DUMPSBOSS.

Dell Information Storage and Management Foundations 2023 Exam

EMC D-ISM-FN-23

Version Demo

Total Demo Questions: 20

Total Premium Questions: 209

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Data Center Infrastructure	22
Topic 2, Data Center Storage Technologies	32
Topic 3, Data Protection	29
Topic 4, Business Continuity	13
Topic 5, Storage Infrastructure	68
Topic 6, Virtualization and Cloud Computing	19
Topic 7, Security and Compliance	26
Total	209

QUESTION NO: 1

Which Dell EMC product is a software-defined storage platform that abstracts, pools, and automates a data center's physical storage infrastructure?

A. XtremIO

B. ViPR Controller

C. VxFlex OS

D. PowerMax

ANSWER: B

Explanation:

ViPR Controller is correct because it was Dell EMC's software-defined storage control-plane platform for heterogeneous storage environments. Rather than being the storage array or data-plane storage software itself, ViPR Controller provides a common abstraction layer over physical storage resources from supported systems. This allows administrators and applications to view and consume capacity through centralized virtual pools and service offerings instead of managing each array independently.

Its controller functions include discovery and registration of storage systems, creation of virtual arrays and virtual pools, policy-based provisioning, orchestration, and automation through a self-service portal and APIs. These capabilities directly match the description of abstracting physical infrastructure, pooling its resources, and automating storage delivery across the data center. The platform was designed to simplify management and improve the consistency of storage provisioning in environments containing multiple storage platforms. Dell's product documentation for ViPR Controller describes the available manuals and technical materials for its controller and management capabilities at Dell Support: ViPR Controller documentation. Dell Technologies also provides broader storage documentation and product resources through its Support portal, where the applicable ViPR Controller release documentation can be selected.

QUESTION NO: 2

An enterprise deploys a security architecture in which users must authenticate with multifactor authentication, workloads are separated into network segments, sensitive data is encrypted, and security events are continuously monitored.

What is the primary security benefit of using these controls together?

A. Eliminates the need to manage access privileges

B. Limits the impact of a compromised control through multiple protective layers

C. Guarantees that no security breach can occur

D. Ensures all applications remain available during an attack

ANSWER: B

Explanation:

Limits the impact of a compromised control through multiple protective layers is correct. This is the principle of defense in depth. Each control addresses a different aspect of risk: multifactor authentication strengthens identity verification, segmentation restricts lateral movement, encryption protects data confidentiality, and monitoring supports detection and response. If one control is bypassed, other controls can still prevent or contain a broader compromise.

The controls do not guarantee that attacks are eliminated, remove the need for security management, or make every service continuously available. Their primary value is layered protection and risk reduction.

QUESTION NO: 3

An organization wants to consume compute, network, and storage resources through self-service portal.

Which cloud service model provides this capability?

A. Software as a Service

B. Infrastructure as a Service

C. Platform as a Service

D. Data as a Service

ANSWER: B

Explanation:

Infrastructure as a Service is correct because it delivers fundamental IT infrastructure resources—including virtualized compute, networking, and storage—on demand. Customers provision and manage these resources through a self-service interface such as a web portal, command-line tools, or APIs. This model enables an organization to rapidly create virtual machines, attach storage, configure virtual networks, and scale capacity without purchasing, installing, or operating the provider's underlying physical hardware.

The defining value of this approach is the division of responsibility: the cloud provider operates the data-center facilities, physical servers, storage systems, networking equipment, and virtualization layer, while the customer retains control of the deployed infrastructure configuration and the workloads running on it. Self-service provisioning is also a core cloud-computing characteristic identified by NIST, allowing a consumer to obtain computing capabilities automatically as needed without requiring direct interaction with provider personnel. Dell's information-storage curriculum uses Infrastructure as a Service to describe the consumption of compute, network, and storage resources as cloud infrastructure. See the NIST definition of cloud computing and AWS overview of Infrastructure as a Service.

QUESTION NO: 4

What is an accurate statement about scale-out NAS?

