

Netapp

Exam NS0-180

NetApp Certified Storage Installation Engineer

Version: 8.0

[Total Questions: 65]

Question No : 1

You are preparing to install a new clustered Data ONTAP system.

Which tool prepares all the relevant customer and configuration information?

- A. AutoSupport
- B. Cluster Config Builder
- C. OnCommand System Manager
- D. Config Advisor

Answer: C

Explanation:

OnCommand System Manager provides fast, simple configuration and management for NetApp FAS storage systems.

Incorrect: Not A: AutoSupport is used in a system that is already installed. It provides simple, effective, proactive monitoring and management of your storage infrastructure. Not D: Config Advisor is run on already installed systems. It performs configuration validation for new NetApp system installation and improve automation for service delivery for assessments and audits.

Reference:<http://www.netapp.com/se/products/management-software/system-manager.aspx>

Question No : 2

Which three types of interface groups are supported on NetApp storage systems? (Choose three.)

- A. single mode
- B. HA interconnect
- C. IPspaces
- D. static multimode
- E. dynamic multimode

Answer: A,D,E

Explanation:

You can create three different types of interface groups on your storage system: single-mode, static multimode, and dynamic multimode interface groups.

Reference: Types of interface groups

<https://library.netapp.com/ecmdocs/ECMP1196907/html/GUID-EFA72201-E035-41E2-AC53-CD81A472B5ED.html>

Question No : 3

What happens to the NTP configuration in clustered Data ONTAP 8.3 when a node joins a cluster?

- A. A node that joins a cluster has a separate NTP service.
- B. A node that joins a cluster automatically adopts the NTP configuration of the cluster.
- C. A node that joins a cluster must be manually configured in the NTP.
- D. A node that joins a cluster must reboot before running the NTP service.

Answer: B

Explanation:

A node that joins a cluster automatically adopts the NTP configuration of the cluster.

Reference: How to configure and troubleshoot NTP on clustered Data ONTAP 8.2 and later using CLI

<https://kb.netapp.com/support/index?page=content&id=1014787>

Question No : 4

Click the exhibit button.

Question **Exhibit**

Onboard Ports for FAS2552 8.3 Clustered Data ONTAP (Click row to view supported cables)

Port Address	Description	Max Data Rate
0a	Serial-Attached SCSI	6 Gbps
0b	Serial-Attached SCSI	6 Gbps
e0a	Ethernet	1 Gbps
e0b	Ethernet	1 Gbps
e0c/0c	Unified Target Adapter 2	16 Gbps
e0d/0d	Unified Target Adapter 2	16 Gbps

Cable Compatibility for Port 0a

SAS 3 /6 Gbps (Cu)		SAS 3 /6 Gbps (Op)			
Mktg Part No	Mfg Part No	Images	Length	End 1	End 2
Copper Cables					
X6557-R6	112-00176	-	0.5m	QSFP	QSFP
X6594-R6	112-00256	-	1m	QSFP	QSFP
X6558-R6	112-00177	-	2m	QSFP	QSFP
X6595-R6	112-00255	-	3m	QSFP	QSFP
X6559-R6	112-00178	-	5m	QSFP	QSFP

Print Preview | Export to PDF | Export to Excel

You are installing a FAS2254 and a DS2246 disk shelf at a customer's site. The controller and disk shelf must be racked in separate rows in the data center due to available space. The customer wants you to use two of the 20-meter copper SAS cables to make connections.

Referring to the exhibit, which statement is correct?

- A. You will need to contract sales to obtain a wireless adapter.
- B. You will need to contract sales to obtain the 20 meter copper cables.
- C. The controller and disk shelf must be closer together if copper cables are to be used.
- D. You will need to contact sales to obtain an Ethernet adapter.

Answer: C

Explanation:

According to the exhibit the maximum copper cable length that can be used is 5 meters.

Question No : 5

Which two statements are true regarding SSDs? (Choose two.)

- A. SSDs are only available in 400 GB sizes.
- B. SSDs must be in different RAID groups than HDDs.
- C. SSDs are available for multi-disk carrier shelves.
- D. The Maintenance Center is not available for SSDs.

Answer: B,D

Explanation:

Incorrect: Not A: There are, for example, 100 GB, 200 GB, and 800 GB SSD sizes. Not C: Disks in multi-disk carriers always have the Data ONTAP disk type of MSATA.

Question No : 6

What is NetApp's best practice regarding Ethernet flow control?

- A. Leave flow control enabled.
- B. Disable flow control on the switch ports.
- C. Disable flow control on the host, switch, and all node ports.
- D. Disable flow control on all 10Gbps ports.

Answer: D

Explanation:

The NetApp Best Practice specifies that Flow Control settings for 10G cluster ports are to be set with Flow Control = None.

Reference:

How to modify the Flow Control setting on clustered Data ONTAP cluster ports

https://kb.netapp.com/support/index?page=content&id=1015245&actp=LIST_RECENT

Question No : 7

According to NetApp, which two scenarios are recommended for a RAID-DP aggregate in clustered Data ONTAP 8.3? (Choose two.)

