

700-905 Dumps

Cisco HyperFlex for Systems Engineers

<https://www.certleader.com/700-905-dumps.html>



NEW QUESTION 1

How many separate VLANs must each HyperFlex node be configured with running ESXi?

- A. 1
- B. 2
- C. 3
- D. 4

Answer: D

Explanation:

The virtual environment has the following characteristics:

- HyperFlex nodes are emulated using VMs running ESXi installations

Server Selection	Chosen Servers (Checkbox)	Server 1, Server 2, Server 3
	Management VLAN	3091
	Storage Traffic VLAN	3092
	vMotion VLAN	3093
	VM Network VLAN	3094

NEW QUESTION 2

What is required to cluster a pair of Fabric Interconnects?

- A. uplink connections to the enterprise network
- B. HXDP 3.5.2 or better
- C. connection between the FI pair using ports L1 and L2
- D. UCS Manager 2.1 or better

Answer: C

Explanation:

You can use a redundant pair of fabric interconnects in a cluster configuration. If one fabric interconnect becomes unavailable, the other takes over.

In addition, a cluster configuration actively enhances failover recovery time for redundant virtual interface connections. When an adapter has an active virtual interface (VIF) connection to one fabric interconnect and a standby VIF connection to the second, the learned MAC addresses of the active VIF are replicated but not installed on the second fabric interconnect. If the active VIF fails, the second fabric interconnect installs the replicated MAC addresses and broadcasts them to the network through gratuitous Address Resolution Protocol (ARP) messages, shortening the switchover time.

The cluster configuration provides redundancy only for the management plane. Data redundancy depends on the user configuration and might require a third-party tool to support data redundancy.

To use the cluster configuration, you must directly connect the two fabric interconnects using Ethernet cables between the L1 (L1-to-L1) and L2 (L2-to-L2) high-availability ports, with no other fabric interconnects in between. Also, you can connect the fabric interconnects directly through a patch panel to allow the two fabric interconnects to continuously monitor the status of each other and quickly know when one has failed.

NEW QUESTION 3

HyperFlex virtual servers differ from regular servers in which two key areas? (Choose two.)

- A. NVMe: Regular servers do not support NVMe drives for high availability.
- B. No RAID is required to consolidate disks into a shared data platform.
- C. CVM: Virtual appliance, which performs reading/writing, caching, deduplication, and compression.
- D. SP: UCS Service Profiles are used to delineate MAC address pools from upstream networks.

E. CCC: Cisco Cloud Center is used for multi-cloud integration and seamless deployment.

Answer: BC

Explanation:

HyperFlex virtual servers differ from **regular** servers in these key areas:

- **No RAID** is required to consolidate disks into a shared data platform.
- **CVM:** Virtual appliance, which performs reading/writing, caching, deduplication, and compression.
- **IOVISOR:** Hypervisor driver, which mounts HyperFlex storage and distributes data.
- **VAAI:** vSphere storage API allowing file-system-level snapshots and cloning.

NEW QUESTION 4

Which three advantages of using the M5 generation of HyperFlex servers over the M4 generation are valid? (Choose three)

- A. Support for Cisco VICs
- B. Multiple GPUs
- C. M.2 SATA drive support for faster disk I/O
- D. DDR3 memory
- E. Microsoft Hyper-V support
- F. NVMe support

Answer: CEF

Explanation:

HyperFlex **M5 generation** servers are configured with these important features:

- HDD or SSD drives for capacity storage.
 - Self-encrypting drive options are available.
- SSD cache drive (SAS, NVMe, or NVMe Optane).
- M.2 SATA drives as boot drives for the hypervisor (ESXi or Hyper-V).
- All nodes use Intel Xeon Scalable CPUs and DDR4 memory.

M5 servers supersede the M4 generation of Cisco UCS servers that was the first to support Cisco HyperFlex. M4 nodes used Intel Xeon processor E5-2600 v4 family CPU. M4 servers did not contain M.2 drives for the hypervisor boot and did not support Microsoft Hyper-V.

NEW QUESTION 5

How many DIMMs are supported per memory channel in the Cisco UCS M5 server?

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

Answer: A

Explanation:

Memory

OS memory is used by the Cisco HyperFlex servers not only to serve the internal hypervisor processes but also to expedite VM-related functions. Its performance has a significant impact on overall system operation.