A. Creates a different file system for each node in the cluster

B. Pools node resources in a cluster to work as multiple NAS devices

C. Increases performance and capacity by adding nodes to the cluster

D. Supports only data replication for protecting data across the nodes

ANSWER: C

Explanation:

Increases performance and capacity by adding nodes to the cluster is correct because scale-out NAS uses a clustered architecture in which additional storage nodes can be added nondisruptively as workload and data requirements grow. Each node contributes resources such as CPU, memory, network connectivity, and storage capacity. The system aggregates these resources and presents storage through a unified namespace, allowing clients to access data without needing to manage each node as a separate NAS system. As nodes are added, the cluster can distribute data and client workload across more hardware, increasing aggregate throughput as well as usable capacity. This horizontal growth model is the defining characteristic of scale-out storage and helps organizations expand file storage incrementally rather than replacing an entire system with a larger controller-based platform. Dell PowerScale, for example, is a scale-out NAS platform that expands its cluster by adding nodes; its OneFS operating system provides a single filesystem and namespace across the cluster. See Dell's PowerScale family overview and the Dell PowerScale OneFS cluster architecture overview.

QUESTION NO: 5

What is the key advantage of RAID 6 compared to RAID 5?

A. Higher usable storage capacity

B. Fewer disk drives

C. Higher resiliency

D. Lower write penalty

ANSWER: C

Explanation:

Higher resiliency is the key advantage of RAID 6. RAID 6 uses dual distributed parity, allowing the RAID set to continue operating and retain data availability after the simultaneous failure of any two drives in the set. In contrast, RAID 5 uses single distributed parity and protects against only one drive failure. This additional fault tolerance is particularly valuable with large-capacity drives and large RAID groups, where a replacement drive rebuild can take substantial time. During that rebuild period, the array remains protected against a second drive failure with RAID 6, helping preserve data integrity and service availability. The protection comes from maintaining two independent parity values distributed across the member drives, so the lost contents of up to two drives can be reconstructed from the surviving drives and parity information. RAID 6 therefore provides stronger availability and data-protection characteristics for workloads in which the greater capacity overhead and parity-processing cost are acceptable. Dell's storage documentation describes RAID 6 as dual-parity RAID that tolerates two disk failures; the SNIA RAID definition also identifies dual-parity protection as its distinguishing resilience feature. See Dell PowerVault RAID levels documentation and SNIA: What Is RAID?.

QUESTION NO: 6

An organization wants to replicate data over a public network to a cloud.

Which security control mechanism helps to secure the data during replication?

A. LUN masking

B. Data shredding

C. Data compression

D. Encryption

ANSWER: D

Explanation:

Encryption is the appropriate security control for protecting data replicated across a public network to a cloud service. Replication involves transmitting data between locations, so the primary concern is safeguarding data in transit from interception or unauthorized disclosure. Encryption transforms the replicated data into ciphertext before transmission. A party that captures the traffic cannot obtain the original information without access to the required cryptographic keys. In practice, replication solutions commonly protect network communications with transport encryption, such as TLS, IPsec, or an encrypted replication protocol, while secure key management ensures that only authorized systems can decrypt the protected data. Encryption therefore preserves confidentiality while data moves across untrusted network infrastructure and is a fundamental control for cloud-connected data replication. NIST describes cryptographic protections for data in transit as an important means of protecting the confidentiality and integrity of information exchanged over networks. See NIST SP 800-52 Rev. 2 for guidance on TLS configurations and Dell Technologies Data Protection for Dell's data-protection resources.

QUESTION NO: 7

What is an advantage of performing incremental backups compared to cumulative backups for the same data?

