

Latest Version: 7.0

Question: 1

Your manager requires you to protect a production SVM called svm01 by using SVM-DR. Your manager insists that all configuration and data LUNs be replicated to the DR cluster. The destination SVM should be called drsvm01. You set up an SVM-DR relationship between the clusters, specifying that the –identity- preserve option is set to true. After the initial replication completes, you discover that the iSCSI LIF configuration has not been replicated to drsvm01.

What should you do to solve this problem?

- A. Run the relationship a second time and check the configuration again.
- B. Manually create and configure the iSCSI LIFs on drsvm01.
- C. Set up a new relationship without using the -identity-preserve option.
- D. Add the FC protocol to the destination and configure the FC LIFs.

Answer: C

Question: 2

Click the Exhibit button.

Error

Using peer-address 192.168.100.153: An introductory RPC to the peer address "192.168.100.153" failed to connect: RPC: Remote system error [from mgwd on node "cluster2-01" (VSID: -1) to xcintro at 192.168.100.153]. Verify that the peer address is correct and try again.

Initiator Cluster
(cluster2)

Initiator Cluster Intercluster LIFs

IPspace Default

(clu...1) 192.168.0.143

Target Cluster

Intercluster LIFs addresses

Target Cluster Intercluster LIF IP Addresses

192.168.100.153

☐ The effective cluster version of the remote cluster is earlier than ONTAP 9.6

You are configuring cluster peering from cluster2 to cluster1 on your ONTAP 9.7 clusters using ONTAP System Manager. You are given the details shown below to set up the cluster peer relationship. You want to determine why the peering has failed.

Cluster2 intercluster LIF IP: 192.168.0.143 Subnet: 255.255.255.0 Port: e0e passphrase: netapp1234

Cluster1 intercluster LIF IP: 192.168.0.153 Subnet: 255.255.255.0 Port: e0e passphrase: netapp1234

Referring to the exhibit, which statement is correct?

- A. The Source Cluster IP is incorrect.
- B. The Target Cluster Passphrase is incorrect.
- C. The Target Cluster IP is incorrect.
- D. The Source Cluster Passphrase is incorrect.

Answer: CD

Question: 3

You are the administrator of your company's ONTAP clusters. You are asked to provide a disaster recovery solution between your primary AFF A400 cluster and your FAS2720 cluster with the highest possible availability at the lowest possible price.

In this scenario, which product will satisfy this requirement?

- A. MetroCluster
- B. SM-S
- C. SnapMirror
- D. SVM-DR

Answer: A

Reference: <https://blog.netapp.com/disaster-recovery-solution-for-enterprise-apps/>

Question: 4

Application data needs to be replicated between two FAS9000 clusters running ONTAP 9.7 for DR purposes only. The application has 50 data volumes to be replicated concurrently. You also require that the application can run from the DR site should the production site suffer a disaster.

Data Protection Cluster Limits		Scope:
		HA Pair
Maximum number of NDMP sessions		36
Maximum number of data protection (DP) mirrors and/or SnapVault relationships		2,000
Maximum number of data protection (DP) mirrors and/or SnapVault relationships for FabricPool		N/A
Maximum number of load sharing (LS) mirrors		N/A
Maximum number of concurrent Snap Mirror or SnapVault transfers		200
Maximum fan-out from source for DP mirror		N/A
Maximum fan-out from source for LS mirror		N/A
Maximum number of clusters that can be peered		255
Maximum Number of constituent volumes in a SnapMirror relationship		1,000
Maximum number of concurrent NDAS relationship transfers		64
Maximum number of SnapMirror Synchronous relationships		40

Referring to the exhibit, what would you use to satisfy the requirements?

- A. SnapVault backups can be used to replicate all the application volumes.
- B. Load sharing can be used to concurrently replicate the data between both systems.
- C. SnapMirror Synchronous can be used to replicate the all the application volumes.
- D. SnapMirror can be used to replicate all the data volumes concurrently.

Answer: A

Question: 5

You experience a node outage as a result of a component failure in a high-availability (HA) pair. You replace the failed component and get the node booted up again, but are unable to perform the giveback operation. The onscreen error is:

"Connected to partner, Takeover is not possible: Storage failover mailbox disk state is invalid, Local node has encountered errors while reading the storage failover partner's mailbox disks. Local node missing partner disks"

What are two ways to solve this problem? (Choose two.)

- A. Correct the shelf cabling errors.
- B. Destroy the root aggregate.
- C. Correct the failover group configuration.
- D. Destroy the mailboxes on both nodes.

Answer: AC