

700-905 Dumps

Cisco HyperFlex for Systems Engineers

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



NEW QUESTION 1

If a GPU card is installed in HyperFlex nodes before a cluster is created, which action can be used to automatically build the service profile in UCS Manager?

- A. Check Run UCS Manager Configuration during the cluster creation process
- B. Check the extended memory option during the cluster creation process
- C. Check the GPU workflow during the cluster creation process
- D. Check the administrative workflow option during the cluster creation process

Answer: C

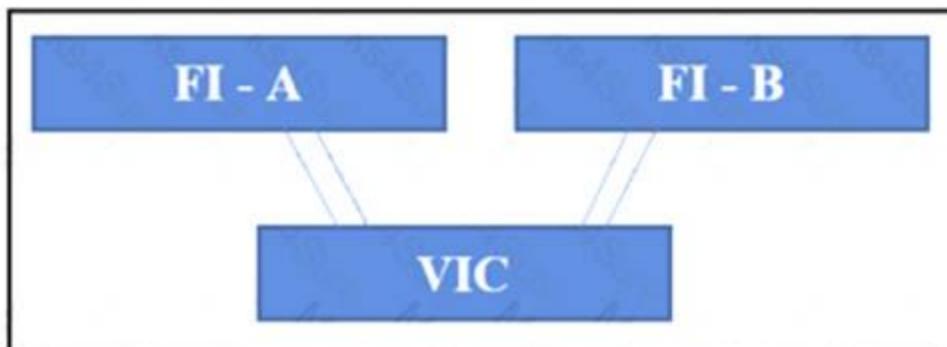
Explanation:

If the **GPU card** is installed before the cluster is created, then, during cluster creation, choose the **Advanced** workflow:

- On the HXDP installer page, choose **I know what I'm doing, let me customize my workflow**.
- Check **Run Cisco UCS Manager Configuration** and click **Continue**. This creates the necessary service profiles for the HyperFlex nodes
- Verify that BIOS Setting by setting **MMIO Above 4-GB** configuration to **Enabled**.
 - If it is not, enable it and you will need to reboot the servers.
- Go back to the **Advanced** workflow on the HX Data Platform Installer page to continue with **Run ESX Configuration, Deploy HX Software, and Create HX Cluster** to complete cluster creation.

NEW QUESTION 2

Refer to the exhibit.



Which VIC model supports two wire connectivity to each Fabric Interconnect?

- A. VIC 1227
- B. VIC 1557
- C. VIC 1387
- D. VIC 1457

Answer: C

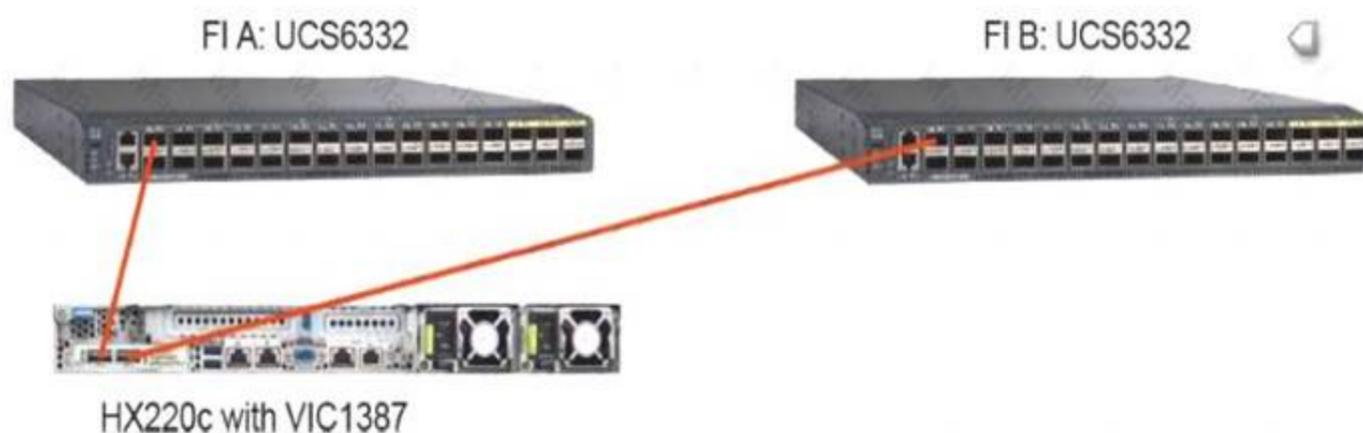
Explanation:

Wiring Cisco HyperFlex Servers to Fabric Interconnects

You connect the Cisco HyperFlex servers to the Fabric Interconnects in the similarly as you wire other rack-mount servers.

