

## 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

There are often disadvantages when using mixed hardware, and uniform systems are highly advised. What are three disadvantages of using validated yet mixed hardware? (Choose three.)

- A. "Hardware only" solutions require less latency for data aggregation.
- B. "Software only" solutions are inherently more secure, adhering to government regulations.
- C. "Software only" HCI solutions often mean that HCI vendors do not have control over the hardware that the consumer is using.
- D. The lack of control over the hardware can result in hardware not being security-compliant (usually the case with white box solutions) or not being optimally configured for the given HCI solution.
- E. Validated yet "mixed" hardware systems may be supported by the specific original equipment manufacturers but software vendors require a master level services agreement.
- F. "Software only" solutions can result in a lack of simplicity for installation and maintenance.
- G. While setup and maintenance of an HCI solution might be simple you still need to configure and maintain the hardware part, which is not orchestrated by the HCI solution.
- H. Validated hardware does not mean smooth performance, which particularly applies to hardware that is seldom chosen for the hardware part of the "software only" solution.

**Answer: CEF**

**Explanation:**

### What Is Cisco HyperFlex?

Cisco HyperFlex is a Cisco interpretation of what a hyperconverged solution should look like. HyperFlex tightly integrates software and hardware, for in an easy to install and easy to operate solution.

Some vendors define hyperconverged solutions as software only, which means that when you buy hyperconverged software, you can use any hardware. But that is not really true, because you must follow the hardware compatibility list (HCL) for your individual solution or you will not have vendor support. There are also often huge disadvantages when using mixed hardware, and uniform systems are highly recommended.

While choice is a good thing, it comes with a cost:

- Validated hardware does not mean smooth performance, which particularly applies to hardware that is seldom chosen for the hardware part of the "software only" solution.
- "Software only" HCI solutions often mean that HCI vendors do not have control over the hardware that the consumer is using. The lack of control over the hardware can result in hardware not being security-compliant (usually the case with white box solutions) or not being optimally configured for the given HCI solution.
- "Software only" solutions can result in a lack of simplicity for installation and maintenance. While setup and maintenance of an HCI solution might be simple you still need to configure and maintain the hardware part, which is not orchestrated by the HCI solution.

#### NEW QUESTION 3

When building a HyperFlex cluster which two recommendations should be followed? (Choose two.)

- A. Use HX 220s for compute nodes and HX 240s for converged nodes

- B. Use B-Series servers to improve converged node scale.
- C. Use the same CPU model but memory configuration can be different.
- D. Use the same server configuration for the cluster.
- E. Use the same server model for the cluster.

**Answer:** DE

#### NEW QUESTION 4

Which two features enable RAID cards striping as well as mirroring and parity? (Choose two.)

- A. Integration with Cisco Intersight for a cloud-based storage management solution.
- B. No load on the system resources, drives seem as one drive to the operating system
- C. On RAID card failure, the RAID onboard concurrent cache assists rebuild cache.
- D. Hot replacement of drives available, depending on configuration
- E. Distributed drives across disparate systems can be in RAID together.

**Answer:** BD

**Explanation:**

RAID cards enable striping as well as **mirroring and parity**, with these features:

- No load on the system resources, drives seem as one drive to the operating system.
- Hot replacement of drives available, depending on configuration.
- Disk replacements require RAID rebuilds, taking a long time.
- On RAID card failure, the RAID card compatibility can be an issue.
- Limited drives in a raid field, depending on solution, limiting scalability.
- Only local drives can be in RAID together.

#### NEW QUESTION 5

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 6

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 7

Which two processes does Disk Failure initiate? (Choose two)

- A. The affected cluster is marked as unhealthy and placed into standby mode
- B. If the replication factor is sufficient for the failure, the system is marked as unhealthy but remains operational
- C. Distributed pooled data is migrated off nodes to master data store.
- D. Performance is almost unaffected Sets 1-minute timer until self-healing starts.
- E. Self-healing mode is initiated and data replication factors applied.

**Answer: BD**

**Explanation:**

## Disk Failure

Disk failure **initiates** this process:

1. If the replication factor is sufficient for the failure, the system is marked as unhealthy but remains operational.
2. VM running on the node is not migrated and the input/output continues from copies.
3. Performance is almost unaffected. Sets 1-minute timer until self-healing starts.
4. After 1 minute, the missing pieces are re-created from the remaining instances.

### NEW QUESTION 8

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 9

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)

