DUMPSARENA

Veeam Certified Engineer Plus v13

Veeam VMCE_v13

Version Demo

Total Demo Questions: 43

Total Premium Questions: 430

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

Which VMware VM backup transport modes are available in Veeam?

A. Direct Storage Access

B. Virtual Appliance

C. Network Mode

D. iSCSI Direct

E. NFS Direct

ANSWER: A B C

Explanation:

Veeam Backup & Replication supports three standard VMware vSphere VM data transport modes: Direct Storage Access, Virtual Appliance, and Network Mode. Direct Storage Access is correct because it allows the Veeam proxy to read VM data directly from the underlying storage infrastructure when the proxy has suitable storage connectivity and permissions. Virtual Appliance is correct because it uses VMware HotAdd behavior, attaching virtual disks of the processed VM to a backup proxy running as a VM in the same VMware environment. Network Mode is correct because it uses the VMware management network path, commonly referred to as NBD/NBDSSL, to move VM disk data through the ESXi host when direct storage or HotAdd processing is not used. These are the transport mode categories Veeam documents for VMware backup proxies, and Veeam can select among them automatically or allow an administrator to configure the preferred mode depending on proxy placement, datastore type, and infrastructure design. See the official Veeam documentation on transport modes and VMware backup proxies.

QUESTION NO: 2

After a network outage, CDP resumes and shows an RPO of 45 minutes instead of 5 seconds. What does this mean?

A. The CDP job must be rebuilt

B. Forty-five minutes of changes are still queued and the RPO should improve as they replicate

C. The target rejected the missing data

D. CDP switched to hourly mode

E. The display will refresh tomorrow

ANSWER: B

Explanation:

Forty-five minutes of changes are still queued and the RPO should improve as they replicate is correct. In Veeam CDP, the displayed RPO is a reflection of how far the replica is behind the source workload at that moment. CDP is designed to continuously capture and transfer I/O changes, but if connectivity is interrupted, changes cannot be delivered to the target site in real time. When the connection is restored, Veeam resumes processing and must work through the accumulated backlog. During that catch-up period, the effective RPO can be much higher than the configured target RPO, because the replica still lacks changes generated during the outage window. As replication traffic drains the queue and the target replica receives the missing data, the reported RPO should decrease until it returns close to the configured objective, assuming the network and target infrastructure can sustain the required throughput. This behavior aligns with Veeam's CDP architecture, where recovery points are created from continuously replicated data and the actual recoverability point depends on successful transfer of recent changes. See Veeam's documentation on Continuous Data Protection and how CDP works.

QUESTION NO: 3

A single file must be recovered from a Windows VM backup. Which recovery method is most efficient?

A. Full VM Restore

B. Instant VM Recovery

C. Guest OS File-Level Recovery

D. Replica failover

ANSWER: C

Explanation:

Guest OS File-Level Recovery is the most efficient recovery method when the requirement is to restore one file from a Windows virtual machine backup. Veeam Backup & Replication is designed to browse the guest file system inside a restore point and restore selected files or folders without bringing back the entire VM, powering on a recovered VM, or performing a failover workflow. For Windows guests, Veeam mounts the backup content to the backup server or mount server and uses the built-in file-level restore browser so the administrator can locate the needed file and restore it back to the original location, another location, or copy it out as needed. This keeps the recovery scope small, reduces restore time, and avoids unnecessary impact on production infrastructure. Veeam's documentation describes guest file recovery as the process for restoring individual guest OS files and folders from image-level backups, which directly matches the scenario of recovering a single file from a Windows VM backup. See the Veeam Help Center for Guest OS file recovery and restoring guest OS files.

QUESTION NO: 4

How can a design minimize AWS egress costs for common restores while still keeping an S3 copy?

A. Keep recent backups on-premises by using Copy mode

B. Use Snowball for every restore

C. Always restore into EC2 first

D. Compress data after upload

E. Use Direct Connect only

ANSWER: A

Explanation:

Keep recent backups on-premises by using Copy mode is correct because it matches the intended Veeam Scale-out Backup Repository design for balancing fast, low-cost restores with cloud durability. In Veeam Backup & Replication, Copy mode for the Capacity Tier copies restore points from the on-premises performance tier to object storage, such as Amazon S3, while the local restore points remain available according to the repository and job retention design. This means the organization still has an S3 copy for resiliency, off-site retention, and cloud-based recovery scenarios, but common restores of recent data can be performed from the local repository instead of reading data back from AWS. Since AWS data transfer out can generate egress charges, restoring frequently needed recent backups from on-premises storage is a practical way to reduce those costs. This approach also preserves restore performance because local backup storage is typically faster and avoids WAN dependency for day-to-day recovery operations. Veeam documents Capacity Tier behavior and the Copy policy for immediately copying backup files to object storage in its official guidance: Veeam Capacity Tier and Copying Backups to Capacity Tier.

QUESTION NO: 5

Which two scenarios are appropriate for backup from storage snapshots?

A. Reducing backup impact on production VMs

B. Eliminating changed-block reads for high-change workloads

C. Offloading snapshot work to the storage array

D. Backing up workloads that cannot use CBT at all

E. Reducing backup file size by itself

ANSWER: A C

Explanation:

Backup from storage snapshots is appropriate for **Reducing backup impact on production VMs** because Veeam can create a short-lived VMware snapshot, trigger a storage snapshot on the supported array, and then remove the VMware snapshot quickly. The backup data is then read from the storage snapshot rather than directly holding the production VM snapshot open for the full backup window, which helps reduce stun time and overall production workload impact. This is especially useful for busy virtual machines where long snapshot commit operations could affect application performance.

Offloading snapshot work to the storage array is also correct because the feature is specifically designed to integrate Veeam Backup & Replication with supported primary storage systems. The storage array performs the snapshot operation and exposes the snapshot data for backup processing, allowing Veeam to leverage array-level capabilities instead of relying only on hypervisor-level snapshot handling. Veeam documents this capability as “backup from storage snapshots” and describes it as a way to reduce impact on the production environment while using supported storage integrations. See the Veeam Help Center pages for Backup from Storage Snapshots and Storage Integration.

QUESTION NO: 6

Which tenant repository settings can a Cloud Connect provider configure?

