



# CV0-003<sup>Q&As</sup>

CompTIA Cloud+ Certification

## Pass CompTIA CV0-003 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/cv0-003.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by CompTIA  
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





### QUESTION 1

A systems administrator wants the VMs on the hypervisor to share CPU resources on the same core when feasible.

Which of the following will BEST achieve this goal?

- A. Configure CPU passthrough
- B. Oversubscribe CPU resources
- C. Switch from a Type 1 to a Type 2 hypervisor
- D. Increase instructions per cycle
- E. Enable simultaneous multithreading

Correct Answer: E

Simultaneous multithreading (SMT) is a type of CPU technology that allows multiple threads to run concurrently on a single CPU core. Enabling SMT can help achieve the goal of having the VMs on the hypervisor share CPU resources on the same core when feasible, as it can increase the CPU utilization and efficiency by executing more instructions per cycle and reducing idle time or wasted cycles. Enabling SMT can also improve performance and throughput, as it can speed up processing and handle increased workload or demand. References: CompTIA Cloud+ Certification Exam Objectives, page 9, section 1.4

---

### QUESTION 2

A company is planning its cloud architecture and wants to use a VPC for each of its three products per environment in two regions, totaling 18 VPCs. The products have interdependences, consuming services between VPCs. Which of the following should the cloud architect use to connect all the VPCs?

- A. MPLS connections
- B. VPC peering
- C. Hub and spoke
- D. VPN connections

Correct Answer: C

The best way to connect all the VPCs for the company that is planning its cloud architecture and wants to use a VPC for each of its three products per environment in two regions, totaling 18 VPCs, is to use a hub and spoke model. A hub and spoke model is a networking model that uses a central hub VPC that connects to multiple spoke VPCs that host the products or workloads. The hub VPC can provide common services and security policies for the spoke VPCs, such as network virtual appliances, DNS servers, firewalls, or VPN gateways. The spoke VPCs can communicate with each other through the hub VPC, using private IP addresses or peering connections. A hub and spoke model can simplify the management and scalability of the network, as well as reduce the costs and complexity of peering multiple VPCs directly. Reference: Hub-and-spoke network topology - Cloud Adoption Framework

---

### QUESTION 3



A cloud administrator is reviewing the annual contracts for all hosted solutions. Upon review of the contract for the hosted mail solution, the administrator notes the monthly subscription rate has increased every year. The provider has been in place for ten years, and there is a large amount of data being hosted. Which of the following is a barrier to switching providers?

- A. Service-level agreement
- B. Vendor lock-in
- C. Memorandum of understanding
- D. Encrypted data

Correct Answer: B

Vendor lock-in is a barrier to switching providers for a hosted mail solution that has increased its monthly subscription rate every year. Vendor lock-in is a situation where a customer becomes dependent on a vendor or provider for a product or service and faces difficulties or costs in switching to another vendor or provider. Vendor lock-in can occur due to various factors, such as proprietary technology, contractual obligations, data migration challenges, compatibility issues, etc. In this case, the customer may face vendor lock-in due to the large amount of data being hosted by the mail provider and the potential challenges or costs of transferring or migrating the data to another provider.

#### QUESTION 4

A company has an SLA of 15ms for storage latency.

	90 days ago	60 days ago	30 days ago	Now
Space usage	60%	70%	75%	80%
Average latency	8ms	10ms	15ms	25ms
Storage CPU utilization	50%	65%	80%	95%

Given the above metrics, which of the following is the MOST appropriate change to the environment?

- A. Add computing nodes and data nodes on the storage side.
- B. Add more storage shelves to the current array.
- C. Enable compression on the storage side.
- D. Enable deduplication on the storage side.

Correct Answer: C

#### QUESTION 5

A systems administrator recently deployed a VDI solution in a cloud environment; however, users are now experiencing poor rendering performance when trying to display 3-D content on their virtual desktops, especially at peak times.

Which of the following actions will MOST likely solve this issue?



- A. Update the quest graphics drivers from the official repository
- B. Add more vGPU licenses to the host
- C. Instruct users to access virtual workstations only on the VLAN
- D. Select vGPU profiles with higher video RAM

Correct Answer: D

A vGPU profile is a configuration option that defines the amount of video RAM (vRAM) and other resources that are allocated to a virtual machine (VM) that uses a virtual graphics processing unit (vGPU). A vGPU profile can affect the rendering performance of a VM, as it determines how much graphics memory and processing power are available for displaying complex graphics content. Selecting vGPU profiles with higher video RAM can most likely solve the issue of poor rendering performance when trying to display 3-D content on virtual desktops, especially at peak times, as it can provide more graphics resources and improve the quality and speed of rendering. References: CompTIA Cloud+ Certification Exam Objectives, page 11, section 1.6

Reference: <https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/desktop-virtualization-solutions-vmware-horizon-view/whitepaper-c11-741606.pdf>

---

#### QUESTION 6

A cloud administrator must ensure all servers are in compliance with the company's security policy. Which of the following should the administrator check FIRST?

- A. The application version
- B. The OS version
- C. Hardened baselines
- D. Password policies

Correct Answer: C

Reference: <https://www.securitymetrics.com/blog/system-hardening-standards-how-comply-pci-requirement-22>

---

#### QUESTION 7

Which of the following cloud services is fully managed?

- A. IaaS
- B. GPU in the cloud
- C. IoT
- D. Serverless compute
- E. SaaS

Correct Answer: E



SaaS (Software as a Service) is a cloud service model that provides fully managed applications to the end users. The users do not have to worry about installing, updating, or maintaining the software, as the cloud provider handles all these tasks. Examples of SaaS are Gmail, Office 365, Salesforce, etc.

