

# Cisco

## Exam Questions 700-905

Cisco HyperFlex for Systems Engineers



**NEW QUESTION 1**

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 2**

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 3**

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 4**

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 5

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 6

How many memory channels does the Cisco UCS M5 server support per CPU?

- A. 1
- B. 2
- C. 6
- D. 8

**Answer: C**

### NEW QUESTION 7

What does the letter W indicate when selecting CPUs for your HX Node (le. 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
<b>HX-CPU-8170M</b>	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 8

HyperFlex virtual servers differ from regular servers in which two key areas? (Choose two.)

- A. NVMe: Regular servers do not support NVMe drives for high availability.
- B. No RAID is required to consolidate disks into a shared data platform.
- C. CVM: Virtual appliance, which performs reading/writing, caching, deduplication, and compression.
- D. SP: UCS Service Profiles are used to delineate MAC address pools from upstream networks.
- E. CCC: Cisco Cloud Center is used for multi-cloud integration and seamless deployment.

**Answer: BC**

**Explanation:**

HyperFlex virtual servers differ from **regular** servers in these key areas:

- **No RAID** is required to consolidate disks into a shared data platform.
- **CVM**: Virtual appliance, which performs reading/writing, caching, deduplication, and compression.
- **IOVISOR**: Hypervisor driver, which mounts HyperFlex storage and distributes data.
- **VAAI**: vSphere storage API allowing file-system-level snapshots and cloning.

### NEW QUESTION 9

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 10

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

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.)

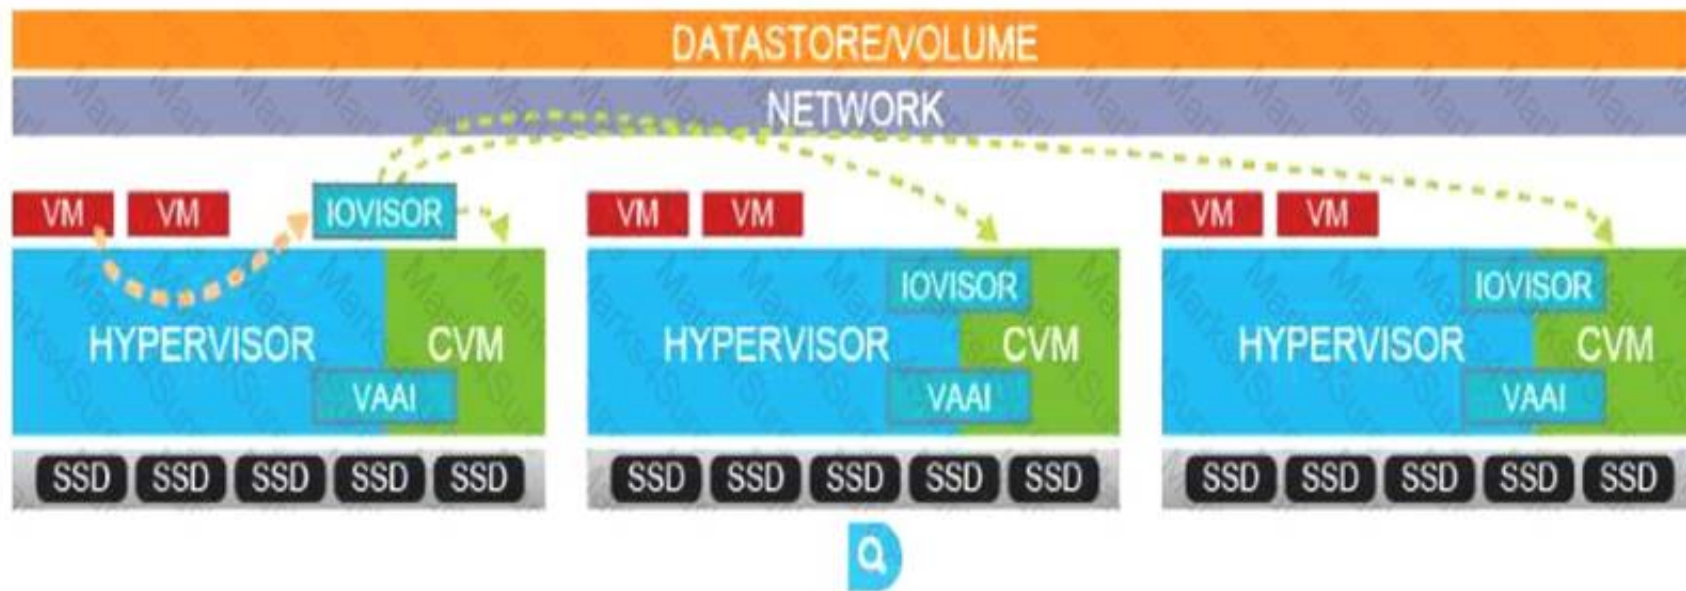
- A. provides redundancy when local CVM fails, offloading data processing to another CVM in the cluster
- B. when an IOVisor fails, the CVM remains active and functional, which enables uninterrupted operation of the system by forwarding IO to another available IOVisor in the HyperFlex cluster
- C. enables 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 11

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

- A. Determine and download recommended installer OVA version required
- B. Complete the pre-installation checklist.
- C. Determine and download recommended hypervisor
- D. Download service profile templates
- E. Determine and download virtual machine OS! required

**Answer:** AB

#### NEW QUESTION 16

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 18

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 23

What is the minimum amount of memory required for an HX node?