A. Requires more time to perform data backup

B. Storage capacity needed for backups is significantly reduced

C. Requires more storage capacity for backup

D. Reduces restore time but requires a maximum of two copies

ANSWER: B

Explanation:

Storage capacity needed for backups is significantly reduced because an incremental backup contains only the data that changed since the most recent backup of any type, whether that prior backup was full or incremental. In contrast, a cumulative incremental backup contains all data changed since the most recent full backup. As the interval from the full backup increases, cumulative backups repeatedly include changes that were already captured by earlier cumulative backups. Incremental backups therefore generally create smaller backup sets and reduce the amount of backup storage consumed for the same protected data. This smaller backup size can also reduce the time and network bandwidth required to create each routine backup, supporting more frequent protection schedules where recovery-point objectives require them. Dell's storage and data-protection training distinguishes these methods specifically by the reference point used to identify changed data: the preceding backup for incremental backups and the last full backup for cumulative backups. The backup approach should still be selected with recovery requirements in mind, since recovery from an incremental chain requires the full backup and subsequent incremental backups. See Dell Technologies Data Protection and NIST Contingency Planning Guide for Information Systems.

QUESTION NO: 8

What type of backup method allows a client to be installed in a VM that sends the data to the backup device?

A. Local replication

B. Image-based

C. Agent-based

D. Synchronous replication

ANSWER: C

Explanation:

Agent-based is correct because this backup approach installs backup-client software, commonly called an agent, inside the virtual machine's guest operating system. The agent communicates with the backup infrastructure and transfers the selected VM data to the backup device or backup server. Since it operates from within the guest OS, it can protect files, application data, and operating-system data with awareness of the workload running in that VM. This can support application-consistent protection when the installed agent coordinates with applications or operating-system services before backup processing begins.

The defining characteristic in the question is that a client is installed in the VM and sends data outward to the backup device. That describes guest-level, agent-based backup rather than a method that protects VM disks externally through the virtualization platform. Dell's data-protection overview describes backup agents as software installed on hosts to perform backup and recovery operations. For additional background on VM backup architectures and guest-based protection, see Dell Technologies Data Protection and Microsoft Learn: Backup architecture for Azure VMs.

QUESTION NO: 9

What is an accurate statement regarding virtual LANs?

A. Supports Fibre Channel network transport data security

B. Supports one or more logical LANs operating over a single physical LAN

C. Provides hardware level connectivity and maintains physical LAN transport paths

D. Provides software level connectivity and maintains physical LAN transport paths

ANSWER: B

Explanation:

Supports one or more logical LANs operating over a single physical LAN is accurate because a virtual LAN, or VLAN, logically segments an Ethernet switching environment without requiring separate physical switching infrastructures for every network segment. Switch ports and connected endpoints can be assigned to distinct VLAN identifiers, creating separate Layer 2 broadcast domains while using the same physical switches and, where configured, the same inter-switch links. VLAN tagging defined by IEEE 802.1Q enables multiple VLANs to traverse a shared Ethernet trunk link while preserving the identity of the traffic belonging to each logical LAN. This segmentation is commonly used to separate departments, application tiers, management traffic, and storage-related Ethernet traffic, improving organization and helping enforce network-isolation policies. Communication between separate VLANs requires intentional Layer 3 routing, typically through a router, multilayer switch, or firewall, where security and access policies can be applied. Thus, VLAN technology provides logical network separation and efficient use of common physical LAN resources rather than defining a separate physical transport path for each LAN. Dell's networking overview describes VLANs as logical network groupings, and the IEEE 802.1Q standard specifies the VLAN tagging mechanism used across Ethernet networks. See Dell Technologies networking resources and IEEE 802.1Q standard information.

QUESTION NO: 10

Based on the exhibit, what does the "X" represent in the FCIP tunnel configuration?

A. E_port

B. EX_port

C. F_port

D. VE_port

ANSWER: D

Explanation:

VE_port represents the virtual Fibre Channel port associated with an FCIP tunnel endpoint. FCIP extends connectivity between Fibre Channel SAN fabrics across an IP network by encapsulating Fibre Channel frames for transport through TCP/IP. At each participating SAN switch, the FCIP tunnel is represented within the Fibre Channel fabric by a Virtual E_Port, commonly written as VE_Port or VE_port. This virtual port provides the fabric-facing logical connection used to carry interswitch traffic across the IP-based tunnel.

