

700-905 Dumps

Cisco HyperFlex for Systems Engineers

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



NEW QUESTION 1

Which two components are automatically configured from the information provided to the HyperFlex installer? (Choose two)

- A. the network
- B. operating system deployment preparation
- C. controller VM configuration
- D. application dependencies
- E. server firmware policy

Answer: AC

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

What does the letter W indicate when selecting CPUs for your HX Node (ie. HX-CPU 8170M)?

- A. support of 1.5 TB/socket of memory
- B. support for all flash drive array
- C. support for NVMe
- D. support for 768 TB/socket of memory

Answer: A

Explanation:

CPU Options

There are several dozens of CPU variants that are available with Cisco HyperFlex M5 servers. The product IDs ending in "M" support 1.5 TB/socket of memory. All other CPU PIDs support 768-Gbps socket memory.

The table lists a few of the many variants, all with product IDs ending in "M". "M" indicates support for 1.5-TB memory per CPU, and up to 3-TB memory in the HyperFlex server (dual CPU.)

| Product ID | Clock Freq (GHz) | Cache Size (MB) | Cores | Highest DDR4 DIMM Clock Support (MHz) |
|---------------------|------------------|-----------------|-------|---------------------------------------|
| HX-CPU-8180M | 2.5 | 38.50 | 28 | 2666 |
| HX-CPU-6142M | 2.6 | 22.00 | 16 | 2666 |
| HX-CPU-6134M | 3.2 | 24.75 | 8 | 2666 |
| HX-CPU-8176M | 2.1 | 38.50 | 28 | 2666 |
| HX-CPU-8170M | 2.1 | 35.75 | 26 | 2666 |
| HX-CPU-8160M | 2.1 | 33.00 | 24 | 2666 |

For a full list of available CPUs, refer to the server specification sheets.

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)

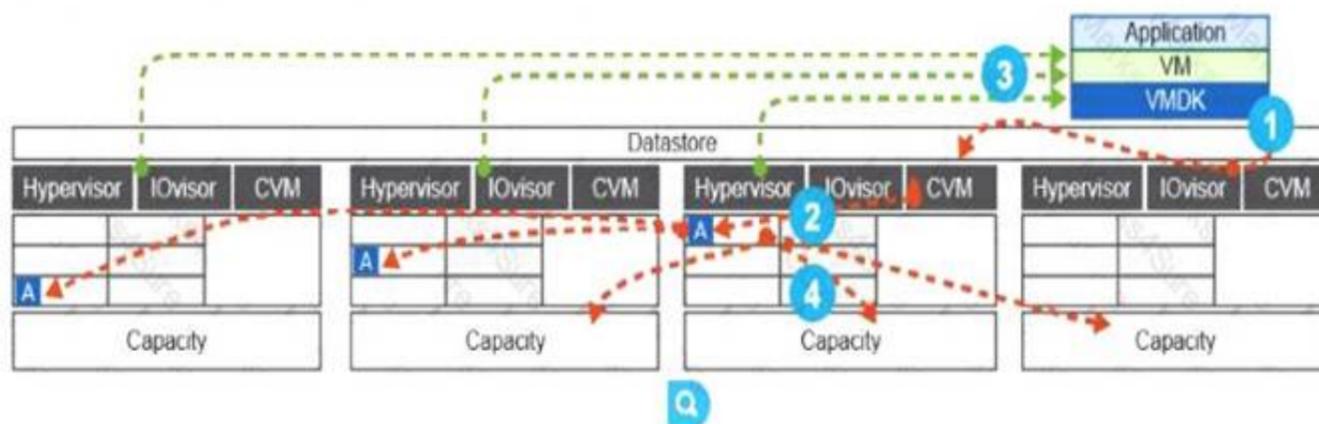
- A. The primary CVM compresses the data, writes it to its cache drive and mirrors it
- B. ACK is sent to the CVM that the write is about to be initiated
- C. On write, the local IOvisor sends the write to the primary CVM for that block
- D. On read/Writ
- E. the distributed VAAI sends the write to the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

