

Question: 1

Exhibit:

Question

Exhibit

```
::>network interface show -fields data-protocol
```

vserver	lif	data-protocol
svm-iscsi1	lif1	nfs,cifs
svm-iscsi1	lif2	nfs,cifs

You are asked to serve iSCSI LUNs in an existing SVM on your AFF A220 using ONTAP 9.5. You verified that the iSCSI license is configured on your cluster and that the iSCSI protocol is enabled. Referring to the exhibit, what is needed to allow hosts to log into the iSCSI target?

- A. Create new LIFs with the iSCSI data protocol.
- B. Create new igroups with the host IQNs.
- C. Add iSCSI to the data protocol of the existing LIFs.
- D. Map LUNs to the igroups.

Answer: D

Question: 2

You are testing iSCSI LUN failover across a 4-node FAS9000 fabric-attached MetroCluster configuration. In this scenario, which front end configuration is required for non-disruptive host LUN failover between sites?

- A. an intercluster LIF
- B. a stretched Layer 2 network
- C. ipv6
- D. a stretched VSAN

Answer: A

Question: 3

Exhibit:

Question**Exhibit**

```
cl11::> ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
01a	0e	cna	target	-	-	online
01a	0f	cna	target	-	-	online
01a	0g	cna	target	-	fc	online
01a	0h	cna	target	-	fc	online
01b	0e	cna	target	-	-	online
01b	0f	cna	target	-	-	online
01b	0g	fc	target	-	-	online
01b	0h	fc	target	-	-	online

Referring to the exhibit, which two pairs of ports are currently configurable as FC SAN LIFs? (Choose two.)

- A. 01b, 0h
- B. 01a, 0g
- C. 01a, 0h
- D. 01b, 0g

Answer: D

Question: 4

You have a 4-node cluster with an AFF A300 HA pair and a FAS8200 HA pair. You plan on using the default storage efficiency settings. With inline data compaction, you estimate that you can save 6% of storage space. AFF A300 volumes that use under 5000 IOPS are moved to a FAS8200 using the volume move command.

In this scenario, what happens to the data after the volume is moved? (Choose two.)

- A. The new written data is not compacted.
- B. The existing data is compacted.
- C. The new written data is compacted.
- D. The existing data is not compacted.


Answer: A

Question: 5

Exhibit:

Question

Exhibit



You have a requirement to serve LUNs with under 200 microsecond latency using local server lass memory. You also are required to use a shared SAN. Using MAX Data as shown in the exhibit, which two LUN access protocols are supported with this solution? (Choose two.)

- A. iSCSI
- B. FCoE
- C. FC
- D. NVMe

Answer: D

Question: 6

A company has 20 ESX hosts. Ten of the hosts are connected using FC, and the rest are connected over iSCSI. The FC host HBA queue depth is set to 32, and the iSCSI LUN queue depth is set to 64. The FC hosts are now experiencing random connectivity issues. In this scenario, what would be the reason for the disconnects?

- A. The FC host has a lower queue depth.
- B. According to the IMT, unsupported HBA firmware is in use.
- C. All of the hosts are part of the same cluster.
- D. The queue depth is not consistent across the cluster.

Answer: A

Question: 7

You want to ensure maximum performance of iSCSI LUNs on an AFF A220.
In this scenario, which two actions will accomplish this task? (Choose two.)

- A. Disable flow control.
- B. Configure virtual LAN (VLAN).
- C. Enable flow control.
- D. Enable jumbo frames.

Answer: C,D

Question: 8

You are provisioning storage to an ESX host that uses iSCSI.
According to NetApp best practice, which three actions accomplish this task? (Choose three.)

- A. Enable the iSCSI service.
- B. Bind the iSCSI ports to the software iSCSI adapter.
- C. Install Virtual Storage Console (VSC) for VMware.
- D. Enable BGP on the switch.
- E. Configure MTU 1500 on the switch.

Answer: B,C,E

Question: 9

Exhibit:

Question**Exhibit**