Connect each HyperFlex server using unified wire to both Fabric Interconnects.

- HX UCS M5 as of HXDP v3.5.1 supports mLOM-based VIC1387 and VIC1457.
 - VIC1457 is supported only for ESXi-based deployments as of HXDP v3.5.1.
 - VIC1457 supports two wire connectivity to each Fabric Interconnect. VIC1387 is single wire to each Fabric Interconnect.
- It is not supported that you use Fabric Extender (FEX) between server and Fabric Interconnects.
- When connecting VIC to Fabric Interconnects, make sure port numbers match.
 - For example, a given server's VIC to port 1/3 on both Fabric Interconnects.
 - If ports do not match, installation will fail.



NEW QUESTION 3

Which two steps should be performed before installing HyperFlex? (Choose two.)

- A. Determine and download recommended hypervisor
- B. Determine and download recommended VCenter required
- C. Download service profile templates
- D. Determine and download recommended UCS firmware required.
- E. Determine and download virtual machine OS' required.

Answer: AD

NEW QUESTION 4

The process of optimizing information is tightly tied to the writing process as it is performed inline as the writing process is being performed The process of data optimization is performed with which two processes? (Choose two)

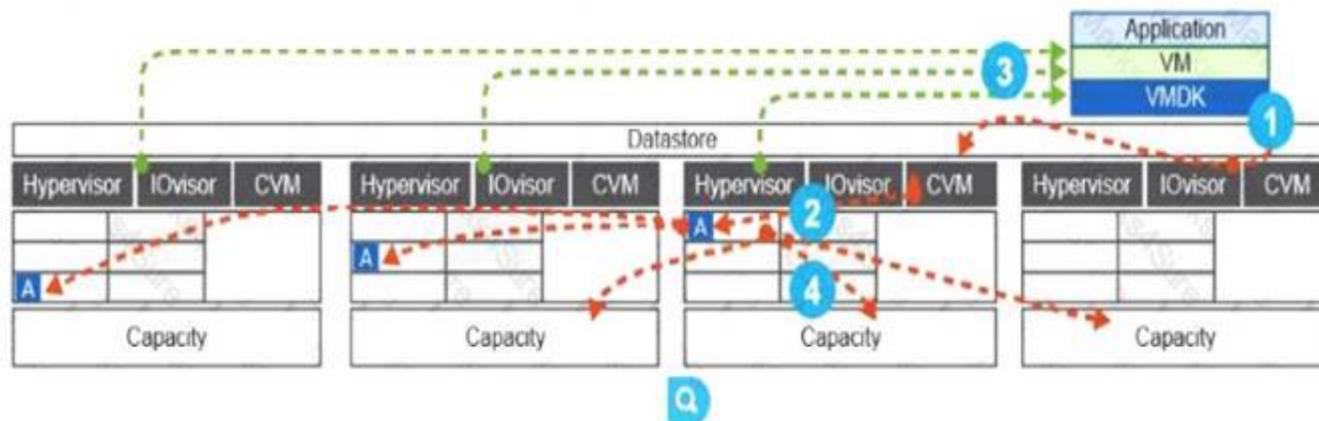
- A. The primary CVM compresses the data, writes it to its cache drive and mirrors it
- B. ACK is sent to the CVM that the write is about to be initiated
- C. On write, the local IOVisor sends the write to the primary CVM for that block
- D. On read/Writ
- E. the distributed VAAI sends the write to the primary CVM for that block
- F. The secondary CVM compresses the data, reads it from its cache drive and mirrors it

Answer: AC

Explanation:

Data Optimization Process and Actual Data Savings

The process of optimizing information is tightly tied to the writing process, as it is performed inline as the writing process is being performed. The system is designed so that the deduplication and compression are done only once by the primary CVM. The IOvisor determines which CVM is primary when the initiated write is intercepted, before it is forwarded to the chosen CVM.



The process of data optimization is performed in this sequence:

1. On write, the local IOvisor sends the write to the primary CVM for that block.
2. The primary CVM compresses the data, writes it to its cache drive and mirrors it.
3. ACK is sent to the virtual machine that the write has been successfully performed.
4. Once the write log is full, a destage is initiated, where the primary CVM performs a best effort deduplication and writes the information across nodes.