- A. 192 GB
- B. 64 GB
- C. 32 GB
- D. 128 GB

**Answer:** D

**Explanation:**

HyperFlex Edge servers have lower hardware requirements than standard HyperFlex servers:

- Cisco Fabric Interconnects are not part of the solution, hardware configured over Cisco IMC.
- Only 1 CPU per server required.
- Minimum 8 RAM sticks per server, up to 12 supported per CPU.
- **128** GB of RAM required, 192 GB recommended.
- 3-8 capacity drives (6-8 on standard HX 220).
- mLOM not required.
- PCIe NICs available with dual 10-G and quad 1-G RJ45 Ethernet connectivity.

### NEW QUESTION 28

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 33

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 37

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 42

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 44

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 49

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 vCPU's, 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 54

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

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

**Answer: C**

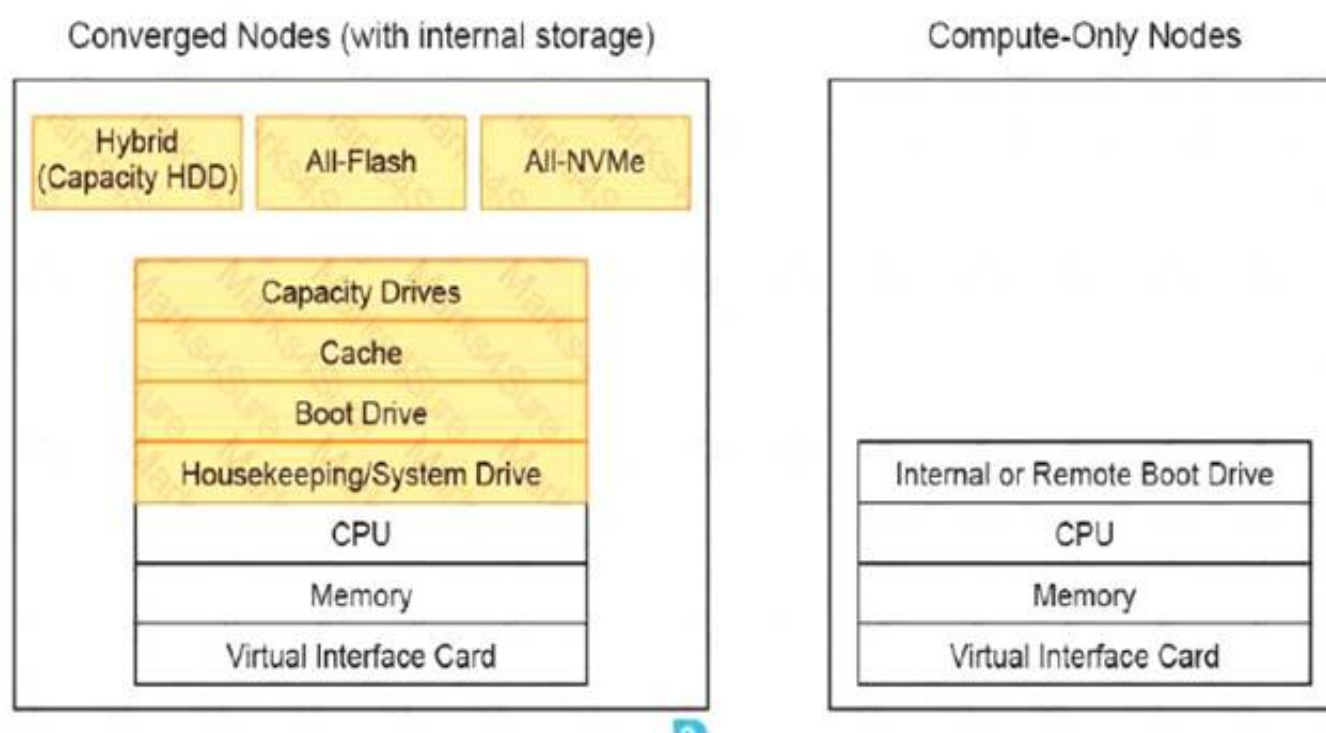
**Explanation:**

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 58

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 63

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### 700-905 Practice Exam Features:

- \* 700-905 Questions and Answers Updated Frequently
- \* 700-905 Practice Questions Verified by Expert Senior Certified Staff
- \* 700-905 Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* 700-905 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The 700-905 Practice Test Here](#)**