Vendor: VMware

Exam Code: 3V0-23.25

Exam Name: Advanced VMware Cloud Foundation 9.0 Storage

Version: DEMO

QUESTION 1
An administrator attempts to enable vSAN Data-at-Rest Encryption on a cluster but receives the following error message:

```
"Key provider <vSphere Native Key Provider> is not available on host."
```

The administrator configured the vSphere Native Key Provider (NKP) using the default settings.

What should the administrator validate before enabling vSAN Data-at-Rest Encryption?

A. Each ESX host has a Trusted Platform Module (TPM) 2.0 device installed.
B. Crypto Safe Mode has been enabled on all ESX hosts.
C. A backup of the vSphere Native Key Provider has been created from the vSphere Client.
D. Secure Boot has been disabled on all ESX hosts.

Answer: C
Explanation:
The administrator must validate that the vSphere Native Key Provider has been backed up. When a Native Key Provider is created, vCenter pushes the key provider configuration to clustered ESX hosts, but the provider is not usable for encryption tasks until it has been backed up. The documentation states that after adding a Native Key Provider, it appears as not backed up, and that it must be backed up before use. The backup process changes the provider status from Not Backed Up, to Warning while vCenter propagates information to hosts, and then to Active when the information has been pushed to all hosts. A TPM 2.0 device is not required when using Native Key Provider with default settings, although TPM can provide enhanced security and can optionally be enforced. Crypto Safe Mode and disabling Secure Boot are not prerequisites for enabling vSAN Data-at-Rest Encryption with NKP. Therefore, the validation step is to confirm that the NKP backup has been created and the provider is active on the hosts.

QUESTION 2
An architect has been tasked with designing a vSAN OSA storage solution for a VMware Cloud Foundation (VCF) Private Cloud. The following requirements and constraints have been gathered from the client:

```
- 2 x 4 TB NVMe disks per host
- 6 hosts
- FTT=1/RAID-5 for all deployed Virtual Machines
- Expected Dedupe/Compression Ratio = 2 (100%)
- Reserve enough capacity to rebuild a host completely in case of a
failure (Host Rebuild Reservation)
- Operational Reserve of 10%
- Expected Overhead for Filesystem, Object, etc. of 25%
```

What is the approximate expected usable capacity of the resulting array?

A. 48 TB
B. 32 TB
C. 40 TB
D. 24 TB

Answer: C
Explanation:
The approximate expected usable capacity is 40 TB. The cluster contains six hosts, each with two 4 TB NVMe disks, giving 48 TB of raw capacity. Because the policy uses FTT=1 with RAID-5

erasure coding, vSAN provides better capacity efficiency than mirroring. RAID-5 with FTT=1 typically uses approximately 1.33x overhead, leaving roughly 75% of raw capacity before additional reserves and data-reduction assumptions. The design also requires Host Rebuild Reservation, which reserves enough capacity to rebuild data after a host failure, and an Operational Reserve of 10% to support internal operations such as resynchronization and policy changes. The expected filesystem and object overhead of 25% further reduces available capacity, while the deduplication/compression ratio of 2 increases the effective usable capacity. After accounting for protection overhead, capacity reserves, operational overhead, and the stated data-reduction assumption, the closest expected usable value from the available choices is 40 TB.

QUESTION 3
An administrator is monitoring a vSAN ESA backed workload domain that is dedicated for running AI inferencing. When the administrator navigates to the Storage Performance dashboard in VMware Cloud Foundation Operations, the performance dashboard shows:

```
- High backend write latency >; 8 ms)
- Low read latency <; 1 ms)
- Normal network throughput
- Disk Group Health = Green
```

Based on the readings above, what would be the explanation?

A. This is caused due to transient commit-queue delays, since the workload is exhibiting random-writes saturating ESA's write buffer.
B. A wrongly sized read cache tier is throttling the write buffer, thus forcing the reads to trespass to the capacity tier.
C. The workload's small-block writes are compressed inline, lowering backend throughput and increasing cache misses.
D. A vSAN network congestion event on the vSAN TCP port 2233 is throttling mirror acknowledgements.

