ExamsLabs

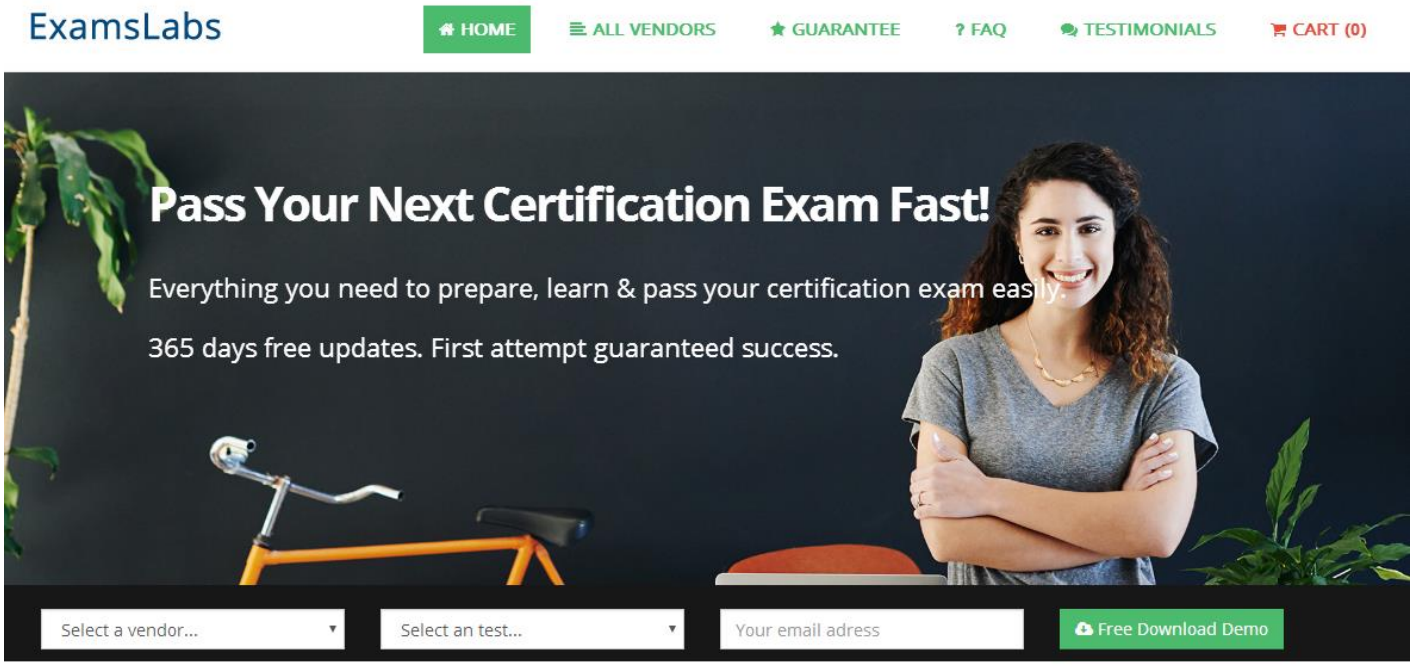

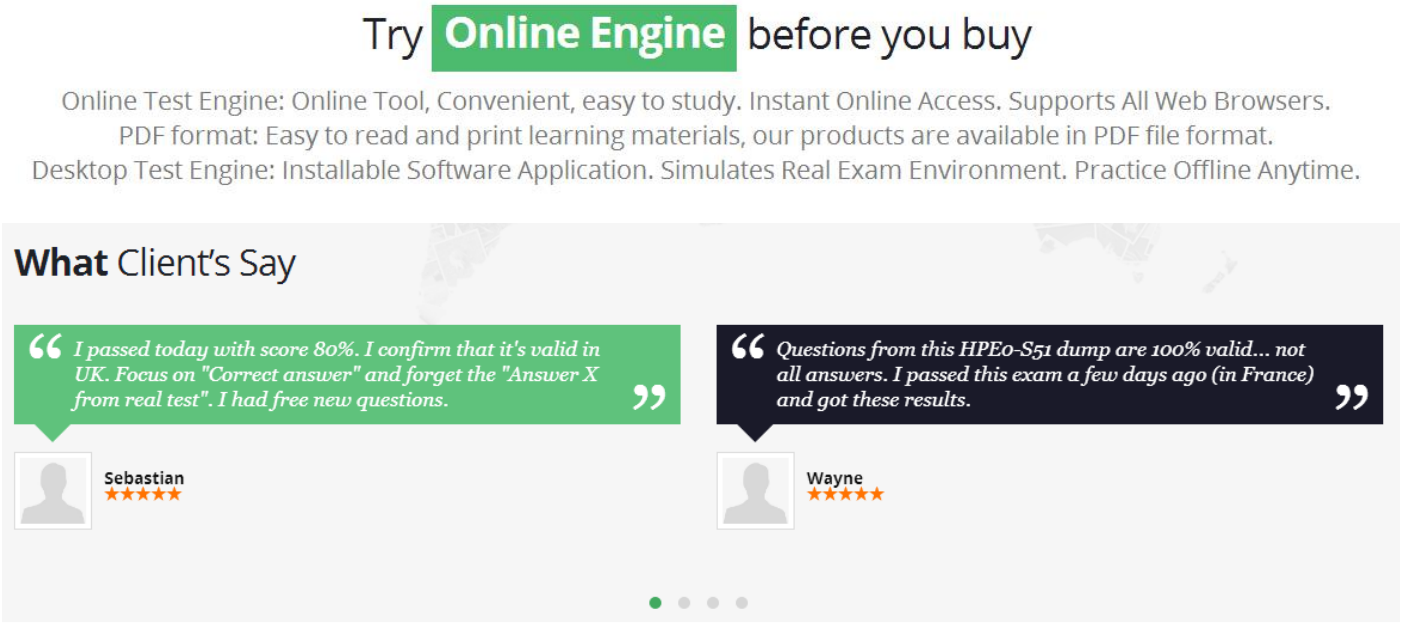

http://www.examslabs.com/

Latest Study Materials, Valid Dumps - ExamsLabs

Exam : **NS0-528**

Title : Implementation Engineer - Data Protection Exam

Vendor : Network Appliance

Version : DEMO

QUESTION NO: 1
After deploying a new on-premises AFF cluster, you have been tasked with backing it up to StorageGRID. You have requested the NetApp ONTAP One licenses.
To successfully deploy this solution, what else needs to be deployed through NetApp ONTAP System Manager?
A. NetApp Cloud Backup service
B. NetApp SnapMirror Cloud API license
C. NetApp BlueXP tiering
D. NetApp BlueXP observability
Answer: B
Explanation:
To back up data from an on-premises AFF cluster to StorageGRID, you must use SnapMirror Cloud, which requires the SnapMirror Cloud API license to be deployed through ONTAP System Manager - even when using ONTAP One. This license enables integration with object storage like StorageGRID for cloud backup.

QUESTION NO: 2
What NetApp tool helps you with the configuration for the cluster switches, when you are expanding an existing 4-node MetroCluster IP setup?
A. Active IQ Config Advisor
B. Interoperability Matrix Tool
C. RCF generator
D. Active IQ Upgrade Advisor
Answer: C
Explanation:
The RCF (Reference Configuration File) generator is the NetApp tool used to create the appropriate switch configuration files for MetroCluster IP setups. When expanding a MetroCluster IP configuration (e.g., from 4 nodes to more), you use the RCF generator to ensure the cluster and interconnect switches are properly configured.

QUESTION NO: 3
What is the minimal interface requirement for a SnapMirror relationship between clusters?
A. one data LIF per SVM
B. one cable directly connected to both clusters
C. one intercluster LIF per node
D. one intercluster LIF per SVM
Answer: C
Explanation:
For a SnapMirror relationship to function properly, each node in a cluster must have at least one intercluster LIF. These LIFs are used exclusively for replication traffic and ensure that all nodes can participate in SnapMirror operations.

QUESTION NO: 4
Which NetApp solution safeguards Kubernetes volumes to object storage in public and

private cloud providers?
A. BlueXP classification
B. BlueXP observability
C. BlueXP backup and recovery
D. Astra Trident
Answer: C
Explanation:
BlueXP backup and recovery provides protection for Kubernetes volumes by backing them up to object storage across public and private cloud providers. It ensures data durability, long-term retention, and recovery for containerized workloads.

