



Vendor: Network Appliance

Exam Code: NS0-522

Exam Name: NetApp Implementation Engineer - SAN
ONTAP Exam

Version: DEMO

QUESTION 1

You are troubleshooting host connectivity issues in a newly deployed ONTAP NVMe/FC SAN environment. Several hosts cannot discover NVMe namespaces, while other hosts on the same FC switch are working correctly. You have verified that the FC switch zoning is correct and NVMe services are running on the ONTAP cluster. Which additional troubleshooting steps should you perform? (Select all that apply)

- A. Verify that the NVMe initiator drivers on the affected hosts are properly installed and up to date
- B. Check the subsystem mapping configuration on ONTAP to ensure the affected host WWNs are mapped to the NVMe subsystem
- C. Confirm that NVMe LIFs are created on the appropriate nodes and are in the UP state
- D. Shut down the FC switch and restart it to reset all configurations automatically
- E. Delete all existing NVMe namespaces and recreate them to resolve discovery issues

Answer: ABC

Explanation:

When troubleshooting selective NVMe discovery failures, you must verify multiple layers: the host-side initiator drivers must be current and properly installed, the ONTAP subsystem mapping must include the affected host WWNs, and the ONTAP NVMe LIFs must exist and be in an UP state. These are standard troubleshooting steps that address the most common causes of connectivity issues. Restarting the FC switch is disruptive and would affect all hosts, not just the affected ones, and is not a standard troubleshooting approach for selective failures. Deleting and recreating namespaces would cause data loss and service disruption without addressing the underlying connectivity problem. The scenario states zoning is already correct, so the issue is in host configuration, driver state, or ONTAP subsystem mapping.

QUESTION 2

You are implementing a Snapshot-based backup strategy for a customer's ONTAP SAN environment with multiple business-critical LUNs. The customer requires point-in-time recovery capability with minimal impact to production workloads. How should ONTAP Snapshots be configured in this scenario?

- A. Create a Snapshot policy with multiple scheduled intervals (e.g., hourly, daily, weekly) and apply it to the volumes containing the LUNs
- B. Create manual Snapshots only when the customer requests them to avoid overhead
- C. Configure Snapshots to run only during maintenance windows to prevent any production impact
- D. Create Snapshots on a separate cluster and replicate all production data every 30 minutes
- E. Disable Snapshots entirely and use only external backup tools for data protection

Answer: A

Explanation:

ONTAP Snapshots are designed to provide data protection with minimal performance impact on production workloads. A well-designed Snapshot policy uses scheduled intervals at multiple granularities (hourly for recent recovery points, daily and weekly for longer retention) applied to volumes containing the LUNs. This provides point-in-time recovery capability without requiring manual intervention and with negligible performance overhead--Snapshots are read-only, space-efficient copies that do not consume significant resources. Manual- only Snapshots are unreliable and lack consistent coverage. Limiting Snapshots only to maintenance windows wastes valuable protection during peak business hours. Replicating all data every 30 minutes to another cluster introduces unnecessary complexity and cost; Snapshots are sufficient for backup without requiring full replication. Disabling Snapshots entirely eliminates a core ONTAP data protection feature and would force reliance on slower external tools.

QUESTION 3

You are designing an ONTAP SAN solution for a customer with 500 TB of storage capacity requirements and strict RPO (Recovery Point Objective) of 4 hours. The customer's budget limits them to a two-node cluster. During the design phase, which factors must you consider regarding storage space utilization and redundancy? (Select all that apply)

- A. In a two-node cluster, storage is mirrored between nodes, so usable capacity is approximately 50% of total raw capacity
- B. RAID-DP protection overhead must be factored in, reducing usable space by additional percentage
- C. Snapshot reserve space should be allocated based on change rate and RPO requirements
- D. A two-node cluster can safely provision 100% of capacity as usable space with no redundancy overhead
- E. The 500 TB requirement means you need 500 TB of raw disk capacity to achieve 500 TB of usable space

Answer: ABC

Explanation:

In a two-node ONTAP cluster, data is mirrored between nodes for high availability, reducing usable capacity to approximately 50% of total raw capacity. Additionally, RAID-DP (ONTAP's default RAID protection) on each node adds further overhead (typically 7-8% depending on drive count and type). Snapshot reserve space must be allocated to maintain the 4-hour RPO--this reserve space reduces available user-provisioning capacity. Therefore, to deliver 500 TB of usable space for LUNs, you must calculate total raw capacity accounting for mirroring, RAID-DP, and Snapshot reserves. A two-node cluster absolutely requires mirroring for availability and fault tolerance--you cannot safely use 100% raw capacity for LUNs. The raw capacity requirement is significantly larger than the usable requirement when factoring in all protection and redundancy mechanisms. Proper capacity planning is essential in the solution design phase.

QUESTION 4

You are performing post-deployment validation testing on a newly implemented ONTAP iSCSI SAN environment. Hosts are connected and can access LUNs, but you notice inconsistent latency and occasional I/O timeouts during peak usage. You have verified that switch bandwidth is sufficient and ONTAP performance metrics show normal CPU and memory utilization. Which additional testing and troubleshooting steps are most relevant to this scenario?

- A. Check iSCSI initiator connection count and verify that I/O sessions are distributed across multiple iSCSI LIFs for load balancing
- B. Verify iSCSI LIF MTU settings match host network configuration and confirm no packet fragmentation is occurring
- C. Review ONTAP event logs and check for any LUN or volume space warnings that could cause I/O delays
- D. Immediately replace all iSCSI LIFs with Fibre Channel LIFs to eliminate protocol latency
- E. Add more nodes to the cluster to automatically redistribute iSCSI traffic without any configuration changes

Answer: ABC

Explanation:

When troubleshooting latency and I/O timeouts in an iSCSI environment with normal cluster metrics, focus on session distribution and network configuration. Verify that iSCSI initiators are connected to multiple iSCSI LIFs for proper load balancing--if all sessions concentrate on one LIF, it can become a bottleneck. MTU (Maximum Transmission Unit) mismatch between ONTAP

LIFs and host network configuration causes packet fragmentation and latency. Space warnings on LUNs or volumes trigger performance safeguards that increase latency. These testing steps target the most common causes of inconsistent performance in properly provisioned environments. Switching protocols or adding nodes without understanding the root cause is inefficient and would not address the actual problem. Cluster expansion does not automatically optimize iSCSI load balancing--explicit LIF configuration and initiator multipathing settings control this behavior.

QUESTION 5

You are planning a NetApp SAN environment for a large financial services customer. The customer has two data centers in different geographic locations and requires that all critical data be mirrored synchronously between sites for disaster recovery purposes. During the assessment phase, you discover that the WAN link between data centers has 500 ms latency and 50 Mbps bandwidth, but the customer's SAN applications require sub-100 ms round-trip response times. Which of the following represents the best approach to address this constraint?

- A. Implement synchronous SnapMirror between the data centers; the 500 ms latency is acceptable for financial services workloads
- B. Recommend asynchronous SnapMirror replication with a defined RPO that meets business requirements, and document the latency constraint as a limiting factor for synchronous mirroring
- C. Deploy a third ONTAP cluster at a third location with lower latency to act as an intermediate replication point
- D. Configure MetroCluster synchronous replication, which is optimized for high-latency WAN links and will automatically compensate for network delays

Answer: B

Explanation:

The correct answer is to recommend asynchronous SnapMirror replication and document the latency constraint. During solution assessment and design, you must evaluate whether the customer's infrastructure can support their stated requirements. Synchronous replication (option A) is inappropriate because synchronous write acknowledgment across 500 ms latency links would violate the sub-100 ms RTT requirement. MetroCluster (option D) actually requires low-latency links and is not suitable for high-latency WAN scenarios. A third cluster (option C) is an unnecessary complication. The correct approach is to align the replication method with actual network capabilities and ensure the customer understands the trade-off between synchronous protection and application performance, documenting this as part of the solution design.

QUESTION 6

A retail customer is implementing a new SAN infrastructure using ONTAP 9.14. They plan to consolidate storage from multiple legacy systems. During the inventory and prerequisites phase, you need to assess their current fabric environment. The customer has existing Cisco MDS 9700 series switches and plans to use both 16 Gbps and 32 Gbps Fibre Channel connections for different host groups. Which two statements are correct regarding FCP fabric planning in this scenario?