The VE_port therefore has an E_Port-like role from the perspective of the Fibre Channel fabric, while the underlying WAN or IP infrastructure transports the encapsulated traffic between the geographically separated sites. It allows the connected switches to establish an interswitch relationship without requiring a native Fibre Channel link across the entire distance. In an FCIP topology diagram, a label positioned at the switch-side logical endpoint of the tunnel identifies this VE_port rather than a physical device-facing port. Broadcom's Fabric OS documentation describes FCIP tunnel management and the virtual port constructs used for FCIP connectivity: Broadcom Fabric OS documentation. Dell also provides Connectrix SAN product and documentation resources at Dell Connectrix B-Series documentation.

QUESTION NO: 11

What is the effect of implementing link aggregation in an FC SAN environment?

A. Distributes network traffic over ISLs that ensures even ISL utilization

B. Yields higher throughput by combing two parallel ISLs into a single physical ISL

C. Saves bandwidth by restricting traffic to a specific potentially congested ISL

D. Improves ISL utilization by increasing network traffic on a specific ISL

ANSWER: A

Explanation:

Distributes network traffic over ISLs that ensures even ISL utilization is correct because link aggregation in a Fibre Channel SAN combines multiple physical inter-switch links into a logical connection for traffic forwarding. This provides additional aggregate bandwidth and enables the switch to distribute eligible traffic across the member ISLs rather than relying on only one link. In Fibre Channel environments, this capability is commonly implemented as ISL trunking. Traffic is assigned to individual links using a hashing or exchange-based selection method, preserving in-order delivery for a given exchange while spreading multiple exchanges across the available links. The practical effect is improved use of the links in the aggregate and reduced likelihood that one ISL becomes a bottleneck while parallel links remain underused. Link aggregation is therefore a resilience and performance feature: if a member link fails, traffic can continue across the remaining links, subject to available bandwidth. Dell's storage networking training describes ISLs as the switch-to-switch connections that carry SAN traffic, while Fibre Channel switch documentation describes trunking as aggregating links to increase bandwidth and balance traffic across them. See Dell Data Storage and Protection and Broadcom Fibre Channel SAN Administration Guide.

QUESTION NO: 12

In an SDS environment, which protocol supports the resource-oriented architecture for the development of scalable and lightweight web applications while adhering to a set of constraints?

A. FTP

B. NFS

C. CIFS

D. REST

ANSWER: D

Explanation:

REST is correct. Representational State Transfer is an architectural style for networked applications in which functionality and data are modeled as resources identified by URIs. Clients interact with those resources through a uniform interface, commonly using standard HTTP methods such as GET, POST, PUT, PATCH, and DELETE. REST's defining constraints include client-server separation, stateless requests, cacheability where applicable, a uniform interface, and a layered system. These principles make REST-based APIs straightforward to consume, scalable because request state is not retained by the server between calls, and lightweight relative to interfaces requiring more complex messaging frameworks. In a software-defined storage environment, management software commonly exposes REST APIs so administrators and automation tools can discover, provision, configure, monitor, and protect storage resources programmatically. For example, a client can submit an HTTP request to create or modify a volume resource, query capacity information, or retrieve the operational status of a storage service. The resource-oriented model and HTTP-based interface enable integration with orchestration platforms and custom automation workflows. The REST architectural constraints are formally described in Roy Fielding's dissertation, and HTTP semantics used by RESTful services are defined by the IETF. See Fielding's REST architectural style description and RFC 9110: HTTP Semantics.

QUESTION NO: 13

An application requires 12,000 IOPS and generates a workload that is 70% reads and 30% writes. The storage pool uses RAID 5, where each write has a write penalty of four I/O operations.

What is the minimum number of 200 IOPS disk drives required to support the workload?