```
scaling::> run -node scaling-01 -command fcp topology show
Switches connected on adapter 0c:
Switch Name:      switch
Switch Vendor:    Cisco Systems, Inc.
Switch Release:   4.2(1)N1(1)
Switch Domain:    75
Switch WWN:       23:e9:00:0d:ec:b4:94:01
```

Port	Port	WWPN	State	Type	Attached WWPN
vfc1	20:00:00:0d:ec:b4:94:3f		Offline	none	
vfc2	20:01:00:0d:ec:b4:94:3f		Offline	none	
vfc3	20:02:00:0d:ec:b4:94:3f		Offline	none	
vfc4	20:03:00:0d:ec:b4:94:3f		Offline	none	
vfc5	20:04:00:0d:ec:b4:94:3f		Offline	none	
2/1	20:41:00:0d:ec:b4:94:00		Offline	none	
2/2	20:42:00:0d:ec:b4:94:00		Online	F-port	50:0a:09:80:00:05:a2:40
2/3	20:43:00:0d:ec:b4:94:00		Online	F-port	50:0a:09:85:8d:9d:bf:ce
2/4	20:44:00:0d:ec:b4:94:00		Online	F-port	50:0a:09:85:8d:ad:c0:02

A customer has an AFF A700 connected to a Cisco switch and has verified connectivity but cannot see any LUNs. The customer issued the run -node scaling-01 -command fcp topology show command to verify the connection.

Referring to the exhibit, what needs to be enabled on the Cisco switch?

- A. Trunking needs to be enabled.
- B. VSAN needs to be enabled.
- C. NPIV needs to be enabled.
- D. VFC needs to be enabled.

Answer: D

Question: 10

Exhibit:

Question		Exhibit					
FC/FCoE and NVMe Adapters							
Edit	Status	Refresh					
WWWNN	Node Name	Slot	Data Protocols	WWPN	Status	Speed	
50:0a:09:80:83:61:11	AFF_1	0e	fcp	50:0a:09:84:80:83:61:11	link not connected	auto	
50:0a:09:80:83:61:11	AFF_1	0f	fcp	50:0a:09:83:80:83:61:11	link not connected	auto	
50:0a:09:80:83:61:11	AFF_1	0g	fcp	50:0a:09:86:80:83:61:11	link not connected	auto	
50:0a:09:80:83:61:11	AFF_1	0h	fcp	50:0a:09:85:80:83:61:11	link not connected	auto	
50:0a:09:80:83:61:11	AFF_1	1a	fcp,fc_nvme	50:0a:09:81:80:83:61:11	online	auto	
50:0a:09:80:83:61:11	AFF_1	1b	fcp,fc_nvme	50:0a:09:82:80:83:61:11	online	auto	
50:0a:09:80:80:d3:61:00	AFF_2	0e	fcp	50:0a:09:82:80:d3:61:00	online	auto	
50:0a:09:80:80:d3:61:00	AFF_2	0f	fcp	50:0a:09:81:80:d3:61:00	link not connected	auto	
50:0a:09:80:80:d3:61:00	AFF_2	0g	fcp	50:0a:09:84:80:d3:61:00	online	auto	
50:0a:09:80:80:d3:61:00	AFF_2	0h	fcp	50:0a:09:83:80:d3:61:00	link not connected	auto	

A customer is using both FC and NVMe in a cluster. The customer needs to do some work on AFF_1.

Referring to the exhibit, what will happen?

- A. The NVMe namespaces on AFF_2 will be available after takeover of AFF_1.
- B. The FC LUNs that use cluster interconnect to access the LUNs on AFF_1 will be unavailable after failover.
- C. The FC LUNs on node AFF_2 will be unavailable after failover of AFF_1.
- D. The NVMe namespaces on AFF_1 will be unavailable after takeover by AFF_2.

Answer: C