- A. The customer should upgrade all connections to 32 Gbps to ensure consistent performance across all hosts
- B. Mixed speed fabrics (16 Gbps and 32 Gbps) can operate together, and hosts using 16 Gbps connections will not negatively impact 32 Gbps hosts if properly configured
- C. SAN fabric best practices allow up to 40% utilization per ISL to maintain headroom for spikes and ensure quality of service
- D. The MDS 9700 series must be completely replaced with newer models to support ONTAP 9.14 cluster connectivity

- E. Zone planning becomes unnecessary when using modern MDS switches with automatic zoning features

Answer: BC

Explanation:

The two correct answers are that mixed-speed fabrics can coexist without negative impact when properly configured (option B), and that SAN fabric best practices recommend maintaining approximately 40% utilization on ISLs (option C). Modern Fibre Channel fabrics support mixed speeds, allowing heterogeneous environments where some hosts connect at 16 Gbps while others use 32 Gbps without interference. Option A is incorrect because upgrading all connections may not be cost-justified and is not necessary for operation. Option D is incorrect because the MDS 9700 series is a modern platform fully compatible with ONTAP 9.14. Option E is incorrect because zoning remains essential regardless of switch capability--it is a fundamental security and isolation practice in SAN design.

QUESTION 7

You are configuring a new SVM (Storage Virtual Machine) for an ONTAP cluster that will serve iSCSI and NVMe/TCP traffic to a virtualization environment. The cluster has two nodes with HA pairs configured, and you need to ensure optimal data path distribution and redundancy. When creating the network interfaces for the SAN protocols, multiple physical interfaces are available on each node. What is the recommended approach for configuring SAN LIFs (Logical Interfaces) across the cluster nodes in this scenario?

- A. Create all SAN LIFs on a single home node to simplify failover and reduce the complexity of multi-path configurations
- B. Create SAN LIFs on both nodes with each LIF assigned a specific home node and home port, distributing LIFs across nodes to balance load and maintain redundancy if a node fails
- C. Create only iSCSI LIFs on one node and NVMe/TCP LIFs on the other node to segregate protocol traffic by node
- D. Create SAN LIFs only on the primary node and do not assign failover locations, allowing ONTAP to handle failover dynamically

Answer: B

Explanation:

The correct answer is to create SAN LIFs on both nodes with each LIF having a defined home node and home port, distributing them for load balancing and redundancy (option B). This is a fundamental SAN configuration best practice. When SAN LIFs are distributed across cluster nodes, traffic is load-balanced, and if one node fails, the HA partner can serve the LIF. Option A is incorrect because concentrating all LIFs on one node creates a single point of failure and wastes cluster capacity. Option C is incorrect because segregating protocols by node gains no operational advantage and complicates the configuration. Option D is incorrect because SAN LIFs require explicit home nodes and failover locations; dynamic failover without assigned locations is not a recommended practice for SAN protocols and can lead to unpredictable behavior.

QUESTION 8

During implementation testing of a new NVMe/FC SAN solution in ONTAP 9.14, you notice that host discovery is successful and connections are established, but I/O performance is significantly lower than expected. After investigation, you discover that the ONTAP cluster has many more namespaces than controllers, and the namespace distribution across controllers is uneven. Some controllers are handling significantly more namespaces than others. Which of the following is the most likely cause of the performance degradation, and what is the recommended corrective action?

