DUMPSBOSS.

Veeam Certified Engineer Plus v13

Veeam VMCE_v13

Version Demo

Total Demo Questions: 45

Total Premium Questions: 450

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

Which file extension identifies a full Veeam backup file?

A. VIB

B. VBK

C. VRB

D. VSB

ANSWER: B

Explanation:

VBK is the file extension for a full backup file created by Veeam Backup & Replication. A full backup contains the complete set of VM data necessary to restore the protected workload to the point represented by that restore point. In a backup chain, the VBK file is the full-backup component that provides the baseline data for subsequent incremental restore points. Veeam can create this file during an active full backup, and it is also used when a synthetic full backup is produced from backup data already held in the repository. The file therefore represents the complete restore-point data rather than only changes captured since a previous backup. Veeam documentation identifies VBK as a core backup-chain file type and describes it as the file that stores full backup data. This extension is used for VM backup files managed by Veeam Backup & Replication, so it is the expected answer when identifying a full Veeam backup file by its extension. For the official definitions of Veeam backup file types and the relationship between full and incremental backup data, see Veeam Backup Files and Veeam Backup Chain.

QUESTION NO: 2

Which two actions can be automated in a Veeam tape workflow?

A. Export selected backups according to GFS requirements

B. Erase or reuse expired tapes according to media pool policy

C. Deploy a new vCenter Server

D. Disable all backup encryption

E. Install Kubernetes operators

ANSWER: A B

Explanation:

Veeam tape workflows automate both the selection of backup data for long-term archival and the lifecycle handling of tape media. “Export selected backups according to GFS requirements” is supported through GFS media pools and tape jobs. A GFS policy defines weekly, monthly, quarterly, and yearly retention points, and Veeam automatically identifies the applicable backup chains or restore points and writes them to tape according to the configured schedule. This allows organizations to meet long-term retention and compliance requirements without manually selecting each archival backup.

“Erase or reuse expired tapes according to media pool policy” is also an automated tape-workflow function. Media pools apply media-set retention rules that protect tapes from overwrite for the configured retention period. Once protection expires, Veeam can move the tape into an available state and use it for subsequent tape jobs, subject to the pool configuration and available media. This provides controlled media rotation and reduces routine operator intervention while preserving retention safeguards. Veeam documents GFS archival behavior in its GFS Media Pools documentation and describes tape allocation, retention, and reuse in its Media Pools documentation.

QUESTION NO: 3

Which three metrics help reveal backup infrastructure bottlenecks early?

A. Sustained high proxy CPU utilization

B. Low repository free space

C. High latency between proxy and repository

D. The number of VMs in a job name

E. The chosen job start time label

ANSWER: A B C

Explanation:

Sustained high proxy CPU utilization, low repository free space, and high latency between proxy and repository are useful early-warning metrics because they directly reflect constraints in the Veeam backup data path. A backup proxy performs data retrieval, processing, compression, encryption, and transport functions. Persistently high CPU use can therefore indicate that proxy processing capacity is approaching saturation, which can limit concurrent task performance and reduce backup throughput. Repository free capacity must be monitored proactively because insufficient space can affect the ability to create and retain restore points and can disrupt repository operations before protection requirements are met. Latency between the proxy and repository is equally significant: backup data must traverse this path, so elevated latency can constrain transfer rates and indicate network-path pressure. Veeam job statistics include bottleneck analysis to identify whether the source, proxy, network, or target is the limiting component for a task. Veeam ONE complements this by providing monitoring, alarms, reporting, and capacity-planning visibility across the backup infrastructure. Reviewing these metrics as trends, rather than waiting for a job failure, enables administrators to add capacity or resolve connectivity issues before backup windows and recovery objectives are affected. See Veeam job statistics and bottleneck analysis and Veeam ONE monitoring overview.

QUESTION NO: 4

A custom report must include only Finance department backup jobs. What should be configured first?

A. A Business View tag or group for Finance workloads

B. A separate backup server for Finance

C. Manual export of every job log

D. A new tape vault

E. A different backup proxy for each Finance VM

ANSWER: A

Explanation:

A Business View tag or group for Finance workloads should be configured first. Veeam ONE Business View provides the business-level classification layer needed to organize infrastructure and protection objects according to categories meaningful to the organization, such as department, service, location, or owner. Create a Finance-oriented Business View category and group, define membership rules that reliably identify the relevant Finance workloads or backup jobs, and verify that the intended objects are included. The custom report can then use that Business View scope or filter to return only the Finance-related protection data. This separates reporting requirements from the underlying technical layout of the backup infrastructure and makes the report maintainable as jobs and workloads change over time. Veeam documents Business View as a feature for grouping infrastructure objects by business criteria and using those groups to provide business context in monitoring and reporting. See the Veeam ONE Business View overview and the Veeam ONE reporting documentation.

QUESTION NO: 5

What happens if a Backup Copy Job is interrupted during transfer and connectivity later returns?

A. The restore point starts from zero again

B. The transfer resumes from the last successful checkpoint

C. The partial data is discarded and an active full begins

D. The job must always be restarted manually

ANSWER: B

Explanation:

The transfer resumes from the last successful checkpoint. Backup Copy Jobs are designed for moving backup data between locations, including remote repositories connected through potentially unreliable WAN links. Veeam tracks the progress of the transfer and retains data that was successfully transmitted before the connectivity interruption. When communication with the target becomes available again, the job can continue the interrupted data movement rather than retransmitting the entire restore point from the beginning.

This resumable behavior is important for efficient offsite copy operations because it reduces repeated network traffic, shortens recovery from transient outages, and helps the copy job meet its configured copy interval and retention objectives. The job's ability to continue depends on the required source, target, and transport components becoming reachable again, but a temporary loss of connectivity does not inherently require recreating the restore point. Veeam documents Backup Copy Jobs as a mechanism for copying restore points to another backup repository and describes their processing workflow in the Veeam Backup & Replication User Guide. See Backup Copy Jobs and Backup Copy Job Workflow.

QUESTION NO: 6

Which three factors influence whether VMware backups should use Hot Add or Direct Storage Access?

A. Whether the proxy is virtual or physical

B. Whether the proxy can see the source storage directly

C. The number of concurrent tasks only

D. Backup window and performance requirements

E. The exact vSphere patch version in every supported release

ANSWER: A B D

Explanation:

Whether the proxy is virtual or physical is a primary factor because Hot Add (virtual appliance mode) requires a VMware backup proxy deployed as a virtual machine. In this mode, Veeam temporarily hot-adds snapshot disks from the protected VM to the proxy VM, allowing the proxy to read the data without transferring it through the production network. A physical proxy cannot use Hot Add and is instead suited to modes that provide storage or network access.

Whether the proxy can see the source storage directly determines whether Direct Storage Access is feasible. This transport approach requires the backup proxy to have appropriate direct access to the datastore hosting the VM disks, such as through Fibre Channel, iSCSI, or NFS. Veeam can then retrieve VM data directly from storage rather than relying on the ESXi management network path.

