

700-905 Dumps

Cisco HyperFlex for Systems Engineers

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



NEW QUESTION 1

When cabling a given HX node to the Fabric Interconnect which three actions are required? (Choose three.)

- A. Connect the node to different port numbers on each of the two Fabric Interconnects.
- B. Connect port 1 on the VIC to Fabric Interconnect A.
- C. Connect server port L1 to Fabric Interconnect port L1.
- D. Connect the node to the same port number on each of the two Fabric Interconnects.
- E. Connect server port L2 to Fabric Interconnect port L2
- F. Connect port 2 on the VIC to Interconnect B.

Answer: BDF

Explanation:

Connect Fabric Interconnect heartbeat: L1-L1 and L2-L2 ports. Optionally connect console management cables to terminal server.

Connect VIC ports on each server to Fabric Interconnects. One port to Fabric Interconnect A, one to Fabric Interconnect B.

Connect uplink both Fabric Interconnects to upstream switch. And connect the IP out-of-band (OOB) management to an access port.

NEW QUESTION 2

Which two statements on Replication Factors are valid? (Choose two.)

- A. RF2 has 100% better protection of data while using 33% more space and is recommended.
- B. RF3 has 100% better protection of data while using 33% more space and is recommended.
- C. LAZ can increase RTO for 16+ node clusters.
- D. When performing rolling upgrades, the upgraded node is down for maintenance, stretching RF2 to the absolute limit.
- E. LAZ can decrease fault tolerance of S+ node clusters to less than 4

Answer: BD

Explanation:

Maximum Tolerable Failures Dependent on Replication Factor

	3-4 Nodes	5 Nodes or More (No LAZ)
Replication Factor 2	1 node or 1 drive failure	1 node failure or 1 drive failure
Replication Factor 3	1 node failure or 2 drive failures	2 node failures or 2 drive failures

Replication factor implies:

- RF3 has 100 percent better protection of data while using 33 percent more space and is recommended.
- When performing rolling upgrades, the upgraded node is down for maintenance, stretching RF2 to the absolute limit.

NEW QUESTION 3

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 4

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 5

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 6

HyperFlex uses file system native snapshots and provides which three features? (Choose three.)

- A. on datastore level, snapshots work the same as vSphere snapshots
- B. automatically deduplicating data of snapshots through StorFS
- C. impact on the VM performance after a lot of writes, requiring future administration
- D. limitations in age and number of snapshot
- E. See user manual for limitations
- F. consolidation of snapshots is still manual but not necessary
- G. automatic update to the golden image when a configuration change is made

Answer: ABD

NEW QUESTION 7

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 8

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 9

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 10

A Controller Virtual Machine (CVM) is an Ubuntu Linux VM that lives outside the converged data platform on the housekeeping drive since it is involved in creating the convergence data platform Which two features for CVMs are valid? (Choose two.)

- A. creates hooks for services related to third-party abstraction applications
- B. does not perform caching, deduplication, and compression of data
- C. an Ubuntu based VM running in the control space of each individual server, having linear access to the server's VMs and networking controls
- D. needs network access to ESX
- E. other CVMs, and management network
- F. is installed automatically by the HyperFlex installer, configured through the installer

Answer: DE

Explanation:

A CVM is an Ubuntu Linux VM that lives outside the converged data platform on the housekeeping drive, since it is involved in creating the converged data platform.

The CVMs have these features:

- An Ubuntu based VM running in the hypervisor of each individual server, having direct access to the server's storage.
- Is installed automatically by the HyperFlex installer, configured through the installer.
- Needs network access to ESXi, other CVMs, and management network.
- Performs caching, deduplication, and compression of data.
- Utilizes IOVisor to distribute data across the HyperFlex cluster.
- Provides HX Connect, HyperFlex CLI, and REST API for management.
- CVMs are responsible for logging.

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 11

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 12

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 13

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 15

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 19

Which two processes does failure on a node initiate? (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. A call-home process is initiated and the failure is reported to TAC
- D. The VMs on the failed node are moved to another node by vSphere high availability
- E. The system is marked unhealthy but remains operational.

Answer: DE

Explanation:

Node Failure

Failure on a node **initiates** the following process:

1. The system is marked unhealthy but remains operational.
2. The VMs on the failed node are moved to another node by vSphere high availability.
3. VMs keep reading from the remaining copies with minimal impact to performance.
4. A 2-hour countdown **initiates** before self-healing process.

NEW QUESTION 24

.....

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>