- A. The uneven namespace distribution is causing controller load imbalance; rebalance namespace ownership so that each controller manages an approximately equal number of namespaces
- B. NVMe/FC protocol has a hard limit of 10 namespaces per controller; delete excess namespaces and recreate them on different controllers
- C. The issue is network-related and requires enabling jumbo frames (MTU 9000) on all switch and host interfaces; namespace distribution is not a performance factor
- D. Enable queue depth optimization on the overloaded controllers; this automatically redistributes namespace I/O requests across all available controllers

Answer: A

Explanation:

The correct answer is that uneven namespace distribution causes controller load imbalance, and rebalancing namespace ownership across controllers is the recommended solution (option A). In NVMe environments, namespaces are the equivalent of LUNs in FCP, and distributing them evenly across controllers ensures balanced load and optimal performance. Option B is incorrect because there is no hard limit of 10 namespaces per controller in ONTAP 9.14 NVMe. Option C conflates network optimization (which may be helpful) with the root cause; jumbo frames alone will not solve a controller load imbalance issue. Option D is incorrect because there is no such feature as 'queue depth optimization' that redistributes namespace ownership; namespace-to-controller assignment is manual or policy-based, not automatic based on I/O patterns.

QUESTION 9

A healthcare customer is migrating SAN-attached workloads from legacy storage to a new ONTAP 9.14 cluster using FCP. The legacy system uses traditional RAID 5 storage with 14-day backup tapes for compliance. The customer wants to reduce backup complexity and improve RTO/RPO using NetApp SnapVault and SnapMirror technologies. During the data protection design phase, you must recommend the appropriate strategy. Which two statements are correct regarding the data protection architecture for this scenario?

- A. Use asynchronous SnapMirror to a secondary ONTAP cluster for near-synchronous replication and RPO measured in seconds, with SnapVault for long-term archive compliance retention
- B. SnapVault can serve as a backup destination and provides deduplication and compression, reducing storage cost for long-term compliance retention
- C. Implement synchronous SnapMirror to the secondary cluster and remove tape backup entirely, as ONTAP replication fully satisfies HIPAA compliance requirements
- D. Legacy RAID 5 parity has better data protection than RAID-TEC used in modern ONTAP systems, so maintain hybrid protection with both technologies
- E. Use local Snapshot copies as the primary backup mechanism without any mirror or vault destination, as Snapshots are immutable and cannot be deleted accidentally

Answer: AB

Explanation:

The two correct answers are that asynchronous SnapMirror to a secondary cluster combined with SnapVault for long-term compliance is an appropriate strategy (option A), and that SnapVault provides deduplication and compression for cost-effective long-term retention (option B). This architecture provides both fast RTO/RPO via SnapMirror and audit-trail compliance via SnapVault with reduced storage footprint. Option C is incorrect because while ONTAP replication is excellent, healthcare compliance often mandates tape or off-site long-term archives as part of a comprehensive DR strategy, and tape is not obsolete. Option D is incorrect because modern RAID-TEC (RAID-6 equivalent) in ONTAP provides superior data protection to legacy RAID-5. Option E is incorrect because while Snapshots are valuable, local-only Snapshots do not provide disaster recovery or off-site compliance copies, and Snapshots can be deleted by administrators.

QUESTION 10

You are troubleshooting a connectivity issue in a newly implemented iSCSI SAN environment on ONTAP 9.14. A subset of hosts can connect to the iSCSI target, but others consistently fail to discover the target or establish sessions. All hosts are running the same OS and iSCSI initiator software. After reviewing the network configuration, you discover that some hosts are on VLAN 100 and others on VLAN 200, and the iSCSI LIFs are only assigned to VLAN 100. What is the root cause of the problem and the recommended resolution?

- A. The iSCSI LIFs are not assigned to VLAN 200, so hosts on that VLAN cannot reach the target; create additional iSCSI LIFs on VLAN 200 or configure multi-VLAN support on existing LIFs
- B. ONTAP iSCSI targets inherently support all VLANs automatically; the issue is that the hosts on VLAN 200 are using outdated iSCSI initiators that do not support modern VLAN discovery
- C. The problem is a duplicate IP address conflict on VLAN 200; create new iSCSI LIFs with different IP addresses but leave them on VLAN 100 only
- D. iSCSI sessions cannot span VLANs; the hosts on VLAN 200 must be physically reconfigured to connect to VLAN 100 infrastructure

Answer: A

Explanation:

The correct answer is that iSCSI LIFs must be assigned to the VLANs where initiator hosts reside, and the solution is to create additional LIFs on VLAN 200 or enable those VLANs on existing LIFs (option A). iSCSI targets are only reachable on VLANs to which their LIFs are assigned. Option B is incorrect because ONTAP does not automatically discover across VLANs; explicit LIF assignments are required. Option C misdiagnoses the issue as a duplicate IP problem, which is not supported by the scenario. Option D is incorrect because hosts do not need to be physically reconfigured; the storage configuration simply needs to accommodate both VLANs through additional LIFs or VLAN tagging.