Memory in HyperFlex M5 nodes provides these benefits:

- Allows up to two **DIMMs** per memory channel.
- Is organized with six memory channels per CPU.
- Comes in 128-, 64-, 32- and 16-GB **DIMMs**.
- Permits 3-TB (3072-GB) maximum memory.

– 2 x 128 GB x 6 channels x 2 CPU = 3072 GB.

NEW QUESTION 6

Which version of HXDP was the first to support multiple VICs on a single server?

- A. HXDP 3.5.1
- B. HXDP 3.0
- C. HXDP 4.0
- D. HXDP 3.5

Answer: A

Explanation:

Network Adapters: **Multi-NIC** Support

Starting with HXDP v3.5.1, **multiple** NICs are supported per server:

- Increases resiliency and enables use cases such as offline streaming and backup.
- Primary, mLOM-placed NIC is still mandatory, other NICs fit into PCIe slots.
- Only supported on fresh installations; no upgrade of existing cluster with additional cards.

NEW QUESTION 7

Which two Cisco UCS Servers support converged nodes in HyperFlex Data Platform (HXDP)? (Choose two.)

- A. HX 220
- B. UCSB200
- C. UCS C480
- D. UCS B480
- E. HX240

Answer: AE

Explanation:

The **converged** nodes can only be HyperFlex rack servers, but the Cisco HyperFlex system also supports expanding the existing data platform with additional compute resources, by integrating compute-only nodes, where M4 and M5 generations of Cisco UCS are supported.

NEW QUESTION 8

Where is the VIC configuration for number type, identify, failover, settings, and bandwidth stored?

- A. in non-volatile memory on the VIC
- B. in UCS Manager service profiles
- C. in VCenter virtual machine image
- D. in UCS Manager server profiles

Answer: D

NEW QUESTION 9

Which three configurations for read caching in Cisco HyperFlex are valid? (Choose three.)

- A. Battery-Initiated Read-back (default): Only read data and most commonly used data are deposited in the Level 4 read-back cache
- B. Write-back (default): Only write information and most commonly used information are deposited in the cache
- C. Write-through (install option for VDI): Only most commonly used data is cached: optimizing VDI performance
- D. No caching (SSD): With all-flash nodes; because there is little difference in read speeds between SSDs
- E. Level 4 cached (SSD): With semi-flash nodes; there is a large difference in read speeds between SSDs
- F. Write-first (default for VDI): Infrequently used data is cached: freeing system resources for VDI performance

Answer: BCD

Explanation:

There are three options for read **caching** in Cisco HyperFlex:

- **Write-back (default):** Only write information and most commonly used information are deposited in the cache
- **Write-through (install option for VDI):** Only most commonly used data is cached, optimizing VDI performance.
- **No **caching** (SSD):** With all-flash nodes, because there is little difference in read speeds between SSDs.

Regular Hybrid
(Write-Through)

VDI Hybrid
(Write-Back)

All-Flash
(No Read Cache)

NEW QUESTION 10

How many PCIe standards-compliant interfaces do Cisco VICs support?

- A. 512
- B. 128
- C. 256
- D. 64

Answer: C

Explanation:

Cisco VICs and Their Benefits

In heavily virtualized environments of modern data center infrastructures, hardware no longer represents the actual topology of a software-defined data center, which is also true for network connectivity. While physical cabling still constructs the physical topology, how individual hardware components are used can be much more flexible. When several virtual machines exist on the same server and in their own network topology, they are still limited by physical network interfaces for communication. However, Cisco VICs allow you to create up to **256** PCIe compliant interfaces that are presented to the hypervisor as individual network interface cards. Allowing for great flexibility when configuring the software-defined network components while maintaining a simple physical topology.

Cisco C-Series VICs resemble regular NICs and use a PCIe slot to connect to the system, while Cisco B-Series VICs use internal mezzanine slots to connect and rely on the Cisco B-Series Chassis to provide physical connectivity through the IOM.

Network Adapters: mLOM

The modular LAN-on-Motherboard (mLOM) slot is used for a Cisco VIC. It incorporates next-generation converged network adapter (CNA) technology from Cisco, providing investment protection for future feature releases.

Important information about Cisco UCS VICs:

- Installed in mLOM slot do not consume a PCIe slot.
- Can present up to **256** PCIe standards-compliant interfaces to the host.
- Available in two variants, for M4 and M5 servers:

NEW QUESTION 10

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 700-905 Exam with Our Prep Materials Via below:

<https://www.certleader.com/700-905-dumps.html>