A. Storage quota

B. Retention enforcement

C. Encryption requirement

D. Tenant backup window

E. Tenant proxy affinity

ANSWER: A B C

Explanation:

Storage quota, Retention enforcement, and Encryption requirement are the correct tenant repository settings a Veeam Cloud Connect provider can configure. In Cloud Connect, the service provider assigns backup storage resources to a tenant and controls how much repository capacity that tenant is allowed to consume, so Storage quota is a core provider-side setting. Providers can also apply retention-related controls for tenant backup data stored in the cloud repository, helping ensure that restore points or deleted backup data are handled according to the provider’s service policy. Encryption requirement is also a provider-controlled security setting: the provider can require tenant backup data placed in the cloud repository to be encrypted, which is a common best-practice control for multi-tenant cloud backup environments. These settings align with the provider’s responsibility to manage shared cloud repository resources, enforce service limits, and maintain security boundaries between tenants. For more detail, see the Veeam Cloud Connect documentation in the Veeam Cloud Connect Guide and the Veeam Backup & Replication documentation on Veeam Help Center.

QUESTION NO: 7

Which two cloud charges commonly apply when restoring data from an object storage capacity tier?

A. Data egress charges

B. Retrieval request charges

C. Normal ingress upload charges

D. Compute instance charges for every restore

E. Extra Veeam license charges per request

ANSWER: A B

Explanation:

Data egress charges and retrieval request charges are the two common cloud costs to expect when restoring data that resides in an object storage capacity tier. In Veeam, the capacity tier stores backup data in object storage as part of a Scale-out Backup Repository, and restore operations may require reading objects back from that cloud storage. Cloud providers typically bill for those read/retrieval operations, especially when the objects are in colder access tiers such as infrequent access, archive, or similar classes. They also commonly bill when restored data leaves the provider's network or region, which is generally categorized as data transfer out or egress. This is why restore planning for object storage should include both the number/volume of retrieval operations and the amount of data transferred out of the cloud provider environment. Veeam documents the capacity tier as object storage used for backup data placement and retrieval, while provider pricing pages such as Amazon S3 pricing explicitly list request/retrieval and data transfer charges as billable items. See Veeam Capacity Tier documentation and Amazon S3 pricing.

QUESTION NO: 8

Which SQL-related backup options are available through application-aware processing?

A. Crash-consistent only with no VSS

B. Truncate transaction logs after successful backup

C. Process transaction logs with the backup job

D. Use tape as the SQL log backup target

E. Disable guest processing but still truncate SQL logs

ANSWER: B C

Explanation:

Application-aware processing in Veeam Backup & Replication integrates with Microsoft VSS and the SQL Server VSS Writer so that SQL databases are backed up in an application-consistent state rather than only at the VM disk level. Within those guest processing settings, SQL-specific log handling can be configured. **Truncate transaction logs after successful backup** is correct because Veeam can request SQL log truncation after the backup completes successfully, helping prevent transaction logs from growing indefinitely when the job is responsible for SQL log maintenance. **Process transaction logs with the backup job** is also correct because Veeam can perform SQL transaction log processing as part of the protected workload's backup configuration, supporting SQL-aware restore scenarios such as replaying logs for more granular recovery points when configured appropriately. These capabilities are part of Veeam's SQL application-aware image processing and are configured in the job's guest processing/application settings. See Veeam's documentation on application-aware processing and Microsoft SQL Server transaction log processing.

QUESTION NO: 9

Which two destinations are valid export targets for Veeam Kasten backups?

A. S3-compatible object storage

B. An NFS share as the primary export target

C. A tape library

D. Microsoft Azure Blob Storage

E. A Veeam Cloud Connect repository

ANSWER: A D

Explanation:

S3-compatible object storage and Microsoft Azure Blob Storage are valid export targets for Veeam Kasten backups because Kasten K10 uses location profiles to define external backup/export repositories, with object storage being the standard target type for portable backup data. In Kasten terminology, exporting a backup means moving application restore points and associated metadata outside the cluster's primary persistent volumes so they can be retained independently, used for disaster recovery, or imported into another cluster. Kasten supports S3-compatible object storage, including AWS S3 and compatible platforms, as a location profile backend. It also supports Microsoft Azure Blob Storage as a location profile provider for storing exported snapshots and backups in Azure. These targets align with Veeam's Kubernetes data protection model, where durable cloud or S3-compatible object storage provides the required scalability, resilience, and portability for backup exports. For reference, see Veeam Kasten's documentation on Location Profiles and the supported storage details in the Kasten storage documentation.

QUESTION NO: 10

Which replication features help adapt VM networking at the target site?

A. Network mapping

B. Re-IP rules

C. Guest file indexing

D. Tape media pools

E. Backup encryption hints

ANSWER: A B

Explanation:

Network mapping and Re-IP rules are the replication features used to adapt VM networking when replicas are created or failed over at a target site. Network mapping lets Veeam Backup & Replication translate source-side virtual networks to the appropriate destination-side networks, such as mapping a production port group to a disaster recovery port group. This is especially useful when the target VMware environment uses different vSwitches, distributed switches, VLAN-backed port groups, or naming conventions. Re-IP rules complement this by changing guest operating system IP configuration during failover, so replica VMs can boot with IP addresses, subnet masks, gateways, and DNS settings that are valid in the disaster recovery network. Together, these features allow a replicated VM to connect correctly in the target site without requiring manual reconfiguration during a failover event. Veeam documents network mapping and replica Re-IP as core replication job settings for handling differences between source and target networking. See the Veeam documentation for network mapping and replica Re-IP.

QUESTION NO: 11

A SOBR should write new backups to local performance storage and move older backups to object storage after 30 days. Which settings are needed?

A. Data Locality placement

B. Capacity Tier using object storage

C. ReFS block cloning

D. Move backups older than the configured age

E. Archive Tier using tape

ANSWER: B D

Explanation:

Capacity Tier using object storage is correct because a scale-out backup repository uses its performance tier for the primary/local backup extents and can attach an object storage repository as a capacity tier. That is the SOBR mechanism that enables backups to be offloaded from local storage to S3-compatible storage, Azure Blob, Google Cloud, or other supported object storage targets. Move backups older than the configured age is also correct because the requirement is age-based movement after 30 days, which maps to Veeam's capacity tier move policy: moving backup files to object storage after they age out of the operational restore window. In practice, the operational restore window would be set to 30 days so newer restore points remain on the local performance tier, while older backup data is moved to the object storage capacity tier. Veeam documents this behavior in its scale-out repository and capacity tier guidance, where the capacity tier extends a SOBR with object storage and the move policy controls when inactive backup chains are moved based on retention age. See the Veeam Help Center pages for Scale-Out Backup Repository and Capacity Tier.