QUESTION NO: 5
You have a large 2PiB flexgroup that will need to be failed over to a DR site for a weekend due to maintenance work at the original site. To speed up time to return to normal operations the Storage Administrator will use the -quick-resync true option when reversing the SnapMirror relationship.
What needs to be taken into consideration when it comes to using quick resync?
A. All storage efficiencies will be preserved
B. All storage efficiencies are done in line with the SnapMirror
C. All storage efficiencies will be ignored
D. All storage efficiencies will be lost on the source
Answer: D
Explanation:
When using the -quick-resync true option in SnapMirror, storage efficiencies such as deduplication and compression are not preserved on the original source during resync. This means that efficiencies will be lost on the source volume, and data will be rehydrated to its full form during transfer.

QUESTION NO: 6
Your customer is running NAS & SAN workloads on an old two node ONTAP Cluster, which reached end of support. As their business is growing, they need a new solution which can provide continuous availability and an automatic switchover for all of his workloads. They already have two data centres that are about 400km apart.
Which solution meets this requirement?
A. MetroCluster IP
B. MetroCluster FC
C. SnapMirror SVM replication
D. SnapMirror active sync
Answer: A
Explanation:
MetroCluster IP is the ideal solution for continuous availability and automatic switchover of both NAS and SAN workloads across geographically separated data centers (up to 700 km apart). It provides synchronous replication with automatic disaster recovery capabilities, making it the best fit for this scenario.

QUESTION NO: 7
You are the administrator of a large NAS environment that includes multiple NetApp and non-NetApp onsite storage systems, as well as multiple cloud storage solutions. One of your responsibilities is to replicate data between these systems.
Which NetApp tool enables you to replicate data between these autonomous systems?
A. BlueXP backup and recovery
B. BlueXP copy and sync
C. BlueXP disaster recovery
D. BlueXP tiering
Answer: B
Explanation:
BlueXP copy and sync enables efficient data replication and synchronization between NetApp and non-NetApp storage systems, both on-premises and in the cloud. It's specifically designed for cross-platform and cross-environment data movement, making it ideal for this scenario.

QUESTION NO: 8
You have set up SnapMirror relationships between two clusters. SnapMirror shows as unhealthy and fails with the following error:
Reason: 'Transfer failed'. Last Transfer Error Was 'The number of Snapshot copies on destination volume "volume_name" has reached the maximum supported count of <count_set_by_policy>.
When you check the EMS logs, you notice the following:
[NetAppProdVault-01: worker_thread_342: sm.xfer.fail.maxlim.snap:alert]: The number of Snapshot copies on destination volume 'SVM_filer02:vflsystems_sm_4' of the relationship with relationship UUID 'ala620fc-a870-lle9-94ef-00a09893ee91' and policy 'filer02_vault_daily_policy' has reached the maximum supported.
[NetAppProdVault-01: sm_logger_main: smc.snapmir.update.fail:error]: Snapmirror update from source volume 'vol1' to destination volume 'vol2' failed with error 'The number of Snapshot copies on destination volume "vol2" has reached the maximum supported count of <count_set_by_policy>.'. Relationship UUID 'ala620fc-a870-lle9-94ef-00a09893ee91'.
What is the cause of this situation, and how can you avoid this in the future? (Choose two.)
A. Modify the SnapMirror policy's -keep parameter.
B. The SnapMirror relationship has reached its maximum number of snapshots, which is 1019.
C. Modify the SnapMirror policy's -snapmirror-label parameter.
D. The SnapMirror relationship has reached its maximum number of snapshots, which is 255.
Answer: A,D
Explanation:
The SnapMirror destination volume has hit the maximum allowed number of Snapshot copies, which is typically 255 for most ONTAP systems. This causes transfers to fail with a "Snapshot copy limit reached" error.
To avoid this issue in the future:
* Adjust the -keep parameter in the SnapMirror policy to control how many Snapshot copies

are retained.
* Implement a retention policy that matches your schedule and storage capacity.

QUESTION NO: 9
A customer is running NetApp ONTAP 9.14.1 software and asks you to protect some data in a way that even disk zeroing cannot delete.
Which command enforces the requirement?
A. volume create -snaplock-type compliance
B. volume create -snaplock-type enterprise
C. aggregate create -snaplock-type compliance
D. aggregate create -snaplock-type enterprise
Answer: A
Explanation:
The SnapLock Compliance mode provides WORM (Write Once, Read Many) protection that is irrevocable, even by administrators, and not removable by disk zeroing or reinitialization. This satisfies stringent regulatory and tamper-proof data protection requirements.

