and C should be used to complete the displayed VPLEX command. The volume information in the exhibit identifies these two storage volumes as the appropriate matching pair for the operation: they provide the required compatible characteristics shown for the command, including the necessary capacity and storage-volume status. VPLEX commands that consume multiple backend volumes require the selected volumes to satisfy the command's eligibility criteria before VPLEX can create or configure the requested object successfully. Selecting B and C ensures that the command operates on the intended eligible volumes represented in the exhibit and allows VPLEX to maintain the expected backend layout for the resulting device configuration. This selection process is important because VPLEX uses the discovered storage-volume properties when presenting storage through its virtual-volume and device layers. Implementation engineers should verify those properties in VPLEX management interfaces or the CLI before issuing a configuration command, particularly when multiple discovered volumes have similar identifiers or capacities. Dell's VPLEX product documentation and support resources provide the applicable CLI and implementation guidance: Dell VPLEX documentation and Dell VPLEX product information.

QUESTION NO: 8

After registering Cluster-2 with ESRS during a VPLEX Metro cluster build, what is the next step to perform before the cluster is successfully built?

A. Complete EZ-Setup Wizard on Cluster-2

B. Configure WAN connections on Cluster-1

C. Complete EZ-Setup Wizard on Cluster-1

D. Configure WAN connections on Cluster-2

ANSWER: C

Explanation:

Complete EZ-Setup Wizard on Cluster-1 is the required next activity in the Metro build workflow. After both VPLEX clusters have been prepared and Cluster-2 has been registered with EMC Secure Remote Services (ESRS), the EZ-Setup workflow is launched from Cluster-1, which acts as the point from which the distributed VPLEX Metro configuration is finalized. The wizard validates the participating clusters and their connectivity, collects the required configuration details, and completes the initialization needed for the two clusters to operate together as a Metro system. Performing EZ-Setup at this stage ensures that remote support registration and the prerequisite cluster-level configuration are in place before the distributed configuration is committed. This sequencing is important because the build process depends on Cluster-1 being able to

discover and establish the intended Metro relationship with the prepared remote cluster. Dell's VPLEX documentation and support resources provide the applicable installation, configuration, and serviceability guidance for VPLEX deployments, including Metro configurations and ESRS connectivity. See the Dell Support documentation for VPLEX and the Dell VPLEX product information.

QUESTION NO: 9

A customer is moving some VMAX devices under a VPLEX Local and is using PowerPath as multipath software.

What is the SCSI product ID assigned to VPLEX volumes?

A. Invista

B. VPLEX

C. PowerPath

D. VMAX

ANSWER: A

Explanation:

Invista is the SCSI product ID presented for VPLEX virtual volumes. Although the back-end storage devices are VMAX devices and PowerPath is the host multipathing software, the host no longer identifies those LUNs using the array's native VMAX identity after they have been claimed and exported by VPLEX. VPLEX presents its own virtualized SCSI device identity to the host. The product identifier *Invista* is retained for VPLEX host presentation compatibility and is therefore the value that host-side discovery and multipathing tools can see in SCSI INQUIRY data. This identity is important during implementation because host configuration, device recognition, and validation must be based on the VPLEX-presented virtual volume rather than the physical array that supplies storage behind VPLEX. PowerPath uses the device information exposed by the VPLEX front end to manage the available paths to the virtual volume. Dell's VPLEX documentation and interoperability resources should be consulted for the supported host, operating-system, and multipathing configuration applicable to the deployment. See the Dell VPLEX Local documentation and the Dell E-Lab Navigator for current implementation guidance.

QUESTION NO: 10

What steps must be performed to encapsulate LUNs into virtual volumes?

A. 2. Create an extent, create a device on the extent, and create a virtual volume on the device

B. 2. Create device

C. 2. Mirror the virtual volume to native volume with data

D. 2. Create extent

ANSWER: A

Explanation:

Create an extent, create a device on the extent, and create a virtual volume on the device is correct because VPLEX virtualizes a back-end storage LUN through these layered objects. After the storage volume has been claimed and is visible to VPLEX, an extent represents the usable storage capacity from that LUN. A device is then constructed from the extent; it is the VPLEX object that provides the internal logical-storage layer and can subsequently support capabilities such as distributed devices or local RAID configurations where applicable. Finally, a virtual volume is created from the device. The virtual volume is the object presented by VPLEX to hosts through a storage view, so it is the required front-end representation of the encapsulated LUN. This ordered relationship—storage volume to extent to device to virtual volume—ensures that VPLEX can manage the physical storage while exposing a virtualized volume consistently to attached hosts.

