PassQueen
a.k.a dumpsqueen.com

# Specialist – Cloud Architect, Cloud Infrastructure Exam

## EMC DES-2T13

**Version Demo**

**Total Demo Questions: 11**

**Total Premium Questions: 114**

**Buy Premium PDF**

https://passqueen.com

support@passqueen.com

# Topic Break Down

| Topic | No. of Questions |
|---|---|
| Topic 1, Cloud Computing | 11 |
| Topic 2, Cloud Infrastructure | 59 |
| Topic 3, Cloud Services | 8 |
| Topic 4, Cloud Management | 22 |
| Topic 5, Cloud Security | 14 |
| Total | 114 |

QUESTION NO: 1

You are designing a cloud backup solution. Consider the following requirements when sizing the backup target:

- Total amount of data to be backed up = 2 TB (2000 GB)
- Full backup once a week
- Retention period for full backups = 2 months
- Daily incremental backups (assuming an average 10% change in data)
- Retention period for incremental backups = 1 month
- No Compression or deduplication
- Number of backup copies = 1

What is the total capacity, in GB, required for the backup target to accommodate the backups?

**A.** 13, 200

**B.** 16, 000

**C.** 21, 200

**D.** 26, 400

**ANSWER: C**

**Explanation:**

**21, 200** is the required backup-target capacity. With a weekly full-backup schedule and a two-month retention period, sizing convention treats two months as eight weeks. Therefore, eight full backups must be retained. Each full backup contains 2,000 GB, so full-backup capacity is 8 × 2,000 GB = 16,000 GB.

Each daily incremental backup contains the estimated changed data only: 10% of 2,000 GB, or 200 GB. A one-month incremental retention period is treated as 26 daily incremental backups for this calculation. Incremental-backup capacity is therefore 26 × 200 GB = 5,200 GB. Adding the concurrently retained full and incremental backups gives 16,000 GB + 5,200 GB = 21,200 GB.

This is raw logical backup capacity because the requirements explicitly state that neither compression nor deduplication is available, and only one backup copy is required. In a production design, capacity planning would normally also account for operational headroom, metadata, growth, and any platform-specific overhead; however, those factors are outside the stated sizing inputs. Dell's data-protection resources are available at Dell Data Protection, and product documentation can be accessed through Dell PowerProtect Data Manager documentation.

**QUESTION NO: 2**

An organization wants to offer a database service through its private cloud catalog. Consumers must be able to request the service without receiving administrative access to the underlying database hosts.

Which cloud characteristic does this requirement support?

**A.** On-demand self-service

**B.** Manual provisioning

**C.** Physical resource isolation

**D.** Direct host administration

**ANSWER: A**

**Explanation:**

**On-demand self-service** is supported because consumers can request and receive a standardized database service through a catalog without direct interaction with the infrastructure administrators or access to the underlying hosts. The cloud-management platform automates the approved deployment workflow while retaining administrative control of the infrastructure and service configuration. NIST defines on-demand self-service as a fundamental cloud characteristic in SP 800-145.

## QUESTION NO: 3

An organization wants to deploy a CMP solution for their private cloud environment. What is required for the CMP to work effectively?

**A.** Programmable infrastructure

**B.** Virtual firewall

**C.** Commodity hardware

**D.** Virtual router

## ANSWER: A

**Explanation:**

Programmable infrastructure is required because a cloud management platform (CMP) must be able to consistently provision, configure, monitor, and retire infrastructure resources through software-driven workflows and APIs. In a private cloud, the CMP provides consumers with a self-service experience while enforcing the organization's policies, approvals, quotas, service definitions, and lifecycle controls. To fulfill a request automatically, it needs underlying compute, storage, network, and security capabilities that can be discovered and controlled programmatically rather than requiring manual configuration by administrators. This programmability enables orchestration across domains, repeatable deployment of service templates, policy-based placement, elastic resource allocation, and accurate metering or chargeback. It also makes the infrastructure consumable as services, which is central to the private-cloud operating model. A CMP can integrate with many infrastructure products and vendors, but its effectiveness depends on those resources exposing manageable interfaces and supporting automation. Dell's cloud-management guidance describes orchestration and automation as foundational capabilities for delivering and managing cloud services, while the NIST cloud definition identifies on-demand self-service and rapid elasticity as essential characteristics enabled by automated resource provisioning. See NIST SP 800-145: The NIST Definition of Cloud Computing and Dell Technologies Cloud resources.