QUESTION NO: 12

Which two changes commonly require a tape infrastructure rescan?

A. A new tape drive is added

B. Tape library firmware is changed

C. Every completed tape job

D. A tape is moved to a vault

E. A media pool retention setting changes

ANSWER: A B

Explanation:

A new tape drive is added and Tape library firmware is changed are correct because Veeam Backup & Replication relies on tape infrastructure discovery to identify the current state and capabilities of connected tape hardware. When a new tape drive is introduced, Veeam must detect the device, enumerate it under the tape infrastructure, and associate it correctly with the tape server and library so it can be used by tape jobs. Similarly, firmware-level changes can alter how a tape library or drive reports its identity, slots, changer, drives, barcodes, or capabilities to the operating system and to Veeam. A rescan refreshes this hardware inventory and helps Veeam synchronize its view with the actual tape environment before jobs are run. Veeam documentation describes rescanning tape infrastructure as the operation used to detect newly added tape devices and update information about connected libraries and drives. For more detail, see the Veeam Help Center topic on rescanning tape infrastructure and the overview of tape devices.

QUESTION NO: 13

Which default port is required when VBR connects to a remote Microsoft SQL Server configuration database?

A. 1433 TCP

B. 443 TCP

C. 3389 TCP

D. 5432 TCP

ANSWER: A

Explanation:

1433 TCP is correct because Microsoft SQL Server's Database Engine listens on TCP port 1433 by default for client connections when using the default instance and standard TCP/IP configuration. When the Veeam Backup & Replication server uses a remote Microsoft SQL Server as its configuration database, the Veeam server must be able to reach that SQL Server listener so it can read and write configuration data, job metadata, infrastructure settings, and other operational

records stored in the VBR configuration database. Veeam's required ports documentation lists Microsoft SQL Server configuration database connectivity using TCP 1433 by default, which aligns with Microsoft's SQL Server networking guidance for the default TCP port. In real deployments, the port can differ if SQL Server is configured to use a named instance, a static non-default port, or dynamic ports, but the exam wording asks for the default port, so the required answer is 1433 TCP. References: Veeam Backup & Replication used ports and Microsoft SQL Server TCP port configuration.

QUESTION NO: 14

What can application-aware image processing do for a SQL Server VM when log processing is configured?

A. Truncate SQL transaction logs

B. Increase backup speed

C. Encrypt SQL data

D. Shut down SQL services

E. Exclude SQL files

ANSWER: A

Explanation:

Truncate SQL transaction logs is correct because Veeam application-aware image processing integrates with Microsoft VSS inside the guest OS to prepare application-consistent backups of workloads such as Microsoft SQL Server. When SQL Server log processing is enabled in the job's guest processing settings, Veeam coordinates the backup with SQL Server and VSS so that transaction logs can be processed safely after a successful backup point is created. This helps prevent SQL transaction logs from growing indefinitely while preserving database consistency and recoverability. In Veeam terms, this is part of application-aware processing and transaction log handling for Microsoft SQL Server VMs, allowing the backup job to create a consistent image-level backup and then perform log truncation according to the configured policy. This behavior is documented in the Veeam User Guide sections covering application-aware processing and Microsoft SQL Server transaction log processing.

QUESTION NO: 15

Which test types can be part of SureBackup verification?

A. VM heartbeat check

B. Application or script-based service test

C. Performance benchmark comparison

D. Network throughput benchmark

E. VSS writer inventory only

ANSWER: A B

Explanation:

VM heartbeat check is correct because SureBackup verification is designed to automatically start backed-up VMs in an isolated Virtual Lab and confirm that they are bootable and responsive. One of the core built-in tests is the heartbeat check, which uses VMware Tools or Hyper-V integration services to confirm that the guest OS has started successfully. Application or script-based service test is also correct because SureBackup can perform application-aware verification, including predefined role checks and custom scripts, to validate that key services inside the guest are available after the VM is powered on. This is the purpose of SureBackup: not just proving that backup files exist, but proving that workloads can be started and that the operating system and applications respond as expected in an isolated environment. Veeam documents these verification tests as part of SureBackup job behavior and application group processing. See Veeam's guidance on SureBackup recovery verification tests and the overview of SureBackup.

QUESTION NO: 16

A company wants to replicate on-premises VMware VMs to AWS for disaster recovery. Which Veeam approach fits?

A. Veeam Backup for AWS protecting only native AWS workloads

B. Veeam Backup & Replication using AWS as a Cloud Connect replication target

C. Installing Veeam Agent for Windows inside every VM as the replication method

D. Using an EC2 instance only as a backup proxy

E. Using Amazon S3 only as a bootable target

ANSWER: B

Explanation:

Veeam Backup & Replication using AWS as a Cloud Connect replication target is the best fit because the requirement is disaster recovery replication of on-premises VMware virtual machines to cloud-hosted recovery infrastructure. Veeam Backup & Replication is the Veeam component that performs image-level VMware VM replication and manages replica failover/failback workflows. With Veeam Cloud Connect Replication, a tenant can replicate production VMs to a cloud host exposed by a service provider, so the replicas are ready to be powered on during a disaster without the customer having to build and maintain a second physical DR site. In an AWS-based design, the target environment can be cloud-hosted infrastructure made available for Veeam replication rather than simply object storage or an AWS-native backup appliance. This aligns with Veeam's documented Cloud Connect Replication model, where Veeam Backup & Replication sends VM replicas to allocated cloud host resources for DR use cases. See Veeam's documentation for Cloud Connect Replication and the broader replication job functionality in Veeam Backup & Replication.

QUESTION NO: 17

A backup file is much larger than expected. Which cause might explain it?