QUESTION NO: 10
You have created a SnapMirror relationship from your production volume serving SMB/CIFS data to your DR cluster. After two days of updates, you do not see any of the expected snapshots on the secondary volume.
Why is SnapMirror failing to replicate snapshots?
A. The SnapMirror label is missing from the SnapMirror policy on the secondary volume.
B. SnapMirror is not properly licensed.
C. SVM peering is not properly configured.
D. The SnapMirror label is missing from the snapshot policy on the secondary volume.
Answer: D
Explanation:
SnapMirror replication is driven by labels that match between the snapshot policy on the source volume and the SnapMirror policy. If the snapshot policy on the source does not include the correct SnapMirror label, SnapMirror will not recognize which snapshots to replicate, resulting in no transfers.

QUESTION NO: 11
An administrator has been tasked to replicate data between two NetApp clusters to ensure that the data is replicated no more than once per hour.
Which NetApp solution meets this requirement?
A. MetroCluster SyncMirror
B. SnapMirror active sync
C. SnapMirror Synchronous
D. SnapMirror Asynchronous
Answer: D
Explanation:
SnapMirror Asynchronous is designed to replicate data based on a scheduled interval, such

as once per hour. It is ideal for disaster recovery scenarios where zero RPO is not required but periodic replication is sufficient.

QUESTION NO: 12
You are using NetApp SnapCenter software to back up several Oracle, SQL, and Exchange instances.
What needs to be created in order to verify application consistency for each application within SnapCenter?
A. a resource group that includes all volumes
B. a consistency group for each application that includes the correct resource groups
C. a resource group for each application that includes the correct consistency groups
D. a consistency group that includes all volumes
Answer: B
Explanation:
In SnapCenter, to ensure application-consistent backups for Oracle, SQL, and Exchange, you must create a consistency group for each application. These groups organize the correct resource groups (which define the backup scope) and ensure all related components are captured in a consistent state during the backup.

QUESTION NO: 13
Your customer has a healthy SnapMirror SVM replication between ClusterA and ClusterB. During maintenance, they want to serve the data from ClusterB.
Which two NetApp ONTAP System Manager steps must be performed before shutting down ClusterA? (Choose two.)
A. Manually resume the relationship.
B. Stop the SVM on ClusterA.
C. Start the SVM on ClusterB.
D. Manually quiesce the relationship.
Answer: B,C
Explanation:
To serve data from the SnapMirror SVM destination (ClusterB), ONTAP requires that the source SVM on ClusterA is stopped, to prevent conflicts, and the destination SVM on ClusterB is started, allowing it to take over serving data.
These steps must be done manually in System Manager or CLI prior to accessing data from the destination SVM.

QUESTION NO: 14
What NetApp data protection solution ensures that data cannot be accessed if a disk shelf is stolen?
A. snapshot volume
B. IPsec
C. disk encryption
D. SnapLock volume
Answer: B

Explanation:
Disk encryption (such as NetApp Storage Encryption or NetApp Volume Encryption) ensures that data stored on disks is inaccessible if a disk or disk shelf is stolen, because the data is encrypted at rest and requires keys that are not stored on the hardware itself.

QUESTION NO: 15
Which native NetApp ONTAP solution protects data by using machine learning to identify abnormal activity?
A. Cloud Insights
B. Autonomous Ransomware Protection
C. external FPolicy engine
D. internal FPolicy engine
Answer: B
Explanation:
Autonomous Ransomware Protection (ARP) is a native NetApp ONTAP feature that uses machine learning to monitor and analyze file activity patterns. It detects anomalies such as sudden spikes in file changes - typical of ransomware behavior - and automatically alerts o r blocks access to limit potential damage.

QUESTION NO: 16
Which two NetApp supported products can be used to protect Kubernetes workloads? (Choose two.)
A. BlueXP backup and recovery
B. Astra Control
C. Acitve IQ Unified Manager
D. SnapCenter
Answer: A,B
Explanation:
BlueXP backup and recovery provides Kubernetes-aware backup and recovery capabilities. Astra Control is a NetApp product specifically designed to manage and protect Kubernetes workloads by providing application-aware backup, snapshot, and replication features.

