DUMPSBOSS.

Associate - Cloud Infrastructure and Services v.3 Exam

EMC DEA-2TT3

Version Demo

Total Demo Questions: 12

Total Premium Questions: 121

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Computing	11
Topic 2, Cloud Services	18
Topic 3, Cloud Infrastructure	30
Topic 4, Cloud Security	18
Topic 5, Cloud Management	38
Topic 6, Cloud Design	6
Total	121

QUESTION NO: 1

To implement a new service desk management tool the capacity of the servers has to be extended. Who is responsible for managing the request for additional capacity?

A. Change Manager

B. Financial Manager

C. Capacity Manager

D. Service Level Manager

ANSWER: C

Explanation:

Capacity Manager is responsible for ensuring that infrastructure and services have sufficient, cost-effective capacity to meet current and anticipated business requirements. Implementing a new service desk management tool creates a changed demand profile: the application may require additional compute, memory, storage, network throughput, or supporting platform capacity. The Capacity Manager assesses that demand, evaluates the existing resource baseline and forecasts, determines the required expansion, and manages the capacity requirement through planning and coordination with the relevant technical teams. This responsibility is not limited to reacting when a resource is exhausted; it includes proactively identifying capacity needs caused by new or changed services and ensuring that the necessary resources are available at the required time and performance level. The resulting implementation work can be controlled through organizational change processes, but the ownership of analyzing and managing the additional-capacity requirement remains with Capacity Management. This aligns with IT service-management practice, in which capacity and performance management balances present and future demand against available service capacity. See the AXELOS ITIL capacity and performance management guidance and the IBM overview of capacity management.

QUESTION NO: 2 - (DRAG DROP)

Match each information security concept with its description.

Select & Place:

ANSWER:

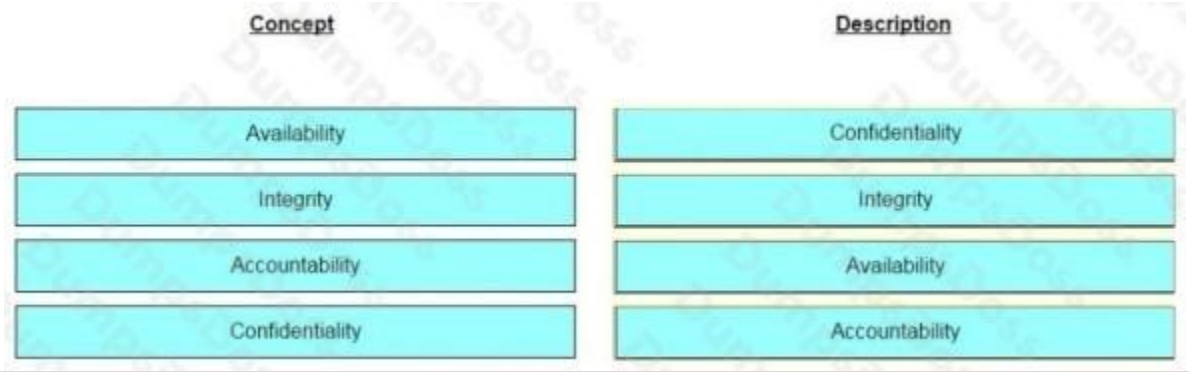

Explanation:

These mappings reflect the core information-security objectives used to protect cloud systems and their data. Confidentiality is the requirement that information is kept secret to the appropriate degree and is disclosed only to authorized entities. In the diagram, this is expressed by ensuring the required secrecy of information. Integrity is the objective of preserving the correctness and trustworthiness of information throughout its lifecycle. A central part of that protection is preventing unauthorized changes, whether those changes are accidental or malicious. Availability ensures that authorized users can access systems, services, and information in a reliable and timely manner when business operations require them. This includes the resilience and operational continuity needed to keep cloud resources accessible. Accountability establishes that user actions are attributable to a specific identity, so users can be held answerable for the actions they execute. Logging, auditing, identity management, and monitoring are common controls that support accountability. Together, confidentiality, integrity, and availability form the familiar CIA security objectives, while accountability adds traceability for activities performed within the environment. These definitions align with the security terminology in NIST's confidentiality glossary entry, NIST's integrity glossary entry, NIST's availability glossary entry, and NIST's accountability glossary entry.