QUESTION 11

You are planning a SAN implementation for a customer with 500 virtual machines requiring 2 TB of shared storage. The customer has two data centers 50 km apart and requires synchronous replication for disaster recovery. During the assessment phase, you discover the customer has legacy 1 Gbps Fibre Channel switches with no plans to upgrade and limited WAN bandwidth of 100 Mbps between sites. Which of the following should you recommend as primary concerns or constraints that will impact the solution design? (Select all that apply)

- A. The 1 Gbps FC switches will limit throughput and may not support the SCSI queue depths required for 500 VMs, requiring either a switch upgrade or workload tiering strategy
- B. Synchronous replication over 100 Mbps WAN will create unacceptable latency and may cause database transaction timeouts; asynchronous replication or disk-to-disk backup should be considered instead
- C. The customer must implement dedicated FCoE converged networks to overcome the bandwidth limitations of legacy FC switches
- D. Two TB of storage is insufficient for 500 VMs and the assessment phase should be halted until the customer agrees to purchase at least 10 TB
- E. The 50 km distance between data centers exceeds the maximum supported distance for FC replication and requires a MetroCluster IP configuration instead

Answer: AB

Explanation:

This question tests assessment and planning skills from the first objective. The correct answers

identify real constraints: legacy 1 Gbps FC switches will bottleneck performance for 500 VMs and may not handle queue depths properly, and 100 Mbps WAN bandwidth is insufficient for synchronous replication, which requires near- zero latency. These are architectural constraints that must be addressed during solution design, either through infrastructure upgrades or workload strategy changes.

The incorrect options contain common misconceptions: FCoE is not a mandatory solution for bandwidth constraints (though it could help), the 2 TB capacity was stated as the customer's requirement so it is not the assessment team's role to reject it, and 50 km is well within FC distance limits and does not require MetroCluster IP (which has no distance restriction anyway).

QUESTION 12

You are configuring iSCSI on a new ONTAP cluster for a bank's transaction processing environment. The customer has four iSCSI initiators and requires 99.99% availability. You create one iSCSI SVM with a single logical interface (LIF) on node A. During implementation testing, you simulate a node A failure and observe that iSCSI connections drop for several minutes before clients reconnect to node B. The customer's RTO is 60 seconds. What is the primary issue with the current configuration, and what should you reconfigure?

- A. The single LIF on node A creates a single point of failure; you should create additional iSCSI LIFs on both node A and node B so that ALUA path failover can redirect I/O within the RTO, and enable ALUA on all initiators
- B. The SVM is not configured with SmartFail; you must enable SmartFail on the iSCSI SVM to automatically redirect sessions to node B within 60 seconds
- C. iSCSI does not support sub-minute failover times; you should recommend that the customer upgrade to FC or NVMe/FC instead to meet the 60-second RTO
- D. The cluster is missing iSCSI enabled on the cluster management LIF; once enabled, failover will happen automatically without needing multiple data LIFs
- E. The initiators must be configured with persistent reservations on node A; add a second LIF on node B and enable iSNS so initiators can discover both paths automatically

Answer: A

Explanation:

This question tests SAN implementation and troubleshooting knowledge. The correct answer identifies that a single iSCSI LIF creates a single point of failure, and the solution is to create multiple LIFs across nodes so that ALUA (Asymmetric Logical Unit Access) can provide active/passive or active/active path failover. With properly configured ALUA paths and multiple LIFs, iSCSI sessions can failover in seconds, well under the 60-second RTO.

The incorrect options use non-existent features (SmartFail for iSCSI), make incorrect assumptions (that iSCSI cannot meet sub-minute RTO, which is false), confuse management and data plane configuration (cluster management LIF is not involved in iSCSI data failover), and misapply concepts like persistent reservations and iSNS. While iSNS can help with discovery, it does not solve the underlying single-LIF bottleneck.

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