Dell's VPLEX documentation and product-support resources describe this object hierarchy and the provisioning workflow: Dell VPLEX documentation and Dell Support knowledge base.

QUESTION NO: 11

A VPLEX Metro distributed device has one device leg at each cluster. What is the purpose of the device legs?

A. They provide the local storage components for the distributed device

B. They provide host initiator definitions for each cluster

C. They store Management Server operating-system logs

D. They replace the WAN-COM connections between clusters

ANSWER: A

Explanation:

They provide the local storage components for the distributed device is correct. A distributed device combines a device from each VPLEX cluster into one logical device that can be accessed from either site. Each local device becomes a device leg and is backed by storage visible to its local cluster. VPLEX maintains consistency between the legs so that the distributed device can be exported as a single distributed virtual volume.

See the Dell VPLEX overview and the VPLEX documentation library for distributed-device configuration information.

QUESTION NO: 12

A storage administrator has completed zoning and LUN masking for a new back-end array that has never been connected to VPLEX.

Which operation should be performed first to make the array known to VPLEX?

A. Array discover

B. Array re-discover

C. Storage-volume resurrect

D. Device detach

ANSWER: A

Explanation:

Array discover is correct. Array discovery is used when VPLEX is first introduced to a back-end storage array. The operation identifies the array through the configured back-end connectivity and adds it to the VPLEX inventory. After the array is discovered, storage volumes presented by that array can be identified and claimed. Array rediscover is used later when the presentation from an array already known to VPLEX has changed.

See the Dell VPLEX documentation library for array-discovery and storage-claim procedures.

QUESTION NO: 13

Which command allows you to navigate vplexcli history by string?

A. CTRL + R string

B. history find string

C. history | grep string

D. UP arrow key

ANSWER: A

Explanation:

`CTRL + R string` is correct because pressing `Ctrl+R` in the VPLEX CLI starts an incremental reverse history search. After entering the key combination, the operator types a text string from a previously issued command, and the CLI searches backward through the command history for the most recent matching entry. Repeated use of `Ctrl+R` continues searching backward for earlier commands containing the same string, allowing an administrator to quickly locate and reuse a command without scrolling through history entries one at a time. This is particularly useful in VPLEX operational work, where commands can be lengthy and often share recognizable object names, device identifiers, or command keywords. Once the intended historical command is displayed, it can be accepted for editing or execution as appropriate. Incremental reverse search is a standard Readline command-history capability, documented as `reverse-search-history`; it searches backward from the current history position using the text supplied by the user. See the GNU Bash manual's Searching for Commands in the History section and the GNU Readline history-search documentation.

QUESTION NO: 14

Refer to the exhibit.

After 6 months of operation, you are monitoring the health of the VPLEX system from the Unisphere Manager for VPLEX.

What is the most likely reason for a state change of the VPLEX Witness?

A. TheESXi host was rebooted.

B. The WAN switch ports failed.

C. The Management Server was restarted.

D. TheMetaVolume was restored.

ANSWER: A

Explanation:

TheESXi host was rebooted. VPLEX Witness is implemented as a virtual appliance running on an ESXi host in a third failure domain. Its purpose is to provide an independent tie-breaker for a VPLEX Metro configuration and to communicate its view of cluster availability to the VPLEX clusters. Rebooting the ESXi host interrupts execution of the Witness virtual machine and temporarily removes its network communication with the VPLEX environment. Unisphere Manager for VPLEX therefore detects that the witness is no longer reachable and reflects that condition as a Witness state change. Once the ESXi host has completed its reboot and the Witness virtual machine is running and reachable again, the Witness can resume normal participation, subject to successful service and network recovery. This is an expected operational effect of maintenance or an unplanned restart of the ESXi platform hosting the Witness appliance. Dell's VPLEX documentation is available from the Dell VPLEX support documentation portal; ESXi host restart behavior is also documented in the Broadcom knowledge base.

QUESTION NO: 15

Which information is required to configure ESRS for VPLEX?

A. B/Top Level Assembly Site ID IP address of the Management Server public port Configuration (1, 2 or 4 Engines)

B. VPLEX model type Level Assembly customer contact IP address of the Management Server public port

C. VPLEX model type Site ID IP address of the Management Server public port Primary Customer contact

ANSWER: A

Explanation:

B/Top Level Assembly Site ID IP address of the Management Server public port Configuration (1, 2 or 4 Engines) is the required information set for VPLEX ESRS configuration. ESRS must be able to associate the call-home configuration with the precise VPLEX hardware identity and its customer installation. The B/Top Level Assembly identifies the VPLEX system at the assembly level, while the Site ID associates it with the customer site registered for support. The public-port IP address of the VPLEX Management Server provides the management-network endpoint used for connectivity and service communication. Finally, the engine configuration identifies the deployed VPLEX topology—one, two, or four engines—which is needed to accurately represent the installed system configuration to Dell support systems. Collecting these details before starting configuration helps ensure that alerts, remote support connectivity, and system identification are registered against the correct VPLEX environment. Dell publishes VPLEX implementation and service documentation through its VPLEX documentation portal; the general Dell support site is available at Dell Support.

