

Cloudera

CDP-5001

CDP Administrator - Public Cloud Certification Exam

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

- 1. Up to Date products, reliable and verified.**
- 2. Questions and Answers in PDF Format.**



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/cdp-5001>

Latest Version: 6.1

Question: 1

You're setting up a CDP environment with a requirement for a private connection between the cloud environment and an on-premises data center. Which of the following solutions should you use?

- A. VPN Gateway
- B. Transit Gateway
- C. VPC peering
- D. Direct Connect (or equivalent service depending on the provider)

Answer: D

Explanation:

A: VPN Gateways create secure connections over the internet but may have bandwidth and latency limitations for a CDP deployment.

B: Transit Gateways simplify the connection between multiple VPCs in the cloud but don't address hybrid cloud connectivity.

C: VPC Peering connects VPCs within the same cloud provider but does not extend to on-premises networks.

Question: 2

Your CDP cluster is running a workload with occasional, unpredictable spikes in resource consumption. Which combination of scaling and instance types would provide the most cost-effective solution?

- A. Vertical Scaling and On-Demand Instances
- B. Horizontal Scaling and Reserved Instances
- C. Horizontal Scaling and Spot Instances
- D. Vertical Scaling and Spot Instances

Answer: C

Explanation:

A: Vertical scaling increases the power of individual instances, potentially incurring higher cost, and might not handle unpredictable spikes well.

B: Reserved instances offer discounts but don't suit unpredictable workloads.

C: Horizontal scaling with spot instances provides flexibility to add more nodes when needed and take advantage of lower costs of unused capacity.

D: Vertical scaling with spot instances might not be as cost-effective for unpredictable work.

Question: 3

You need to select the most appropriate compute instance type for the CDP Master Node. Which of the following considerations are the most important? (Choose two)

- A. High-performance storage (like NVMe SSD)
- B. Maximum network throughput
- C. Consistent and reliable memory availability
- D. Largest possible number of CPU cores

Answer: B,C

Explanation:

- A: While fast storage is beneficial, it's less critical for the Master Node than other components.
- B: Network throughput is crucial for Master Node communication with workers and other cluster services.
- C: Master Node holds critical metadata, requiring reliable memory.
- D: The number of cores is less important than memory and network performance for the Master Node.

Question: 4

You are deploying a CDP cluster that must access data stored in an object storage service (e.g., S3). To allow this access, which of the following should you configure and consider? (Select two.)

- A. An IAM role granting the required object storage permissions to the CDP cluster instances.
- B. A route within the VPC route table to allow traffic to flow to the object storage endpoint.
- C. An S3 Gateway endpoint, providing private connectivity to S3 from within the VPC.
- D. An Internet Gateway attached to your VPC to route traffic to the public S3 endpoints.

Answer: A,D

Explanation:

- A: IAM roles are critical for authorizing services to interact with other cloud resources.
- B: A route table may be needed, but not specifically for reaching public service endpoints like S3.
- C: An S3 gateway offers enhanced connectivity but may not be strictly necessary.
- D: An Internet Gateway enables access to services outside the VPC, including S3.

Question: 5

Which of the following scenarios would typically NOT be a reason to create multiple, separate VPCs for your CDP deployments?

- A. Isolating environments based on their purpose (production, development, testing).
- B. Applying different security policies and network controls.
- C. Implementing overlapping IP address ranges between environments.
- D. Simplifying network management and troubleshooting.

Answer: D

Explanation:

A, B, C: All these are valid reasons for using separate VPCs.

D: Separate VPCs can sometimes complicate network management due to additional configurations required (peering, route tables, etc.).

Compute Architecture

Question: 6

You anticipate a CDP workload involving data processing jobs that benefit significantly from GPU acceleration. Which type of compute instances should you consider?

- A. Memory-optimized Instances
- B. Instances with enhanced network capabilities
- C. GPU-accelerated Instances
- D. General purpose instances

Answer: C

Explanation:

Data processing jobs involving machine learning or other GPU-intensive tasks will see the most benefit from GPU-accelerated instances.

Question: 7

Your data engineers need to run frequent, short-lived batch jobs for data transformations. To optimize for cost-savings in this scenario, which of the following is the most appropriate strategy?

- A. Use reserved instances with a 3-year commitment.
- B. Use on-demand instances to ensure instant availability.
- C. Implement auto-scaling with a mix of on-demand and spot instances.
- D. Vertically scale compute instances as needed.

Answer: C

Explanation:

A: Reserved instances don't suit short-lived workloads.

B: On-demand instances can be more expensive for frequent jobs.

C: Auto-scaling with spot instances offers cost-effectiveness for intermittent, short-lived workloads.

D: Vertical scaling is less flexible for fluctuating resource needs.

Question: 8

You need to ensure that critical CDP cluster metadata is stored persistently in case of node failures. Which storage option would be the most suitable for this?

- A. Instance Store
- B. Network File System (NFS)
- C. Amazon Elastic Block Store (EBS) (or equivalent depending on the cloud provider)
- D. Object Storage (like S3)

Answer: C

Explanation:

- A: Instance stores are temporary.
- B: NFS offers file-level access, potentially overkill for cluster metadata.
- C: EBS provides persistent block storage, ideal for critical data requiring resilience.
- D: Object storage is less suitable for file-system-like operations, often used for archival or less structured data.

Question: 9

Your CDP cluster requires low-latency, high-bandwidth communication between compute nodes. You're considering different networking options within your cloud provider. Which of the following should you likely choose?

- A. Standard network configuration within a VPC.
- B. Transit Gateway to connect the VPC to another region.
- C. VPC Peering to connect the VPC to on-prem infrastructure
- D. Placement groups to cluster instances within an Availability Zone

Answer: D

Explanation:

- A: Standard network setups might not guarantee low latency.
- B: Transit Gateways simplify multi-VPC connections but don't directly optimize low latency within a single VPC.
- C: VPC Peering connects VPCs or hybrid environments, not expressly for low latency within a VPC.
- D: Placement groups ensure instances are physically close, minimizing network latency.

Question: 10

You need to collect detailed network performance metrics from your CDP cluster's instances for troubleshooting. Which of the following would provide the most granular level of network data?

- A. Basic Cloud Monitoring
- B. Custom metrics written to object storage
- C. VPC Flow Logs
- D. Enhanced Networking metrics (if supported by your cloud provider)

Answer: D

Explanation:

A: Basic cloud monitoring gives a general overview but may lack fine-grained network data.

B: Custom metrics require implementation effort and might not capture all network metrics.

C: VPC Flow Logs track network traffic at a flow level but may not have all the detailed metrics needed.

D: Enhanced Networking metrics provide the most in-depth network performance information on supported instance types.

Thank You for Trying Our Product
Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>