DUMPSBOSS.

Veeam Certified Architect 2022

Veeam VMCA2022

Version Demo

Total Demo Questions: 12

Total Premium Questions: 126

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

While looking through the requirements regarding database servers, you see that Veeam should manage the database backups and logs. Which configurations should you check for that are unsupported base on the technical requirements?

A. Physical windows server with Oracle.

B. Oracle cluster configurations.

C. Virtual machines with physical raw device mapping.

D. Virtual machines with Oracle Automatic Storage Management storage.

ANSWER: A

Explanation:

Physical windows server with Oracle is the unsupported configuration when the requirement is for Veeam to manage Oracle database backups and archive logs through the Veeam Plug-in for Oracle RMAN. The plug-in integrates with Oracle Recovery Manager to create application-consistent database backups, process archived redo logs, and coordinate backup metadata with Veeam Backup & Replication. Its supported operating-system platforms for Oracle database hosts are Linux distributions, Solaris, and IBM AIX; Microsoft Windows is not a supported platform for the Oracle RMAN plug-in. Therefore, an Oracle database installed on a physical Windows server cannot meet this particular technical requirement through that Veeam-managed RMAN integration. This assessment concerns the database-host operating system and the supported platform matrix, rather than whether the host is physical or virtual. During architecture design, confirm the Oracle host OS early, because it determines whether the Veeam Plug-in for Oracle RMAN can be deployed and used to protect database data and archived logs. Veeam documents the plug-in's requirements and supported platforms in its Veeam Plug-in for Oracle RMAN system requirements and describes the integration workflow in the Veeam Plug-in for Oracle RMAN documentation.

QUESTION NO: 2

While going through the discovery data for the NAS environment, you determine several key metrics are missing for later deign and sizing. Which of the following should you collect from the customer about the data stored on the on the NAS per site? (Choose 3)

A. Retention requirements

B. Total number of files (in millions) to be backed up

C. Amount of source data before dedupe and compression

D. Number of shares and compressed source data

E. Large file size

ANSWER: A B C

Explanation:

Retention requirements, **Total number of files (in millions) to be backed up**, and **Amount of source data before dedupe and compression** are essential discovery inputs for NAS backup design at each site. Retention requirements establish how long restore points must be retained and therefore drive the capacity required in the backup repository, including any immutable or archival storage tier. The total number of files is equally important because NAS protection processes file-system objects and their metadata; a large file count can materially affect scanning, indexing, change detection, processing time, and the infrastructure needed to meet the backup window. Source capacity before deduplication and compression provides the baseline data-volume measurement needed for forecasting backup traffic and repository consumption. Although storage-reduction ratios can be estimated later, they vary based on file types, data change rate, and repository configuration, so uncompressed logical source size is the appropriate discovery starting point. Together, these metrics let an architect estimate capacity, performance requirements, and operational feasibility per location. Veeam's NAS backup documentation describes file-share protection and its processing behavior, while its sizing guidance emphasizes

collecting workload and retention characteristics before infrastructure design. See Veeam NAS Backup documentation and Veeam sizing resources.

QUESTION NO: 3

Veeam Financial Services has expressed concern that the backups for their virtualized stock trading application meet all stated technical and business requirements. Which Veeam backup method should be

used?

A. Backup the virtual machines using a Veeam agent.

B. Separate the virtual machine into its own job and enable application aware processing.

C. Run pre-freeze and post-thaw scripts on the virtual machines running the trading application.

D. Create a replication job at each site with the corresponding remote site as the target.

ANSWER: B

Explanation:

Separate the virtual machine into its own job and enable application aware processing is the appropriate method because a stock trading application is likely transaction-sensitive and requires application-consistent recovery, not merely a crash-consistent virtual-machine image. Application-aware processing uses Microsoft Volume Shadow Copy Service (VSS) for supported Windows applications to quiesce application activity and create a consistent backup state. Veeam can also perform application-specific processing, including transaction-log handling where applicable, so that recovery objectives, retention requirements, and restore operations can be managed appropriately.

Placing the virtual machine in its own backup job provides the required control and isolation. The architect can apply settings specifically suited to this business-critical workload, such as the application-aware processing configuration, guest credentials, backup schedule, retention policy, repository selection, encryption, and validation or testing requirements. This avoids inheriting settings intended for less critical workloads and makes it easier to demonstrate that the application's technical and business requirements are addressed. Veeam documents application-aware processing and its use of guest processing to create recoverable application-consistent backups in the Veeam Backup & Replication User Guide. Veeam's guidance for protecting transaction-sensitive workloads is also covered in its application-aware processing overview.