A. High source change rate

B. Compression disabled or ineffective

C. An active full backup was created

D. CBT was reset

E. All of the above

ANSWER: E

Explanation:

All of the above

is correct because an unexpectedly large Veeam backup file can be caused by several normal product behaviors and environmental conditions. A high source change rate means an incremental restore point must store more changed blocks, so the resulting file can be much larger than a typical daily increment. If compression is disabled or the source data is already compressed, encrypted, or otherwise poorly compressible, Veeam cannot reduce the stored data as much as expected, increasing the backup file size. An active full backup also creates a new full restore point by reading the entire source workload, so it is naturally much larger than an incremental backup. Finally, if Changed Block Tracking is reset or cannot be trusted, Veeam may need to read and process more blocks than usual to ensure backup consistency, which can make the restore point larger than expected. These behaviors align with Veeam's documented backup methods, data reduction, and CBT-based incremental processing concepts. See the Veeam documentation on backup methods and Changed Block Tracking.

QUESTION NO: 18

Which requirements support granular Active Directory object recovery with Veeam Explorer?

A. Application-aware backup

B. Domain controller must be online

C. Backup contains domain controller/system state data

D. Forest functional level must be 2016

E. Schema Admin rights are always required

ANSWER: A C

Explanation:

Granular Active Directory object recovery with Veeam Explorer depends on having a usable, consistent backup of a domain controller that contains the Active Directory database and related system-state components. “Application-aware backup” is correct because Veeam application-aware processing uses Microsoft VSS to quiesce the domain controller workload and create a transactionally consistent restore point. That consistency is important when Veeam Explorer mounts the backup and reads Active Directory data for item-level browse, search, compare and restore operations. “Backup contains domain controller/system state data” is also correct because Veeam Explorer for Microsoft Active Directory needs the domain controller data set, including the directory database, to expose users, groups, organizational units, Group Policy objects and other AD items for granular recovery. In practice, the backup must be of a machine that actually hosts Active Directory Domain Services, not just any domain-joined server. Veeam documents application-aware processing as the mechanism for producing consistent backups of applications such as Microsoft Active Directory, and Veeam Explorer for Microsoft Active Directory is specifically designed to work with AD data from such restore points. See the Veeam documentation for application-aware processing and Veeam Explorer for Microsoft Active Directory.

QUESTION NO: 19

Which requirements fit Quick Migration of backup data within a SOBR?

A. Source and target locations are extents in the same SOBR

B. The repository must be Windows SMB

C. Backup files must be physically moved manually

D. Configuration metadata must be updated

E. The repository must be offline

ANSWER: A C D

Explanation:

Quick Migration for backup data in a scale-out backup repository is used when the backup placement is being changed between extents that belong to the same SOBR. In this workflow, the backup files are already relocated by manual or external file/storage operations, and Veeam Backup & Replication then performs the “quick” part by registering the new location rather than copying the backup chain itself. That is why “Source and target locations are extents in the same SOBR,” “Backup files must be physically moved manually,” and “Configuration metadata must be updated” are the requirements that fit this scenario. The metadata update is essential because Veeam must know which extent now owns the backup chain so future operations such as restore, retention, health check, and backup chain processing address the correct storage location. This behavior aligns with Veeam’s SOBR model, where extents are managed as parts of one logical repository, and backup file movement must be reflected in the Veeam configuration database. See Veeam’s documentation for Scale-Out Backup Repositories and moving backup files.

QUESTION NO: 20

Which statements correctly describe the Guest Interaction Proxy?

A. It is used for application-aware processing and guest file indexing

B. It replaces the backup proxy for data transport

C. It can communicate with guests through VMware Tools or the network

D. It is needed only for Linux guests

E. It stores the final backup files

ANSWER: A C

Explanation:

The correct statements are “It is used for application-aware processing and guest file indexing” and “It can communicate with guests through VMware Tools or the network.” In Veeam Backup & Replication, the Guest Interaction Proxy is the component that coordinates guest OS processing tasks inside protected machines. This includes operations required for application-aware image processing, such as preparing applications for consistent backup and collecting guest-related metadata, as well as guest file system indexing when that feature is enabled. Veeam can use a Guest Interaction Proxy to reach the guest OS and perform these operations without making the backup proxy itself responsible for guest-level communication. For VMware workloads, guest interaction can be performed through VMware Tools, or by connecting to the guest OS over the network when appropriate credentials and connectivity are available. This design helps separate data transport from guest processing and allows Veeam to select a suitable server for guest interaction tasks. Veeam documents this role in its Guest Interaction Proxy guidance and related guest processing documentation: Veeam Help Center: Guest Interaction Proxy and Veeam Help Center: Guest Processing.

QUESTION NO: 21

Which file extension identifies a full Veeam backup file?

A. VIB

B. VBK

C. VRB

D. VSB

ANSWER: B

Explanation:

VBK is the correct file extension for a full Veeam backup file. In Veeam Backup & Replication, a full backup restore point contains the complete set of protected VM data required to restore the workload at that point in time. Veeam stores that full backup restore point in a file with the .vbk extension. This applies to standalone full backups as well as full restore points that form the base of an incremental backup chain. Because the full backup file contains the complete data set, it is the anchor point that later restore points depend on in forward incremental backup chains. Veeam documentation identifies .vbk as the full backup file format and describes it as one of the core backup chain file types used by backup jobs. For exam purposes, when the question asks which file extension identifies a full Veeam backup file, the expected answer is VBK. See Veeam’s documentation on backup files and backup chains for the official file type definitions: Veeam Backup Files and Veeam Backup Chain.

QUESTION NO: 22

Which report categories are available natively in Enterprise Manager?

A. Backup success or failure summary

B. Storage capacity usage

C. Retention compliance

D. Per-job bandwidth detail

E. Application performance metrics

ANSWER: A B C

Explanation:

Enterprise Manager includes native reporting intended to give administrators a centralized view of protection status, repository consumption, and policy/SLA adherence across connected Veeam Backup & Replication servers. **Backup success or failure summary** is correct because Enterprise Manager aggregates job session information and presents job status reporting so teams can quickly identify whether protected workloads are being backed up successfully. **Storage capacity usage** is also correct because Enterprise Manager exposes repository and backup storage consumption information, helping administrators understand used space, available capacity, and growth trends from a central console. **Retention compliance** is correct because Enterprise Manager includes compliance-oriented views and reporting around whether workloads meet configured protection and retention expectations, which is especially useful when managing multiple backup servers or delegating restore/reporting access. These capabilities align with Enterprise Manager's role as a centralized management, reporting, and self-service portal rather than a detailed application monitoring or network analytics tool. For reference, see the Veeam Backup Enterprise Manager overview in the Veeam Help Center and the Enterprise Manager reporting documentation at Veeam Help Center reports section.

