TestKingfree

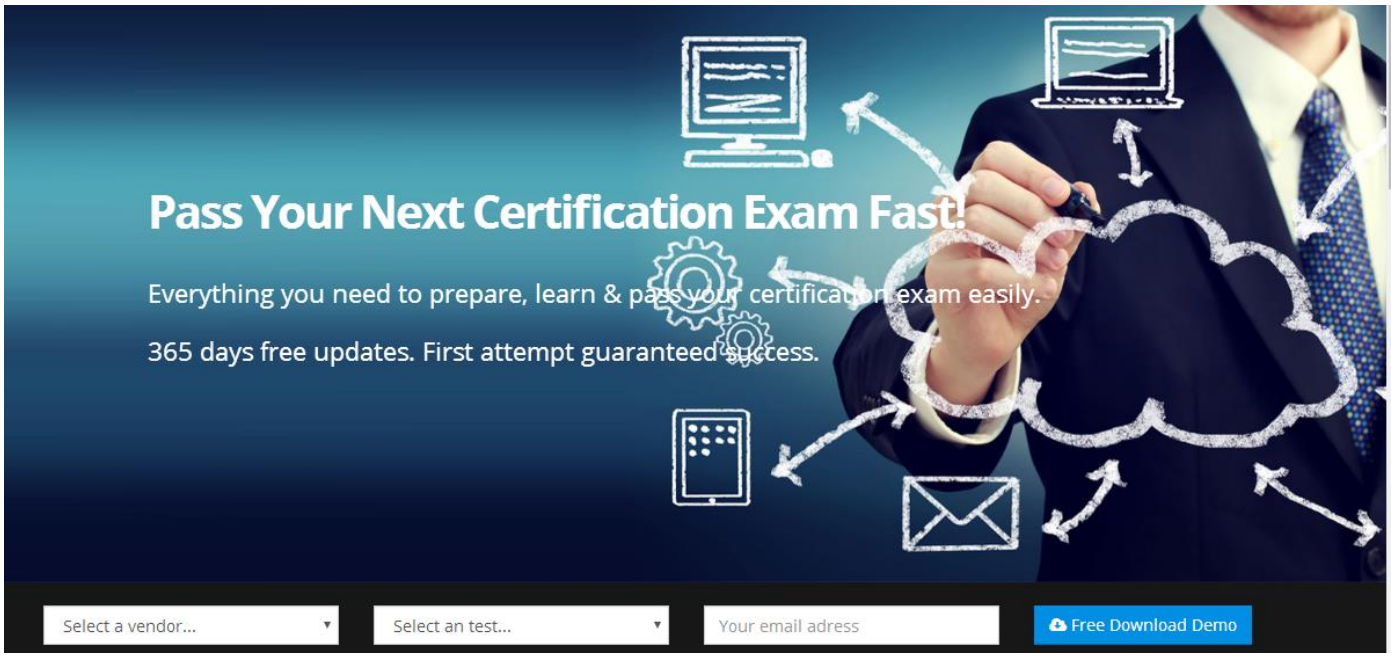

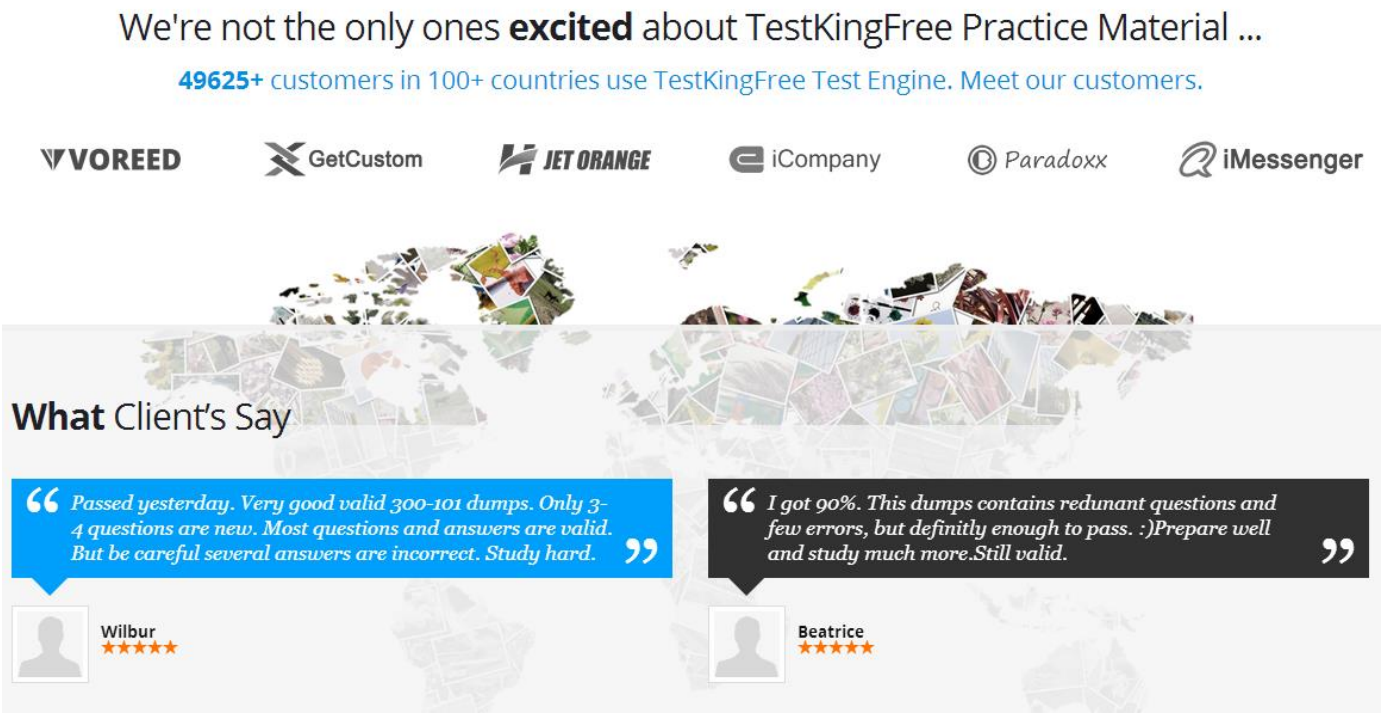

http://www.testkingfree.com/

Pass For sure Certification Exam Guide and Exam Dumps - TestKingFree

Exam : **3V0-23.25**

Title : Advanced VMware Cloud Foundation 9.0 Storage

Vendor : VMware

Version : DEMO

NO.1 A cache drive failed on one of the vSAN OSA nodes in the cluster.
When the drive failed, vSAN started a resync to ensure the health of the data, and all objects are showing a healthy and compliant state.
The vSAN administrator needs to replace the failed cache drive.
Which set of steps should the vSAN administrator take?

A. Place the disk group into maintenance mode, and select Full Data Migration. Then, physically replace the failed cache device. Afterwards, vSAN will rebuild the disk group automatically.

B. Physically replace the failed cache device, and vSAN will automatically allocate the storage.

C. Remove the existing vSAN disk group, and physically replace the device. Then, check to verify that the ESX host automatically detects the new device. Afterwards, manually recreate the Disk Group.

D. Physically replace the failed cache device, and vSAN will automatically create a new disk group. Then, remove the disk group with the failed device.

Answer: C
Explanation:
In vSAN Original Storage Architecture, a disk group consists of one flash cache device and one or more capacity devices. The cache device is a required component of the disk group, so a failed cache device cannot be replaced as an isolated drive while preserving the same disk group. The supported operational approach is to remove the affected disk group, physically replace the failed cache device, verify that ESX detects the new device, and then manually recreate the disk group using the replacement cache device and the appropriate capacity devices. The scenario states that vSAN has already completed resynchronization and that all objects are healthy and compliant, so removing the failed disk group no longer risks object noncompliance. Selecting Full Data Migration on a failed cache device is not the correct workflow because the cache failure affects the entire disk group. Simply inserting a replacement drive does not cause vSAN OSA to automatically rebuild the disk group.
Reference topics: vSAN OSA Disk Groups, Replace a Cache Device, Remove Disk Group, Recreate Disk Group.

NO.2 An administrator notices alerts triggering IOPS and Disk Throughput storage performance problems in the Fibre Channel datastore in a VMware Cloud Foundation (VCF) Workload Domain.
What can the administrator review to identify which Virtual Machines (VMs) may be experiencing storage IOPS and disk throughput performance issues?