## QUESTION NO: 4

An organization uses Microsoft Active Directory (AD) for service authentication in their private cloud. They want to use the same authentication source for services in the public cloud. External employees access services in both clouds and need to authenticate with AD.

The organization wants to minimize and secure network traffic. Which solution will address these requirements?

**A.** Connect the two clouds through IPsec VPNPlace AD controllers in the private cloud only

**B.** Connect the two clouds through IPsec VPNPlace AD controllers in the private and public clouds

**C.** Configure perimeter firewalls to allow AD trafficPlace AD controllers in the private and public clouds

**D.** External users connect with IPsec VPNPlace AD controllers in the private cloud only

## ANSWER: B

**Explanation:**

**Connect the two clouds through IPsec VPNPlace AD controllers in the private and public clouds** is the appropriate design because it combines protected inter-cloud connectivity with local Active Directory Domain Services authentication

capacity in each cloud. An IPsec site-to-site VPN encrypts traffic traversing the public network and provides a controlled private path for essential AD replication and service communications. This meets the security requirement without exposing directory traffic directly to the internet.

Deploying domain controllers in both the private and public cloud locations lets workloads and users authenticate against a nearby domain controller. Active Directory uses sites, subnets, and replication topology information to direct clients toward local authentication services where possible. As a result, routine logon, Kerberos, LDAP, and directory lookup traffic does not need to cross the WAN link for every request. The encrypted VPN is then primarily used for directory replication and any required cross-site communications, reducing latency and minimizing inter-cloud network traffic while retaining one AD domain and a consistent identity source.

Microsoft's guidance on AD DS site design and replication is available at Selecting a site design for Active Directory Domain Services and Understanding Active Directory replication topology.

## QUESTION NO: 5

Case 1

Refer to the exhibit.