Answer: A
Explanation:
The readings indicate a write-path pressure condition rather than a read-cache, network, or device-health issue. In vSAN ESA, storage devices are organized into a storage pool architecture where NVMe devices contribute both capacity and performance. The dashboard shows low read latency and normal network throughput, so the read path and vSAN network are not the likely bottlenecks. Disk Group Health being green also rules out an obvious device or pool-health failure. AI inferencing environments can still create bursts of small random writes from logs, metadata updates, temporary files, checkpoints, and application-side write activity. When backend write latency rises while reads remain low, the most likely explanation is transient write-buffer or commit-queue pressure. ESA does not use the traditional OSA read-cache tier model, so option B is incorrect. ESA also does not rely on OSA-style inline deduplication and compression behavior for the stated cause, making option C incorrect. Normal network throughput makes option D unlikely.

QUESTION 4
A multinational corporation is deploying a VMware Cloud Foundation (VCF) Workload Domain cluster to host tier-1 business applications with the following requirements:

```
- Deliver high performance to all workloads.
- Integration with automated VM placement.
```

```
- Reuse of existing infrastructure including multiple Fibre Channel
(FC) arrays partitioned by business unit.
- The storage architecture must avoid vSAN.
- Meet organizational separation requirements.
```

Which two deployment approaches meet these requirements? (Choose two.)

A. Configure dedicated VMFS datastores on assigned FC arrays as principal storage.
B. Configure dedicated VMFS datastores on assigned FC arrays as supplemental storage.
C. Configure a shared VMFS datastore on an assigned FC array as principal storage.
D. Configure shared VMFS datastores on an assigned FC array as supplemental storage.
E. Configure shared VMFS datastores on an assigned FC array as principal storage.

Answer: AB
Explanation:
Dedicated VMFS datastores on assigned Fibre Channel arrays meet both the performance and separation requirements. VCF supports Fibre Channel-backed VMFS as principal storage for workload domains, making it valid as the primary datastore type for the tier-1 business application cluster. Fibre Channel provides low-latency block storage, resilient pathing, SAN zoning, LUN masking, and mature enterprise performance characteristics. Because the existing FC arrays are partitioned by business unit, dedicated datastores preserve organizational separation and avoid mixing business- unit workloads on shared storage resources. The same dedicated FC-backed VMFS approach can also be used as supplemental storage after the workload domain exists, allowing extra capacity or specialized tiers while maintaining separation. Shared VMFS datastores do not satisfy the organizational separation requirement because they allow multiple business units to consume the same datastore boundary. vSAN is explicitly excluded, so the design must rely on supported non- vSAN storage. Automated VM placement can be supported by VMFS datastore clusters and Storage DRS where datastores are compatible and similarly managed.

QUESTION 5
An agency is designing its secure private cloud on VMware Cloud Foundation with the following requirements:

```
- Strict data segregation between the management and workload domains.
- Company policy prevents using vSAN as a storage solution.
- Data encryption at rest is mandatory for both the management and
workload domains.
- Data encryption in transit is mandatory for the workload domains.
- Data-at-rest encryption must be performed by the storage array and
not rely on VMware native or vSAN-specific mechanisms.
- Allow for automated VM placement, operational integrity with VCF
Operations, and assurance that file-based workloads scale efficiently.
```

Which storage architecture fulfills these technical and regulatory requirements?

A. Deploy metro clusters backed by Self-Encrypting Disk (SED) using iSCSI for both domains and configure encrypted VLANs for data-in-transit between VCF Operations and hosts.
B. Configure dedicated VMFS datastores on separate Fibre Channel arrays for management and workloads, with array-based encryption enabled and SPBM storage policies assigned. Use NFS v4.1 shares with Kerberos-based encryption for in-transit data for file workloads in the workload domain, restricting cross-domain datastore access.
C. Enable VMware VM-level encryption using vSphere Native Key Provider (NKP) on local VMFS disks for all environments, implementing manual file workload allocation through generic SMB

shares.

D. Create a unified NAS platform offering both NFS and SMB exports shared across management and workload domains, then enable application-level encryption for sensitive workloads and restrict access via firewall rules.

Answer: B
Explanation:
Option B satisfies every requirement because it separates management and workload storage onto dedicated Fibre Channel-backed VMFS datastores while relying on array-based encryption at rest. This avoids VMware native encryption and vSAN-specific encryption mechanisms, matching the regulatory requirement. Fibre Channel is a supported non-vSAN storage architecture for VCF and provides enterprise block storage with multipathing, LUN masking, zoning, and consistent datastore presentation. SPBM storage policies allow storage placement and compliance to align with array and LUN characteristics, supporting automated placement and operational consistency. For scalable file- based workloads, NFS v4.1 is appropriate as supplemental file storage, and Kerberos privacy mode provides authentication, integrity, and encryption in transit for NFS traffic. Restricting cross-domain datastore access preserves strict management/workload segregation.