A. vSAN Health dashboard.

B. vSphere Storage Inventory dashboard.

C. Live! vSphere Heavy Hitter VM dashboard.

D. Storage Operations page.

Answer: C
Explanation:
The Live! vSphere Heavy Hitter VM dashboard is the correct place to identify which virtual machines are driving storage IOPS and disk throughput problems. In VCF Operations, heavy hitter dashboards are designed to isolate the consumers creating the highest demand on shared infrastructure resources.
For storage troubleshooting, this dashboard identifies VMs producing significant I/O load, including IOPS and throughput. This is especially useful when a Fibre Channel datastore is reporting performance alerts because the datastore itself may show contention, but the root cause is often one or more high-demand VMs consuming disproportionate backend storage resources. The vSAN Health

dashboard is not correct because the affected datastore is Fibre Channel, not vSAN. The vSphere Storage Inventory dashboard is useful for topology and inventory visibility, but it is not the best view for identifying the VM-level consumers causing IOPS and throughput pressure. The Storage Operations page is broader and less targeted than the Live heavy hitter view. Reference topics: VCF Operations Performance Dashboards, Storage Heavy Hitters, VM IOPS and Throughput Troubleshooting.

NO.3 Select the storage capabilities supported for use with persistent volumes in the VMware Kubernetes Service (VKS) in VMware Cloud Foundation (VCF).
Drag and drop the five supported capabilities from the vSphere Storage Capabilities list on the left and place them into the Supported Storage Capabilities list on the right in any order. (Choose five.)

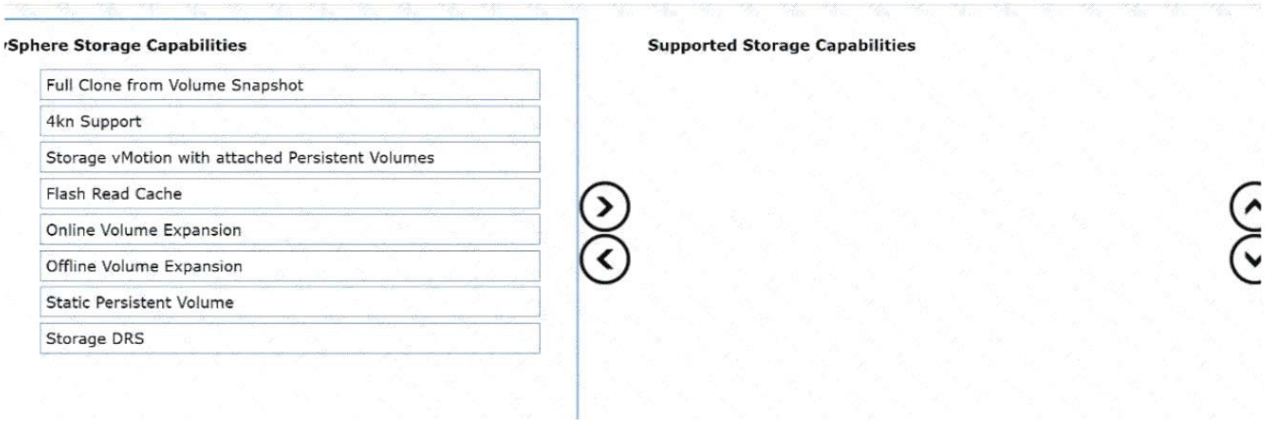

Answer:

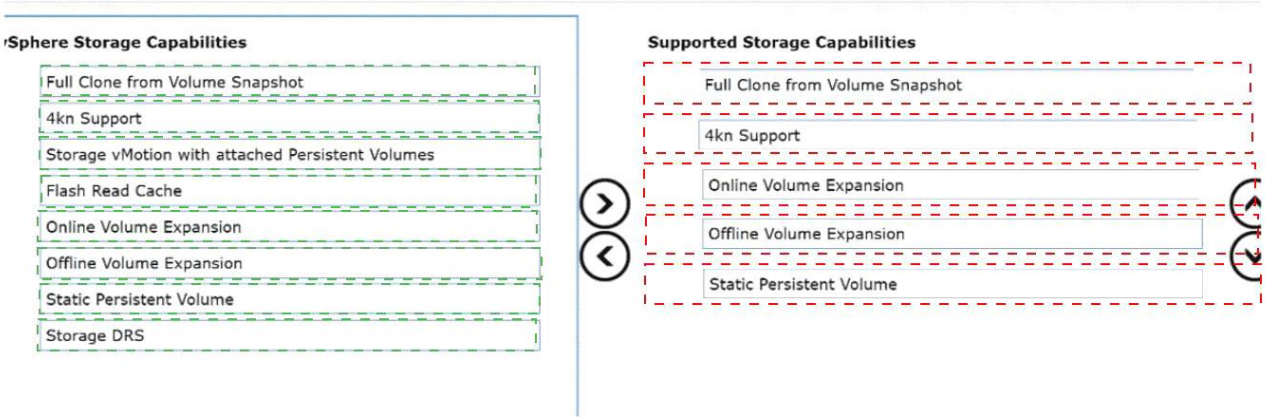

Explanation:
Full Clone from Volume Snapshot
4kn Support
Online Volume Expansion
Offline Volume Expansion
Static Persistent Volume
VMware Kubernetes Service persistent volumes use vSphere storage integrations such as CNS and pvCSI to provide Kubernetes storage services backed by vSphere datastores and storage policies. The supported persistent volume capabilities include Static Persistent Volume, which allows pre-existing storage to be consumed by Kubernetes workloads; Offline Volume Expansion and Online Volume Expansion, which allow persistent volumes to be resized depending on workload and filesystem state; Full Clone from Volume Snapshot, which enables a new volume to be created from an existing snapshot; and 4kn Support, which supports modern 4K native storage devices where the underlying

stack supports them. Storage vMotion with attached Persistent Volumes is not supported because moving an attached Kubernetes persistent volume while it is in active use can break attachment and consistency assumptions. Storage DRS is also not a supported persistent volume capability in this VKS context. Flash Read Cache is not part of the supported VKS persistent-volume capability set.
Reference topics: VMware Kubernetes Service Storage, CNS, pvCSI, Persistent Volume Capabilities, Volume Expansion, Volume Snapshot Clone.

NO.4 An administrator is tasked with deploying a VMware Cloud Foundation (VCF) Workload Domain that meets the following requirements:
* vSAN ESA as principal storage
* RAID-6 with FTT=2
* Support for Storage Traffic Separation
The administrator is provided the following hardware to perform the task:
* Four ESX hosts, each host contains:
* 24 CPU cores
* 96 GB memory
* Two 25GbE network NICs
* 12 NVMe devices 4 TB each, connected to a single SATA/SAS/NVMe Tri-mode controller What four changes must the administrator make to the hardware before deploying the new Workload Domain? (Choose four.)

A. Increase the ESX host count to a minimum of seven.

B. Increase the ESX host count to a minimum of six.

C. Increase the CPU quantity on each host to a minimum 32.

D. Increase the Tri-mode controller quantity on each host to two, with six NVMe devices connected to each.

E. Increase the network NICs on each host to minimum of four 25 GbE network NICs.

F. Replace the network NICs on each host to a minimum of two 100 GbE network NICs.

G. Increase the memory on each host to a minimum 128 GB.

H. Replace the Tri-mode controller on each host with a dedicated NVMe controller.

Answer: B E G H
Explanation:
The design requires four hardware corrections. First, RAID-6 with FTT=2 requires six fault domains or hosts because RAID-6 erasure coding must place data and parity components across enough independent hosts to tolerate two failures. Therefore, four hosts are insufficient and the host count must increase to at least six.
Second, Storage Traffic Separation requires sufficient physical NICs to support separate storage and non- storage traffic paths. With two distributed switches and redundant uplinks, the design must increase to four 25 GbE NICs per host. Third, vSAN ESA requires at least 128 GB of host memory; the provided 96 GB is below the minimum. Fourth, NVMe devices attached to a SATA/SAS/NVMe Tri-mode controller are not supported for vSAN storage pools. The documentation states that vSAN supports SAS and SATA devices on Tri-mode controllers, but NVMe devices must be directly connected to PCIe or use a supported NVMe design. The CPU count does not need to increase because 24 cores satisfies the stated vSAN ESA minimum in the referenced design guidance.
Reference topics: vSAN ESA Hardware Requirements, RAID-6 FTT=2, Storage Traffic Separation, NVMe Controller Support.

NO.5 A storage administrator is being presented with the following VMware Cloud Foundation (VCF) architectural details:
* The applications data will require 2.5 PB of capacity.
* The production applications will be hosting archival solutions and gateways.
* There will be some applications deployed for the purpose of testing and development.
What is the optimal principal storage that the administrator can recommend?