- A. three or more disks if it is a data aggregate
- B. three or more disks if it is a root aggregate
- C. two or more disks if it is a root aggregate
- D. five or more disks if it is a data aggregate

Answer: B,D

Explanation:

The minimum number of disks in a RAID-DP group is three: at least one data disk, one regular parity disk, and one double-parity (dParity) disk. However, for non-root aggregates with only one RAID group, you must have at least 5 disks (three data disks and two parity disks).

Reference:

Clustered Data ONTAP 8.3, Physical Storage Management Guide, page 108

https://library.netapp.com/ecm/ecm_download_file/ECMP1636022

Question No : 8

Multiple customers want to access a cluster from different domains that may contain duplicate IP addresses.

In this situation, which feature should you use?

- A. IPspaces
- B. interface groups
- C. VLAN
- D. logical interfaces

Answer: A

Explanation:

IPspaces enable you to configure a single Data ONTAP cluster so that it can be accessed by clients from more than one administratively separate network domain, even if those clients are using the same IP address subnet range Reference: Clustered Data ONTAP 8.3, Network Management Guide, page 32

Question No : 9

You are planning to install a new clustered Data ONTAP system for a customer.

Which action should be performed before configuring the cluster network switches?

- A. Confirm IP address and configuration information for Brocade switches before configuring the switches.
- B. Use the Cluster Network and Management Network Compatibility Matrix URL before configuring the switches.
- C. Use the Inventory Collect Tool before configuring the switches.
- D. Download the Fibre Channel switches software before configuring the switches.

Answer: B

Explanation:

The NetApp Interoperability Matrix defines infrastructure components and versions that can be used for configurations that NetApp corporation support end to end.

Incorrect: Not C: Inventory Collect Tool enables you to create an inventory report that contains information about 7-Mode storage systems, hosts, and host applications. The inventory report can then be imported to the 7-Mode Transition Tool for transition assessment.

Reference:

<http://www.netapp.com/us/technology/interop.aspx>

Question No : 10

A customer learns that a certain NetApp tool allows remote data collection, intelligent core file handling, and notification of down storage controllers for technical support analysis and troubleshooting.

What is the name of this tool?

- A. RSA
- B. HWU
- C. Synergy
- D. HW Assist

Answer: A

Explanation:

Remote Support Agent (RSA) is a remote diagnostics data collector that is embedded directly into the storage controller's remote management device firmware. Remote Support Agent enables a NetApp support engineer to remotely request an automated upload of log files, core files, and other diagnostic information stored in the storage controller and to remotely trigger an on-demand AutoSupport collection.

Incorrect: Not B: The Hardware Universe (HWU) application is NetApp's system configuration solution. Not C: NetApp Synergy is a complete design, deployment and documentation tool. The program enables you to build a very accurate, detailed model of a NetApp storage system. You can represent most aspects of a storage deployment including the hardware configuration and the sizing of all storage containers, including aggregates, volumes, qtrees and LUNs.

Reference: Clustered Data ONTAP 8.2, Remote Support Agent Configuration Guide, Page 5

Question No : 11

Regarding clustered Data ONTAP 8.3, what happens to the roles of the on board ports, HBAs, and NICs?

- A. These roles are assigned when LIFs are created.
- B. These roles are assigned based on the controller type.
- C. These roles are predefined and must be changed before creating your LIFs.
- D. The roles are divided into two categories: Data Management and Node Management.

Answer: A

Explanation:

In earlier releases of Data ONTAP, default roles were assigned to each Ethernet network port. The roles included data, cluster, cluster management, intercluster, and node management. In Data ONTAP 8.3, these roles are assigned when LIFs are created.

Reference: Clustered Data ONTAP 8.3, Network Management Guide, page 7

Question No : 12

You want to create a data volume on aggr0 and you receive a message with the following warning:

"Warning: You are about to create a volume on a root aggregate. This may cause severe performance or stability problems and therefore is not recommended. Do you want to proceed?"

What is the reason for this severe performance or stability problem?

- A. Controller failover and storage failover are separate processes that must occur in parallel, otherwise problems will arise.
- B. The performance load generated by the data volume can be very strong, so it should not be shared with vol0 on aggr0, requiring a separate aggregate for it.
- C. The performance load generated by the vol0 cannot be shared with any other volume on aggr0, requiring a dedicated aggregate for it.
- D. Controller failover and storage failover are separate processes that must occur at different times, otherwise problems will arise.

Answer: B

Explanation:

One possibly cause is Disk I/O contention on the data volume.

The root volume in Cluster-Mode is used to store and update various tables of the replicated database. Crucial information regarding the locations of LIFs, volumes, aggregates, and different jobs required to run in the cluster are stored in these tables. If a root aggregate has very busy data volumes, the disks in the aggregate will experience higher latency. When a node is unable to update its copy of the replicated database fast enough, it will consider itself unhealthy and stop serving all the data until it can catch up. This is extremely disruptive and affects all the volumes on the node, even if the cause is related to the data volumes stored on the root aggregate only.

Reference:

Why is a warning displayed when attempting to create one or more data volumes in the root aggregate in Data ONTAP Cluster-Mode?

https://kb.netapp.com/support/index?id=3013563&page=content&locale=en_US

Question No : 13