QUESTION NO: 23

What is the key difference between a backup job and a replication job?

A. Replication creates a bootable target VM

B. Replication always compresses more

C. Replication never uses snapshots

D. Replication is always slower

E. Replication cannot be scheduled

ANSWER: A

Explanation:

Replication creates a bootable target VM is correct because a Veeam replication job is designed to maintain a ready-to-start copy of a production virtual machine on another VMware vSphere or Microsoft Hyper-V host. The replica is registered as a VM at the target location and includes restore points, so during an outage an administrator can perform failover and power on the replica to resume service quickly. This is the key operational distinction from a backup job: backups are stored as backup files in a repository and are primarily used for restore operations, while replicas are VM copies prepared for rapid failover. Veeam documentation describes replication as creating an exact copy of a VM in its native format on a standby host, with subsequent job runs transferring only changes to keep that replica up to date. See Veeam's guidance on replication and failover to a VM replica for the official workflow and behavior.

QUESTION NO: 24

Which component reduces traffic for offsite backup transfer across a poor WAN link?

A. Guest Interaction Proxy

B. WAN Accelerator

C. Veeam Catalog Service

D. SQL Server Browser

ANSWER: B

Explanation:

WAN Accelerator is correct because it is the Veeam component specifically designed to optimize data transfer over slow, high-latency, or unreliable WAN connections. In Veeam Backup & Replication, WAN acceleration is used primarily with backup copy jobs and replication scenarios where data must be transferred to an offsite location. The source and target WAN accelerators reduce the amount of data sent across the link by using global data caching, variable-block deduplication, compression, and traffic optimization techniques. Instead of repeatedly transferring identical blocks over the WAN, Veeam can reuse data already stored in the WAN accelerator cache at the remote site, significantly lowering bandwidth consumption and improving transfer performance across poor links.

This makes WAN Accelerator the best-fit component when the requirement is to reduce offsite backup transfer traffic, especially between geographically separated locations. Veeam documentation describes WAN accelerators as optional infrastructure components that help optimize data transfer over WAN and slow connections. See the official Veeam guidance on WAN acceleration and WAN acceleration for backup copy jobs.

QUESTION NO: 25

After a clean restore, what security step is critical before reconnecting the VM to production?

A. Change passwords for accounts that existed on the VM

B. Reinstall the operating system every time

C. Only update VMware Tools

D. Enable CBT

E. Skip validation and reconnect quickly

ANSWER: A

Explanation:

Change passwords for accounts that existed on the VM is correct because a clean restore only returns the workload to a known-good state; it does not automatically invalidate credentials that may have been exposed before or during the ransomware incident. If attackers obtained local administrator, service account, application, or cached domain credentials, reconnecting the restored VM with the same passwords can give them an immediate path back into the environment. Credential rotation before production reconnection helps break persistence, reduces the chance of reinfection, and supports a controlled recovery process. This aligns with ransomware recovery guidance that emphasizes validating systems and securing identity before returning services to normal operation. Veeam's recovery approach includes secure restore and validation capabilities to help ensure restored machines are safe before use, as described in Veeam Secure Restore documentation. Broader incident-response guidance from CISA's StopRansomware Guide also stresses securing compromised accounts and credentials as part of recovery. In practice, password changes should cover accounts that existed on the VM, especially privileged and service accounts, before the machine is trusted again in production.

QUESTION NO: 26

Where does Veeam ONE store collected monitoring and reporting data?

A. A Microsoft SQL Server database

B. Only in vCenter inventory

C. Only in Windows Event Logs

D. Inside backup files on the repository

E. In a separate CSV for each alarm

ANSWER: A

Explanation:

A Microsoft SQL Server database is correct because Veeam ONE uses its own database to retain configuration, monitoring, performance, alarm, and reporting data collected from the protected virtual and backup infrastructure. In a standard Veeam ONE deployment, the Veeam ONE Server components write collected data to the Veeam ONE database hosted on Microsoft SQL Server, which is then used by Veeam ONE Client and Veeam ONE Web Client for monitoring dashboards, historical performance views, alarms, capacity planning, and reports. This centralized database is essential because Veeam ONE is designed not only to show current infrastructure state, but also to provide historical analytics and reporting over time. Official Veeam deployment documentation lists Microsoft SQL Server as the database platform for Veeam ONE and describes the Veeam ONE database as a required component of the product architecture. See the Veeam ONE system requirements in the Veeam ONE Deployment Guide and the product architecture overview in Veeam ONE Architecture.

QUESTION NO: 27

A healthcare organization requires backups encrypted at rest and controlled with on-premises keys. Which two settings fit?

A. Enable backup file encryption with passwords stored in the on-premises Veeam credential store

B. Use only TLS for Veeam component communication

C. Encrypt repository volumes with BitLocker and on-premises key management

D. Use provider-managed cloud encryption keys

E. Use agent backup without repository encryption

ANSWER: A C

Explanation:

Backup file encryption with passwords stored in the on-premises Veeam credential store is correct because Veeam Backup & Replication can encrypt backup data before it is written to the target repository. The encryption password is managed within the Veeam configuration, so the organization retains control of the encryption secret instead of relying on a cloud provider-managed key. This directly satisfies the requirement for backup data to be encrypted at rest while keeping key ownership on premises. See Veeam's documentation on data encryption.

Encrypt repository volumes with BitLocker and on-premises key management is also correct because volume-level encryption protects the storage that holds backup files. When BitLocker keys are managed internally, the healthcare organization maintains local administrative control over the keys used to unlock the repository volumes. This adds another at-rest protection layer for backup storage and aligns with security and compliance expectations for regulated environments. Microsoft documents BitLocker as a volume encryption technology designed to protect data at rest; see BitLocker overview.

QUESTION NO: 28

What are the two primary Veeam ONE components?

