

Cisco and NetApp FlexPod Design Specialist

Netapp NS0-175

Version Demo

Total Demo Questions: 8

Total Premium Questions: 86

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Solution Design	20
Topic 2, Storage Design	22
Topic 3, Network Design	26
Topic 4, Compute Design	16
Topic 5, Mix Questions	2
Total	86

QUESTION NO: 1

A customer is deploying a FlexPod Datacenter and planning to boot from SAN each their UCS blades using the Fibre Channel Protocol.

In this scenario, which three actions would accomplish their objective? (Choose three.)

- A. Use only a pair of Cisco Nexus 9318YC-FX switches
- B. Create an igroup on the NetApp ONTAP storage system with the appropriate IQNs
- C. Use only a pair of Cisco Nexus 9336C-FX2 switches
- D. Use a pair of Cisco MDS 9148T and a pair of Cisco 9336C-FX2 switches together
- E. Create an igroup on the NetApp ONTAP storage system with the appropriate WWPNS
- F. Create a boot LUN on the ONTAP storage system and map it to the WWPNS-based igroup

ANSWER: D E F

Explanation:

Fibre Channel SAN boot requires an end-to-end FC storage path, FC host identification, and a bootable LUN presented to the host. "Use a pair of Cisco MDS 9148T and a pair of Cisco 9336C-FX2 switches together" provides the required separation of SAN and Ethernet functions in this FlexPod design: the MDS switches provide the Fibre Channel SAN fabric, while the Nexus switches provide the Ethernet network infrastructure. Redundant FC fabrics are essential so that each UCS blade has resilient paths to its boot storage.

"Create an igroup on the NetApp ONTAP storage system with the appropriate WWPNS" is required because FC initiators are identified by their World Wide Port Names. Adding the UCS host adapter WWPNS to an ONTAP FC igroup authorizes the appropriate hosts to access the boot LUN through the SAN fabric.

"Create a boot LUN on the ONTAP storage system and map it to the WWPNS-based igroup" completes the storage-presentation requirement. ONTAP must map the LUN to the FC igroup so that the UCS service profile can select the presented LUN as its SAN boot target. NetApp documents FC LUN access control through igroups and LUN mapping in its ONTAP SAN administration guidance: [ONTAP igroups overview](#) and [Map a LUN to an igroup](#).

QUESTION NO: 2

Your environment currently has a 12-node NetApp cluster serving iSCSI. You need to add at least four more nodes to serve FC to your environment.

In this scenario, what should you do?

- A. You can safely add up to 4 more storage controllers in the current cluster
- B. You can safely add up to 6 more storage controllers in the current cluster
- C. You must create a new cluster with the new storage controllers
- D. You can safely add up to 12 more storage controllers in the current cluster

ANSWER: C

Explanation:

You must create a new cluster with the new storage controllers is correct because the existing cluster is already at the supported node-scale limit for a SAN-serving ONTAP cluster. iSCSI and Fibre Channel are both SAN protocols, so adding FC capability does not change the fact that the cluster remains a SAN configuration for scalability purposes. A 12-node SAN cluster has reached its supported maximum; therefore, it cannot be expanded with the four additional FC-serving nodes requested in this scenario.

The appropriate design is to deploy the new controllers as a separate ONTAP cluster and configure the required FC SAN services there. Hosts can access storage from both clusters through the normal supported SAN design mechanisms, including appropriate zoning, LIF configuration, multipathing, and host-side connection management. This approach keeps each cluster within its validated ONTAP SAN scale boundary while allowing the environment to add FC capacity and connectivity.

NetApp documents cluster scale and platform limitations in its ONTAP documentation; designers should also validate the exact ONTAP release and hardware combination in the Hardware Universe before implementation. See the [ONTAP cluster limitations reference](#) and the [NetApp Hardware Universe](#).

QUESTION NO: 3

Which two support account levels are minimally required when you deploy a FlexPod Datacenter solution? (Choose two.)