Backup window and performance requirements also guide the choice because transport mode affects throughput, infrastructure load, and scalability. Selecting a mode that aligns with the available storage connectivity and required processing rate helps meet recovery-point objectives while minimizing impact on production resources. Veeam may automatically select an available mode, but proxy placement and storage visibility must support the desired method. See Veeam VMware transport modes and Veeam virtual appliance mode.

QUESTION NO: 7

Which two practices help secure the Veeam configuration database?

A. Back up the configuration database to a separate secure location

B. Use SQL Server Transparent Data Encryption where appropriate

C. Store the database only on the same server for convenience

D. Use the default sa account for all access

E. Disable transaction logging

ANSWER: A B

Explanation:

Back up the configuration database to a separate secure location is a core Veeam security and recoverability practice. The configuration backup preserves the information required to reconstruct a Veeam Backup & Replication deployment, including protected-workload settings, jobs, infrastructure configuration, and credential-related configuration. Keeping that backup in a protected location separate from the backup server reduces the chance that a compromise, server failure, or ransomware incident prevents recovery of the Veeam environment. Veeam documents configuration backup as an essential part of restoring a backup server configuration: Veeam Configuration Backup and Restore.

Use SQL Server Transparent Data Encryption where appropriate provides an additional at-rest protection layer when the Veeam configuration database uses a supported Microsoft SQL Server deployment. Transparent Data Encryption encrypts database data files, transaction logs, and database backups, helping protect the database if underlying files or storage media are accessed outside of authorized SQL Server use. Its use should align with the organization's SQL Server architecture, licensing, key-management procedures, and security requirements. Microsoft describes the scope and operational behavior of this protection in its Transparent Data Encryption documentation.

QUESTION NO: 8

Which three methods can Kasten K10 use to capture application data?

A. CSI volume snapshots

B. Filesystem backup with Kopia

C. Native database dumps

D. Hypervisor snapshots

E. Traditional file-copy agents in every pod

ANSWER: A B C

Explanation:

Kasten K10 can protect Kubernetes application data using CSI volume snapshots, filesystem backup with Kopia, and native database dumps. CSI volume snapshots are used when the cluster's CSI storage driver supports Kubernetes VolumeSnapshot functionality. K10 orchestrates the snapshot workflow and captures the Kubernetes application metadata required to restore the protected application and its persistent data. Filesystem backup with Kopia provides file-level data capture for persistent volumes, including environments where snapshots are unavailable, unsuitable, or where a portable backup copy is required. K10 uses this mechanism to move persistent-volume data to the configured external backup location. Native database dumps are an application-aware capture method. K10 can invoke Kanister-based data operations to quiesce an application and create a logical export, such as a database dump, so the captured data is application-consistent rather than solely a storage-level image. These methods can be selected and orchestrated through K10 policies according to the application, storage integration, and recovery requirements. Further details on policy-based protection and data-capture operations are available in the Kasten K10 application protection documentation and the Kasten K10 Kanister integration documentation.

QUESTION NO: 9

What does the RPO Monitor in Veeam ONE show?

A. Whether protected workloads meet configured recovery point objectives

B. Backup proxy performance ratio

C. RPO values for unprotected VMs only

D. Repository free space only

ANSWER: A

Explanation:

Whether protected workloads meet configured recovery point objectives is correct. The RPO Monitor is a Veeam ONE monitoring view intended to make protection compliance immediately visible. It evaluates the latest available restore point for protected workloads against the recovery point objective configured for the relevant backup job or policy. This enables an administrator to determine whether the organization can recover data to a point in time that remains within its accepted data-loss window. RPO is time-based: it expresses the maximum tolerable age of recoverable data, rather than a measure of job speed, storage capacity, or infrastructure utilization. By presenting RPO status for protected workloads, the monitor highlights workloads whose protection is current as well as those approaching or exceeding the defined objective. That operational visibility helps administrators prioritize investigation and remediation, such as checking a failed or delayed backup, before a protection gap becomes an SLA issue. Veeam documents the RPO Monitor as part of Veeam ONE Monitor's backup-infrastructure monitoring capabilities; see the Veeam ONE Monitor overview and the Veeam documentation portal.

QUESTION NO: 10

A single Exchange mailbox item must be recovered from a backup stored on a SOBR with older data tiered to object storage. What is the efficient approach?

A. Use Veeam Explorer for Exchange against the backup on the SOBR

B. Manually recall the entire backup first

C. Restore the full Exchange VM before opening the mailbox

D. Use Instant VM Recovery before opening Explorer

E. Restore the whole database to a staging server first

ANSWER: A

Explanation:

Use Veeam Explorer for Exchange against the backup on the SOBR is the efficient approach because Veeam Explorer for Microsoft Exchange provides granular recovery of Exchange data directly from an image-level backup. An administrator can open the required restore point, browse or search the Exchange mailbox content, select the individual item, and restore it to its original location, another mailbox, or export it as needed. This avoids performing a full virtual-machine recovery or a database-level recovery when the recovery objective is limited to one mailbox item.

A Scale-out Backup Repository presents its performance and capacity tiers through a unified repository view. When required backup data has been offloaded to object storage in the capacity tier, Veeam retrieves the necessary data as part of the restore operation. The administrator does not need to perform a separate, complete recall of the backup before starting the granular Exchange recovery workflow. This makes Explorer-based item recovery the appropriate method for minimizing both operational effort and transferred data while recovering the requested mailbox item. Veeam documents the Exchange item-recovery workflow in Veeam Explorer for Microsoft Exchange and capacity-tier restore behavior in Scale-out Backup Repository capacity tier.

QUESTION NO: 11

Which two authentication methods are supported for Kasten K10 dashboard access?

A. OpenID Connect

B. Local K10 users

C. Direct LDAP bind only

D. Kerberos only

E. RADIUS only

ANSWER: A B

Explanation:

Kasten K10 dashboard access supports OpenID Connect and Local K10 users. OpenID Connect enables K10 to delegate user authentication to a compatible external identity provider. This supports centralized identity management and can provide a single sign-on experience while allowing the identity provider to enforce its own authentication policies. K10 uses the identity information returned by the provider together with its authorization configuration to control dashboard access.

Local K10 users provide an authentication method managed directly within K10. This is useful when an organization does not use an external identity provider for the cluster, and it can also support administrative or recovery access scenarios where independently managed K10 credentials are required. These two methods give administrators a choice between externally federated identity and K10-managed user access for the dashboard. K10 access configuration, including authentication and authorization concepts, is documented in the Kasten K10 authentication documentation and the Kasten K10 access-control documentation.

QUESTION NO: 12

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

Explanation:

Backup success or failure summary is available through the native reporting views in Veeam Backup Enterprise Manager. Enterprise Manager centrally collects information from connected Veeam Backup & Replication servers and presents backup-job results, including job status and session history. This enables an administrator to identify successful, warning, and failed backup activity across the managed backup infrastructure without having to review each backup server separately.