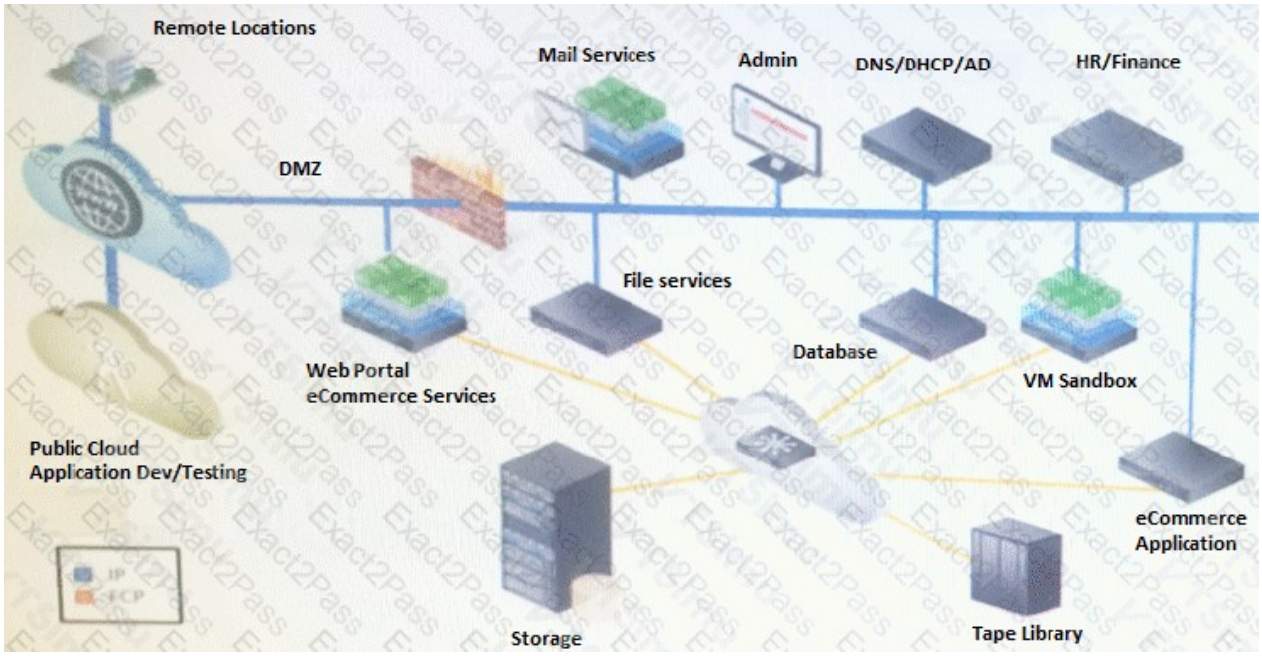


Your company recently negotiated several contracts with new customers to use its proprietary eCommerce system. As a result, the customer base has grown rapidly, which has introduced several challenges and requirements.

Business requirements:

Applications need to scale quickly

Metrics and reporting to anticipate business trends and efficient use of resources

External customers are looking to license the eCommerce application

Other considerations:

Applications need to be recorded for the public cloud

Address security vulnerabilities for customer data in the cloud

Refer to Case 1.

You are designing a transformation from Platform 2.0 to cloud. You review the existing infrastructure and

business requirements. You must define the method to connect the on-premises network directly to the public cloud provider.

Which technology should you select?

**A.** KMS

**B.** IDPS

**C.** MPLS

**D.** IPSec VPN

**ANSWER: C**

**Explanation:**

MPLS is the appropriate technology because it provides a carrier-managed private wide-area network that can extend an organization's on-premises network to a cloud provider through a service-provider connection. This approach supports the requirement for a direct, controlled path to public-cloud resources rather than relying on best-effort public Internet transport. MPLS VPN services logically isolate customer traffic and can provide predictable routing, traffic engineering, and service-level commitments that are valuable for an expanding eCommerce environment handling customer data and rapidly changing application demand.

In practice, the MPLS service is commonly integrated with the cloud provider's private-connect offering at a provider exchange or partner location. The organization retains a private WAN model while gaining access to cloud virtual networks, enabling more consistent performance and operational governance than an Internet-only design. Cloud providers also describe private connectivity as a way to establish dedicated network connections from on-premises environments to cloud services; for example, see AWS Direct Connect documentation and Microsoft Azure ExpressRoute documentation. MPLS therefore best matches the requested direct on-premises-to-public-cloud connectivity method.

**QUESTION NO: 6 - (DRAG DROP)**

You are tasked with designing a new cloud infrastructure. Match each compute related requirement with its proper category.

**Select and Place:**


**ANSWER:**

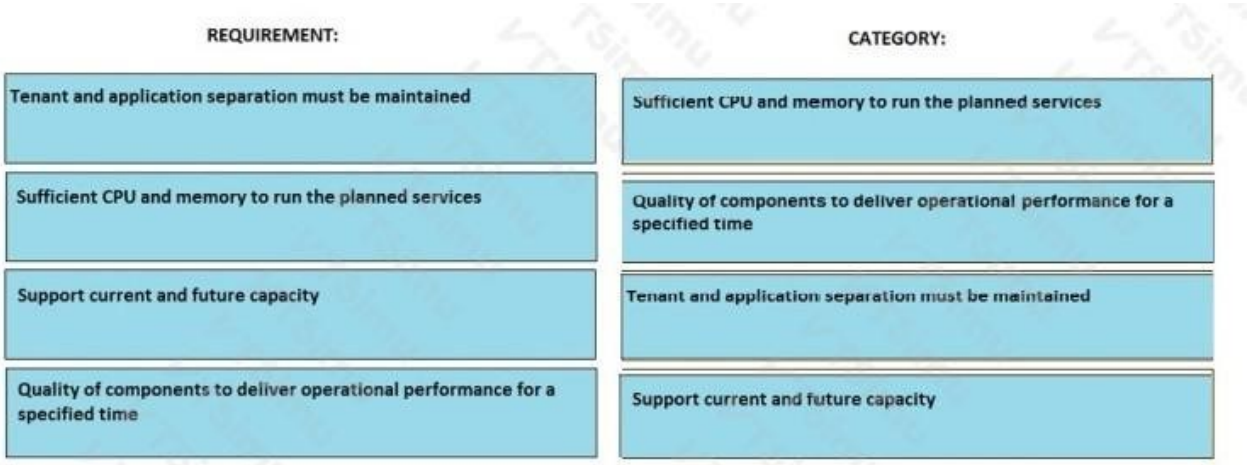


**Explanation:**

Capacity addresses the amount of compute resource needed to operate the intended workload. Therefore, sufficient CPU and memory for planned services belongs to Capacity, since processor and memory sizing directly determine whether the infrastructure can run those services at the required level. Scalability concerns the infrastructure's ability to accommodate changing demand over time, so supporting both current and future capacity is a scalability requirement. This includes designing compute resources so they can be expanded or adjusted as workloads grow.