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 6

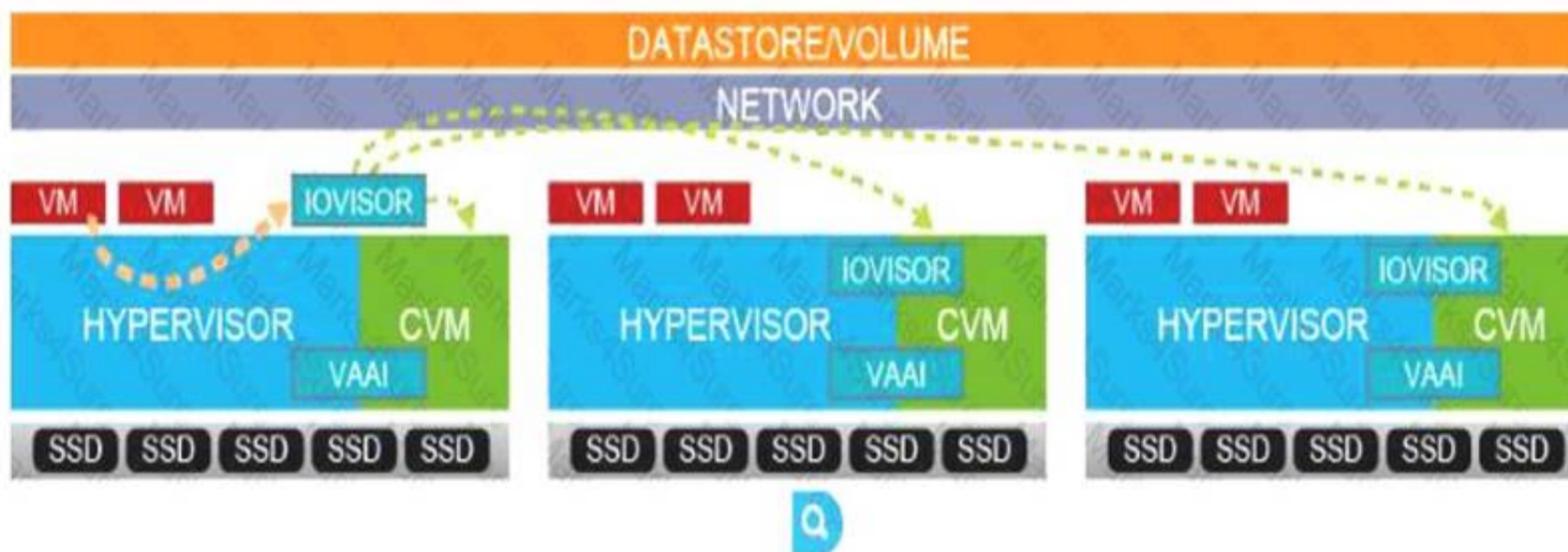
When local writing or reading is performed, the IOVisor intercepts the read/write requests and forwards them to CVMs across the cluster. This action allows non-local CVMs to be aware of the input/output requests so that they can perform the appropriate input/output action. IOVisor provides which two 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 7

How much memory is reserved for the controller VM in the HX220c?

- A. 48 GB
- B. 12 GB
- C. 24 GB
- D. 78 GB

Answer: A

Explanation:

CPU and Memory Guidelines

When selecting the most appropriate CPU for your cluster, you should consider the overhead consumed by the Controller VM and RAM support limits.

Consider these facts when choosing hardware:

- These resources are reserved for the Controller VM:
 - 8 vCPUs, shared.
 - 10,8-GHz of CPU power.
 - 48-GB memory on each HX220c, reserved.
 - 72-GB memory on each HX240c, reserved.
 - 78-GB memory on each HX240c LFF, reserved.

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

Which two results are expected when you replace a node or expand a cluster? (Choose two.)

- A. Distributed pooled data is migrated off nodes to master data store.
- B. Affected node is marked as unhealthy and placed into standby mode
- C. vSphere DRS migrates the virtual machines to the new node to balance the load
- D. On node replace, the self-healing must finish for the cluster to be healthy
- E. The cluster profile is updated and RAID takes care of rebalancing the load.

Answer: CD

Explanation:

Expansion and Hardware Replacement

When you replace a node or **expand** a cluster, the following happens:

1. vSphere DRS migrates the virtual machines to the new node to balance the load.
2. On node replace, the self-healing has to finish for the cluster to be healthy.
3. The new node is already used for writing, but the old data is not migrated until the rebalance process.
4. Rebalance is initiated daily at 5:15 AM or can be executed manually with the **stcli cluster rebalance** command.

NEW QUESTION 10

Drag the server type from the left onto the maximum number of capacity drives on the right.

Drag the server type from the left onto the maximum number of capacity drives on the right.

| | | |
|-------------|-------|------|
| HX220c-M5SX | _____ | 6-12 |
| HX240c-M5SX | _____ | 6-8 |
| HX240c-M5L | _____ | 6-23 |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

HX220c-M5SX_6-8 HX240c-M5SX_6-23 HX240c-M5L_6-12

Capacity Drive options

| Server | Drives | Capacity Drive Type |
|--------------------------------|--------|--|
| HX220c-M5SX Hybrid | 6-8 | 1.8-TB or 1.2-TB SFF HDDs |
| HX220c-M5SX Hybrid with SED | | 1.2-TB SED SFF HDDs |
| HX220c-M5SX All-Flash | | 3.8-TB or 960-GB SSDs |
| HX220c-M5SX All-Flash with SED | | 3.8-TB, 960-GB, or 800-GB SED SFF SSDs |
| HX220c-M5SX All-NVMe | | 4-TB or 1-TB NVMe SSD |
| HX240c-M5SX Hybrid | 6-23 | 1.8-TB or 1.2-TB SFF HDDs |
| HX240c-M5SX Hybrid with SED | | 1.2-TB SED SFF HDDs |
| HX240c-M5SX All-Flash | | 3.8-TB or 960-GB SSDs |
| HX240c-M5SX All-Flash with SED | | 3.8-TB, 960-GB, or 800-GB SED SFF SSDs |
| HX240c-M5L Hybrid | 6-12 | 8-TB or 6-TB LFF HDDs |

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 11

Cisco HyperFlex All-NVMe nodes are expected to be supported beginning in which HXDP version'?