A. vSAN ESA Storage Clusters

B. vSAN ESA

C. vSAN OSA - All Flash

D. vSAN OSA - Hybrid

Answer: A

Explanation:
vSAN ESA Storage Clusters are the optimal choice because the workload profile is capacity-heavy and benefits from disaggregated storage. The requirement for 2.5 PB of application data, archival workloads, gateways, and development/test applications indicates a need to scale storage capacity independently from compute resources. VMware Cloud Foundation describes a vSAN storage cluster as a disaggregated storage solution based on vSAN ESA. It provides storage resources but not compute resources, and its datastore can be mounted by vSAN compute clusters or vSAN HCI clusters. This model is designed for storage-dense environments where capacity and performance can be expanded without adding unnecessary compute. A standard vSAN ESA HCI cluster is excellent for high- performance workloads, but it scales compute and storage together. vSAN OSA all-flash and hybrid use the older disk-group architecture and are less suitable for this large-capacity, storage-centric design. Therefore, the principal storage recommendation should be vSAN ESA Storage Clusters to meet scalability, performance, and capacity growth requirements. Reference topics: vSAN Storage Cluster, Disaggregated Storage, vSAN ESA, Storage Models.

NO.6 A vSAN ESA cluster experienced a host NVMe failure, causing several objects to become non-compliant.
Select the steps vSAN ESA follows to detect, evaluate and repair the non-compliant objects.
Drag and drop the four correct options from the Options list on the left and place them into the Required Options on the right in any order. (Choose four.)

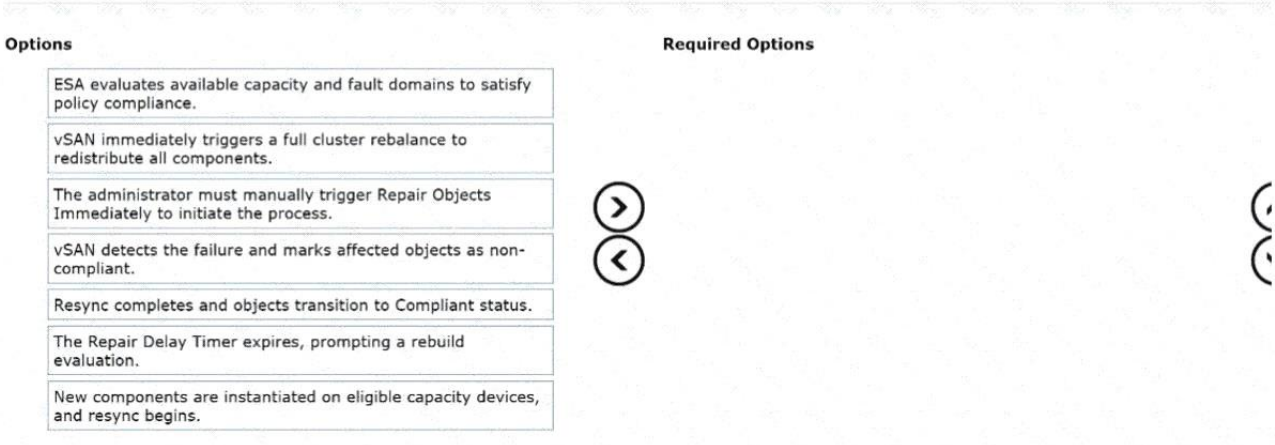

Answer:

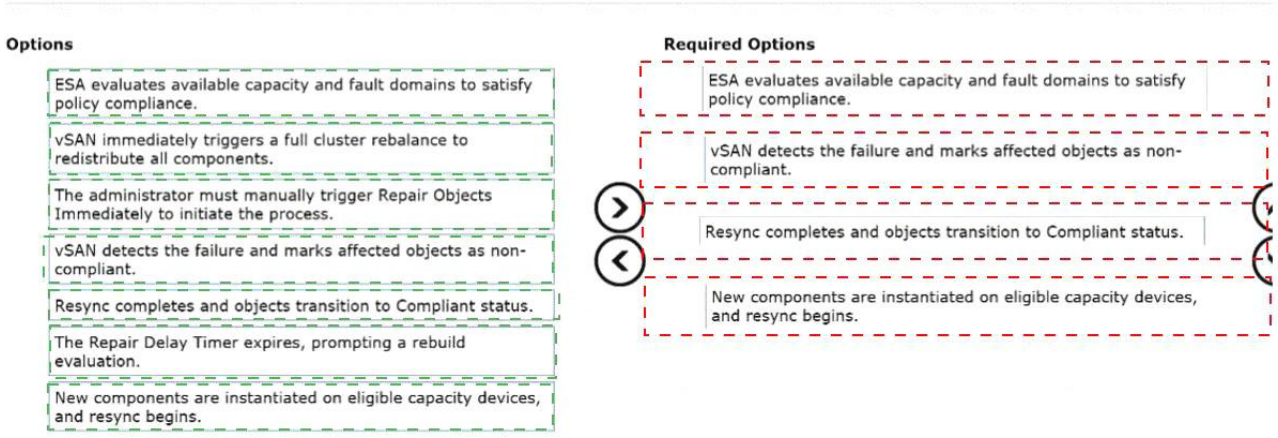

Explanation:
ESA evaluates available capacity and fault domains to satisfy policy compliance.
vSAN detects the failure and marks affected objects as non-compliant.
Resync completes and objects transition to Compliant status.
New components are instantiated on eligible capacity devices, and resync begins.
A failed NVMe device in a vSAN ESA cluster is handled as a storage pool device failure. vSAN first detects the failure and marks affected objects as non-compliant when their assigned storage policy requirements are no longer satisfied. ESA then evaluates available capacity and fault domains to determine whether replacement components can be placed in a way that restores compliance. If sufficient placement resources exist, vSAN creates new components on eligible capacity devices and begins resynchronization. After the resync completes, the affected objects transition back to Compliant status. The Repair Delay Timer is primarily associated with absent components, such as a host or device that may return, and is not the required step for a permanent NVMe failure. The administrator does not have to manually trigger Repair Objects Immediately for the normal automatic repair process. A full cluster rebalance is separate from repair and is not performed before rebuilding missing or degraded components. Reference topics: vSAN ESA Storage Pool Device Failure, Object Compliance, Automatic Reprotection, Resynchronization.

NO.7 A six-node vSAN ESA cluster contains multiple Virtual Machines (VMs), and a vSAN storage policy with the rule "Failures to tolerate" set to "1 failure - RAID-5 (Erasure Coding)" is assigned. A vSAN administrator has changed the rule in the assigned policy to "2 failures - RAID-6 (Erasure Coding)." What is the result of this change?

A. No changes occur until the policy is reapplied.

B. The policy change is rejected immediately.

C. The updated policy is serially applied to the Virtual Machines.

D. The changes are queued for 60 minutes.

Answer: A

Explanation:
The policy edit is valid for a six-node vSAN ESA cluster because RAID-6 with Failures to tolerate set to 2 requires at least six ESX hosts in the vSphere cluster. However, editing a VM storage policy that is already associated with virtual machine objects does not automatically communicate the new requirements to those objects. vSphere marks the compliance status as Out of Date, indicating that the policy has been edited but the new requirements have not yet been communicated to the datastore where the VM objects reside. The administrator must reapply the VM Storage Policy to the

affected objects before vSAN evaluates the new RAID-6 requirement and begins any layout changes or resynchronization needed to bring objects into compliance. The change is not rejected because the cluster satisfies the minimum host count. It is not automatically applied serially to the VMs, and the 60- minute timer is related to repair-delay behavior after failures, not normal storage policy changes. Reference topics: Reapply Virtual Machine Storage Policy, Out-of-Date Compliance, vSAN ESA RAID-5, vSAN ESA RAID-6.

NO.8 An administrator is tasked with enabling vSAN Data Protection.
Which action is required to enable vSAN Data Protection?

A. Enable vSAN advanced options

B. Deploy VMware Live Recovery (VLR)

C. Deploy Data Services Manager

D. Deploy vSphere Replication

Answer: B

Explanation:
The required action is to deploy the VMware Live Recovery appliance. In VMware Cloud Foundation 9.0, vSAN Data Protection is powered by VMware Live Recovery and provides local virtual machine protection and remote virtual machine replication for supported vSAN ESA environments. The vSAN Services configuration workflow explicitly states that before vSAN Data Protection can be used, the VMware Live Recovery appliance must be deployed. This appliance provides the data protection control plane required for protection groups, native snapshots, recovery workflows, and replication-based protection. Enabling vSAN advanced options is not sufficient because those settings control cluster behaviors such as repair timers, read locality, thin swap, unmap, and rebalance. Data Services Manager is used for database/data-service lifecycle use cases, not vSAN VM data protection. vSphere Replication can be part of broader site-recovery scenarios, but it is not the required action to enable vSAN Data Protection from the vSAN services workflow. Reference topics: vSAN Data Protection, VMware Live Recovery Appliance, vSAN Services Configuration, Local VM Protection and Replication.