QUESTION 6
A vSAN ESA solution is configured using the following requirements:

```
Seven ESX Hosts, each host contains:

32 CPU

256 GB memory

25 GbE network

12 storage devices 4 TB each

One storage pool using the 12 storage devices

RAID-6 with FTT=2
```

If a storage device on a single host fails, what percentage of that host's capacity is impacted?

A. 50%
B. 25%
C. 8.3%
D. 0%

Answer: C
Explanation:
In vSAN ESA, each host uses a storage pool instead of the older OSA disk-group model. The question asks what percentage of that host's local capacity is impacted by the failure of one storage device. Each host has 12 equal-capacity storage devices, and all 12 devices participate in one storage pool. Therefore, one failed device represents one-twelfth of that host's local pool capacity. One divided by twelve equals approximately 8.3%. RAID-6 with FTT=2 determines object resilience across hosts and fault domains, but it does not change the simple local-capacity fraction represented by a single failed disk on one host. With sufficient cluster resources and policy compliance, vSAN ESA can continue serving data and begin repair or reprotection using remaining capacity. However, the impacted local host capacity is still the failed device's share of

the host storage pool. The correct percentage is therefore 8.3%, not 25%, 50%, or 0%.

QUESTION 7
An administrator needs to quickly test a possibly destructive change to a Virtual Machine (VM) in production. The VM is currently protected by vSAN Data Protection. Which feature of vSAN Data Protection can be leveraged to achieve this objective?

A. Immutable snapshots
B. Multiple snapshot schedules
C. Protection group
D. Linked clone
E. Replication

Answer: D
Explanation:
Linked clone is the correct feature because the administrator needs to test a potentially destructive change without directly modifying the production VM. vSAN Data Protection uses native vSAN ESA snapshot technology to protect virtual machines through protection groups and scheduled snapshots. Beyond simple restore operations, vSAN Data Protection supports clone workflows that allow a VM to be created from a protected snapshot for development, validation, testing, or recovery use cases. A linked clone is especially appropriate because it can be created quickly from an existing snapshot while maintaining dependency on the base snapshot, avoiding the time and capacity impact of a full independent copy. Immutable snapshots protect recovery points from deletion or modification, but they do not themselves provide an isolated test VM. Multiple schedules only control snapshot timing. A protection group defines which VMs are protected. Replication is used for site-level protection and disaster recovery, not for rapidly testing a destructive change against a local production snapshot.

QUESTION 8
An administrator has been tasked with deploying a new VMware Cloud Foundation (VCF) environment that includes a Management Domain and two workload domains. Compliance regulations require that production and non-production workloads reside in separate failure domains, with the production workload environment using low-latency block storage and the non-production environment relying on high-capacity file-based storage. Which combination of supported non-vSAN storage solutions should the administrator recommend to meet these requirements?

A. Management domain on vVols over NFS, production workload domain on VMFS over iSCSI, and non-production workload domain on SMB
B. Management domain on local VMFS datastores, production workload domain on iSCSI, and non-production workload domain on NFS v3
C. Management domain on vSAN ESA, production workload domain on vSAN HCI Mesh, and non-production workload domain on vSAN File Services
D. Management domain on VMFS over Fibre Channel, production workload domain on VMFS over NVMe/FC, and non-production workload domain on NFS v4.1

Answer: D
Explanation:
Option D is the only design that uses supported non-vSAN storage models while matching the workload characteristics. VMFS over Fibre Channel is a supported enterprise block-storage option for the Management Domain and provides shared, resilient storage with SAN fabric pathing, zoning, LUN masking, and multipathing. Production requires low-latency block storage,

and VMFS over NVMe/FC aligns with that requirement by using NVMe over Fibre Channel transport for high-performance block access. Non-production requires high-capacity file-based storage, and NFS v4.1 is the supported file-based option listed in the design.

QUESTION 9
A financial organization successfully deployed a new Workload Domain in VMware Cloud Foundation (VCF) using NFS storage. All ESX hosts are commissioned and connected to the same storage network. The system administrator is tasked to configure a new NFS datastore in the existing cluster. Which two steps are correct? (Choose two.)