- A. HXDP 4.0.1
- B. HXDP 3.5.1
- C. HXDP 3 5.2
- D. HXDP 4.0

Answer: D

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 15

What is the maximum size of an HXDP cluster running 3.5.1?

- A. 64 nodes
- B. 8 nodes
- C. 16 nodes
- D. 32 nodes

Answer: A

Explanation:

Cisco HyperFlex is a scalable system:

- As of HXDP v3.5.1, maximum size of standard ESXi-based cluster is 64 servers.
 - Cluster, with exception of stretched cluster, cannot be a part of more than one Cisco UCS domain.
 - You can only achieve cluster of this size with Cisco UCS 6296, other fabric interconnects do not have enough ports.
 - An alternative is to have a stretch cluster where servers are split across two Cisco UCS domains.
- If you want to connect Fibre Channel storage to the same Cisco UCS domain, consider that all Fabric Interconnects, except Cisco UCS 6332, support unified ports.

NEW QUESTION 20

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 24

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 28

HyperFlex compute nodes contribute what percentage of the overall disk storage capacity?

- A. 5%
- B. 20%
- C. 0%
- D. 10%

Answer: C

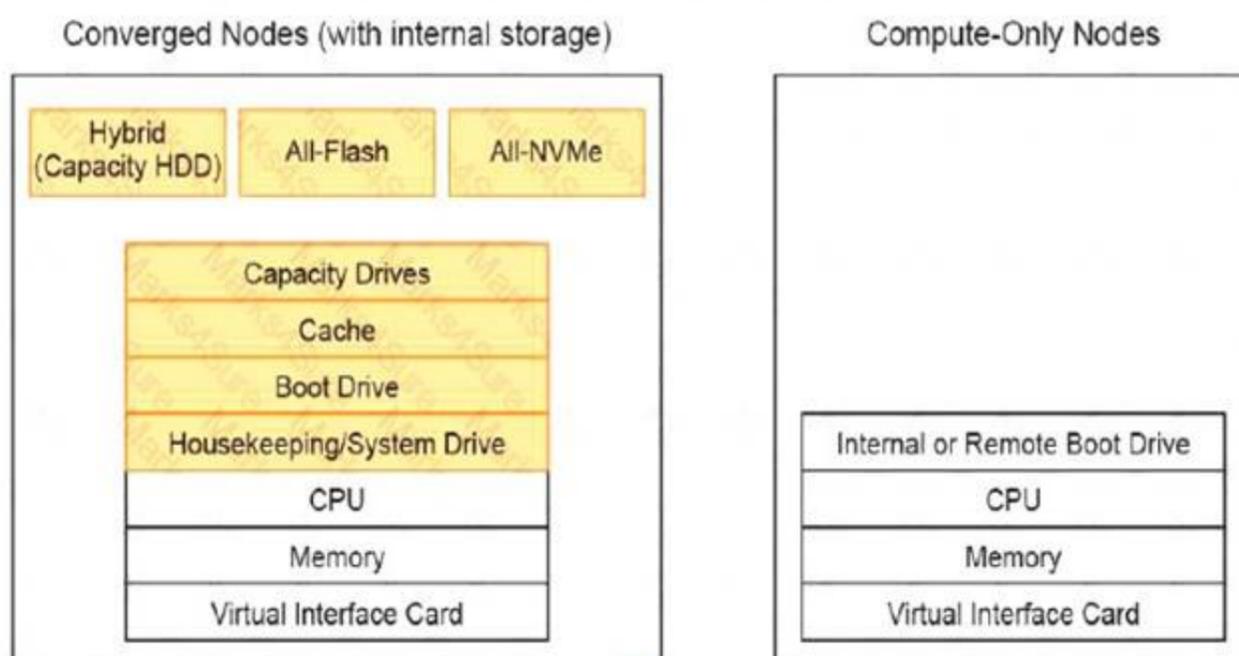
Explanation:

Compute-only nodes are part of the same vSphere cluster as the converged nodes. Since **compute-only nodes do not have storage, they utilize resources available from the converged nodes.**

Storage Components of Cisco HyperFlex Converged Nodes

Cisco HyperFlex converged nodes differ from compute-only nodes by the internal storage resources that they contribute to the overall storage pool. These storage resources include the capacity drives and cache drives.

The figure illustrates a high-level diagram of hardware components of HyperFlex servers.



NEW QUESTION 31

.....

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>