QUESTION NO: 16

What is an EMC best practice for fault tolerance when zoning a VPLEX Metro together over Fibre Channel for successful communication between the sites?

A. Zone each WAN-COM port to all WAN-COM ports in the other fabric.

B. Zone all WAN-COM ports to each other in both SAN fabrics within a single zone.

C. Zone each WAN-COM port to all WAN-COM ports in the other fabric.
Zone one Local-COM port to all WAN-COM ports in the other fabric for VPLEX Witness.

D. Zone all WAN-COM ports to each other in both SAN fabrics within a single zone.
Zone one Local-COM port to all WAN-COM ports in the other fabric for VPLEX Witness.

ANSWER: A

Explanation:

Zone each WAN-COM port to all WAN-COM ports in the other fabric. This is the EMC-recommended fault-tolerant Fibre Channel connectivity model for VPLEX Metro inter-cluster communication. WAN-COM ports carry the distributed-cache and

cluster communication traffic between the two VPLEX clusters. Providing full WAN-COM connectivity across the corresponding fabrics creates multiple available communication paths between sites rather than relying on an individual port-to-port relationship. As a result, inter-cluster communication can continue through the remaining paths when a WAN-COM port, cable, switch port, or an entire SAN-fabric path is unavailable. This redundancy is essential to maintaining the high-availability design expected of a VPLEX Metro configuration and supports continued distributed-volume operation during infrastructure faults. The zoning design must preserve fabric separation: WAN-COM connectivity is established through each corresponding independent SAN fabric, while each WAN-COM port has visibility to all applicable remote WAN-COM ports in that fabric. Dell's VPLEX product information describes VPLEX Metro as a solution designed for continuous availability across geographically separated sites: Dell VPLEX family. Dell also provides VPLEX technical documentation, including installation and configuration guidance, through its VPLEX support documentation.

QUESTION NO: 17

Refer to the exhibit.

You are preparing to perform an upgrade on an unconfigured VPLEX Metro. When you run the ndu pre-check command, it indicates that it has failed and there are no first upgraders. In your troubleshooting, you run the version -a command.

Based on the output, what is the reason for the failure?

A. health-check --full was not run

B. Directors are at incompatible code levels

C. ndu pre-check --force flag was not used

D. Directors are not connected.

ANSWER: D

Explanation:

Directors are not connected. is correct because a VPLEX NDU relies on director-to-director communication to assess director state and select an eligible first upgrader. The `version -a` output provides version and connectivity information for the directors. When that output indicates that the directors are disconnected, the NDU framework cannot establish the coordinated director relationship needed to choose the first director to be upgraded. Consequently, `ndu pre-check` reports that no first upgraders are available. This must be resolved before beginning the upgrade: confirm that the directors are powered and healthy, verify the required internal management and communication cabling, and ensure that each director is visible and connected in the VPLEX management environment. Once communication is restored and the directors report a connected state, rerun the applicable health validation and the NDU pre-check. Dell's VPLEX documentation portal provides the platform's installation, service, and upgrade documentation, including NDU prerequisites and director connectivity requirements: Dell VPLEX Series documentation. Additional Dell support resources are available through the Dell Technologies Support portal.

QUESTION NO: 18

What three elements are part of the VPLEX Storage View?

A. VPLEX FE Port
Top-Level Device
Host IP address

B. VPLEX FE Port
Initiator Port
Virtual Volume

C. VPLEX BE Port
VPLEX FE Port
Virtual Volume

D. VPLEX BE Port
Initiator Port
Virtual Volume

ANSWER: B

Explanation:

VPLEX FE Port, Initiator Port, and Virtual Volume is correct because a VPLEX storage view is the access-control and presentation object that associates a host with the storage it is permitted to use through VPLEX. The initiator port represents the host-side Fibre Channel initiator WWN that logs in to VPLEX. The VPLEX front-end port is the target-facing VPLEX port through which that host accesses exported storage. The virtual volume is the VPLEX-managed volume presented to the host as a LUN. Creating the storage view establishes this specific relationship, allowing VPLEX to mask and present the selected virtual volumes only to the defined host initiators through the selected front-end ports. This design supports controlled host access and is fundamental to VPLEX provisioning workflows. Back-end ports are used for connectivity between VPLEX and storage arrays rather than host presentation, so they are not part of the storage-view definition. For VPLEX product documentation and available implementation guides, see the Dell VPLEX documentation portal and the Dell VPLEX family overview.