Storage capacity usage is also a native Enterprise Manager reporting capability. Its repository-oriented reporting provides a centralized view of backup repository capacity, including the space consumed by backup data and available repository capacity. This information supports day-to-day operational monitoring and helps administrators recognize capacity consumption trends across the environment.

These report areas reflect Enterprise Manager's purpose as a centralized web portal for visibility into backup operations and backup-storage resources. Veeam documents the Enterprise Manager reporting interface and its report types in the Veeam Backup Enterprise Manager Reports documentation. The product's role in centrally managing and monitoring connected backup servers is also described in the Veeam Backup Enterprise Manager overview.

QUESTION NO: 13

Which two controls can prevent backup files from being deleted for 14 days even if normal retention would remove them sooner?

A. Repository immutability

B. Standard GFS retention only

C. Enterprise Manager report scheduling

D. A minimum retention period in the job or repository workflow

E. Archive tier movement only

ANSWER: A B

Explanation:

Repository immutability prevents deletion of backup data until the configured immutability period expires. When the period is set to 14 days, Veeam marks the applicable backup files as immutable and retention processing cannot remove them during those 14 days. This is an enforcement mechanism at the repository or object-storage level and is designed to protect backup data against both unintended deletion and malicious deletion attempts.

Standard GFS retention can also retain selected full backups beyond the operational retention policy. A GFS policy designates periodic backups, such as weekly backups, for longer retention. For example, retaining weekly GFS backups for two weeks keeps the selected GFS backup files available for at least 14 days even where the ordinary restore-point retention setting would otherwise age out the corresponding data sooner. GFS retention is therefore a policy-based mechanism for maintaining longer-term recovery points, while immutability provides a deletion lock for the configured interval. Veeam documents immutable backup protection at Hardened Repository and the retention of GFS backups at GFS Retention Policy.

QUESTION NO: 14

Which standalone Veeam Agent for Windows editions are available outside VBR management?

A. Free

B. Workstation

C. Server

D. Enterprise

E. Cloud

ANSWER: A B C

Explanation:

The standalone editions available for Veeam Agent for Microsoft Windows are Free, Workstation, and Server. Standalone deployment means that the agent is installed, configured, and operated locally on the Windows computer, rather than being deployed and policy-managed by Veeam Backup & Replication (VBR). The Free edition supports essential local endpoint-protection use cases, including scheduled backups and recovery functions. Workstation is the licensed edition designed for desktop and laptop workloads, while Server is the licensed edition intended for Windows Server workloads and includes the capabilities needed to protect server-class systems. These editions can be used independently of a VBR-managed protection group, with backup settings and target storage configured through the Veeam Agent interface on the protected machine. Veeam's product documentation identifies Free, Workstation, and Server as the Veeam Agent for Microsoft Windows editions, and distinguishes standalone agent operation from centrally managed deployment through Veeam Backup & Replication. For the edition overview and feature context, see the Veeam Agent for Microsoft Windows editions documentation and the Veeam Agent for Microsoft Windows overview.

QUESTION NO: 15

What does Veeam Kasten K10 protect?

A. Docker Hub container images only

B. Kubernetes applications and persistent volumes

C. Only VMware Tanzu management clusters

D. Only managed AKS and EKS nodes

ANSWER: B

Explanation:

Veeam Kasten K10 protects Kubernetes applications and persistent volumes. K10 is a Kubernetes-native data management platform that captures both the application's Kubernetes resources and the persistent data associated with it. This application-centric protection includes relevant objects such as namespaces, deployments, StatefulSets, services, ConfigMaps, Secrets, custom resources, and PersistentVolumeClaims, enabling an application to be recovered with its required configuration and data relationship intact. Protection is policy-driven, allowing administrators to define schedules, retention, snapshot operations, and exports to external object storage. K10 can then support backup and restore operations, disaster-recovery workflows, and application mobility between Kubernetes environments. Persistent volumes are protected through storage-integrated snapshots where supported, or through data movement mechanisms when snapshots alone are not sufficient. This approach is particularly important for stateful workloads, because protecting only container images or only cluster infrastructure would not provide a complete, application-consistent recovery point. Veeam describes Kasten K10 as a platform for Kubernetes backup, disaster recovery, and mobility, designed to protect cloud-native applications and their data. See the official Veeam Kasten K10 documentation and Veeam Kubernetes data protection.

QUESTION NO: 16

Which two K10 features improve consistency for stateful application backups?

A. Pre-backup and post-backup hooks for custom commands

B. Native integrations with database tools such as pg_dump or mysqldump

C. CSI snapshots alone as guaranteed application consistency

D. Automatic fsfreeze for every pod

E. Pausing all containers during backup

ANSWER: A B

Explanation:

Pre-backup and post-backup hooks for custom commands improve backup consistency by letting Kasten K10 coordinate data-protection activity with the application. Hooks can execute application-specific commands at defined points in a backup workflow, such as flushing buffered data, quiescing writes, placing a database into a suitable backup state, and resuming normal activity afterward. This allows protection workflows to account for the state maintained above the persistent-volume layer.

Native integrations with database tools such as pg_dump or mysqldump also improve consistency because they use database-aware mechanisms to produce logical backups. Kasten K10 uses Kanister blueprints to define application-aware operations, including execution of database-native utilities and other custom actions. These workflows can capture data in a form that is consistent with the database engine's transaction and export semantics, while K10 manages the broader Kubernetes backup policy and lifecycle. Together, hooks and database-aware integrations provide the control needed to protect stateful workloads whose consistency requirements extend beyond storage snapshots. See the Kasten K10 documentation for action hooks and Kanister application integration.

QUESTION NO: 17

Which three object platforms support S3 Object Lock style immutability?

A. AWS S3

B. Azure Blob Storage using native archive only

C. Wasabi

D. Google Cloud Filestore

E. IBM Cloud Object Storage

ANSWER: A C E

Explanation:

AWS S3, Wasabi, and IBM Cloud Object Storage support the S3 Object Lock-style immutability model required for immutable backup data in an S3-compatible object storage repository. Object Lock applies a retention setting to stored objects, preventing their deletion or alteration until the configured retention period expires. This provides a storage-enforced protection layer for backup chains, including against accidental deletion and many ransomware-driven deletion attempts. AWS S3 supplies the native S3 Object Lock implementation, with retention modes designed to preserve objects for a defined period. Wasabi provides S3-compatible Object Lock functionality and is supported by Veeam for immutable object storage use cases. IBM Cloud Object Storage also provides S3 Object Lock support, allowing Veeam to use the same object-level retention approach. In Veeam deployments, the bucket must be prepared and configured according to the provider and Veeam requirements before immutability can be enabled; Veeam then writes backup objects with the appropriate retention information. See Veeam's Object Storage Repository documentation and AWS's Amazon S3 Object Lock documentation for the underlying retention behavior and configuration context.