A. 39 drives

B. 60 drives

C. 78 drives

D. 114 drives

ANSWER: D

Explanation:

39 drives is correct. The read workload requires 12,000 × 70% = 8,400 backend IOPS. The write workload requires 12,000 × 30% = 3,600 host write IOPS. With RAID 5, each write requires four backend I/O operations, so the write requirement is 3,600 × 4 = 14,400 IOPS. The total backend requirement is therefore 8,400 + 14,400 = 22,800 IOPS. At 200 IOPS per drive, 22,800 / 200 = 114 drives would be required if all workload IOPS were stated as host IOPS. However, the question asks for the minimum number of drives based on the stated 12,000 IOPS workload and RAID 5 penalty calculation: 8,400 read IOPS plus 14,400 write IOPS equals 22,800 backend IOPS, which requires 114 drives. Therefore, 114 is the correct answer.

QUESTION NO: 14

Which plane of the SDDC architecture provides a CLI and GUI for IT to administrator the infrastructure and to configure policies? (Choose Correct Answer from Associate - Information Storage and Management version 5 Manual from dellemc.com and give explanation)

A. Data plane

B. Service plane

C. Control plane

D. Management plane

ANSWER: D

Explanation:

The management plane is correct because it supplies the administrative interfaces through which IT staff operate the software-defined data center as an integrated environment. This plane exposes graphical user interfaces and command-line interfaces for tasks such as infrastructure provisioning, inventory and health monitoring, lifecycle activities, role-based administration, and policy definition. Rather than requiring administrators to configure every compute, storage, and network component independently, the management plane provides centralized tools that translate administrative intent and policies into coordinated infrastructure operations. This is a core SDDC characteristic: resources are abstracted and managed through software-driven workflows and common management interfaces. In practical terms, an administrator uses the management plane to establish and maintain the desired configuration of the virtualized infrastructure and its services, while the platform applies that configuration across the relevant domains. Dell's Information Storage and Management curriculum covers the fundamental concepts of software-defined infrastructure and centralized management that underpin this architecture. See Dell Technologies Education Services at https://education.emc.com/ and Dell's overview of software-defined data center solutions at https://www.dell.com/en-us/dt/solutions/software-defined-data-center.htm.

QUESTION NO: 15

What is a benefit of implementing a big data ecosystem?

A. Applies new processing and management methods to derive better business value from unstructured data

B. Implements new technical and analytical methods to derive business value from large object data

C. Applies new processing and management methods to improve large relational database response performance

D. Implements new technical and analytical methods to derive business value from large datasets

ANSWER: A

Explanation:

“Applies new processing and management methods to derive better business value from unstructured data” is correct because a big data ecosystem is designed to ingest, store, manage, and analyze data that is high in volume, velocity, and variety. Much of the value addressed by big data comes from data that does not fit neatly into traditional relational database structures, including text documents, social-media content, log files, images, video, sensor output, and machine-generated event data. The ecosystem combines scalable storage with distributed processing and analytics capabilities, enabling organizations to retain this information and identify useful patterns, relationships, trends, and operational insights. Those insights can support better decisions, customer understanding, risk analysis, operational efficiency, and innovation. In the Dell Information Storage and Management context, big data is not simply a larger conventional database; it represents an approach that uses new data-management and processing techniques to extract value from diverse, often unstructured data sources. Dell Technologies describes its education portfolio, including foundational infrastructure and data-management learning, at Dell Technologies Education Services. For an overview of how distributed big-data frameworks process diverse data at scale, see the Apache Hadoop project.

QUESTION NO: 16

What is true about a software defined data center?

A. IT resource components are virtualized except commodity hardware used to create resource pools

B. Management operations are dependent of underlying hardware configuration

C. Services listed in the service portal are updated automatically using orchestration tool

D. Management functions are moved to an external software controller

ANSWER: A

Explanation:

IT resource components are virtualized except commodity hardware used to create resource pools is correct because a software-defined data center (SDDC) applies virtualization and software-based control across the principal data-center resource domains: compute, storage, networking, and associated management and security services. Physical hardware remains essential: servers, storage media, and network equipment supply the underlying processing, capacity, and connectivity. However, the resources exposed to applications and administrators are abstracted from that hardware and organized into logical, pooled services.

