KTest

更に上のクオリティ　更に上のサービス

問題集

http://www.ktest.jp

1年で無料進級することに提供する

Exam : **3V0-23.25**

Title : Advanced VMware Cloud Foundation 9.0 Storage

Version : DEMO

1.An L3 Support Engineer is troubleshooting a severe performance degradation and partial VM unavailability in a VCF Workload Domain configured with a vSAN Stretched Cluster. The cluster utilizes the "Witness Traffic Separation" (WTS) feature.
The engineer pulls the vmkernel.log from host-sec-01 in the Secondary fault domain:
```
2026-05-12T14:15:22.456Z INFO vsanmgmt - Entering Stretched Cluster Health Check
2026-05-12T14:15:30.112Z WARN vsan-network [vmk2:vSAN-Data] Failed to ping Preferred-Gateway: Destination unreachable
2026-05-12T14:15:35.889Z INFO vsan-network [vmk3:Witness-Traffic] Ping to Witness-Appliance (10.50.1.10) Successful
2026-05-12T14:15:36.001Z ERROR cmmds - Cluster partition detected. Secondary site isolated from Preferred site.
2026-05-12T14:15:38.220Z INFO clom - Object 5543505c-xxxx entering DEGRADED state.
2026-05-12T14:15:40.500Z WARN vsan-network [vmk2:vSAN-Data] High congestion detected on ISL. TxQueue=100%
2026-05-12T14:15:45.000Z ERROR vobd - [vSAN] Node uuid-sec-01 has lost communication with Witness node uuid-wit-01 via WTS network.
2026-05-12T14:15:45.500Z ERROR cmmds - Component state for Object 5543505c changed to INACCESSIBLE.
```
Based on the logs and the integration between Fault Domains and Witness Traffic Separation, which of the following statements explain the root cause and system behavior? (Select all that apply.)
A. The Secondary fault domain successfully maintained quorum via vmk3 at 14:15:35, preventing the VMs from entering an inaccessible state at that moment.
B. The congestion on vmk2 indicates that storage I/O traffic is incorrectly being routed over the Witness Traffic Separation network.
C. The Dual Site Mirroring policy requires the Secondary site to maintain connectivity to either the Preferred site OR the Witness to keep objects accessible.
D. A subsequent failure of the Witness Traffic Separation network (vmk3) at 14:15:45 caused the Secondary site to lose its tie-breaker vote.
Answer: A, C, D

2.A Network Administrator is troubleshooting a newly deployed vSAN Witness Appliance that cannot join the Stretched Cluster CMMDS network.
The administrator queries the Witness Appliance network adapters via SSH:
```
[root@witness-01:~] vim-cmd hostsvc/net/vnic_info
vmk0: 10.10.1.15 (Traffic: Management)
vmk1: 172.16.50.15 (Traffic: vSAN Witness)
[root@witness-01:~] esxcfg-route -l
Network          Netmask          Gateway          Interface
default          0.0.0.0          10.10.1.1        vmk0
```
The ESXi data hosts exist on the 192.168.100.0/24 subnet. Pings from vmk1 to the data hosts fail.

What is the specific missing configuration causing this network partition?
A. A static route is missing; because vmk1 is on a different subnet than the data hosts, the Witness is trying to route the vSAN traffic through the vmk0 default gateway (Management), which violates network isolation.
B. The administrator failed to tag vmk1 with the "vMotion" traffic type, which is required for Witness replication.
C. The ESXi data hosts must be configured with "vSAN Direct" to bypass the gateway and establish a Layer 2 tunnel to vmk1.
D. The Witness Appliance requires dual vmk adapters for vSAN traffic configured in an Active/Active LACP bond to process heartbeats.
Answer: A