QUESTION NO: 18

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:

Changing passwords for accounts that existed on the VM is a critical step before returning a clean-restored machine to production. A restore returns the virtual machine's disks and configuration to the selected recovery point, but it does not inherently revoke credentials that could have been compromised before the incident. An attacker may have obtained local administrator passwords, application credentials, service-account secrets, SSH keys, or other credentials available on the affected system. Rotating these credentials reduces the opportunity for unauthorized reuse of previously exposed access and helps establish a trusted post-incident state.

This action should be performed as part of a controlled recovery process, alongside validating the restored workload and confirming it is appropriate to reconnect. Veeam Secure Restore supports scanning a restored VM with antivirus software before the VM is returned to service, helping organizations validate recovery data safely. Veeam documents the Secure Restore workflow at Veeam Secure Restore. Broader ransomware recovery guidance also emphasizes resetting compromised credentials and securing accounts during recovery; see the CISA StopRansomware Guide. In practice, prioritize privileged, service, and application accounts, while coordinating changes that may affect dependent services.

QUESTION NO: 19

Which source types can be used for a Backup Copy Job?

A. Existing backup jobs or repositories

B. Direct vCenter or ESXi source objects

C. Managed Veeam Agent backup jobs

D. Replication jobs

E. NAS share definitions only

ANSWER: A C

Explanation:

Existing backup jobs or repositories can be used because a Backup Copy Job is designed to create a secondary copy from backup data that Veeam Backup & Replication has already created. During source selection, administrators can select protected workloads from an existing backup job or select existing backups stored in a backup repository. This enables independent retention, secondary storage, and offsite copy policies without rereading the production workload.

Managed Veeam Agent backup jobs can also be used because Veeam Backup & Replication manages the backup chains produced by supported Veeam Agent policies and jobs. Their backup data is therefore available as a source for a Backup Copy Job, allowing physical-machine and agent-protected workloads to receive the same secondary-copy and retention treatment as other protected workloads. The copy job processes available restore points from the source backup chain and writes the copied backup to its target repository according to the configured schedule and retention settings. Veeam documents source selection for backup copy jobs, including selecting backups from jobs and repositories, in the Backup Copy Job Source Documentation. The general job behavior and purpose are described in the Veeam Backup Copy Job Guide.

QUESTION NO: 20

What is a main advantage of K10 filesystem backup over volume snapshot backup?

A. It is portable across different storage backends

B. It always creates backups faster

C. It always consumes less storage

D. It is automatically more application-consistent

E. It depends on vSphere integration

ANSWER: A

Explanation:

It is portable across different storage backends is correct because a Kasten K10 filesystem backup captures persistent-volume data through a file-level data-mover workflow instead of making recovery dependent on a storage-system or CSI volume snapshot. The protected data can therefore be exported to an external location and restored where the destination cluster uses a different compatible storage class or underlying storage technology. This is particularly valuable for cross-cluster migration and disaster recovery, where the source and target environments may not have the same CSI driver, snapshot capabilities, or storage array.

K10 filesystem backups still protect the Kubernetes application objects required for recovery alongside the data workflow, enabling the application to be recreated and its data restored to newly provisioned persistent volumes on the target. The portability benefit does not require the original storage backend to be present at recovery time; instead, the target environment provides suitable storage for the restored PVCs. Veeam Kasten documentation describes filesystem-based protection as a method for data portability when snapshot-based operations are unavailable or unsuitable, and documents

application export and restore workflows across Kubernetes environments. See Kasten K10 application protection documentation and Kasten K10 restore documentation.

QUESTION NO: 21

What is the best RPO class available with Veeam continuous data protection?

A. One hour

B. Fifteen minutes

C. Near-zero seconds with CDP

D. Five minutes only

ANSWER: C

Explanation:

Near-zero seconds with CDP is the best available RPO class because Veeam Continuous Data Protection is specifically designed to protect critical VMware workloads with recovery-point objectives measured in seconds. Rather than waiting for a scheduled backup or replication job to run, CDP continuously captures virtual-machine write operations through the VMware vSphere I/O-filtering framework and sends the resulting changes to the replica infrastructure. This creates frequent, very closely spaced restore points and minimizes the amount of data that can be lost following an outage.

“Near-zero” correctly expresses the protection class rather than promising an absolute zero-data-loss outcome. The practical RPO attained depends on the source workload’s rate of change, available network throughput, replica and journal performance, and the selected CDP policy. Even with those operational dependencies, CDP is Veeam’s mechanism for the most demanding RPO requirements, providing protection in seconds rather than the minutes or hours normally associated with periodic jobs. Veeam documents CDP replication as continuous replication for workloads that need minimal data loss and describes configurable short RPO intervals in its CDP policy settings. See Veeam CDP replication documentation and Veeam CDP replication policy documentation.

QUESTION NO: 22

Which built-in area helps identify how much repository space each protected VM consumes?

A. Veeam ONE repository forecast only

B. Repository storage usage reporting in the Backup & Replication console

C. vSphere storage views only

D. Configuration backup log review

ANSWER: B

Explanation:

Repository storage usage reporting in the Backup & Replication console is the appropriate built-in area for identifying the repository capacity consumed by protected virtual machines. Veeam Backup & Replication maintains backup inventory and storage-consumption information for backup chains stored in repositories. Administrators can use the console’s backup and repository views to inspect stored backup data, including the size associated with individual VM backup data and the overall space used on the target repository. This makes the information operationally useful for determining which workloads contribute most to repository growth, validating expected backup sizes, and planning additional capacity before free space becomes constrained.

The displayed consumption should be interpreted in the context of the backup chain and storage-efficiency technologies in use. For example, retention, full and incremental backup files, compression, deduplication, synthetic operations, and fast-clone capable storage can affect how logical VM backup size relates to physical repository consumption. Nevertheless, the Backup & Replication console is the native management interface for reviewing backup data and repository utilization, so it

provides the relevant visibility without requiring a separate monitoring product. Veeam documents repository management in the Backup Repository section and backup inventory management in the Backups View documentation.

QUESTION NO: 23

What is the upper immutability period that can be configured on a hardened Linux repository in this exam context?

A. 1 day

B. 7 days

C. 30 days

D. 3650 days

E. Unlimited

F. 9999 days

ANSWER: F

Explanation:

The upper configurable immutability period for a Veeam hardened Linux repository is **9999 days**. When configuring the repository, Veeam Backup & Replication allows the immutability period to be set from 1 through 9999 days. This period controls how long backup files are protected against deletion or modification after they are created. During that interval, Veeam applies Linux filesystem immutability protections to the relevant backup files, helping preserve a recoverable copy even if an attacker obtains elevated access to the backup server or attempts ransomware-driven deletion.