QUESTION NO: 3

A specific intrusion detection and prevention system (IDPS) scans events against known attack patterns. Which technique is the IDPS using?

A. Plug-in-based detection

B. Signature-based detection

C. Device-based detection

D. Anomaly-based detection

ANSWER: B

Explanation:

Signature-based detection is correct because it compares observed network traffic, host activity, log events, or other monitored data with a database of predefined signatures that represent known attacks or malicious behaviors. A signature may identify a distinctive byte sequence, packet characteristic, protocol misuse, command pattern, exploit payload, or other recognizable indicator associated with a previously identified threat. When the IDPS finds a match, it can generate an alert and, in prevention mode, take an action such as dropping a packet, terminating a connection, or blocking the source according to its configured policy.

This approach is highly effective for detecting threats with known and well-defined patterns, provided that the signature set is current and accurately tuned for the environment. The wording "known attack patterns" directly describes the matching process used by signatures. NIST identifies signature-based IDPS detection as a method that compares monitored events with patterns of known malicious activity; see NIST SP 800-94, Guide to Intrusion Detection and Prevention Systems. For further security-operations context on detection technologies and signature updates, see CISA Cyber Threats and Advisories.

QUESTION NO: 4

Which product provides key management and a data-at-rest encryption mechanism for private, public, and hybrid clouds?

A. RSASecurID

B. Dell EMC CloudArray

C. RSA Archer eGRC

D. Dell EMC CloudLink

ANSWER: D

Explanation:

Dell EMC CloudLink is the correct product because it was designed to provide centralized encryption key management and data-at-rest encryption for cloud-based workloads. CloudLink protects data stored in virtual machines and cloud environments by encrypting data at rest while separating key management from the protected workload. This approach enables organizations to retain control of encryption keys and apply consistent security controls as workloads operate across private, public, and hybrid cloud deployments. CloudLink also supports policy-based security and access controls, helping address data-protection and compliance requirements when applications and their data move between environments. Dell EMC describes CloudLink as delivering data-at-rest encryption and centralized key management for cloud infrastructures, making it the product that directly satisfies both parts of the requirement in the question. Further product background is available in Dell Technologies' data protection resources and in the archived product overview available through the Dell Technologies Support knowledge base.

QUESTION NO: 5

Which storage-defined controller function automates the storage provisioning task and delivers virtual storage resources?

A. Storage component management

B. Service Provisioning

C. Resource abstraction and pooling

D. Storage discovery

ANSWER: B

Explanation:

Service Provisioning is the storage-defined controller function responsible for automating the delivery of storage services to consumers. It translates a requested service level—such as capacity, performance, protection, and availability requirements—into provisioned virtual storage resources. Rather than requiring administrators to manually perform each storage configuration activity, this function uses predefined policies, templates, and workflows to create and present the required storage in a consistent, repeatable manner. This is central to a software-defined storage model because consumers can request storage as a service while the controller coordinates the underlying resource allocation and configuration. Provisioning automation also supports faster delivery, reduces operational effort, and helps ensure that the delivered resource conforms to the organization's defined service requirements. Dell's software-defined storage portfolio emphasizes policy-driven automation and abstraction of underlying infrastructure to provide storage resources efficiently. See Dell Technologies' Cloud Infrastructure overview and the Dell PowerFlex overview for examples of software-defined infrastructure and automated storage-service delivery.

QUESTION NO: 6

What is the purpose of a bitmap in snapshot creation?