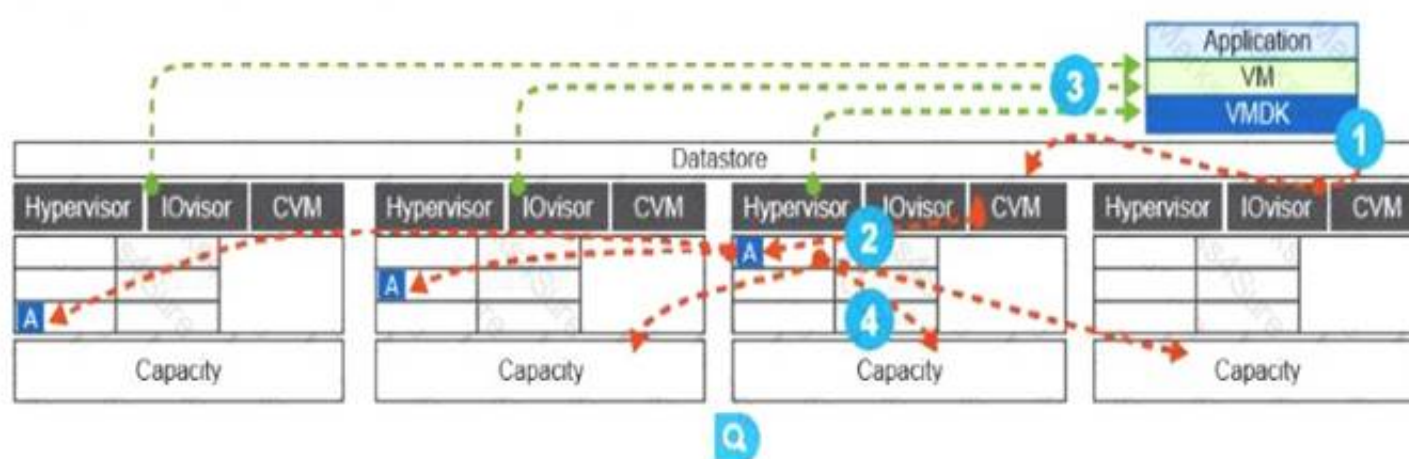
- A. The primary CVM compresses the data, writes it to its cache drive and mirrors it
- B. ACK is sent to the CVM that the write is about to be initiated
- C. On write, the local IOvisor sends the write to the primary CVM for that block
- D. On read/Writ
- E. the distributed VAAI sends the write to the 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 10

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 10

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 15

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 20**

What is the maximum number of cores supported in the Cisco UCS M5 server?

- A. 28
- B. 22
- C. 12
- D. 8

**Answer:** A

**NEW QUESTION 22**

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 24**

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 29

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 31

How many vCPUs does the HXDP controller VM require?

- A. 8
- B. 6
- C. 2
- D. 4

**Answer:** A

**Explanation:**



## CPU and Memory Guidelines

When selecting the most appropriate CPU for your cluster, you should consider the overhead consumed by 1 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 I FF reserved.

### NEW QUESTION 34

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

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

**Answer: C**

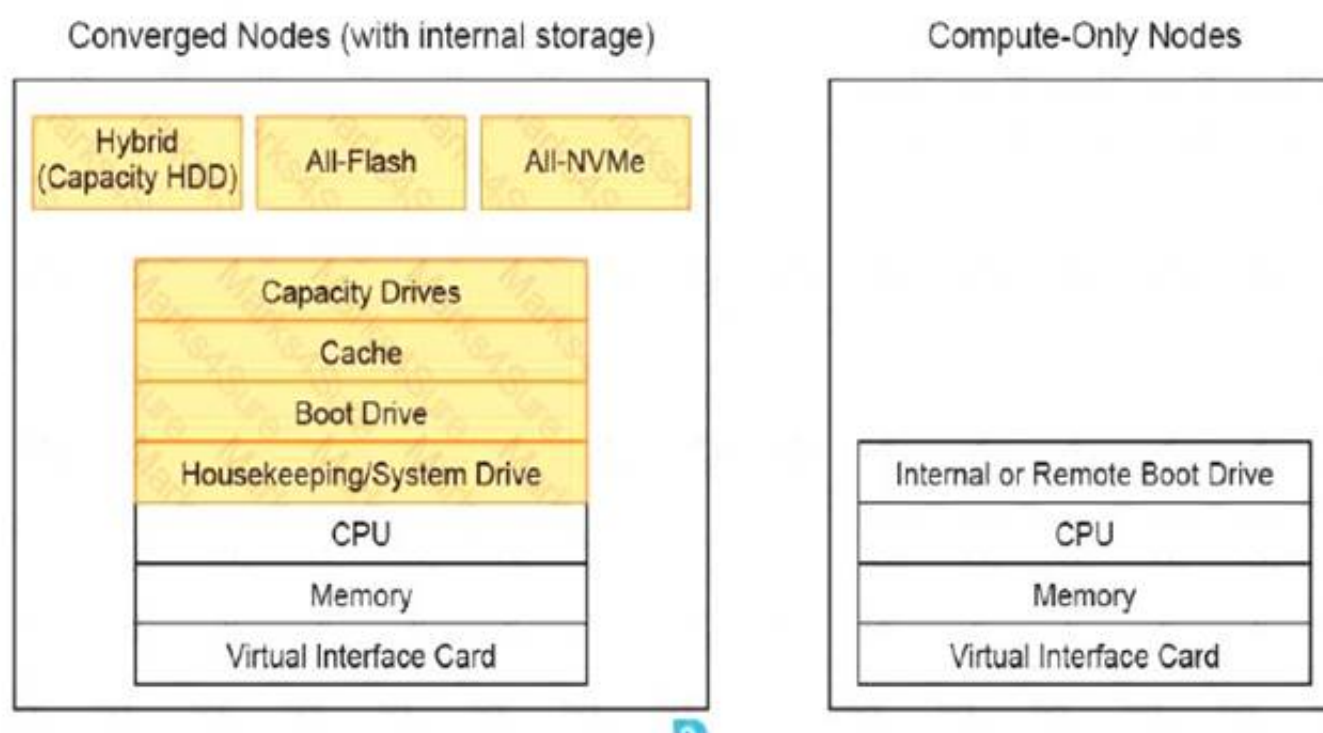
**Explanation:**

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 37

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 41

.....

## 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>