Security includes isolation controls that protect tenants and applications from one another in a shared cloud environment. Maintaining tenant and application separation is therefore a security requirement because it enforces logical boundaries, limits unauthorized cross-tenant access, and supports multitenancy. NIST describes cloud security and isolation considerations in its Guidelines on Security and Privacy in Public Cloud Computing.

Availability is concerned with delivering the required operational service over a defined period. The quality and dependability of infrastructure components affect whether compute services remain operational for the stated time, making this an availability requirement. Availability engineering commonly relies on dependable components, resilient design, and appropriate maintenance practices. The relationship between dependable operation and availability is covered in the NIST/SEMATECH e-Handbook of Statistical Methods. Together, these mappings cover resource sufficiency, growth, isolation, and continued service operation in a cloud compute design.

## QUESTION NO: 7

You are a cloud architect responsible for designing a multiple availability zone configuration for the

organization's geographically dispersed clouds. You must provide the ability for services to survive a failure at the data center or site level. The primary requirement is to provide for automated site recovery.

What is a key consideration for automating the setup, failover, and recovery of services at the remote site?

**A.** processes

**B.** The number of storage volumes in the DR site must be identical in number and size to the primary site.

**C.** site.

**D.** The DR site must contain the same vendor hardware as the primary site.

**E.** Use orchestration runbooks that account for application dependencies and automate setup, failover, validation, and failback at the recovery site.

## ANSWER: E

**Explanation:**

Use orchestration runbooks that model application dependencies and automate the full recovery workflow, including service setup, startup sequencing, failover, validation, and failback. A site-recovery solution must recover more than replicated storage or virtual machines: it must bring up the dependent infrastructure and application components in an order that results

in a usable service. For example, network services, identity services, databases, middleware, and application tiers may each have prerequisites and required readiness checks. Capturing these dependencies in tested recovery plans removes error-prone manual activity during a site outage and makes recovery behavior repeatable.

Automated workflows should also define the recovery target, recovery order, network mappings, protection groups, and post-recovery verification steps. They should be exercised regularly through non-disruptive tests where possible, so that the documented recovery point and recovery time objectives can be demonstrated before an actual disaster. This aligns with disaster-recovery guidance that emphasizes documented, tested recovery procedures and coordinated restoration of system components. See NIST SP 800-34 Rev. 1, Contingency Planning Guide and VMware Site Recovery Manager documentation.

## QUESTION NO: 8 - (SIMULATION)

As part of the planning process, a cloud architect may use tools to gather data from customer systems. Match each tool with its associated description.

## ANSWER: See the explanation for the answer

**Explanation:**

**Correct matches:**

**VMware Capacity Planner** — Performs agentless discovery of servers and gathers hardware inventory, resource-utilization, and capacity data. It is used to size/right-size a virtual or cloud infrastructure.
**VMware vCenter Operations Manager (vC Ops/vROps)** — Collects and analyzes operational data from the vSphere environment, including performance, capacity, health, and workload trends.
**EMC ViPR SRM / EMC Storage Resource Management Suite** — Discovers and reports on multi-vendor storage infrastructure, including storage configuration, capacity utilization, performance, and chargeback/showback information.
**EMC Data Protection Advisor (DPA)** — Collects backup and data-protection information, such as protection policies, job success/failure, backup capacity, and recovery compliance.
**EMC CloudScape** — Provides discovery and assessment information used to understand application/infrastructure dependencies and assess workloads for cloud adoption or migration.

In short, use `Capacity Planner` for server sizing, `vC Ops` for virtual-infrastructure operations and trends, `ViPR SRM` for storage, `DPA` for data protection, and `CloudScape` for cloud-readiness/dependency assessment.