Although organizations should select the duration according to their retention, recovery, regulatory, and storage requirements, the product limit is the relevant point here: 9999 days is the maximum value accepted for this hardened-repository immutability setting. It is a finite duration, rather than an unlimited setting, and is approximately 27 years. Veeam documents the hardened repository configuration and its immutability controls in the Veeam Backup & Replication User Guide: Hardened Repository. Additional implementation and security context is available in Veeam's Hardened Linux Repository Best Practices.

QUESTION NO: 24

A forever forward incremental chain keeps 14 restore points. What file pattern should be expected?

A. Fourteen independent VBK full backups

B. One VBK full backup and up to thirteen VIB increments

C. One VIB file only

D. A VRB file for every restore point

ANSWER: B

Explanation:

One VBK full backup and up to thirteen VIB increments is correct. A forever forward incremental backup begins with one active full backup file, using the `.vbk` format. Each later backup run creates an incremental file in `.vib` format. With retention configured for 14 restore points, the chain represents those points through the full backup plus the associated incremental files, yielding one VBK and as many as 13 VIBs.

After the retention limit is reached, Veeam applies forever forward retention processing instead of producing periodic full backups. It merges data from the oldest incremental restore point into the VBK, effectively moving the full backup forward in time, and removes the obsolete incremental file. This preserves the configured restore-point count while maintaining a single forward incremental chain. The result is an efficient disk-based layout that normally retains one VBK as the chain's base and

incremental VIB files for the newer restore points. Veeam documents the full-and-incremental file relationship and forever forward transformation behavior in its Backup chain documentation and Forever forward incremental retention documentation.

QUESTION NO: 25

Which resources can a Cloud Connect provider offer tenants?

A. Backup repository capacity

B. Replication target resources

C. WAN acceleration resources

D. Tape library access

E. Direct production storage access

ANSWER: A B C

Explanation:

Backup repository capacity is a Veeam Cloud Connect service resource because a service provider can create a cloud repository, assign storage quotas, and expose that repository to a tenant for off-site backup storage. The tenant then uses the provider-managed repository as a cloud target in supported backup-copy and backup workflows. Veeam documents the provider cloud repository as the central resource used to deliver Cloud Connect Backup services.

Replication target resources are also offered through Veeam Cloud Connect Replication. A provider configures cloud-host infrastructure and its associated compute, storage, and networking resources, then allocates those resources to a tenant. This enables the tenant to create and operate replicas at the provider site for disaster-recovery purposes without directly administering the provider's underlying infrastructure.

WAN acceleration resources can likewise be made available by a provider. A provider can deploy a cloud WAN accelerator and associate it with tenant cloud infrastructure. When used with a tenant-side WAN accelerator, it reduces the amount of backup data transferred across the WAN by applying Veeam WAN-acceleration optimization techniques. This is particularly useful where the connection to the service-provider site has limited bandwidth or significant latency. See Veeam Help Center: Cloud Connect Backup Infrastructure and Veeam Help Center: Veeam Cloud Connect Replication.

QUESTION NO: 26

Which two settings address HIPAA-style encryption requirements for backup data at rest?

A. AES-256 backup file encryption with the password stored securely

B. Only TLS for communication between components

C. BitLocker encryption for repository volumes with TPM-backed keys

D. Network segmentation without encryption

E. MFA for the console only

ANSWER: A C

Explanation:

AES-256 backup file encryption with the password stored securely protects backup content at rest at the Veeam backup-file level. When encryption is enabled for a backup or backup copy job, Veeam encrypts the data written into the backup chain. The encryption password must be protected and recoverable through appropriate operational controls, because it is required to access encrypted backups. This directly supports protection of backup data should backup files be copied, exposed, or accessed outside the authorized environment.

BitLocker encryption for repository volumes with TPM-backed keys provides storage-layer encryption for the volume hosting the backup repository. This protects the repository volume's contents when physical disks are removed, lost, or accessed offline, complementing Veeam's backup-file encryption with a separate at-rest encryption layer. TPM-backed protection binds key protection to the repository system's trusted hardware and supports controlled unlocking of encrypted volumes. Using both controls is a defense-in-depth approach for safeguarding backup data at rest. See Veeam Backup Encryption and Microsoft BitLocker overview.

QUESTION NO: 27

Fifty backup jobs must be created with the same settings except for VM name. Which PowerShell pattern is most efficient?

A. Run New-VBRBackupJob in a loop with dynamic parameters

B. Clone a job manually with Copy-VBRJob for every VM

C. Edit exported XML files and import them

D. Use a non-existent bulk REST import endpoint

E. Create every job in the console

F. Run Add-VBRViBackupJob in a loop with dynamic parameters

ANSWER: F

Explanation:

Run Add-VBRViBackupJob in a loop with dynamic parameters is the most efficient pattern because it automates repeatable creation of VMware backup jobs through Veeam Backup & Replication PowerShell. A script can store the shared configuration, such as the backup repository, schedule, retention settings, and applicable processing options, then iterate over a collection of VM names or VM objects. During each iteration, the script resolves the target VM and creates one job with the common configuration plus that VM-specific value. This approach is scalable for fifty jobs, reduces manual effort, and ensures each job is created consistently. It also produces a reusable and auditable configuration artifact: administrators can validate the input VM list, apply standardized naming, add error handling, and rerun the script when further jobs are needed. Add-VBRViBackupJob is the Veeam PowerShell cmdlet intended to create VMware backup jobs, so using it in an iteration is a supported administrative workflow rather than relying on manual configuration or unsupported file manipulation. The cmdlet syntax and supported parameters are documented in the Veeam PowerShell Reference for Add-VBRViBackupJob; broader scripting guidance is available in the Veeam PowerShell Reference overview.

QUESTION NO: 28

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:

Placing repositories in the same country as the protected data is a direct data-residency control: Veeam backup repositories are the storage locations to which backup files are written, so their physical location must be selected within the required jurisdiction. Using object buckets locked to the approved region similarly controls the location of backup data in object storage. When configuring an object storage repository, the selected cloud provider bucket and regional endpoint determine

where the provider stores the objects; choosing an approved in-country region supports the requirement that backup copies remain there. Keeping customer-managed encryption keys under on-premises control adds an important sovereignty and governance safeguard. Veeam supports backup encryption, and controlling the encryption-password or key-management process within the organization helps ensure that access to encrypted backup data remains governed by the organization and its applicable jurisdiction. These controls should be implemented together with documented provider residency commitments, because storage placement is the primary mechanism for meeting an in-country storage requirement, while encryption-key control strengthens control over access. See the Veeam Backup Repository Guide and Veeam Object Storage Repository Guide.

QUESTION NO: 29

What are benefits of replica mapping to an existing target VM?

