PassQueen
a.k.a dumpsqueen.com

IBM Spectrum Storage Solution Architect V2

IBM C1000-088

Version Demo

Total Demo Questions: 8

Total Premium Questions: 87

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Storage Requirements	15
Topic 2, Solution Planning	53
Topic 3, Solution Design	14
Topic 4, Solution Implementation	5
Total	87

QUESTION NO: 1

Which IBM storage software product can be used to migrate data non-disruptively from one storage array to another storage array?

A. IBM Spectrum Protect

B. IBM Spectrum Virtualize

C. IBM Spectrum Discover

D. IBM Spectrum Control

ANSWER: B

Explanation:

IBM Spectrum Virtualize is correct because it provides storage virtualization and data-migration capabilities that allow host-accessible volumes to be moved between managed storage systems without interrupting application access. It virtualizes internal and externally attached storage behind a common storage pool, so a volume can initially remain presented to hosts while its extents are copied from the source array to a destination managed disk or storage pool. After synchronization, the system transparently switches the volume's underlying storage location while preserving the volume's identity and host mappings. This approach supports technology refreshes, storage consolidation, and replacement of an existing array with minimal operational disruption. IBM documents this capability as volume migration, including migration from external storage systems managed by Spectrum Virtualize. The feature is central to the platform's value as a storage-virtualization layer: applications continue using their existing logical volumes while the administrator changes the physical storage supplying capacity. See IBM's Volume migration documentation and the IBM Storage Virtualize product page.

QUESTION NO: 2

What are the two licensing options offered through IBM's Storage Suite for IBM Cloud

Paks?

A. Standard Edition

B. Capacity Edition

C. Data Management Edition

D. Advanced Edition

E. Entry Edition

ANSWER: A B

Explanation:

IBM Storage Suite for IBM Cloud Paks is offered in the two licensing editions named **Standard Edition** and **Capacity Edition**. These editions provide customers with alternative consumption and entitlement approaches for deploying IBM storage software services in Red Hat OpenShift environments. Standard Edition is the suite's standard software licensing offering, while Capacity Edition is designed for licensing based on managed storage capacity. Providing both editions lets an organization select the model that aligns with its expected storage footprint, operational requirements, and procurement approach while using the Storage Suite software portfolio.

IBM identifies these offerings by their full product names, IBM Storage Suite for IBM Cloud Paks Standard Edition and IBM Storage Suite for IBM Cloud Paks Capacity Edition, in its product and licensing materials. The editions are therefore the applicable choices when the question asks specifically for the two licensing options offered through the suite, rather than for feature tiers or data-management functions. See IBM's Storage Suite for IBM Cloud Paks product page and the IBM Storage Suite for IBM Cloud Paks documentation for product information and planning guidance.

QUESTION NO: 3

Which two benefits can IBM Storage Scale Active File Management provide for a globally distributed file environment?

A. Local caching of data from a remote fileset

B. Asynchronous replication for disaster recovery

C. Synchronous block replication between volumes

D. Object storage authentication for S3 clients

E. Deduplication of backup client data

ANSWER: A B

Explanation:

Local caching of data from a remote fileset and **asynchronous replication for disaster recovery** are correct. AFM can cache files from a home fileset at a remote site, helping users access data with lower latency. AFM disaster recovery can asynchronously replicate a fileset to an independent recovery fileset at another site.

AFM is not a synchronous block replication technology and does not replace object storage protocols. IBM describes AFM caching and disaster recovery capabilities in the IBM Storage Scale Active File Management documentation.

QUESTION NO: 4

Which IBM Storage Virtualize feature reduces the amount of physical capacity allocated to a volume by allocating capacity only as host data is written?

A. Thin provisioning

B. Volume mirroring

C. Data compression

D. Safeguarded Copy

ANSWER: A

Explanation:

Thin provisioning is correct because it allows a volume to have a logical size larger than the physical capacity currently allocated from its storage pool. Capacity is allocated as hosts write data, rather than reserving all capacity when the volume is created. This improves utilization but requires administrators to monitor pool capacity carefully.

IBM describes thin-provisioned volumes and their capacity-allocation behavior in the IBM Storage Virtualize documentation.

QUESTION NO: 5

Which two items should be collected when sizing an IBM Storage Protect environment for a new backup workload?

A. Daily data change rate

B. Required backup window

C. Number of storage administrators

D. Number of tape-library doors

E. Desktop screen resolution

ANSWER: A B

Explanation:

Daily data change rate and **required backup window** are correct. The daily change rate helps determine the expected amount of data that must be processed during incremental backup operations. The available backup window establishes the throughput needed for clients, network connectivity, servers, and storage pools to complete protection on time.

Total protected capacity is also important, but storage design requires workload behavior and service-level requirements rather than capacity alone. IBM Storage Protect planning information is available in the IBM Storage Protect documentation.

QUESTION NO: 6

Which two Information Dispersal Algorithm (IDA) thresholds are used by IBM Cloud Object Storage to determine object availability?

A. Read threshold

B. Write threshold

C. Compression threshold

D. Snapshot threshold

E. Deduplication threshold

ANSWER: A B

Explanation:

Read threshold is correct because it specifies the minimum number of object slices that must be available to reconstruct and read an object. This enables IBM Cloud Object Storage to continue serving data when some slices are unavailable.

Write threshold is also correct because it specifies the number of successful slice writes required before the system acknowledges a write operation. These thresholds, together with the IDA width, define the resilience and operational behavior of an IBM Cloud Object Storage deployment. See the IBM Cloud Object Storage documentation on the Information Dispersal Algorithm.

QUESTION NO: 7

Which IBM Storage Scale capability allows applications to access a file system through NFS, SMB, and Object protocols?

A. Cluster Export Services

B. Transparent Cloud Tiering

C. Dispersed Data Protection

D. Storage Scale RAID

ANSWER: A

Explanation:

Cluster Export Services is correct because CES provides the protocol-serving layer for IBM Storage Scale. Protocol nodes configured for CES can export file systems through NFS and SMB and can provide object access when the applicable protocol services are configured. CES uses highly available service IP addresses that can move between eligible protocol nodes.

IBM documents CES and protocol-node functions in the IBM Storage Scale Cluster Export Services documentation.

QUESTION NO: 8

How can information on storage workloads to obtained?

A. IBM Cloud Repository Assessment Portal

B. IBM Client Storage Assessment

C. IBM Data Estimator

D. IBM Workload Scheduler Assessment

ANSWER: B

Explanation:

IBM Client Storage Assessment is the appropriate mechanism for obtaining information about storage workloads. A client storage assessment is used to gather and document the characteristics that drive a storage-protection design, such as the client systems involved, protected data volumes, data types, change rates, backup and recovery requirements, retention expectations, and expected growth. This workload information gives the solution architect a factual basis for sizing and selecting IBM Storage Protect and associated storage resources, rather than relying solely on high-level capacity figures.

Workload collection is especially important because the required architecture is influenced by more than the amount of protected data. Ingestion windows, concurrency, performance objectives, deduplication behavior, retention policies, and recovery expectations can all affect component sizing and configuration. IBM's Storage Protect documentation describes the product's role in protecting client data and managing policy-based backup and recovery, while IBM's implementation guidance emphasizes collecting environment and workload requirements before designing the solution. These sources provide the technical context for using a client-focused assessment to capture the required workload data: IBM Storage Protect documentation and IBM Spectrum Protect implementation guidance.