3.A Compliance Auditor is tracking the success of an automated "Shallow Rekey" task scheduled across a massive VCF 9.0 multi-cluster environment. The task failed on a specific vSAN Stretched Cluster.
```
[Skyline Health > vSAN > Encryption Health]
Status: Warning
Message: "KMS Server unreachable on Host esx-04. Rekey Aborted."
[Architecture Details]
esx-04 is part of the Secondary Site. The Inter-Site Link is currently DOWN (Partition).
```
How does the vSAN encryption architecture prevent data loss and split-brain when a Rekey operation hits a partitioned cluster? (Choose 2.)
A. esx-04 will instantly cryptographically shred its local drives to prevent data compromise during the network partition.
B. Even though the Rekey failed, the virtual machines on the surviving site remain fully operational because the ESXi hosts maintain the *current* KEK cached in their secure RAM, requiring no active KMS connection to serve I/O.
C. esx-04 will automatically fallback to the local Witness appliance to generate a temporary KEK until the network is restored.
D. The DOM Client forces esx-04 to perform a Deep Rekey using the vSphere TPM chip to bypass the KMS outage.
E. The Shallow Rekey operation is strictly an atomic transaction; if esx-04 cannot reach the KMS to receive the new Key Encryption Key (KEK), the vCenter master node rolls back the KEK on all other hosts to ensure cluster-wide key consistency.
Answer: B, E

4.A VCF Deployment Specialist is investigating a localized physical drive failure in two separate VCF domains: Domain A (vSAN OSA with Dedupe Enabled) and Domain B (vSAN ESA).
In both domains, a single 3.84 TB Capacity SSD/NVMe drive has suffered a "Permanent Device Loss" (PDL).
```
[RVC Output: vsan.disks_stats Domain A (OSA)]
Failed: naa.500A... (Capacity Tier)
```

[RVC Output: vsan.disks_stats Domain B (ESA)]
Failed: naa.500B... (Storage Pool)
```

Based on the architectural implementation of deduplication and the filesystem structure, which TWO statements accurately contrast the failure blast radius in these environments? (Choose 2.)
A. Domain B (ESA) will suffer a total host failure because the Log-Structured filesystem cannot isolate single NVMe failures.
B. In Domain A (OSA), the loss of a single capacity drive in a deduped disk group invalidates the entire deduplication hash table; vSAN must fail the ENTIRE disk group (including the cache and all other healthy capacity drives) and rebuild the data across the network.
C. Domain A (OSA) will rebuild faster because deduplication pointers are automatically remapped to the remaining capacity drives in the group without requiring network resynchronization.
D. In Domain B (ESA), because Deduplication is eliminated and the Storage Pool is flat, the loss of the single NVMe drive ONLY affects the components physically stored on that specific drive; the other drives remain active.
E. Both architectures experience the exact same failure domain (loss of 3.84 TB), as deduplication state is maintained independently inside the ESXi RAM.
**Answer:** B, D

5.An Infrastructure Manager is planning to scale out a massive VCF 9.0 Workload Domain. The target host utilizes a dense storage configuration with multiple physical storage controllers.
[Architecture Diagram: Dense Storage Host]
- Controller 1: Adaptec SmartRAID 3154 (Pass-Through) -> 8x SAS HDD
- Controller 2: Broadcom 3908 (Hardware RAID-0) -> 8x SATA SSD
- Direct PCIe Bus: 2x NVMe Drives
Which of the following statements correctly evaluate how the vSAN validation logic processes this specific dense hardware topology during the SDDC Manager commissioning phase? (Select all that apply.)
A. If the host is targeted for vSAN ESA, the 8x SAS HDD and 8x SATA SSD drives will be completely ignored, and SDDC Manager will only validate the 2x NVMe drives connected to the PCIe bus.
B. The presence of SAS HDDs permanently disqualifies the host from SDDC Manager, as VCF requires all storage to be 100% flash.
C. Validation will fail because vSAN strictly prohibits mixing different storage controller vendors (Adaptec and Broadcom) inside the same ESXi host.
D. SDDC Manager will intelligently pair the 2x NVMe drives as Cache and the 8x SATA SSDs as Capacity to form an All-Flash vSAN OSA cluster, ignoring the HDD controller.
E. Validation will immediately fail on Controller 2 because vSAN strictly requires Host Bus Adapters (HBAs) to run in Pass-Through (JBOD) mode; Hardware RAID-0 creates a false single-disk abstraction that blinds vSAN to physical disk health.
**Answer:** A, D, E
```