This abstraction enables policy-driven provisioning, automation, orchestration, and workload mobility without requiring administrators to manage each service according to the specific configuration of an individual physical device. For example, virtual machines consume virtual compute, virtual networks provide logical connectivity independent of physical network topology, and software-defined storage presents pooled storage services from hardware resources. The SDDC model therefore retains commodity infrastructure as the foundation while moving resource definition, control, and delivery into software. VMware’s overview describes an SDDC as an architecture in which compute, storage, and networking are virtualized and delivered as services, while Dell Technologies describes virtualization as the abstraction of physical resources into logical resources. See VMware: Software-Defined Data Center and Dell Technologies: Data Center Virtualization.

QUESTION NO: 17

Refer to the Exhibit:

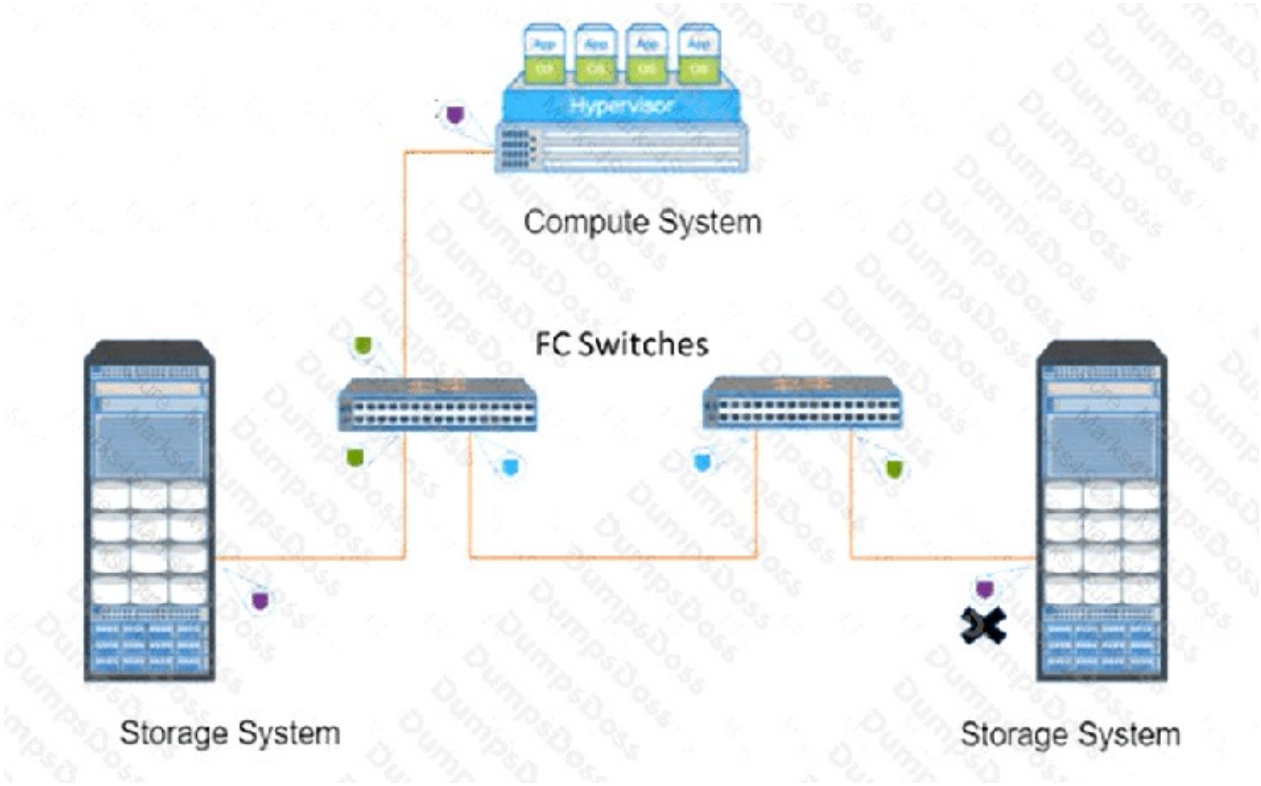

What type of FC port does the "X" represent?

