

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

Which three features do Managed Deployments provide? (Choose three.)

- A. Great for managing large deployments, scalability, and oversight of the UCS servers.
- B. Consistent deployment by replicating working configurations from development labs to the production deployment
- C. High availability of the management system and connectivity when using 2 Cisco Fabric Interconnects
- D. Individual configuration of each redundant fabric or global configuration.
- E. Increased operating overhead, raising Operating Expenses (OpEx)
- F. Decentralized yet complex management of an entire UCS domain.

Answer: ACD

Explanation:

Managed deployments provide these features:

- Centralized and simplified profile-based management of the entire Cisco UCS domain.
- Individual configuration of each redundant fabric or global configuration.
- High availability of the management system and connectivity when using two Cisco Fabric Interconnects.
- Great for managing large deployments, scalability, and oversight of the Cisco UCS servers.
- Reduced operating overhead, lowering operating expenses (OpEx).

In the context of Cisco HyperFlex, the centralized management platform for the entire cluster allows the HyperFlex installation to configure the servers automatically. The installation is therefore much simpler than if you had to configure the BIOS, disk drives, networking, and other hardware related features yourself.

NEW QUESTION 3

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 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

Which three features for NVMe transfer protocol are valid? (Choose three.)

- A. Uses PCIe interface
- B. Streamlined commands for fewer CPU instructions
- C. More and deeper queues
- D. Requires 10 controller
- E. Increases memory speed for more IOPs
- F. Improves SAS and SATA speed

Answer: ABC

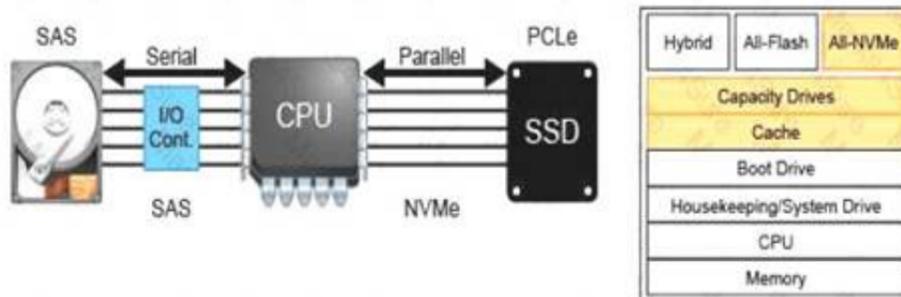
Explanation:

Transfer Protocol: NVMe

Earlier technologies (such as SAS and SATA) that were architected for hard disk drives are unable to take the full advantage of SSDs' potential. This problem warranted the need for a new architecture: Non-Volatile Memory Express (NVMe).

NVMe is a transfer protocol with these features:

- Uses the PCIe interface.
- Does not need an I/O controller and communicates directly with the CPU.
- Uses many more and deeper queues for command submission than SAS/SATA.
- Streamlines the command set to generate fewer CPU instructions.
- In HyperFlex:
 - NVMe (regular or Optane) disks can be used for cache in All-Flash version.
 - Regular NVMe is capacity and Optane is cache in All-NVMe version.



NEW QUESTION 6

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 7

Which three additional management tools are included in HXDP to configure HyperFlex clusters? (Choose three.)

- A. UCS Manager
- B. Storage CLI
- C. Data Center Network Manager
- D. Cisco IMC 13
- E. REST API
- F. HyperFlex Connect

Answer: BEF

Explanation:

Three management tools cover a similar configuration scope related to native HyperFlex features: HyperFlex Connect, Storage CLI, and REST API. HX Connect is an HTML5-based web interface; stcli is CLI-based and lends itself very well to troubleshooting. REST API offers the optimal solution when you integrate the HyperFlex system with RESTful orchestration tools.

NEW QUESTION 8

How can the maximum IOPS 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

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 11

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 16

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 17

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 19

.....

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>