NEW QUESTION 5

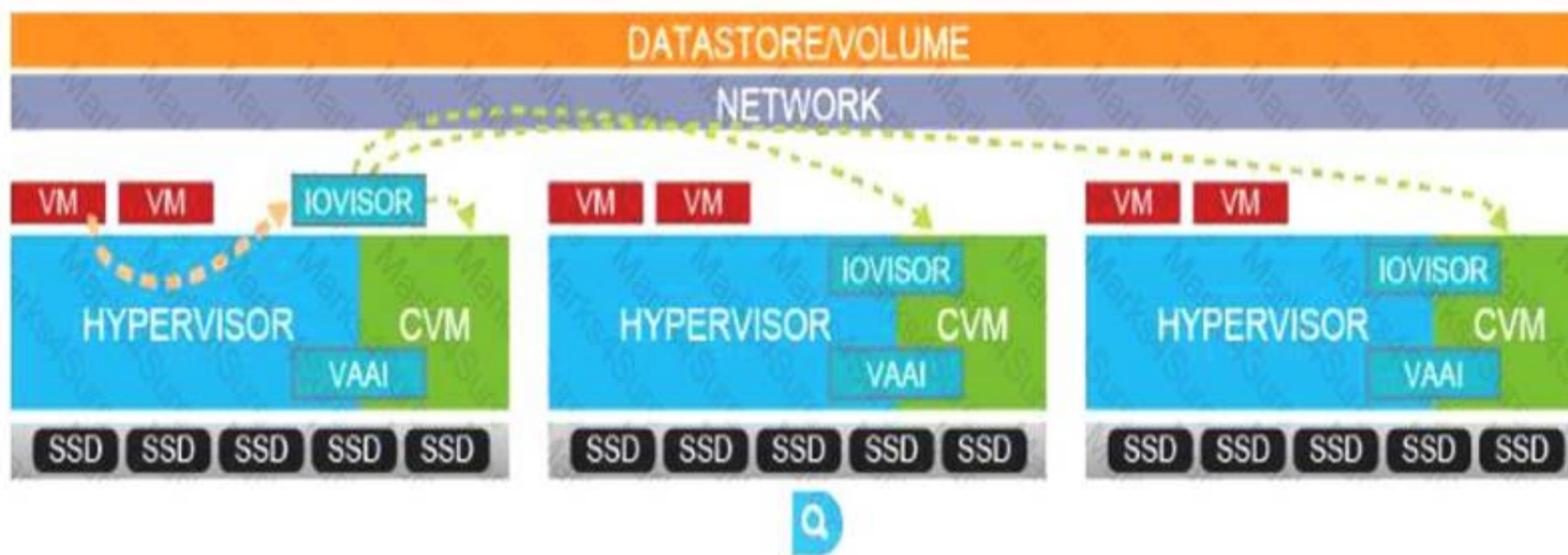
When local writing or reading is performed, the IOvisor intercepts the read/write requests and forwards them to CVMs across the cluster. This action allows non-local CVMs to be aware of the input/output requests so that they can perform the appropriate input/output action. IOvisor provides which two additional functionalities? (Choose two.)

- A. provides redundancy when local CVM fails, offloading data processing to another CVM in the cluster
- B. when an IOvisor fails, the CVM remains active and functional, which enables uninterrupted operation of the system by forwarding IO to another available IOvisor in the HyperFlex cluster
- C. enables asynchronous replication of data across individual HyperFlex nodes with sub-second re-convergence
- D. integration point for deployment of cloud-based SaaS offerings from eco-system partners
- E. intercepts local virtual machines' reads and writes and distributes them across the network eliminating hotspots

Answer: AE

Explanation:

IOVisor provides these functionalities:



- Intercepts local virtual machines' reads and writes and distributes them across the network, eliminating hotspots.
- Provides redundancy when local CVM fails, offloading data processing to another CVM in the cluster.
- Enables synchronous replication of data across individual HyperFlex nodes according to replication factor.

When local writing or reading is performed, the IOVisor intercepts the read/write requests and forwards them to CVMs across the cluster. This action allows non-local CVMs to be aware of the input/output requests so that they can perform the appropriate input/output action. This feature enables the entire cluster to function as one coherent storage using the network.

NEW QUESTION 6

Which statement about Standalone Cisco UCS Server Deployments is valid?