- A. NetApp SupportEdge Expert NBD
- B. Cisco SmartNet Total Care 24x7x4
- C. NetApp SupportEdge Advisor 24x7x4
- D. Cisco Solution Support 24x7x4

ANSWER: B C

Explanation:

A FlexPod Datacenter deployment requires vendor support coverage for both the Cisco and NetApp portions of the converged infrastructure. **Cisco SmartNet Total Care 24x7x4** provides the required around-the-clock Cisco technical-support entitlement and four-hour hardware-replacement service level for Cisco infrastructure components. This ensures that Cisco TAC assistance, software access, and hardware replacement are available at the service level expected for a production FlexPod design.

NetApp SupportEdge Advisor 24x7x4 is the corresponding minimum NetApp support level. It provides 24-hour support and four-hour replacement coverage for the NetApp storage components, while enabling the support relationship and proactive operational capabilities used to maintain the storage environment. Together, these support levels establish the minimum separate vendor-support coverage for the validated Cisco and NetApp components in FlexPod Datacenter. Cisco describes Smart Net Total Care service features and coverage at [Cisco Smart Net Total Care](#). NetApp describes its support-service portfolio, including SupportEdge offerings, at [NetApp Support Services](#).

QUESTION NO: 4

Exhibit:

You are building a new FlexPod system booting from SAN.

Referring to the exhibit, how many paths will there be to the target LUN?

- A. 1
- B. 2
- C. 8
- D. 4

ANSWER: D

Explanation:

4 is correct. A resilient FlexPod SAN-boot configuration presents the boot LUN through redundant, independent SAN paths. The server uses its redundant Fibre Channel connectivity across both fabrics, while ONTAP provides the LUN through

appropriately configured target paths. This design lets the host maintain access to its boot LUN if an individual fabric component, FC link, adapter, or storage path is unavailable. In practice, the multipathing software discovers and manages these distinct paths as a single LUN device, selecting an active optimized path where appropriate and failing over non-disruptively when necessary. Four paths reflects the standard highly available design objective of maintaining end-to-end redundancy between the compute initiators and the storage target interfaces, rather than relying on a single path or a single fabric. NetApp SAN documentation describes multipathing as the mechanism that provides multiple paths between a host and storage system and supports continuous access during path failures. FlexPod design guidance likewise uses redundant SAN fabrics and connections for boot-from-SAN deployments. See the [NetApp FC host configuration documentation](#) and the [Cisco FlexPod solution overview](#).

QUESTION NO: 5

You are asked to design a new FlexPod solution that will host VMware virtual machines using NFS datastores, file services using SMB, and NFS-mounted home directories. You need to ensure that the network traffic of these three use cases are isolated?

In this scenario, which supported end-to-end network topology would be used to accomplish this task?

- A. VSAN
- B. VXLAN
- C. VLAN
- D. vPC

ANSWER: C D

Explanation:

VLAN and vPC provide the supported end-to-end FlexPod network design for this requirement. VLANs create separate Layer 2 broadcast domains for VMware NFS datastore traffic, SMB file-service traffic, and NFS home-directory traffic. Each workload can be assigned a dedicated VLAN and corresponding IP subnet, allowing the design to maintain traffic separation across the Cisco switching infrastructure and the NetApp storage network interfaces. This segmentation also supports clear policy application, controlled routing between networks where needed, and predictable operational troubleshooting.

vPC complements this VLAN-based separation by connecting the FlexPod components redundantly to a pair of Cisco Nexus switches while presenting port-channel connectivity. A vPC topology permits the attached devices to use both switches concurrently, providing resilient paths without relying on a single active uplink. The VLANs are carried consistently across the vPC port channels, so the separate storage and file-service networks retain their intended end-to-end connectivity while benefiting from link and switch resiliency. This combination is the standard FlexPod approach for isolated Ethernet traffic classes with highly available network attachment. See the [NetApp FlexPod design guide](#) and Cisco's [vPC configuration documentation](#).

QUESTION NO: 6

A customer is unhappy with their current data center infrastructure. They want to purchase a FlexPod system to replace it. The customer asks you to help them understand the benefits of Solution Support to determine which support model would suit them best.

In this scenario, which three statements are correct? (Choose three.)