## QUESTION NO: 9

An organization is migrating workloads to a private cloud. The cloud architect discovers that an application has a dependency on a physical USB security device attached to its existing server.

Which workload characteristic will most affect the migration decision?

**A.** Dependency on physical hardware

**B.** Use of IP addressing

**C.** Use of block storage

**D.** Requirement for monitoring

## ANSWER: A

**Explanation:**

**Dependency on physical hardware** will most affect the migration decision because the application depends on a device that may not be available or supported in the target virtualized cloud environment. Before migration, the architect must determine whether the device can be securely redirected, replaced with a virtual or network-based alternative, or retained through a supported passthrough configuration. If no suitable approach exists, the workload may need to remain on dedicated physical infrastructure.

Cloud workload assessment must identify application dependencies in addition to compute, storage, and network requirements. Dependencies can include hardware devices, operating system versions, middleware, identity services, network services, and external applications. NIST discusses the importance of understanding cloud migration and interoperability considerations in SP 800-146, Cloud Computing Synopsis and Recommendations.

## QUESTION NO: 10

Case 1

Refer to the exhibit.

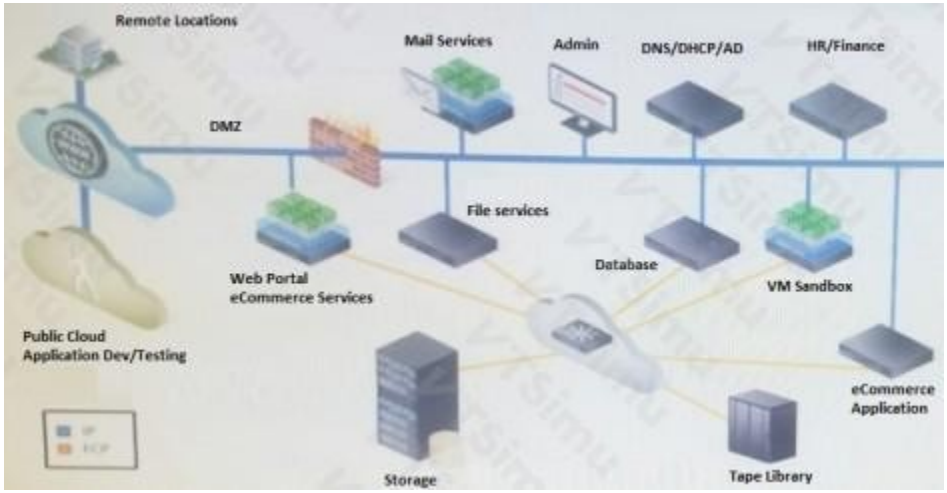


Your company recently negotiated several contracts with new customers to use its proprietary eCommerce system. As a result, the customer base has grown rapidly, which has introduced several challenges and requirements.

Business requirements:

- Applications need to scale quickly
- Metrics and reporting to anticipate business trends and efficient use of resources
- External customers are looking to license the eCommerce application

Other considerations:

- Applications need to be recorded for the public cloud
- Address security vulnerabilities for customer data in the cloud

Refer to Case 1.

You are designing the application development platform for the cloud and are considering using a PaaS Solution. You must ensure that monitoring is included in the platform. Which key metrics need to be available from the PaaS Solution that you select?

**A.** Queue length and storage utilization

**B.** Storage utilization and account login

**C.** Account login and response time

**D.** Response time and queue length

## ANSWER: D

**Explanation:**

**Response time and queue length** are the key monitoring metrics because together they show both the customer-facing performance of the eCommerce service and the demand pressure building behind it. Response time measures how long the platform takes to complete a request from a user or calling service. It is therefore a direct indicator of whether the application is meeting its expected service experience as demand increases. Queue length measures the number of requests or jobs waiting to be processed. A sustained increase indicates that incoming workload exceeds available processing capacity, making it a useful leading signal for scaling application instances or dependent services before customers experience unacceptable delays.

These metrics support the stated need for rapid scaling and for reporting that anticipates business trends. Monitoring response time identifies current service degradation, while queue depth helps forecast imminent capacity issues and supports automated or operational scale-out decisions. Major cloud platform monitoring services expose equivalent latency and queue-depth measures for this purpose; see Azure Monitor metric aggregation documentation and Amazon SQS CloudWatch metrics documentation.