- A. They require Cisco Fabric Interconnects to operate, which reduces the Operating Expenses (OpEx) associated with the deployment
- B. They do not require Cisco Fabric Interconnects to operate, which reduces the Operating Expenses (OpEx) associated with the deployment
- C. They do not require Cisco Fabric Interconnects to operate, which reduces the Capital Expenses (CapEx) associated with the deployment
- D. They require Cisco Fabric Interconnects to operate, which reduces the Capital Expenses (CapEx) associated with the deployment

Answer: C

Explanation:

Standalone deployments have these features:

- Reduced upfront cost, but increased management overhead.
- Good for single deployments or small environments, but do not scale well.
- You are always able to integrate a single deployment into a centrally managed infrastructure.

Standalone deployments of servers do not require Cisco Fabric Interconnects to operate, which reduces the Capital Expenses (CapEx) associated with the deployment. It does not mean that the long-term total cost of ownership (TCO) is better in standalone deployment scenarios, because management overhead is much greater than in a managed deployment scenario, especially in larger deployments.

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

How can the maximum 10 performance be achieved?

- A. Use the HX 220 with all flash drives
- B. Use the HX 240 with all flash drives
- C. Use the HX 220 with all SAS drives
- D. Use the HX 240 with all SAS drives

Answer: B

NEW QUESTION 9

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 10

In all HX server types, where are capacity drives installed?

- A. side
- B. top
- C. back
- D. front

Answer: D

Explanation:

Identifying Capacity Drives

In all server types, the capacity drives are installed on the **front**.

Capacity drives are installed in:

- All HX220c (hybrid/all-flash/all-NVMe):
 - **Front** slots 3-10.
 - First two slots used by housekeeping and cache drives.
- HX240c-M5SX (hybrid/all-flash):
 - **Front** slots 2-24.
 - First slot used by housekeeping drive.
- HX240c-M5L (hybrid only):
 - **Front** slots 1-12.

NEW QUESTION 10

With which three components must every HyperFlex cluster be equipped with in regard to disks? (Choose three.)

- A. NVMe drives
- B. there are no specific requirements
- C. same type of cache drives
- D. same type and size of capacity of drives
- E. same number of capacity drives
- F. SAS drives

Answer: CDE

Explanation:

Drive Selection Rules

Similar to the limitations about mixing different nodes in a cluster, you must follow these guidelines when selecting drives for each node within a cluster:

Every node in Cisco HyperFlex cluster must be equipped with:

- The same type and size of capacity drives:
 - **HDD:** 1.2, 1.8, 6, or 8 TB.
 - **SSD:** 960 GB or 3.8 TB.
 - **NVMe SSD:** 1 or 4 TB.
- The same number of capacity drives
 - 6–8 in HX220 (all types).
 - 6–23 in HX240c-M5SX.
 - 6–12 in HX240c-M5L.
- The same type of cache drive:
 - SAS SSD, NVMe SSD, or NVMe Optane SSD.
 - Size does not matter; the same amount of space is used no matter the disk size.

NEW QUESTION 13

Which Cisco UCS Server running HXDP supports the largest storage pool?

- A. UCS B200
- B. HX220
- C. HX 240
- D. UCS B480

Answer: C

Explanation:

When you evaluate the servers that are most appropriate for your environment, consider these general guidelines:

- Choose HX240 servers to maximize the storage pool.
- Choose HX220 servers to ensure high compute power (relative to storage).
- Choose all-flash platforms to increase IO performance.
- For environments where storage performance is crucial, use All-NVMe nodes once HyperFlex 4.0 is released.

NEW QUESTION 16

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 20

Which two ways does Cisco HyperFlex upgrade the traditional RAID? (Choose two.)

- A. HyperFlex enables stretched RAID arrays spanning multiple geographic sites.
- B. Hardware replacement initiates self-healing with minimal impact
- C. Limiting the number of drives locally, which are not a part of the shared datastore.
- D. Distributing data locally, not just across the hosts in HyperFlex cluster
- E. Eliminating the need for additional hardware cards, while maintaining high performance.

Answer: BE

Explanation:

Cisco HyperFlex upgrades the **traditional** RAID by:

- Not limiting the number of drives, which are a part of the shared datastore.
- Distributing data across the hosts in HyperFlex cluster, not just locally.
- Hardware replacement initiates self-healing with minimal impact.
- Eliminating the need for additional hardware cards, while maintaining high performance.

NEW QUESTION 22

.....

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>