QUESTION NO: 4

During Veeam deployment, Veeam Financial Services has determined that the latency encountered during backup of the virtualized stock trading application servers running on VSAN is too high. What is the

recommended step Veeam Financial Services can take to improve this?

A. Use backup from storage snapshot for the virtualized stock trading application servers.

B. Implement Veeam Agent on the virtualized stock trading application servers.

C. Use Veeam Replication to create replicas of the virtualized stock trading application servers at the remote site, then back up the replicas.

D. Turn off application-aware processing for the virtualized stock trading application servers.

E. Use a virtual backup proxy configured for the HotAdd transport mode for the virtualized stock trading application servers.

ANSWER: E

Explanation:

Use a virtual backup proxy configured for the HotAdd transport mode for the virtualized stock trading application servers. This is the recommended approach for VMware vSAN workloads because the proxy can access the protected virtual disks

through the VMware virtual infrastructure rather than requiring backup traffic to traverse the production network between a separate proxy and the ESXi hosts. Keeping the proxy virtual machine on the vSAN datastore also enables efficient access to vSAN-resident workload data and helps reduce the backup-related network load that can contribute to latency on performance-sensitive applications. Veeam can automatically select an available transport mode, but defining HotAdd as the preferred mode for suitably sized virtual proxies provides a predictable architecture for vSAN protection. The backup proxy must have appropriate CPU, memory, and disk-controller capacity, and proxy placement should be planned so that it does not compete excessively with the trading workloads for host resources. Veeam's VMware proxy guidance documents HotAdd as a transport mode for virtual backup proxies, while VMware's vSAN guidance describes the importance of designing workload and infrastructure placement to control performance impact. See Veeam Backup Proxy documentation and VMware vSAN Design Guide.

QUESTION NO: 5

During further discussions, the 3-2-1 rule was recommended. What is the most efficient way to protect primary and secondary sites following the 3-2-1 rule?

A. From existing backups at the primary site, create a backup copy to the secondary site. Offload backup data to object storage using Scale-out Backup Repository capacity tier.

B. Send replicas directly from the primary to the secondary site. Backup replicas on the secondary site and tier this to object storage using Scale-out Backup Repository capacity tier.

C. Send backups directly from the primary to the secondary site, avoiding duplicate data storage. Tier backup data to object storage using Scale-out Backup Repository capacity tier.

D. Use storage snapshot support to backup a virtual machine using the secondary array. Backup snapshots on the secondary site and tier this to object storage using Scale-out Backup Repository capacity tier.

ANSWER: A

Explanation:

From existing backups at the primary site, create a backup copy to the secondary site. Offload backup data to object storage using Scale-out Backup Repository capacity tier. This design efficiently implements the core 3-2-1 principles by retaining the production data as one copy, maintaining a backup copy locally, and creating a separate backup copy at the secondary site. The secondary site provides geographic separation and therefore protects recoverability if the primary location becomes unavailable. Using a Veeam backup copy job is efficient because it uses the existing backup chain as its source rather than requiring another full production backup workflow. It can also be configured to transfer only the required restore-point data across the WAN and to maintain a defined retention policy at the remote location. Offloading data through the Scale-out Backup Repository capacity tier adds a distinct object-storage layer for longer-term retention and scalable capacity. Veeam supports moving or copying backup data from performance-tier extents to object storage according to capacity-tier policy, helping balance operational recovery performance with lower-cost, durable storage. This layered approach gives both site-level resilience and an additional storage medium while using Veeam-native copy and tiering mechanisms. See the Veeam Backup Copy documentation and Capacity Tier documentation.

QUESTION NO: 6

A customer requires proof that ransomware recovery procedures remain effective over time. Backup files are retained immutably, but the customer also needs operational evidence that critical workloads can be recovered and that the recovery environment is not exposed to production services. Which design measures should be included? (Choose 3)

A. Create SureBackup jobs that use an isolated virtual lab.

B. Configure application groups to start dependent virtual machines in the required order.

C. Use heartbeat, ping, application tests or custom scripts to validate restored services.

D. Perform every verification by failing over the production virtual machines.

E. Rely on immutable retention as evidence that applications can be restored.

F. Allow virtual-lab virtual machines to use the production network without isolation.

ANSWER: A B C

Explanation:

SureBackup jobs running in an isolated virtual lab provide automated verification that workloads can start from backup data without connecting the recovered machines to production. Application groups provide controlled boot order for dependent services, and application tests or scripts can validate that the restored services are usable. Immutability protects backup data but does not itself prove application recovery. Production failover and manual console review do not provide the same isolated and repeatable verification process.