QUESTION NO: 19

Starting in Geosyncrony version 5.1 what are the minimum password requirements?

A. Fourteen characters, including numbers, upper-case and lower-case letters and special characters

B. Seven characters, including numbers

C. Eight characters, including numbers, and special characters

D. Ten characters, including numbers, upper-case and lower-case letters, and special characters

ANSWER: A

Explanation:

Starting with GeoSynchrony 5.1, VPLEX management-account password policy was strengthened to require a minimum length of fourteen characters and password complexity from all four character groups: numeric characters, uppercase letters, lowercase letters, and special characters. Therefore, “Fourteen characters, including numbers, upper-case and lower-case letters and special characters” states the required minimum accurately. This policy is intended to reduce the likelihood that privileged VPLEX credentials can be guessed or compromised through password attacks. It applies the complexity requirement at the point where passwords are created or changed, so an implementation engineer should plan credentials that meet the full requirement before performing installation, upgrade, or administrative tasks. Password-policy requirements should also be confirmed against the documentation matching the installed GeoSynchrony release, because procedures and security controls can vary between releases. Dell’s VPLEX documentation portal is available at Dell VPLEX Product Support Documentation. Dell’s support search can be used to locate the applicable GeoSynchrony 5.1 security and installation documentation at Dell Support Search.

QUESTION NO: 20

When creating logging and metadata volumes, which LUN requirement applies to both?

A. Array LUNs must be zeroed-out.

B. Array LUNs must be RAID 1.

C. Array LUNs must be thin.

D. Array LUNs must be RAID 5.

ANSWER: B

Explanation:

Array LUNs must be RAID 1. VPLEX uses metadata volumes to retain critical distributed-system configuration and device metadata, while logging volumes provide protected write logging used by VPLEX consistency mechanisms. Both types therefore require mirrored RAID 1 storage protection. Mirroring maintains a second copy of each block on an independent member, allowing the volume to remain available after the loss of a single underlying drive or RAID member. This resilience is essential because loss or corruption of either volume type can affect VPLEX recovery operations and the availability of distributed devices. RAID 1 also provides predictable low-latency write behavior, which is appropriate for the small but operationally critical I/O generated by metadata and logging functions. The required capacity of a logging volume and a metadata volume can differ, but their protection requirement is the same: they must be provisioned as RAID 1 LUNs before being claimed and configured by VPLEX. Dell publishes VPLEX product documentation, including implementation and configuration guidance, through its VPLEX documentation portal. See also Dell's VPLEX manuals page for the applicable installation and planning documentation.

QUESTION NO: 21

A customer is performing maintenance on their network that might affect their VPLEX Witness VM.

What can be done to ensure there are no adverse effects due to a Witness outage?

A. Disable the cluster's management server ipsec.

B. Disable the VPLEX Witness.

C. Shut down the VPLEX Witness VM.

D. Shut down the cluster's management server.

ANSWER: B

Explanation:

Disable the VPLEX Witness before the planned network maintenance. This is the supported administrative action for a known, intentional Witness interruption: it tells the VPLEX clusters not to use Witness participation while connectivity to the Witness VM may be unavailable. Disabling it in advance avoids treating the planned loss of communication as an unexpected Witness failure and keeps the distributed VPLEX environment in a controlled operational state throughout the maintenance window.

VPLEX Witness provides an independent arbitration mechanism for VPLEX Metro configurations, particularly during site or inter-cluster communication failure scenarios. Because its availability can influence cluster behavior during fault handling, planned work affecting the Witness VM or its network path should be coordinated by explicitly disabling Witness first. After the network work is complete and stable connectivity between both VPLEX clusters and the Witness VM has been verified, Witness can be enabled again so normal arbitration protection is restored. The operational procedure and required health checks should follow the VPLEX documentation for the installed software release. See the Dell VPLEX documentation portal and the Dell support knowledge base for VPLEX operational guidance.

QUESTION NO: 22

A customer wants to add a mirror to a VPLEX Local RAID-0 device.

Which RAID type is created after the command "device attach-mirror"?