---

#### QUESTION 8

An organization is currently deploying a private cloud model. All devices should receive the time from the local environment with the least administrative effort. Which of the following ports needs to be opened to fulfill this requirement?

- A. 53
- B. 67
- C. 123
- D. 161

Correct Answer: C

Port 123 is what needs to be opened to ensure all devices receive the time from the local environment with the least administrative effort in a private cloud model. Port 123 is the port used by NTP (Network Time Protocol), which is a protocol that synchronizes the clocks of network devices and systems. NTP can help to ensure accurate and consistent time across different devices and systems in a cloud environment, which can facilitate coordination, communication, logging, auditing, etc.

Reference: <https://www.sciencedirect.com/topics/computer-science/network-time-protocol#:~:text=NTP%20is%20a%20built%20on,for%20example%2C%20a%20desktop>.

---

#### QUESTION 9

A company has hired a security firm to perform a vulnerability assessment of its environment. In the first phase, an engineer needs to scan the network services exposed by the hosts. Which of the following will help achieve this with the LEAST privileges?

- A. An agent-based scan
- B. A credentialed scan
- C. A network-based scan
- D. An application scan

Correct Answer: C

A network-based scan is a type of vulnerability assessment that scans the network services exposed by the hosts without requiring any credentials or agents. This type of scan will help achieve the objective of scanning the network services with the least privileges, as it does not need any access to the hosts or their internal configurations. A network-based scan can identify open ports, running services, and potential vulnerabilities on the hosts. Reference: CompTIA Cloud+ Certification Exam Objectives, Domain 2.0 Security, Objective 2.4 Given a scenario, implement security automation and orchestration in a cloud environment.

---



### QUESTION 10

A systems administrator needs to configure SSO authentication in a hybrid cloud environment.

Which of the following is the BEST technique to use?

- A. Access controls
- B. Federation
- C. Multifactor authentication
- D. Certificate authentication

Correct Answer: B

Federation is a type of authentication mechanism that allows users to access multiple systems or applications across different domains or organizations with a single login credential. Federation can help configure SSO authentication in a hybrid cloud environment, as it can enable seamless and secure access to cloud-based and on-premises resources using the same identity provider and authentication method. Federation can also improve user convenience, productivity, and security, as it can simplify the login process, reduce login errors, and enhance password management. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

---

### QUESTION 11

Company A has acquired Company B and is in the process of integrating their cloud resources. Company B needs access to Company A's cloud resources while retaining its IAM solution. Which of the following should be implemented?

- A. Multifactor authentication
- B. Single sign-on
- C. Identity federation
- D. Directory service

Correct Answer: C

Identity federation is a type of authentication mechanism that allows users to access multiple systems or applications across different domains or organizations with a single login credential. Identity federation can help integrate the cloud resources of Company A and Company B after Company A has acquired Company B, as it can enable seamless and secure access to both companies' cloud resources using the same IAM solution. Identity federation can also improve user convenience, productivity, and security, as it can simplify the login process, reduce login errors, and enhance password management. References: CompTIA Cloud+ Certification Exam Objectives, page 14, section 2.7

Reference: <https://medium.com/@dinika.15/identity-federation-a-brief-introduction-f2f823f8795a>

---

### QUESTION 12

Different healthcare organizations have agreed to collaborate and build a cloud infrastructure that should minimize compliance costs and provide a high degree of security and privacy, as per regulatory requirements. This is an example of a:



- A. private cloud.
- B. community cloud.
- C. hybrid cloud.
- D. public cloud.

Correct Answer: B

Reference: <https://www.sciencedirect.com/topics/computer-science/community-cloud>

---

### QUESTION 13

A cloud administrator configured a local cloud-resource pool to offer 64GB of memory, 64 cores, and 640GB of storage. Thirty-two machines with identical resource allocations are started, but one machine is unable to handle requests. Which of the following is the MOST likely cause?

- A. Insufficient guest bandwidth
- B. Overwhelmed vCPU
- C. A storage error on the guest
- D. Incorrect VLAN assignment
- E. Inadequate memory allocation

Correct Answer: E

---

### QUESTION 14

A systems administrator would like to reduce the network delay between two servers.

Which of the following will reduce the network delay without taxing other system resources?

- A. Decrease the MTU size on both servers
- B. Adjust the CPU resources on both servers
- C. Enable compression between the servers
- D. Configure a VPN tunnel between the servers

Correct Answer: A

The maximum transmission unit (MTU) is the largest size of a packet or frame that can be sent over a network. Decreasing the MTU size on both servers can reduce the network delay between them, as it can reduce the fragmentation and reassembly of packets, improve the transmission efficiency, and avoid packet loss or errors. Decreasing the MTU size can also avoid taxing other system resources, as it does not require additional CPU, memory, or disk resources. References: CompTIA Cloud+ Certification Exam Objectives, page 16, section 3.2



Reference: <https://cseweb.ucsd.edu/~calder/papers/HPDC-01-DynComp.pdf>

---

### QUESTION 15

A systems administrator is troubleshooting network throughput issues following a deployment. The network is currently being overwhelmed by the amount of traffic between the database and the web servers in the environment.

Which of the following should the administrator do to resolve this issue?

- A. Set up affinity rules to keep web and database servers on the same hypervisor
- B. Enable jumbo frames on the gateway
- C. Move the web and database servers onto the same VXLAN
- D. Move the servers onto thick-provisioned storage

Correct Answer: A

[Latest CV0-003 Dumps](#)

[CV0-003 Study Guide](#)

[CV0-003 Exam Questions](#)