QUESTION NO: 7

What information related to the virtual infrastructure is missing and must be collected during the discovery phase) (Choose 2)

A. Number of vSphere clusters

B. Backup window

C. Recovery time objective

D. Total of virtual machines

E. Number of scale-out Backup Repository extents currently used

ANSWER: A D

Explanation:

Number of vSphere clusters and **Total of virtual machines** are essential virtual-infrastructure inventory details to collect during discovery. The cluster count identifies the structure and distribution of the VMware estate: it helps the architect understand where workloads reside, plan connectivity to vCenter Server, assess likely proxy placement, and distribute data-mover load appropriately across sites or compute domains. Cluster topology also affects resilience and traffic-path decisions when designing the Veeam Backup & Replication architecture.

The total virtual-machine count establishes the scale of the protected environment. It is a fundamental input for estimating concurrent tasks, proxy and repository throughput requirements, configuration-database growth, operational management overhead, and the capacity needed as the environment grows. Combined with workload sizing and change-rate information gathered elsewhere in discovery, the VM count supports realistic estimates of backup duration, repository consumption, and infrastructure scalability. Veeam's sizing guidance uses environment-scale and workload characteristics to select suitable backup-server, proxy, and repository resources, while the Veeam calculator provides a structured way to model those design inputs. See the Veeam Backup & Replication Sizing Estimator and Veeam Backup & Replication system requirements.

QUESTION NO: 8

What additional information is necessary for sizing the backup solution and must be gathered

during further discovery? (Choose 2)

A. The available bandwidth per site

B. The amount of virtual machines to be protected

C. The total amount of production storage provisioned space

D. Repository storage file system

E. The replication schedule

ANSWER: B C

Explanation:

The amount of virtual machines to be protected and **the total amount of production storage provisioned space** are essential discovery inputs for sizing a Veeam backup solution. The protected virtual-machine count establishes the expected scale of the environment and helps the architect plan job organization, concurrent processing, and the compute resources required for backup proxies and related data-movement components. It also provides necessary context for operational design, such as how many workloads must be processed within available backup windows.

Provisioned production storage is a primary capacity input because it establishes the initial protected-data footprint from which backup repository capacity can be estimated. During detailed design, this value is evaluated together with expected data-change rates, retention requirements, backup method, compression and deduplication efficiency, and any required backup copies to calculate usable repository capacity and growth allowance. Accurate collection of both workload quantity and protected-storage volume allows sizing to be based on the customer's actual protection scope rather than assumptions. Veeam's sizing guidance emphasizes gathering workload and data-volume characteristics before selecting infrastructure resources and repository capacity. See Veeam Backup Repository Capacity Calculator and Veeam Backup & Replication system requirements.

QUESTION NO: 9

What information related to the virtual infrastructure is missing and must be collected during the discovery phase? (Choose 2)

A. Local area network speed.

B. Recovery point objective.

C. Backup window.

D. Size of the source data set.

E. number of Microsoft Exchange mailboxes.

ANSWER: A D

Explanation:

Local area network speed and Size of the source data set are essential discovery inputs because they define the technical constraints on which a Veeam architecture is sized. Local area network speed establishes the available bandwidth between production workloads, backup proxies, repositories, and any remote components. This information is needed to estimate achievable data-transfer rates, identify potential traffic bottlenecks, and determine whether the design requires proxy placement, network traffic rules, throttling, or dedicated backup connectivity. Veeam's architecture uses data movers on proxies and repositories, so the performance of the network paths between those roles directly affects backup and restore processing.

Size of the source data set is equally necessary for capacity and performance planning. It provides the baseline for estimating the initial full-backup volume, repository capacity, projected retention consumption, and the processing resources needed to complete protection operations. Together with expected change rates and retention requirements collected elsewhere in discovery, the source-data size enables an architect to calculate usable storage requirements and validate that the infrastructure can support the intended protection design. Veeam documents the backup infrastructure components and their data-transfer roles in the Veeam Backup & Replication User Guide, and provides architecture-planning guidance in its VMware vSphere Best Practices Guide.

QUESTION NO: 10

While looking through the requirements regarding database servers, you see native Oracle utilities must be available for backups and restores. To make sure the native tools can interface directly with Veeam, which option most closely matches the requirements?