A. Keeps track of blocks that have changed in the filesystem

B. Indicates the exact address from which data has to be read

C. Indicates the exact address to where data has to be written

D. Keeps track of the number of blocks that are highly utilized

ANSWER: A

Explanation:

Keeps track of blocks that have changed in the filesystem is correct. A snapshot represents data at a specific point in time without necessarily making a full, independent copy of every block at the instant the snapshot is created. To preserve that point-in-time view efficiently, the storage system needs to identify which data blocks are modified after snapshot

creation. A bitmap provides this tracking mechanism: each bit corresponds to a block or region and records whether that area has changed. When a write occurs to a block associated with an active snapshot, the system can use this change information to preserve the pre-change data as required by its copy-on-write or redirect-on-write implementation. The snapshot can therefore continue presenting the earlier version of changed data, while unchanged blocks can still be referenced directly from the source volume. This block-level tracking reduces unnecessary copying, supports efficient snapshot operation, and is fundamental to maintaining a consistent point-in-time image. Dell storage platforms document snapshot behavior and management in their product documentation; see the Dell PowerStore documentation portal. For a broader description of point-in-time shadow-copy architecture and its coordinated preservation of prior data versions, see the Microsoft Volume Shadow Copy Service documentation.

QUESTION NO: 7

What is an organizational benefit of adopting modern applications?

A. Enhances competitiveness

B. Removes dependency on IT completely

C. Eliminates CAPEX

D. Removes dependency on service provider

ANSWER: A

Explanation:

Modern applications enhance competitiveness by enabling an organization to respond more rapidly to changing customer expectations, market conditions, and business priorities. Architectures such as cloud-native, microservices-based, and API-driven applications support faster development and deployment cycles, allowing teams to deliver new capabilities and improvements more frequently. They also help organizations scale services efficiently and use automation to improve reliability and operational consistency. These capabilities make it easier to innovate, differentiate digital offerings, and bring products or features to market sooner—key contributors to maintaining a competitive position. Dell Technologies describes modern applications as an approach that helps organizations improve agility, accelerate innovation, and deliver better business outcomes. Cloud-native practices similarly emphasize rapid, reliable delivery through automation and scalable application design. The organizational value is therefore not merely technical modernization; it is the ability to compete more effectively through increased speed, flexibility, and responsiveness. See Dell Technologies: Modern Applications and Cloud Native Computing Foundation.

QUESTION NO: 8

What is a benefit of the active-active cluster implementation?

A. Enhances performance and availability of a service

B. Lower total cost of ownership for providing services

C. Improves security of applications running on the cluster

D. Ability to run multiple services on one physical compute system

ANSWER: A

Explanation:

Enhances performance and availability of a service is the benefit of an active-active cluster implementation. In this design, two or more cluster members actively provide the same service at the same time. Client requests and workloads can be distributed across the available active members, which increases aggregate processing capacity and can improve response times versus directing all work to a single active node. Because service capacity is already operating on multiple nodes, the service can continue through the remaining active member or members if one member becomes unavailable. This provides high availability while also making productive use of the compute resources in normal operation, rather than reserving a

node solely for standby use. Effective active-active designs use load balancing, health monitoring, and appropriate data or session coordination so that clients are directed only to healthy service instances. Dell EMC infrastructure supports highly available clustered architectures intended to keep applications and services available during component or node failures. Microsoft similarly describes failover clustering as a high-availability technology that uses multiple servers to provide continued service when a failure occurs. See Dell Technologies Cloud Infrastructure and Microsoft Failover Clustering overview.

QUESTION NO: 9

Which management process ensures business continuity by eliminating single points of failure in a data center?

A. Availability management

B. Security management

C. Capacity management

D. Performance management

ANSWER: A

Explanation:

Availability management is correct because it focuses on ensuring that IT services and supporting infrastructure remain accessible when needed, at agreed service levels. A core availability-management practice is identifying and removing single points of failure through resilient design. In a data center, this commonly includes redundant power feeds and UPS systems, multiple network paths and switches, clustered compute hosts, multipathed storage connectivity, and replicated components or services. If one component, link, or device fails, an alternative component or path can continue providing the required service. This fault-tolerant approach reduces unplanned downtime and directly supports business continuity objectives. Availability management also evaluates service availability requirements, designs appropriate redundancy, monitors availability, and drives improvements based on failures and risks. Dell Technologies infrastructure guidance similarly emphasizes highly available architectures that eliminate single points of failure across compute, network, storage, and power layers. These controls help maintain operations during component failures and provide the foundation on which broader continuity and recovery plans operate. See the IBM overview of availability management and the Dell Technologies cloud infrastructure resources.