A. Veeam ONE Monitor and Veeam ONE Reporter

B. Veeam Proxy and Veeam Repository

C. Veeam Explorer and SureBackup

D. Cloud Gateway and Tape Server

ANSWER: A

Explanation:

Veeam ONE Monitor and Veeam ONE Reporter is correct because these are the two main functional areas administrators use in Veeam ONE for operational visibility and reporting. Veeam ONE Monitor provides real-time monitoring, alerting, performance analysis, and infrastructure visibility for protected virtual and backup environments. It is used for day-to-day health checks, troubleshooting, alarms, dashboards, and tracking the status of workloads and Veeam Backup & Replication activity. Veeam ONE Reporter provides the reporting and planning side of the product, including predefined and customizable reports, capacity planning, audit information, chargeback/showback-style visibility, and compliance-oriented reporting. Together, they represent the monitoring and reporting capabilities that Veeam ONE is designed to deliver. This aligns with Veeam's product documentation, which describes Veeam ONE as a monitoring, reporting, and analytics solution for virtual and backup infrastructures, with monitoring and reporting delivered through its client and web/reporting interfaces. See the official Veeam ONE documentation for more context: Veeam ONE Architecture and Veeam Monitoring & Analytics.

QUESTION NO: 29

What is the main idea behind Veeam Universal License?

A. A license tied only to physical CPU sockets

B. A portable workload-based license pool for different protected workload types

C. A license used only for Microsoft 365 mailboxes

D. A free license that covers unlimited agents

ANSWER: B

Explanation:

"A portable workload-based license pool for different protected workload types" is correct because Veeam Universal License is designed around workload portability rather than being locked to one infrastructure type. The core idea is that an organization buys a pool of license units and consumes them across supported protected workloads, such as virtual machines, cloud instances, physical servers protected with agents, NAS capacity, and other Veeam-supported workloads depending on the product and edition. This model is especially useful in modern environments where workloads may move between on-premises virtualization, public cloud, physical systems, and unstructured data platforms over time. Instead of purchasing separate, rigid license types for each platform, Veeam Universal License provides operational flexibility and simplifies license management through a common consumption model. Veeam describes this approach as portable licensing for protecting diverse workloads with Veeam Data Platform capabilities. See Veeam's official licensing information at Veeam Licensing & Pricing and the Veeam Backup & Replication licensing documentation at Veeam Help Center.

QUESTION NO: 30

During restore from a Cloud Connect repository, which components can participate in the data path?

A. Tenant local backup proxy

B. Provider Cloud Gateway

C. Tenant VBR server as data mover

D. Provider production vCenter

E. Tenant WAN Accelerator if configured

ANSWER: A B E

Explanation:

During a restore from a Veeam Cloud Connect repository, the Provider Cloud Gateway is part of the transport path because it is the service provider-side network entry point that tunnels tenant traffic to and from the cloud repository. A Tenant local backup proxy can also participate when the restore writes VM data back to the tenant's local virtualization environment; in that case, the proxy runs the transport/data mover component that receives restored blocks and writes them to the target

datastore or host using the selected transport mode. Tenant WAN Accelerator if configured is also correct because WAN acceleration can be used for Cloud Connect backup copy and restore traffic to reduce data sent across the WAN, with the tenant-side accelerator participating in the optimized transfer path. In short, Cloud Connect restore traffic is carried through the provider gateway, can be optimized through WAN acceleration, and, for VM image restores to local infrastructure, is handled on the tenant side by a backup proxy/data mover. See Veeam's Cloud Connect and WAN acceleration documentation for the supported architecture and roles: Veeam Cloud Connect and WAN Acceleration.

QUESTION NO: 31

A production VM was deleted yesterday, and the latest backup is 24 hours old. Which restore method minimizes downtime?

A. Full VM Restore

B. Instant VM Recovery

C. Guest file restore

D. Virtual disk restore only

E. Instant disk recovery only

ANSWER: B

Explanation:

Instant VM Recovery is correct because it is designed specifically to bring a failed or deleted production virtual machine back online as quickly as possible from an existing restore point. Instead of waiting for the entire VM to be copied back to production storage before users can access it, Veeam publishes and powers on the VM directly from the backup repository. This allows service availability to be restored within minutes, depending on infrastructure performance, while the VM is still running from backup storage. After the VM is accessible again, it can be migrated back to production storage using the appropriate migration method, such as Quick Migration or storage migration, completing the recovery without requiring a long initial outage. For a deleted VM where the latest available backup is 24 hours old, the priority in the question is minimizing downtime, and Instant VM Recovery is the Veeam restore workflow intended for that objective. Veeam describes this capability in its Instant Recovery documentation: Veeam Instant Recovery. Additional Veeam recovery guidance is available in the user guide: Veeam Restore from Backup.

QUESTION NO: 32

A company wants AWS as a DR target for on-premises VMware with 15-minute RPO and about 1-hour RTO. Which Veeam design fits best?

A. Backup copy to S3 only

B. Basic replication job directly to EC2 without provider design

C. CDP directly to AWS

D. Agent backup to AWS storage

E. Cloud Connect Replication through an AWS-based provider

ANSWER: E

Explanation:

Cloud Connect Replication through an AWS-based provider is the best fit because it is designed for disaster recovery as a service using provider-hosted replica resources. In this design, the on-premises VMware workloads are replicated to a cloud host supplied by a Veeam Cloud Connect service provider, giving the tenant a ready-to-start replica environment rather than only off-site backup data. That aligns well with a 15-minute recovery point objective when replication jobs are scheduled frequently, and it also supports a roughly 1-hour recovery time objective because failover can be performed to already-created replicas in the provider environment. Veeam Cloud Connect Replication also includes DR-oriented capabilities such

as cloud failover, failback, replica resources, and network extension scenarios, which are important when AWS is being used through a provider architecture as the recovery site. Veeam's documentation describes Cloud Connect as the mechanism for sending replicas to a service provider cloud and using those replicas for failover operations. See the Veeam documentation for Veeam Cloud Connect and replication to cloud.

QUESTION NO: 33

A restore fails because Veeam cannot mount a backup file. Which two checks are most relevant?