A. Mount the NFS datastore on each ESX host in the cluster.
B. Attach NFS as a secondary datastore.
C. Create an NFS export on the storage array.
D. Configure vSAN on all hosts.
E. Mount the NFS datastore on one ESX host in the cluster.

Answer: AC
Explanation:
The correct steps are to create an NFS export on the storage array and mount the NFS datastore on each ESX host in the cluster. NFS datastore creation begins on the storage system, where the storage administrator creates and exports the NFS volume or directory with the required access permissions for the ESX hosts. After the export exists, vSphere mounts that export as a datastore. For a clustered VCF workload domain, the datastore must be mounted consistently across all ESX hosts in the cluster so that HA, DRS, vMotion, and VM placement can operate correctly. Mounting the datastore on only one host would create inconsistent datastore accessibility and could prevent workloads from moving or restarting on other hosts. "Attach NFS as a secondary datastore" is not the correct procedural step; the datastore is created and mounted as a vSphere datastore. Configuring vSAN is irrelevant because the workload domain already uses NFS storage and the task is to add another NFS datastore.

QUESTION 10
An administrator wants to protect a site from disaster and re-protect after fail-over. Industry regulations require to change the primary site every six months. Which two options are required during the inventory mapping and while running the recovery plans to fulfill this requirement? (Choose two.)

A. Once fail-over is successful, reconfigure vSAN data protection to reverse replication.
B. Use "Prepare reverse mappings" during inventory mapping creation.
C. Use a test recovery plan.
D. Once the fail-over is successful, re-protect from VMware Live Site Recovery.
E. Extend the network between both sites.

Answer: BD
Explanation:
The requirement is to alternate the primary site every six months, which requires a disaster recovery design that can reverse protection direction after failover. Inventory mappings define how resources such as compute resources, folders, networks, and datastores map between the protected and recovery sites. Using "Prepare reverse mappings" during inventory mapping creation ensures that the same mapping structure can be used when the recovery site later becomes the protected site. After a failover or planned migration succeeds, VMware Live Site Recovery uses the re-protect operation to reverse the protection relationship. This makes the formerly recovered site the new protected source and prepares replication and recovery plans for

failback or for the next scheduled site rotation. A test recovery plan validates recovery without changing the primary site. Extending the network can simplify workload mobility but is not required for reprotection. Reconfiguring vSAN Data Protection manually is not the Live Site Recovery workflow for reversing protected and recovery site roles.

QUESTION 11
An administrator is tasked with deploying a vSAN ESA Stretched Cluster for a VMware Cloud Foundation (VCF) workload domain. The stretched cluster must meet the following requirements:

```
- Ensure the cluster supports a storage policy with a secondary level
of resilience of FTT=2 using RAID-6 erasure coding.
- Allow vSAN to regain its prescribed level of resilience in the event
of a sustained host outage.
```

What is the minimum number of hosts required in each availability zone?

A. 6
B. 14
C. 7
D. 12

Answer: C
Explanation:
Seven hosts per availability zone is required for this design. In a vSAN ESA stretched cluster, data components are distributed across two data sites, while a witness site provides quorum. The first requirement is a secondary level of resilience of FTT=2 using RAID-6 erasure coding. RAID-6 requires enough fault domains to place data and parity components while tolerating two failures. The second requirement is more restrictive: vSAN must be able to regain its prescribed level of resilience after a sustained host outage. This means the site must have enough remaining placement capacity to rebuild components and return objects to policy compliance after a host is lost. A six-host site can satisfy the RAID-6 placement minimum, but it does not provide the additional host needed to rebuild cleanly after a sustained host outage. Therefore, the design requires seven hosts in each availability zone. This produces fourteen data hosts across the stretched cluster, plus the witness.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

★ More than **99,900** Satisfied Customers Worldwide.

★ Average **99.9%** Success Rate.

★ **Free Update** to match latest and real exam scenarios.

★ **Instant Download** Access! No Setup required.

★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.

★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.

★ **100%** Guaranteed Success or **100%** Money Back Guarantee.

★ **Fast**, helpful support **24x7**.

View list of all certification exams: http://www.lead2pass.com/all-products.html

10% Discount Coupon Code: **ASTR14**