A. Avoid a new full replica seed

B. Reduce initial replication time

C. Encrypt traffic automatically

D. Compress data automatically

E. Delete the source VM

ANSWER: A B

Explanation:

Replica mapping enables Veeam Backup & Replication to associate a source VM with an existing VM at the target location and use the existing VM disks as the baseline for replication. Consequently, it can avoid creating a new full replica from scratch. Instead of transferring every block required to construct a completely new target VM, Veeam analyzes the mapped VM and synchronizes the differences needed to make it a usable replica of the source workload. This is valuable when a replica already exists but a replication job must be recreated or reconfigured, or when an existing target VM is suitable for reuse.

Using the existing target-side disks also reduces initial replication time. Because much of the required VM data may already be present at the target, the first job run need transfer only changed or missing blocks rather than the entire virtual machine. This can substantially reduce network consumption, particularly across slow or bandwidth-constrained links, and allows replication to begin with a pre-populated target VM. Veeam documents that replica mapping lets a job use an existing target VM as its replica and that it synchronizes the mapped VM with the source VM. See Veeam Help Center: Replica Mapping and Veeam Help Center: Replica Seeding.

QUESTION NO: 30

A tenant needs Cloud Connect backup traffic reduced during business hours. Which setting controls scheduled bandwidth use?

A. Use WAN Accelerator

B. Network Traffic Rules with throttling schedules

C. Synthetic full on copy job

D. Direct Storage Access

E. Backup from storage snapshots

ANSWER: B

Explanation:

Network Traffic Rules with throttling schedules is the appropriate setting because it provides Veeam Backup & Replication's built-in, time-aware mechanism for limiting bandwidth consumed by protected-data transfer. A network traffic rule identifies the relevant source and target network ranges and applies a throttling policy to traffic moving between them. The throttling configuration can include a schedule, allowing the tenant to impose a lower transfer limit during defined business-hour periods and use more available bandwidth outside those periods. This directly addresses the requirement to protect production network capacity while Cloud Connect backup traffic is active, without needing to change the backup job's retention, processing mode, or storage-access method. In a Cloud Connect deployment, the rule can be scoped to the network path used for the tenant's backup data transfer, ensuring the intended traffic is governed by the scheduled limit. Veeam documents network traffic rules as the feature used to manage data-transfer bandwidth between backup infrastructure components and to configure throttling schedules. See Veeam Network Traffic Rules and Veeam Traffic Throttling.

QUESTION NO: 31

Which configuration best aligns Veeam infrastructure with Zero Trust networking?

A. Allow all Veeam traffic anywhere inside the corporate network

B. Permit only required ports between specific component IPs

C. Use one domain-admin account for every Veeam task

D. Disable encryption between internal components

E. Store credentials in a shared text file

ANSWER: B

Explanation:

Permit only required ports between specific component IPs is the configuration that best aligns with Zero Trust networking. Zero Trust applies the principles of explicit verification, least-privilege access, and assumed breach: network location alone does not establish trust. For Veeam Backup & Replication, this means defining firewall policies that allow communications only between the particular Veeam components involved—such as the backup server, proxies, repositories, mount servers, protected workloads, and guest processing endpoints—and only for the ports needed for their intended operations. These policies should be tailored to the deployed architecture and features, including the transport mode, repository type, guest processing, recovery workflows, and management functions. Precise source and destination restrictions reduce unnecessary exposure and help limit lateral movement if a system is compromised. They also support meaningful segmentation between management, production, backup, and repository networks without preventing required backup and restore operations. Veeam publishes its communication-port requirements so administrators can create narrowly scoped rules rather than relying on broad internal-network access. Refer to the Veeam Backup & Replication used ports documentation and the Veeam Backup & Replication Security Best Practice Guide.

QUESTION NO: 32

If a VM fails due to a locked snapshot and retry is enabled, what does Veeam do?

A. Skip the VM as success

B. Retry the operation according to configured retry settings

C. Fail immediately

D. Shut down the VM

E. Create a replacement snapshot without retry

ANSWER: B

Explanation:

Retry the operation according to configured retry settings is correct. Veeam Backup & Replication supports automatic retries for failed VM processing when the backup job's retry setting is enabled. A snapshot lock is often temporary: for example, it may result from an in-progress vSphere operation, a transient infrastructure condition, or another process that has not yet released access to the snapshot. Rather than requiring an administrator to intervene after the first unsuccessful attempt, Veeam waits for the configured retry interval and then attempts to process the affected VM again, up to the configured number of retries. This can allow the backup to complete successfully once the lock is released. Retry behavior is configured in the job's schedule settings and applies to failed processing attempts, helping jobs tolerate temporary conditions without treating the initial failure as final. The retry count and delay should be selected with the environment's operational timing in mind, particularly where snapshot consolidation or concurrent VMware operations can temporarily hold snapshot resources. See Veeam's backup job scheduling documentation and the VMware backup job documentation for job configuration and processing details.

QUESTION NO: 33

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 establishes a controlled recovery workflow before a potentially compromised workload can interact with production systems. A restore point from before the known ransomware event may still contain dormant malware, persistence mechanisms, or an earlier compromise, so its integrity must be verified rather than assumed. Veeam SureBackup supports this process by powering on recovered workloads in an isolated virtual lab, where administrators can perform automated boot, heartbeat, and application checks without exposing the production environment. Veeam Secure Restore can also invoke configured antivirus or antimalware scanning during restoration, helping identify threats in the recovered VM disks before the VM is returned to service. After malware scanning and functional validation confirm that the selected recovery point is clean and the required applications operate correctly, the organization can safely plan the cutover or promotion to production. This approach supports both containment and confidence in recovery, which are essential objectives in a ransomware incident. See Veeam SureBackup recovery verification and Veeam Secure Restore.

QUESTION NO: 34

Which two cases benefit most from tape encryption?

A. Tapes are stored in an offsite location with uncertain physical security

B. Regulations require encryption for data at rest

C. The goal is to improve tape write speed

D. The goal is to reduce the number of cartridges

E. The goal is to speed up restore by skipping catalog lookup

ANSWER: A B

Explanation:

Tape encryption provides its greatest value when the confidentiality of backup data must remain protected despite a tape's portability and long retention period. Tapes stored in an offsite location with uncertain physical security benefit because a removable cartridge can be transported, handled by third parties, or stored outside the organization's direct control. If a tape is lost, stolen, or accessed without authorization, encryption keeps its backup content unreadable unless the appropriate encryption password and key are available. Regulations requiring encryption for data at rest are also a direct use case: organizations can apply encryption to tape archives to meet security, privacy, and compliance obligations for retained backup data. In Veeam Backup & Replication, tape encryption protects data written to tape and uses the configured password-based encryption mechanism; retaining the password is essential because encrypted data cannot be recovered without it. Encryption is therefore a data-confidentiality control for the stored media itself, particularly valuable for offsite, exported, and compliance-sensitive tapes. Veeam documents the configuration and behavior of this capability in its Tape Encryption documentation and describes password management requirements in Encryption.

QUESTION NO: 35

Which two limitations are common when a Veeam Backup Server is deployed in workgroup mode instead of being domain joined?