A. Repository connectivity and permissions

B. Mount Server services, resources, and drivers

C. Report schedule count

D. Tape vault barcode format

E. Business View category

ANSWER: A B

Explanation:

The most relevant checks are **Repository connectivity and permissions** and **Mount Server services, resources, and drivers**. During restore workflows that require browsing or exposing backup contents, Veeam must be able to access the backup files on the repository and use the mount server associated with that repository. If the repository path is unavailable, credentials are invalid, permissions do not allow reading the backup chain, or network connectivity to the repository is interrupted, the mount operation can fail before restore data is exposed. Veeam also relies on the mount server to perform key restore mounting functions, so its services must be running and it must have sufficient CPU, RAM, disk space, and required components or drivers available. In Veeam documentation, the mount server is specifically tied to backup repositories and restore operations that mount backup content for browsing or recovery. For more detail, see Veeam's guidance on mount servers and backup repositories.

QUESTION NO: 34

Which three configurations help satisfy data sovereignty rules that require backups to remain in-country?

A. Place repositories in the same country as the protected data

B. Use object buckets locked to the approved region

C. Keep customer-managed encryption keys under on-premises control

D. Copy backups to a different country for disaster recovery

E. Use a service provider in another jurisdiction

ANSWER: A B C

Explanation:

Place repositories in the same country as the protected data is correct because Veeam backup repositories define where backup files are physically stored, so selecting repository infrastructure located inside the required jurisdiction directly supports data residency and sovereignty requirements. Use object buckets locked to the approved region is also correct because object storage used by Veeam is tied to a provider region or endpoint, and choosing an approved in-country region helps ensure backup data written to object storage remains within the required geographic boundary. Keep customer-managed encryption keys under on-premises control is correct because sovereignty programs commonly combine residency controls with customer-controlled encryption, ensuring the organization retains control over access to protected backup data and cryptographic material rather than delegating that control to an external provider in another jurisdiction. Together, these configurations align storage placement, cloud-region selection, and encryption-key governance with the requirement that backup data remain under the approved country's legal and operational control. See Veeam guidance on backup

repositories and object storage repositories in the Veeam Backup Repository documentation and Veeam Object Storage Repository documentation.

QUESTION NO: 35

After a VM is migrated to Azure, what is the recommended Veeam protection method?

A. Install only a generic in-guest Veeam Agent

B. Protect the native Azure VM with Veeam Backup for Microsoft Azure

C. Replicate it back to VMware as the primary backup

D. Use Azure only as a VBR repository for that VM

E. Use Azure Backup as a Veeam-integrated product

ANSWER: B

Explanation:

Protect the native Azure VM with Veeam Backup for Microsoft Azure is correct because, after migration, the workload is no longer primarily a vSphere or Hyper-V object; it is an Azure IaaS virtual machine and should be protected using Veeam's cloud-native Azure data protection capabilities. Veeam Backup for Microsoft Azure is designed specifically for Azure workloads, using Azure-native snapshot orchestration, policy-based protection, application-aware processing where supported, and recovery options that understand Azure constructs such as subscriptions, resource groups, regions, virtual networks, and managed disks. This aligns the protection method with the workload's new operating environment instead of treating Azure merely as storage or trying to continue using an on-premises-style protection model. It also integrates with Veeam Backup & Replication, allowing centralized management, backup copy, retention, and recovery workflows across hybrid and cloud environments. Veeam's documentation describes Veeam Backup for Microsoft Azure as the product intended to protect Azure VMs and other Azure workloads with native backup and restore capabilities. See the Veeam Backup for Microsoft Azure User Guide and Veeam's Azure backup product page for details.

QUESTION NO: 36

Which authentication methods are supported for Cloud Connect tenant access?

A. Local Veeam user accounts

B. Active Directory integration

C. SAML SSO

D. RADIUS

E. Certificate-only authentication

ANSWER: A B

Explanation:

Cloud Connect tenant access is based on tenant credentials defined and managed by the service provider in the Cloud Connect infrastructure. Local Veeam user accounts is correct because a provider can create standalone tenant accounts in Veeam Backup & Replication and assign those accounts the required cloud repository, cloud host, quota, and other tenant resources. The tenant then uses those credentials when adding the service provider in their Veeam console, and those credentials control access to the provider-side Cloud Connect resources.

Active Directory integration is also correct because Cloud Connect tenant access can be tied to directory-based identities where the provider uses Active Directory users or groups for tenant authentication and management. This lets service providers align Cloud Connect access with existing identity administration while still applying Veeam-specific tenant resource assignments and quotas. Veeam documents Cloud Connect as a provider-managed infrastructure where tenants connect

with assigned credentials and consume allocated resources; see the Veeam Cloud Connect Guide and the Veeam Backup & Replication Cloud Connect documentation.

QUESTION NO: 37

A VM must be restored to a point before ransomware infection. Which approach is safest before returning it to production?

A. Restore the newest point without checking it

B. Restore the oldest point and accept the data loss

C. Only scan one restore point during restore

D. Assume a weekly GFS point is clean

E. Restore in an isolated network, scan and validate it, then promote it to production

ANSWER: E

Explanation:

Restore in an isolated network, scan and validate it, then promote it to production is the safest approach because it separates recovery from business operations until the restore point has been proven trustworthy. In a ransomware scenario, the key risk is not just restoring data, but accidentally reintroducing malware, persistence mechanisms, encrypted files, or compromised services back into the production environment. Veeam supports this type of controlled validation through capabilities such as SureBackup, which can start restored workloads in an isolated virtual lab for verification, and Secure Restore, which can scan restored machine disks with antivirus or antimalware tools before the workload is placed back into service. This gives administrators a safer workflow: restore the VM away from production, perform malware scanning and application or OS validation, confirm the selected restore point meets recovery requirements, and only then reconnect or promote the workload to production. This aligns with Veeam ransomware recovery best practices because it reduces reinfection risk while still enabling administrators to choose the cleanest usable restore point. References: Veeam SureBackup recovery verification and Veeam Secure Restore.

QUESTION NO: 38

Which settings are part of the Storage step in a backup job wizard?