## QUESTION NO: 11

SPECIAL INSTRUCTIONS REMINDER

The remaining questions are associated with two cases.

Each case has an introductory, descriptive passage (case text) that describes background information and requirements for a specific cloud based design. And each case has a set of six related, randomized questions. You will need to read the case text in order to answer the related questions correctly.

Given that there are two cases with six scored questions each, the total number of case-based scored questions is 12.

When you display each questions you can easily review the related case text in a pop-up window by clicking on the “Case” icon (shown below) Case 4

Refer to the exhibit.

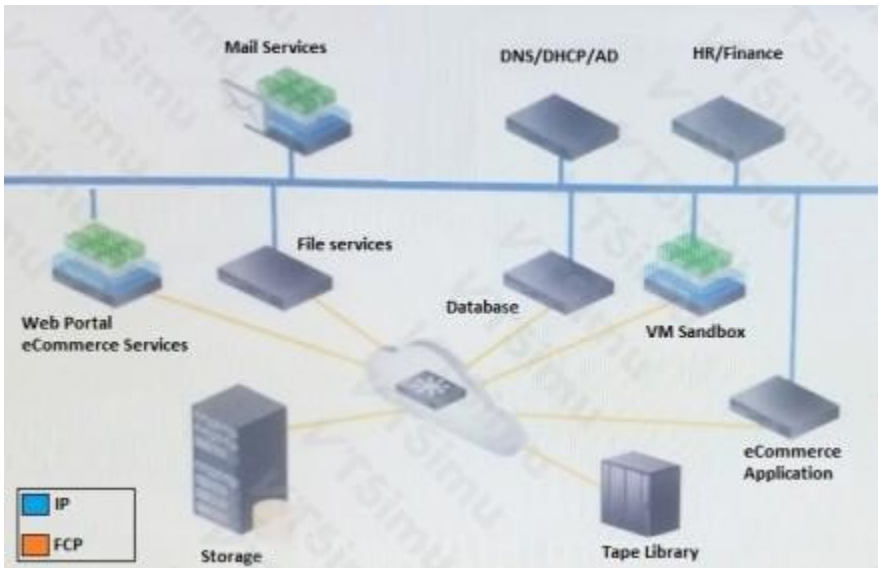


Your company is currently using a traditional SAN storage network. The server infrastructure is partially virtualized. There is a public cloud in place that is used to test development and application migration to cloud.

Business requirements:

- Consolidate the network, storage, and compute resources
- Simplify network management and storage provisioning
- Maintain application performance
- Need to share resources to reduce cost and improve efficiency

Other considerations:

- Monitoring of protection SLAs
- All transactions must be logged for auditing review
- Users are validated using internal credentials Refer to Case 4.

You are designing a brownfield cloud. You review the requirements and must determine the best, low-cost solution. Which components must you include to meet the business requirements?

**A.** AD, AppDev platform, and Service Catalog

**B.** CI infrastructure, CMP, and AppDev platform

**C.** HCI infrastructure, CMP, and AD

**D.** CMP, AppDev platform, and Service Catalog

**ANSWER: C**

**Explanation:**

**HCI infrastructure, CMP, and AD** is the appropriate low-cost brownfield-cloud combination because it directly covers the infrastructure, management, and identity capabilities stated in the case. Hyperconverged infrastructure combines compute and software-defined storage in a scale-out platform, substantially reducing the separate infrastructure domains that must be deployed and administered. It can be expanded incrementally as demand grows, supporting resource consolidation and sharing while preserving performance through clustered, policy-driven operations. Dell describes VxRail as an HCI system that integrates compute, storage, networking, and virtualization management into a single operational model: Dell VxRail.

A cloud management platform supplies the cloud-control layer needed to simplify provisioning and ongoing operations. It provides automated, policy-based delivery, metering and governance capabilities, and centralized visibility appropriate for monitoring protection service levels and retaining operational records for audit review. Finally, Active Directory provides the existing internal directory-based identity source required to validate users. Its domain-based authentication and authorization model allows the cloud-management environment to use established enterprise credentials rather than introducing a separate identity repository. Microsoft documents Active Directory Domain Services identity and access functions at Active Directory Domain Services overview.