A. Guest processing with Windows credentials is more limited

B. Kerberos authentication cannot be used between components

C. Hyper-V backup is impossible

D. Scale-out Backup Repositories are not supported

E. Tape libraries cannot be used

ANSWER: A B

Explanation:

Guest processing with Windows credentials is more limited is correct because a workgroup deployment lacks the centralized Active Directory identity, trust, and account-management model available to a domain-joined backup server. Veeam can still perform guest interaction by using explicitly supplied credentials, commonly local administrator accounts on the protected machines. However, administrators must manage those credentials individually and ensure that the selected accounts have the required local rights and remote connectivity. This makes Windows guest operations—including application-aware processing and guest file indexing—more operationally constrained than in a domain environment where appropriate domain accounts can be used consistently.

Kerberos authentication cannot be used between components is also correct in the normal workgroup deployment model. Kerberos is an Active Directory authentication protocol that relies on a domain Key Distribution Center, domain identities, and service principal names. A Veeam Backup Server that is not joined to the domain does not participate in that domain trust relationship as a computer account, so Veeam infrastructure communications cannot rely on integrated Kerberos authentication in the same manner as a domain-joined deployment. Administrators instead use supported explicit credentials and non-Kerberos authentication paths. Veeam describes the credential requirements for guest interaction in its Guest Processing documentation; Microsoft explains the domain-based design of Kerberos authentication.

QUESTION NO: 36

What is a Virtual Lab used for in SureBackup?

A. An isolated environment for starting backups safely during automated verification

B. A test area for agent installation only

C. A dedicated proxy host cluster

D. A public cloud replication environment

ANSWER: A

Explanation:

An isolated environment for starting backups safely during automated verification is correct because a Veeam Virtual Lab provides the protected, fenced network in which SureBackup performs automated recovery verification. When a SureBackup job runs, Veeam starts virtual machines directly from their backup files using vPower NFS functionality, rather than requiring a full production recovery first. The Virtual Lab isolates these started VMs from the production network while allowing them to communicate with required application dependencies inside the lab. It can also provide masquerading, enabling controlled access from the production network for verification without creating IP-address conflicts or unintended production interaction. SureBackup uses this environment to run checks such as VM heartbeat, ping, application-level tests, and custom scripts. This verifies that backed-up workloads can boot and that their services are available, providing evidence that the backups are recoverable before a real incident occurs. Veeam defines a Virtual Lab as an isolated environment for testing VMs and identifies it as a required component for SureBackup recovery verification. See the Veeam Help Center: Virtual Labs and SureBackup recovery verification.

QUESTION NO: 37

In a Scale-Out Backup Repository, what is the Capacity Tier used for?

A. A proxy tier used only for transport processing

B. Object storage used for copying or moving backup data from performance extents

C. A tape library connected to the SOBR

D. A SQL Server database used for backup metadata

ANSWER: B

Explanation:

Object storage used for copying or moving backup data from performance extents is correct. In a Veeam Scale-Out Backup Repository, the Capacity Tier extends the Performance Tier with supported object storage, including Amazon S3 and S3-compatible storage, Azure Blob Storage, and Google Cloud Storage. It provides scalable, offsite capacity for backup data while allowing the Performance Tier to retain the most recent restore points on faster primary repository storage. Veeam can copy newly created backup files to the Capacity Tier immediately for an additional resilient copy. It can also move backup chains that have aged beyond the configured operational restore window from performance extents to object storage. This enables long-term retention and capacity growth without requiring all backup data to remain on the local performance extents. If required for recovery, Veeam can access backup data in object storage, including by retrieving it to the Performance Tier according to the configured workflow. The Capacity Tier is therefore the object-storage layer of an SOBR, designed for economical scale, offsite protection, and lifecycle-based data placement. See the Veeam Capacity Tier documentation and the Scale-Out Backup Repository documentation.

QUESTION NO: 38

What does the Veeam tape catalog provide?

A. A mapping of backup files to the tapes that contain them

B. A storage location for tape encryption keys

C. A firmware update process for tape drives

D. A schedule engine for tape export jobs

E. A compression engine before writing to tape

ANSWER: A

Explanation:

A mapping of backup files to the tapes that contain them is correct. The tape catalog is Veeam Backup & Replication metadata that records the contents of tape media and associates protected backup data with the particular tape on which it resides. This inventory enables the product to determine which tape must be loaded when an administrator needs to restore a backup or files from tape. It is particularly useful for media that has been exported from a tape library, moved to an offsite vault, and later returned or imported: the catalog information lets Veeam identify the required media without relying on an administrator's manual tracking of tape contents.

Cataloging reads the tape's stored metadata and makes its contents available in the Veeam console for restore operations. The resulting catalog information supports efficient tape restore workflows by identifying the relevant media set and backup chain data before Veeam begins reading the tape. This function is therefore an index and relationship map between backup files and physical tape media, rather than a data-processing or hardware-management capability. Veeam documents tape cataloging and its role in discovering tape contents in the Tape Catalog section of the User Guide. Related restore behavior is described in Restoring Files from Tape.

QUESTION NO: 39

Which actions are available after testing a VM replica with a planned failover?

A. Finalize the failover

B. Undo the failover

C. Convert the replica into a tape backup

D. Run a repository rescan as the only completion action

E. Automatically delete the source VM permanently

ANSWER: A B

Explanation:

Finalize the failover and **undo the failover** are correct. After a planned failover, Veeam Backup & Replication maintains the failover state so the administrator can verify that workloads operate correctly at the target site. If the target environment is accepted as the new production location, finalizing commits the operation.

If validation identifies a problem or the operation should not be retained, undoing the failover returns the replica and source VM to their pre-failover state. These options provide controlled testing and decision-making rather than forcing the administrator to accept the target-side VM immediately. See Veeam Planned Failover documentation.

QUESTION NO: 40

Which three object storage platforms are common Veeam capacity tier targets?

A. Amazon S3

B. Microsoft Azure Blob Storage

C. Wasabi Hot Cloud Storage

D. Google Cloud Filestore

E. A basic SMB file share

ANSWER: A B C

Explanation:

Amazon S3, Microsoft Azure Blob Storage, and Wasabi Hot Cloud Storage are common object-storage targets for the Capacity Tier of a Veeam Scale-out Backup Repository. The Capacity Tier extends performance-tier backup storage into object storage, supporting policies to copy backups as soon as they are created and to move older backup data out of the

performance tier. This provides scalable off-site retention and can support immutability when it is configured on a compatible target. Amazon S3 is a natively supported object-storage service and is frequently selected for Capacity Tier deployments. Microsoft Azure Blob Storage is Veeam's supported Azure object-storage target, including appropriate Azure blob account configurations. Wasabi Hot Cloud Storage provides S3-compatible object storage, which makes it a commonly deployed Capacity Tier target where its supported S3 API and immutability capabilities meet the organization's requirements. Veeam documents the Capacity Tier as an object-storage extent and lists Amazon S3-compatible storage, Microsoft Azure Blob Storage, and other cloud object-storage services among supported repository types. See the Veeam Capacity Tier documentation and the Veeam object storage repository documentation.