A. Backup Proxy

B. Backup Repository

C. Retention Policy

D. Guest Processing

E. Schedule

ANSWER: A B C

Explanation:

Backup Proxy, Backup Repository, and Retention Policy are all settings handled in the Storage step of the Veeam backup job wizard. In this step, Veeam lets the administrator decide how the backup data will be transported, where it will be stored, and how long restore points should be kept. Backup Proxy selection determines which proxy components are used for data processing and transport, either automatically or through explicit selection. Backup Repository selection defines the target repository where backup files are written. Retention Policy is also configured here because it controls the number of restore points or retention period that Veeam maintains for the job on the selected storage target. These settings are grouped together because they directly affect backup data flow, capacity planning, restore point availability, and repository consumption. Veeam's documentation for the backup job wizard describes the Storage step as the place to configure repository, retention, proxy, and related storage settings. See the official Veeam guidance for the Storage step and the broader backup job wizard.

QUESTION NO: 39

What is the difference between Active Full and Synthetic Full backup creation?

A. Active Full reads source workload data again; Synthetic Full builds a full from repository backup files

B. Active Full never uses production storage; Synthetic Full always rereads the VM

C. Both methods always read every block from production

D. Synthetic Full is available only for tape jobs

ANSWER: A

Explanation:

Active Full reads source workload data again; Synthetic Full builds a full from repository backup files is correct because it captures the essential operational difference between the two full-backup creation methods in Veeam. An active full backup is created by reading the required data directly from the protected production workload, such as a VM, physical machine, or other source, and writing a new full backup file to the repository. This produces an independent full restore point based on a fresh read of production data. A synthetic full backup, however, is assembled on the backup repository by combining data already stored in the previous full backup and subsequent incremental restore points. In other words, it creates a new full backup file without rereading the entire protected workload from production storage. This can reduce load on production systems and networks, while shifting the processing and I/O activity to the backup repository. Veeam documents this distinction in its backup chain behavior and full backup method descriptions, including active full and synthetic full creation. See the Veeam User Guide sections on full backup methods and synthetic full backup.

QUESTION NO: 40

What happens in a planned failover of a replicated VM?

A. The source VM is gracefully stopped after a final sync and the replica is started

B. The replica starts without any source coordination

C. Only an isolated test is created

D. The source VM is backed up before the replica starts

ANSWER: A

Explanation:

The source VM is gracefully stopped after a final sync and the replica is started is correct because Veeam planned failover is designed for controlled migration or anticipated outage scenarios where the production VM is still available. During this process, Veeam Backup & Replication first performs a final incremental synchronization to transfer the latest changes from the source VM to the replica. It then shuts down the original production VM in an orderly way to prevent additional writes and to preserve application and data consistency as much as possible. After the final state is applied to the replica, Veeam powers on the replica VM at the target location so workloads can continue running with minimal or no data loss. This differs from an emergency failover because planned failover coordinates with the source VM and intentionally brings it down before starting the replica, making it the preferred approach for planned maintenance, data center migration, or avoiding an expected site outage. Veeam documents planned failover as a controlled operation that includes synchronization, source VM shutdown, and replica startup. See the official Veeam User Guide for Planned Failover and related Replica Failover behavior.

QUESTION NO: 41

What is a key advantage of a synthetic full backup over an active full backup?

A. It builds the full backup from existing repository data instead of rereading the whole source workload

B. It removes the need for incrementals

C. It requires the VM to be powered off

D. It disables compression

E. It writes only to tape

ANSWER: A

Explanation:

It builds the full backup from existing repository data instead of rereading the whole source workload is correct because a synthetic full backup is constructed on the backup repository from the previous full backup and subsequent incremental restore points. In Veeam, this produces a new full restore point without requiring the backup job to read the entire protected workload from production storage again. That is the key operational benefit compared with an active full backup, which retrieves all source data from the production environment to create a new full backup file. By assembling the full backup from existing backup-chain data, synthetic full backup activity shifts much of the workload to the repository, helping reduce production storage I/O and network traffic during the backup window. This is especially valuable for large virtual machines, busy production datastores, or environments where minimizing impact on source systems is important. Veeam documents synthetic full backups as full backups created from existing backup files, while active full backups are created by reading VM data from source storage. See the Veeam documentation for synthetic full backup and active full backup.

QUESTION NO: 42

What does WORM tape support provide?

A. Data written to WORM cartridges cannot be overwritten or erased

B. Tape cartridges become automatically encrypted by Active Directory

C. Tape jobs run only once per month

D. A tape library no longer needs media pools

ANSWER: A

Explanation:

Data written to WORM cartridges cannot be overwritten or erased is correct because WORM stands for “Write Once, Read Many.” In the context of Veeam tape support, WORM media is used when an organization needs backup data placed on tape in a non-rewriteable, non-erasable state for long-term archival retention, regulatory compliance, or ransomware-resilient offline storage. Once data has been successfully written to a WORM tape cartridge, the media’s behavior prevents that recorded data from being modified or deleted through normal tape operations. This gives the tape copy immutability characteristics: it can be read many times for restore or audit purposes, but the protected written blocks cannot be overwritten. Veeam Backup & Replication supports tape workflows for long-term retention and archival use cases, and WORM-capable tape media extends that by enforcing write-once protection at the media level. This is especially useful for compliance-driven retention policies where backup records must remain unchanged for a defined period. See Veeam’s tape device and tape support documentation in the Veeam Backup & Replication User Guide, and the general WORM media concept described by IBM tape documentation.

QUESTION NO: 43

Which three resources can be included in a Veeam Kasten application backup?

A. Persistent volume data

B. ConfigMaps and Secrets

C. Deployment or service definitions

D. Container runtime logs as the main target

E. The host node operating system

ANSWER: A B C

Explanation:

Veeam Kasten protects Kubernetes applications as a collection of related data and Kubernetes objects, so **Persistent volume data**, **ConfigMaps and Secrets**, and **Deployment or service definitions** can all be included in an application backup. This is central to Kasten's application-centric model: it does not just copy storage blocks, but captures the state needed to restore the application in a usable form. Persistent volume data represents the application's stateful data, while ConfigMaps and Secrets provide configuration and sensitive runtime information required by workloads. Deployment and service definitions are Kubernetes API resources that describe how workloads should run and how they are exposed or reached inside the cluster. Together, these resources allow Kasten to restore both the data and the Kubernetes context of the application, which is especially important for migration, disaster recovery, and rollback scenarios. Kasten documentation describes applications as being made up of Kubernetes resources and associated persistent data, and protection policies can snapshot/export this application state for recovery. See the Veeam Kasten documentation on applications and protecting applications for more detail.