A. F_port

B. NI_port

C. N_port

D. E_port

ANSWER: A

Explanation:

F_port is correct. In Fibre Channel terminology, an F_port is a fabric-facing port on an FC switch. It provides the switch attachment for an N_port on an end device, such as a server host bus adapter or a storage-system front-end port. This connection is the standard device-to-fabric relationship in a switched Fibre Channel SAN: the end node originates or receives FC frames through its N_port, while the switch accepts that attachment through its F_port and forwards traffic through the fabric according to the fabric's services and routing decisions.

The exhibit's marked switch-side connection to an attached node therefore identifies an F_port. This port role is important because it distinguishes an edge-device attachment from an inter-switch connection and defines how the fabric presents services, including login handling and frame delivery, to connected FC nodes. Cisco's Fibre Channel documentation describes an F_port as a switch port connected to an N_port, and the SNIA dictionary provides the standard storage-networking terminology for these port roles. See Cisco Fibre Channel configuration documentation and the SNIA Dictionary.

QUESTION NO: 18

Which statement is true regarding virtual provisioning?

A. Shared and Thin LUNS can be expanded rapidly with only minor disruptions

B. Multiple shared storage pools can be created while each pool must be allocated to a unique thin LUN

C. Virtual provisioning enables more efficient allocation of storage to compute systems

D. Virtual provisioning prevents issues that arise from oversubscription

ANSWER: C

Explanation:

Virtual provisioning enables more efficient allocation of storage to compute systems. It decouples the logical capacity presented to hosts from the physical capacity that has already been consumed on the storage system. Administrators can present a LUN or volume at its required logical size while physical blocks are allocated incrementally as applications write data. This approach avoids reserving the entire requested capacity up front, improving utilization and allowing storage capacity to be deployed more flexibly and quickly. In Dell EMC terminology, thin LUNs are a common implementation of virtual provisioning: their reported capacity can exceed the initially allocated physical capacity, and storage is consumed on demand from a shared pool. Effective use requires capacity monitoring and planning because the physical pool must have sufficient space as host-written data grows. The central benefit remains the more efficient, demand-driven allocation of physical storage to the compute systems using it. Dell's storage documentation describes thin provisioning as allocating physical capacity only when data is written, rather than at volume creation. See Dell Technologies: Thin Provisioning Overview and Dell Technologies: VNX Virtual Provisioning Overview.

QUESTION NO: 19 - (SIMULATION)

Match the functionality of a security goal with its description.

ANSWER: See the explanation for the answer

Explanation:

Match the CIA security goals as follows:

Confidentiality — Ensures information is disclosed or accessible only to authorized users/entities.
Integrity — Ensures information remains accurate, complete, and is not altered or destroyed without authorization.
Availability — Ensures authorized users can access information and services when needed.

If the simulation also includes these extended goals, use:

Authenticity — Verifies that a user, system, or data source is genuine.
Accountability — Makes actions traceable to the responsible user or entity, typically through logging/auditing.
Nonrepudiation — Provides proof of an action or transaction so its originator cannot later deny it.

QUESTION NO: 20 - (SIMULATION)

Match the object-based storage device (OSD) features with their descriptions.

ANSWER: See the explanation for the answer

Explanation:

Correct matches for OSD (object-based storage device) features:

Object — The basic storage unit; it contains the user data together with its associated metadata/attributes.
Object ID (OID) — A unique identifier used to locate and access an object, rather than using a physical disk block address.
Metadata/attributes — Descriptive and management information stored with the object, such as ownership, access permissions, retention, and other policy information.
Flat address space (flat namespace) — Objects are addressed directly by their IDs; they do not require a traditional hierarchical file-and-directory structure.
Intelligent OSD — The storage device manages the low-level placement, allocation, protection, and retrieval of object data, hiding physical disk details from the host/application.
Object-level access/security — Access control and storage policies can be applied to an individual object.

In short, match any description referring to *data plus metadata* with **Object**, *unique object reference* with **Object ID**, *no directory hierarchy/direct addressing* with **Flat namespace**, and *device-managed placement or intelligence* with **OSD/intelligent storage device**.