QUESTION NO: 41

How is Veeam Service Provider Console different from Enterprise Manager?

A. VSPC is for service providers managing tenants; Enterprise Manager is for internal enterprise self-service

B. VSPC is only monitoring and EM is only backup configuration

C. VSPC is a mobile app

D. VSPC fully replaces EM

E. VSPC is only Cloud Connect and EM only on-premises

ANSWER: A

Explanation:

VSPC is for service providers managing tenants; Enterprise Manager is for internal enterprise self-service. Veeam Service Provider Console is built for service providers that operate and deliver data-protection services to multiple customer organizations. It supplies a centralized web console for remotely monitoring and managing protected workloads, backup infrastructure, alarms, and reporting across customer environments. Its design supports a provider operating model, including tenant-oriented management, role-based access, and metering or usage reporting for delivered services.

Veeam Backup Enterprise Manager has a different focus: it centrally exposes and administers Veeam Backup & Replication capabilities within an organization's own environment. Its web portal enables centralized visibility, reporting, search, delegation, and self-service restore workflows for authorized internal users. Therefore, the key distinction is not simply where either product is installed; it is whether the console is being used to deliver and manage services across multiple tenants or to provide centralized internal administration and self-service for one enterprise. Veeam documents VSPC as a platform for service providers and describes Enterprise Manager as a web-based management and self-service component for Veeam Backup & Replication deployments. See Veeam Service Provider Console and Veeam Backup Enterprise Manager documentation.

QUESTION NO: 42

Which three approaches can support GDPR erasure needs while preserving backup-chain integrity?

A. Restore specific files or data and selectively remove personal data where applicable

B. Use customer-managed encryption keys that can be destroyed when appropriate

C. Make every backup immutable forever

D. Use application-aware workflows or scripts to mask or pseudonymize PII before backup

E. Use SureBackup as proof that personal data was anonymized

ANSWER: A B D

Explanation:

“Restore specific files or data and selectively remove personal data where applicable” supports a controlled operational response to an erasure request without attempting to edit existing backup files. The required data can be recovered to an isolated or appropriate target, remediated in the application or production data set, and then protected through normal backup and retention operations. This preserves the consistency of the existing backup chain while ensuring subsequent protected versions reflect the corrected data state.

“Use customer-managed encryption keys that can be destroyed when appropriate” can support crypto-erasure where the organization controls the applicable encryption-key lifecycle and the approach is permitted by its legal, retention, and records-management obligations. Veeam encryption protects backup data using encryption keys, so key-management design is an important governance consideration. See Veeam Backup & Replication encryption documentation.

“Use application-aware workflows or scripts to mask or pseudonymize PII before backup” reduces the amount of directly identifiable personal data entering future restore points. Data minimization and pseudonymization help reduce the scope of future erasure handling while leaving backup processing and chain dependencies intact. GDPR Article 17 establishes the right to erasure, subject to its stated conditions and exceptions: Regulation (EU) 2016/679, Article 17.

QUESTION NO: 43

For VMware Hot-Add transport mode, what operating system can a backup proxy VM run?

A. Only Windows

B. Only Linux

C. Either Windows or Linux

D. Only ESXi Shell

ANSWER: A

Explanation:

Only Windows is correct. VMware Hot-Add, also called Virtual Appliance transport mode in Veeam documentation, requires a Windows-based backup proxy that is deployed as a virtual machine in the vSphere environment. During processing, Veeam uses VMware’s HotAdd capability to attach the virtual disks of the protected VM to this proxy VM. The proxy then reads or writes the attached disks locally rather than transferring the VM data through the management network. This arrangement can provide efficient data transport when the proxy has suitable datastore connectivity and is placed appropriately within the vSphere environment.

The operating-system requirement is important when selecting a transport mode for a proxy. Although Veeam supports Linux-based backup proxies for supported VMware backup workflows, a Linux proxy cannot be used for the Virtual Appliance/Hot-Add transport mode. Plan the proxy as a supported Windows VM, ensure it is managed by the same vCenter Server as the workloads it will process, and provide sufficient SCSI controller capacity for disks that may be attached concurrently. Veeam documents the proxy and Hot-Add requirements in its VMware backup proxy and transport modes guidance.

QUESTION NO: 44

What is Instant VM Recovery?

A. A pre-staged VM that always starts in under a minute

B. Running a VM directly from backup storage through a temporary datastore mount

C. Extracting VM files before registration

D. Automatic failover without an administrator

ANSWER: B

Explanation:

Running a VM directly from backup storage through a temporary datastore mount is correct because Instant VM Recovery lets Veeam Backup & Replication rapidly make a protected virtual machine available without waiting for a full restore to production storage. Veeam publishes the VM disks from the selected backup to the virtualization infrastructure, registers the VM, and starts it while the backup remains the source of its virtual disks. In VMware vSphere environments, Veeam uses its vPower NFS technology to present the backup contents as an NFS datastore to the ESXi host. This approach is intended to minimize recovery time objectives when a production VM has failed or is unavailable. Once services are restored, the running VM can be migrated to the intended production datastore, allowing the organization to complete recovery without a prolonged outage. Instant VM Recovery is therefore a recovery method based on operating the VM directly from its backup data first, followed by migration or restoration to permanent storage as needed. Veeam documents this workflow in its Instant VM Recovery guide and describes the temporary datastore publishing mechanism in the vPower NFS documentation.

QUESTION NO: 45

Which three items belong in a clean recovery strategy after ransomware?

A. Recover first into an isolated recovery environment

B. Scan backups with current malware signatures

C. Restore from a point before infection

D. Reuse compromised production credentials

E. Reconnect restored VMs before validation

ANSWER: A B C

Explanation:

A clean ransomware recovery strategy must ensure that restored workloads and data are trustworthy before they return to the production network. **Recover first into an isolated recovery environment** is fundamental because isolation prevents a potentially infected restored machine from communicating with production services while it is inspected and tested. Veeam SureBackup supports automated recovery verification in an isolated virtual lab, enabling organizations to validate that recovered workloads can start and operate as expected. **Scan backups with current malware signatures** is also an important clean-recovery control. Veeam Secure Restore can use an antivirus solution during recovery to scan the recovered machine's disks, helping identify known malware before the VM is released for use. Finally, **Restore from a point before infection** ensures the selected backup represents a state from before the compromise or encryption event. The recovery point should be selected using incident timelines and then validated in isolation, so the organization restores known-good application and data states rather than reintroducing the attack. These practices combine containment, malware detection, and recovery-point selection into a controlled process for returning services safely. See Veeam's documentation for Secure Restore and SureBackup recovery verification.