QUESTION NO: 17
A customer wants to replicate between their two NetApp ONTAP clusters but is concerned that replication traffic might slow down normal data traffic. They still have two unused network ports on every node. You recommend using dedicated ports for the Intercluster LIFs, and you are tasked with creating the Intercluster LIFs so that they are on dedicated ports at all times. Which command must you use before creating InterCluster LIFs?
A. system services firewall policy create
B. network interface failover-groups create
C. network ipspace create -ipspace
D. network interface service-policy create
Answer: D
Explanation:

To ensure Intercluster LIFs are bound to dedicated ports and used only for replication traffic, you must first create a service policy that defines the LIF's role (e.g., intercluster). The network interface service-policy create command allows you to specify that the LIF is intended strictly for intercluster services, ensuring it's isolated from data traffic.

QUESTION NO: 18
Your customer wants to set up a local retention time of at least four months. They also want to protect a week of retention from being deleted by ransomware or rogue administrators.
Which two settings create the correct Snapshot policy to achieve this? (Choose two.)
A. Enable Autonomous Ransomware Protection (ARP) on the volume.
B. Create a Snapshot policy with the weekly Snapshot copy retention time set to 122 days.
C. Create a Snapshot policy with the maximum daily snapshots set to 122.
D. Set the SnapLock retention period for the daily Snapshot copies to 7 days.
Answer: B,D
Explanation:
To meet the requirement of 122 days ≈ 4 months, so setting the weekly Snapshot copy retention to 122 days ensures local backups are retained for that duration.
To protect a week of retention from tampering, apply SnapLock to the daily Snapshot copies with a 7-day retention, which makes them tamper-proof and resistant to deletion by ransomware or rogue administrators.

QUESTION NO: 19
You are testing a MetroCluster disaster recovery scenario for your customer, and you get an error messaging stating that the test failed.
Which command should you use to view detailed error information?
A. metrocluster interconnect mirror show
B. metrocluster operation show
C. metrocluster show
D. metrocluster check config-replication show
Answer: B
Explanation:
The metrocluster operation show command provides detailed status and error information about MetroCluster operations, including switchover, switchback, and testing activities. It's the correct command to diagnose a failed disaster recovery test.

QUESTION NO: 20
You need to migrate volumes, export policies, and data LIFs to a new NetApp ONTAP cluster.
Which NetApp feature should be used?
A. BlueXP copy and sync
B. SnapMirror SVM replication
C. SnapMirror active sync
D. Foreign LUN Import
Answer: B

Explanation:
SnapMirror SVM replication enables the replication of an entire SVM (Storage Virtual Machine), including volumes, export policies, LIFs, and configuration, from one ONTAP cluster to another. It is the correct feature for full SVM migration scenarios.

QUESTION NO: 21
A company is using SnapMirror SVM replication, with identity preserve set to false to protect an SMB-only dataset in case of a disaster.
What needs to be manually configured to allow SMB to have access to the data on the destination?
A. SMB server
B. home directories
C. SMB shares
D. name mappings
Answer: C
Explanation:
When SnapMirror SVM replication is configured with identity-preserve=false, network and protocol configurations (like SMB shares) are not replicated. Therefore, to allow SMB access on the destination SVM, you must manually create SMB shares on the destination after replication is initialized.

QUESTION NO: 22
An administrator is deploying NetApp BlueXP backup and recovery for VMs to protect both Windows and Linux systems. The site has two instances of VMware Cloud on AWS that need to be backed up.
How will this goal be accomplished? (Choose two.)
A. Pair each instance of VMware Cloud on AWS with a unique instance of BlueXP backup and recovery.
B. Deploy one instance of BlueXP backup and recovery on a single Linux VM.
C. Pair both instances of VMware Cloud on AWS to a single instance of BlueXP backup and recovery.
D. Deploy two instances of BlueXP backup and recovery on separate Linux VMs.
Answer: C,D
Explanation:
BlueXP backup and recovery supports multiple VMware Cloud on AWS instances with a single BlueXP backup and recovery deployment, as long as proper pairing is done. Alternatively, separate deployments can also be used if desired, each on its own Linux VM. This provides deployment flexibility depending on isolation, scalability, or administrative needs.

QUESTION NO: 23
You are at a customer site and assisting with testing a sudden site disaster scenario with their MetroCluster solution. All the equipment at the production site is suddenly powered off, and the command metrocluster switchover -forced-on-disaster true is issued. After the switchover to the disaster site is complete, the customer notices that one of their database

applications is offline. All other data is accessible from the disaster site. You notice the following in the volume show output:

```
vserver volume in-nvfailed-state,state,total,used,type
vserver   volume     state     total    used     type   in-nvfailed-state
--------  -------    --------  -------  -------  -----  --------------------
SVM1-mc   APP02_DB   online    800 GB   785 GB   RW     true
```

What can you do to help bring the application online while in switchover?