A. RAID-0

B. RAID-S

C. RAID-1

D. RAID-C

ANSWER: C

Explanation:

RAID-1 is created when `device attach-mirror` is used on a VPLEX Local RAID-0 device. A RAID-0 device initially presents a single device leg without redundancy at the VPLEX virtual-volume layer. The attach-mirror operation adds a second leg and establishes a mirror relationship between the original leg and the newly attached storage. VPLEX then synchronizes the contents of the added leg with the existing device so that both legs contain equivalent data. Once this synchronization is established, the device has the mirrored, two-leg layout represented by RAID-1. This operation is commonly used to introduce VPLEX-level redundancy or to move a leg onto different backing storage while keeping the virtual volume available. The resulting RAID-1 geometry allows either mirror leg to continue serving data if the other leg is unavailable, subject to the remaining VPLEX and backend-storage configuration. Dell's VPLEX documentation and product support resources provide the applicable CLI procedures and device-geometry concepts: Dell VPLEX documentation and Dell Technologies Support.

QUESTION NO: 23

After an expansion of a RAID-0 device in VPLEX, the customer notices a slight degradation of performance. What could cause this?

A. After the expansion operation is complete, the original mapped device has been converted into type RAID-C.

B. After the expansion operation is complete, the original mapped device has been converted into type RAID-S.

C. The device is striped across mismatched array types.

D. The device is striped across RAID-S devices.

ANSWER: A

Explanation:

After the expansion operation is complete, the original mapped device has been converted into type RAID-C. This is expected VPLEX behavior for this type of expansion. VPLEX preserves the original RAID-0 structure and incorporates the added capacity through a RAID-C (concatenation) layer rather than redistributing all pre-existing data across the newly available capacity. The resulting virtual-device layout can therefore introduce an additional mapping layer and can leave I/O distributed according to where the data already resides. I/O directed to the original portion continues through its existing structure, while I/O to newly allocated portions follows the added component. Depending on workload locality, this can create a modest performance difference after expansion, particularly when the original and added storage portions have different response characteristics or when the workload remains concentrated on existing data. This behavior avoids a disruptive full data restriping operation during the capacity expansion, but the concatenated layout is a valid reason for the slight degradation described. Dell's VPLEX family information is available at Dell VPLEX Family, and VPLEX product documentation can be accessed through Dell Support: VPLEX Series Documentation.

QUESTION NO: 24

You need to claim storage that is presented to VPLEX from a VNX. What is required in order to create the mapping file needed by the VPLEX Claim Storage Wizard?

A. A host with solutions enabler installed and zoned to SPA and SPB

B. The user name, password and management address of the VNX

C. A host with WinSCP installed, with an IP connection to the VPLEX Management Server

D. The username and password of the VPLEX Management Server

ANSWER: B

Explanation:

The user name, password and management address of the VNX is required because the mapping file is generated by connecting to the VNX management environment and collecting the VNX configuration information that identifies the presented storage. In particular, the process needs authenticated access to obtain the relevant LUN and storage-group mapping details that let VPLEX associate discovered back-end volumes with their intended VNX presentation information during the Claim Storage workflow. The VNX management address specifies the target array, while the user name and password authorize the mapping-file generation utility or workflow to query that array. This information is therefore the essential input for creating a valid VNX mapping file before the volumes are claimed in VPLEX. Dell documentation for VPLEX is available through the Dell VPLEX documentation portal, and VNX family product documentation can be accessed from the Dell VNX documentation portal.

QUESTION NO: 25

Which VPLEX volume contains virtual-to-physical mapping information?

A. Logging volume

B. Metadata volume

C. Virtual volume

D. Repository volume

ANSWER: B

Explanation:

Metadata volume is correct because VPLEX uses metadata volumes to persist the system configuration and mapping data that define how VPLEX virtual volumes are presented and how they relate to underlying storage resources. This includes the metadata needed for VPLEX to maintain the virtual-to-physical relationship between its exported virtual volumes and the back-end storage volumes/devices from which they are constructed. The metadata is essential to consistent operation: it enables the VPLEX directors to recognize the configured devices and reconstruct the virtualization configuration after service events or director restarts. In a distributed VPLEX configuration, this persistent metadata also supports the coordinated view of the virtualized storage environment across clusters. Metadata volumes are therefore a required internal VPLEX component and must be provisioned according to the applicable VPLEX implementation guidance, including the required availability and protection characteristics. Dell's VPLEX documentation library contains the platform planning, installation, and administration manuals that describe metadata volumes and their role in VPLEX configuration: Dell VPLEX Series Manuals and Documents. Additional Dell technical support resources for VPLEX are available through the Dell VPLEX Series Support page.