QUESTION NO: 10 - (SIMULATION)

Match each virtual network type with its characteristics

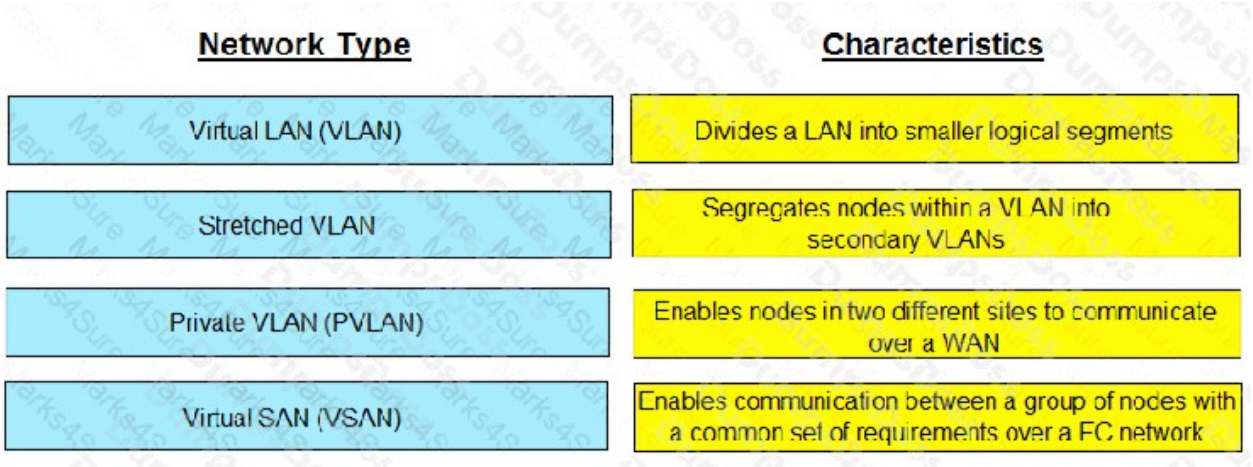

ANSWER: See the explanation for the answer

Explanation:

Correct matches:

QUESTION NO: 11

Which security control converts plaintext into unreadable ciphertext to protect confidential data?

A. Hashing

B. Encryption

C. Masking

D. Compression

ANSWER: B

Explanation:

Encryption is correct because it transforms plaintext data into ciphertext using a cryptographic algorithm and key. Authorized users or systems with the appropriate key can decrypt the ciphertext and recover the original information. Encryption can protect data at rest on storage media and data in transit across networks, helping preserve confidentiality if data is accessed or intercepted by an unauthorized party. Effective encryption also requires secure generation, storage, rotation, and retirement of cryptographic keys. See the NIST Cryptographic Standards and Guidelines.

QUESTION NO: 12

A cloud service provider is concerned about the security risks due to velocity-of-attack.

What can the cloud service provider deploy to mitigate the impact of this attack?

A. Shredding data-at-rest

B. Role-based access control

C. Multi-factor authentication

D. Containment mechanisms

ANSWER: D

Explanation:

Containment mechanisms are appropriate because velocity-of-attack describes how rapidly a threat can spread, affect multiple workloads, or exploit the scale and connectivity of a cloud environment. The key defensive objective is therefore to limit the attack's blast radius immediately after suspicious or malicious activity is detected. Containment can include isolating affected virtual machines or containers, segmenting networks, restricting east-west traffic, quarantining compromised accounts or workloads, and applying automated policy controls that prevent further propagation. These measures allow the provider to preserve service availability and protect unaffected tenants while investigation, remediation, and recovery occur. In a multitenant cloud, rapid isolation is especially important because shared infrastructure and highly automated provisioning can otherwise enable an incident to affect many resources in a short time. This approach aligns with the NIST incident-response lifecycle, which identifies containment as a core activity for limiting an incident's scope and impact. NIST's cloud-computing definition also emphasizes the pooled and rapidly provisioned nature of cloud resources, characteristics that make prompt containment operationally essential. See NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide and NIST SP 800-145, The NIST Definition of Cloud Computing.