A. Set up RMAN to backup locality to a dedicated disk and use Veeam Agent File Folder backups to send the backups to a Veeam Backup & Replication repository.

B. Create a script to kick off virtual machine backups and initiate Oracle explorer exports in a RMAN format.

C. Set up RMAN to backups to a NAS share that gets backed up by Veeam NAS backups.

D. Install the Veeam RMAN Plug-in to allow native RMAN backups and restores to leverage Veeam repository directly from the Oracle servers.

ANSWER: D

Explanation:

Install the Veeam RMAN Plug-in to allow native RMAN backups and restores to leverage Veeam repository directly from the Oracle servers. This meets the requirement because Oracle Recovery Manager (RMAN) remains the utility used by the database administrator to initiate, manage, and restore Oracle database backups. The plug-in provides the integration layer between RMAN and Veeam Backup & Replication, allowing RMAN backup pieces to be written directly to a Veeam backup repository rather than first requiring a separate local disk, file-share staging area, or a secondary file-level protection workflow.

This design preserves the expected native Oracle operational model: administrators use familiar RMAN commands and catalog/control-file metadata for backup and recovery operations, while Veeam supplies repository-based storage and its associated backup-storage capabilities. It is therefore appropriate when the requirement specifically calls for native Oracle backup and restore utilities to interface directly with Veeam. The Veeam Plug-in for Oracle RMAN is designed for this purpose and supports RMAN-driven database backup and restore workflows. Veeam documents the architecture, deployment requirements, and supported operations in its Veeam Plug-in for Oracle RMAN documentation; Oracle describes RMAN as its native database backup and recovery manager in the Oracle Database Backup and Recovery User's Guide.

QUESTION NO: 11

When deciding on the design of the primary backup repository, which option best fits the requirements in the case study?

A. Dedupe appliance leveraging vendor API for access

B. Public cloud storage with S3 object-lock for immutability

C. Windows repository using ReFS integration, single-use credential and persistent VSS snapshot

D. Linux Repository using XFS integration, single-use credentials, and immutability

ANSWER: D

Explanation:

Linux Repository using XFS integration, single-use credentials, and immutability is the best fit because it combines performance, storage efficiency, and ransomware resilience in the primary backup tier. Veeam can use XFS fast-clone capabilities on a supported Linux repository to create synthetic full backups and transform backup chains without physically copying all backup blocks. This substantially reduces the time and storage I/O needed for these operations while retaining efficient backup-chain management. When the repository is deployed as a Veeam Hardened Repository, immutability prevents backup files from being changed or deleted for the configured retention period, including by an attacker who gains access to the backup server. The hardened-repository deployment model also uses a dedicated non-root account and supports the single-use credentials approach during setup, limiting credential exposure. Together, these capabilities provide a practical on-premises primary repository design that supports reliable restores, efficient synthetic operations, and a strong recovery posture against ransomware. Veeam documents Linux hardened repositories and their immutability controls in the Veeam Backup & Replication User Guide; XFS fast-clone requirements and behavior are covered in the block cloning documentation.

QUESTION NO: 12

Which of the following areas would benefit from additional analysis on areas mentioned in the case study? (Choose 3)

A. VMware

B. MSSQL

C. NAS

D. On-premises Exchange

E. PostgreSQL

F. Hyper-V

ANSWER: B C E

Explanation:

MSSQL, NAS, and PostgreSQL warrant additional analysis because each workload has data-consistency, recovery, performance, and capacity requirements that materially affect a Veeam architecture. For MSSQL, the design must establish application-aware processing and transaction-log backup requirements, including log-retention and point-in-time recovery objectives. These decisions affect backup frequency, repository consumption, restore procedures, and the credentials and permissions required for processing SQL Server workloads. Veeam documents the application-aware SQL Server processing considerations at Veeam Help Center: SQL Server.

NAS analysis is needed to identify the file-share protocol, file count, change rate, permissions, unstructured-data growth, and required restore granularity. Those inputs are essential for sizing NAS backup proxies, repositories, and backup windows, as well as selecting suitable storage integration and change-tracking methods. Veeam's NAS backup design and processing guidance is available at Veeam Help Center: File Share Backup.

PostgreSQL also requires workload-specific planning for database-consistent protection, log handling, recovery-point objectives, and validated restore workflows. Where scripts are used to coordinate database quiescence with image-level protection, the pre-freeze and post-thaw behavior must be tested with the database team. Analyzing these three areas ensures that the resulting design protects both the data and the applications' recovery requirements.