A. Advise the customer to reboot their application server.

B. Increase the available space using the volume size command.

C. Modify the -in-nvfailed-state option using the volume modify command.

D. Reboot the NetApp nodes with the -ignore-strict-sync-warnings true option.

Answer: C

Explanation:

The in-nvfailed-state: true indicates that the volume is in Nonvolatile RAM (NVRAM) failed state, which happens when write operations could not be committed properly during a disaster event. This causes the volume to go into a read-only protected mode to avoid data corruption.

To restore access, you must manually override this state with:

volume modify -vserver SVM1-mc -volume APP02_DB -in-nvfailed-state false This will bring the volume out of failed state and allow the database application to function again during the switchover.

QUESTION NO: 24

What two types of Compliance Clocks are available on NetApp ONTAP software? (Choose two.)

A. volume

B. SVM

C. NTP

D. system

Answer: A,D

Explanation:

NetApp ONTAP supports two types of Compliance Clocks for write once, read many (WORM) and SnapLock compliance:

The system Compliance Clock is maintained at the cluster level and ensures tamperproof retention.

The volume Compliance Clock is specific to SnapLock volumes and is used to enforce retention periods at the volume level.

QUESTION NO: 25

You have production WORM data that is locked by SnapLock Compliance with a 14-year retention. You need to protect the compliance data against hardware failures and site-level disasters.

Which three NetApp products would you use to preserve the compliance metadata? (Choose three.)

A. SnapMirror Asynchronous

B. NDMP backups
C. SnapMirror SVM
D. BlueXP copy and sync
E. MetroCluster
Answer: A,C,E
Explanation:
To protect SnapLock Compliance WORM data with long-term retention (14 years) against hardware and site-level failures, you must use NetApp solutions that preserve both the data and its compliance metadata:
* SnapMirror Asynchronous supports replication of SnapLock volumes, preserving WORM and compliance metadata.
* SnapMirror SVM replicates the entire SVM, including configuration and compliance settings .
* MetroCluster provides synchronous replication and automatic site failover with full preservation of SnapLock compliance attributes.

QUESTION NO: 26
An administrator needs to maintain 10 snapshots per volume on a NetApp AFF system and replicate to a destination with a snapshot retention of 120.
Which NetApp feature should the administrator use?
A. SnapMirror active sync
B. XCP
C. MetroCluster
D. SnapMirror
Answer: D
Explanation:
SnapMirror supports asymmetric snapshot retention, allowing the source volume (on the AFF system) to retain 10 snapshots while the destination retains up to 120 snapshots. This is useful for long-term retention or compliance at the destination without consuming extra capacity on the source.

QUESTION NO: 27
You have a requirement to keep multiple yearly Snapshot copies on your SnapMirror destination volume. However, you do not want to have the Snapshot copies locked on your source volume because you do not have the required capacity.
Where do you set the schedule option to create the required Snapshot copies?
A. primary SnapMirror policy
B. secondary Snapshot policy
C. primary Snapshot policy
D. secondary SnapMirror policy
Answer: B
Explanation:
To retain multiple yearly Snapshot copies only on the SnapMirror destination, you must apply a Snapshot policy on the secondary volume. This avoids consuming space on the source

volume, as these Snapshots are created and retained independently on the destination.

QUESTION NO: 28
You have set up a SnapMirror relationship between volumes in ClusterA and ClusterB. Both the clusters have the same NetApp ONTAP software release and platform. You have enabled Temperature Sensitive Storage Efficiency (TSSE) on the source volume.
How does TSSE affect the efficiency on the destination volume?
A. TSSE settings are inherited by SnapMirror destination volumes by default.
B. TSSE cannot be enabled for volumes in a SnapMirror relationship.
C. TSSE must be manually configured on the destination volumes, otherwise the SnapMirror transfers will fail.
D. TSSE settings are inherited by SnapMirror destination volumes only if the destination volume is less than 50% full.
Answer: A
Explanation:
When Temperature Sensitive Storage Efficiency (TSSE) is enabled on a SnapMirror source volume, the settings are automatically inherited by the destination volume during SnapMirror initialization. This ensures consistent storage efficiency behavior across the replication relationship.