- A. You have a single point of contact for all cases
- B. Solution Support is a lower-cost support option
- C. There is no need to diagnose the problem before you call support
- D. Customers do not need to open tickets; all issues are automatically reported
- E. Issues are solved by solution experts

ANSWER: A C E**Explanation:**

Solution Support provides a single point of contact for all cases, which simplifies support for a FlexPod environment consisting of integrated technologies from multiple vendors. Instead of the customer having to determine whether a reported symptom originates in storage, networking, compute, virtualization, or another solution component, NetApp coordinates case management and works with the appropriate technology partners as needed.

There is no need to diagnose the problem before you call support because Solution Support is intended to remove the burden of initial cross-vendor fault isolation from the customer. The customer can report the observed solution issue to NetApp, which manages the investigation and engagement process through resolution. This is particularly valuable for converged infrastructure, where a symptom may cross component boundaries.

Issues are solved by solution experts because Solution Support uses personnel with solution-level expertise to coordinate troubleshooting and resolution across the supported FlexPod stack. This integrated ownership helps reduce support complexity and accelerates progress on incidents requiring collaboration among vendors. NetApp describes the service and its single-contact, multivendor coordination model at [NetApp Solution Support](#). FlexPod support and lifecycle resources are also available from [FlexPod Support](#).

QUESTION NO: 7

A customer wants to use a Cisco Nexus 93180YC-FX switch to support the FC protocol for booting from SAN without having to deploy additional Cisco 900 Series MDS switches in their FlexPod Datacenter configuration.

In this scenario, what is required for this solution?

- A. The customer must enable the FC license on the 93180YC-FX switch
- B. The customer must deploy the 93180YC-FX switch as part of a Cisco ACI fabric
- C. The customer must directly connect the NetApp AFF system to the Cisco UCS Fabric Interconnects
- D. The customer must add an expansion FC module to the 93180YC-FX switch

ANSWER: C**Explanation:**

The customer must directly connect the NetApp AFF system to the Cisco UCS Fabric Interconnects. This is the supported design approach when Fibre Channel SAN boot is required without a dedicated Cisco MDS Fibre Channel switching layer. Cisco UCS Fabric Interconnects provide the Fibre Channel connectivity and SAN switching functions for the UCS domain, while the NetApp AFF system supplies FC target ports for the storage LUNs that host the server boot volumes. The Fabric Interconnects must be configured with the appropriate unified FC ports and VSAN configuration, and the NetApp FC target interfaces must be connected redundantly to both Fabric Interconnects. UCS service profiles can then define the SAN boot policy, including the storage target and boot LUN, so each server has resilient paths to its boot LUN. This design keeps Ethernet switching on the Nexus switches while using the Fabric Interconnects for the direct-attached FC storage connectivity. Cisco's FlexPod design guidance documents direct-attached storage as an option for deployments that do not use MDS switches, and NetApp documentation describes the required FC SAN connectivity and LUN access configuration. See the [Cisco FlexPod Datacenter design guide](#) and [NetApp ONTAP FC configuration documentation](#).

QUESTION NO: 8

A customer wants to upgrade their code fabric between each site with a modular chassis switch that supports up to 400GbE per port to accommodate more business transaction volume.

In this scenario, which device would help them accomplish their objective?

- A. Cisco Nexus 9504
- B. Cisco Nexus 9332C

C. Cisco MDS 9706

D. Cisco Catalyst 9400

ANSWER: A

Explanation:

Cisco Nexus 9504 is the appropriate device because it is a modular Nexus 9000-series chassis designed for high-capacity data-center fabrics. Its modular architecture accepts Nexus 9500 line cards, enabling the platform to be configured for high-density Ethernet connectivity, including interfaces that operate at up to 400GbE per port when supported line cards and optics are selected. This makes it suitable for upgrading an intersite fabric where increased transaction volume requires substantially greater bandwidth, port density, and scalable forwarding capacity.

The chassis design is particularly relevant to the requirement: rather than being limited to a fixed port layout, the Cisco Nexus 9504 can be expanded with line cards as traffic demands grow. It also supports the data-center network capabilities commonly used in scalable fabric designs, including VXLAN EVPN-based deployments and high-speed spine or aggregation roles. Cisco identifies the Nexus 9500 modular switches as supporting 400G-capable line-card options for large-scale data-center networks. See the [Cisco Nexus 9500 Series Switches data sheet](#) and Cisco's [400G data-center networking overview](#).