PassQueen
a.k.a dumpsqueen.com

Veeam Certified Engineer 2021

Veeam VMCE2021

Version Demo

Total Demo Questions: 29

Total Premium Questions: 294

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

What is the purpose of scheduling a virtual Synthesized full backup for tape jobs?

A. To be able to restore files and folders from tape

B. To be able to create the synthetic full .vbk backup directly on tape

C. To help save additional disk spate on tape

D. To avoid copying a full .vbk file to tape from the source reverse incremental backup Job

ANSWER: D

Explanation:

The purpose is to avoid copying the complete .vbk file from a source reverse-incremental backup job to tape every time a new full tape backup is needed. A reverse-incremental job maintains its most recent restore point as a full backup file, while older restore points are represented by reverse increments. If a tape job performed a conventional full backup, it would need to read and write that entire current .vbk file to tape for each full tape cycle.

With a virtual synthesized full backup, Veeam uses the backup data already available in the tape backup chain together with the required incremental changes to produce a new full restore point on tape. This lets the tape job create a full backup set without repeatedly transferring the source job's entire full backup file. The approach can reduce the load on the backup repository and the amount of data that must be read from it during recurring tape fulls, while retaining a full backup point on tape for the scheduled cycle. Veeam documents tape jobs and their full-backup processing in the Veeam Backup & Replication User Guide; see also the product's Veeam Backup & Replication User Guide PDF.

QUESTION NO: 2

A Veeam administrator has been tasked to create a backup that will automatically relocate data after a specified period of time to different media type, while also being off site,without creating additional backup copy jobs. The administrator determined that Scale-out backup Repository can ben used as part of the solution . How can this be accomplished? -

A. Add a capacity tier to the Scale-out Backup Repository and enable "GFS Tiering"

B. Add a archive tier to the Scale-out Backup Repository and enable "Copy" option

C. Add a capacity tier to the Scale-out Backup Repository and enable "Copy" option

D. Add a capacity tier to the Scale-out Backup Repository and enable "Move" option

ANSWER: D

Explanation:

Add a capacity tier to the Scale-out Backup Repository and enable "Move" option is correct because the Scale-out Backup Repository's capacity tier is designed to extend local performance-tier storage to object storage, which can be located off site. The Move policy offloads backup chains that have aged beyond the administrator-defined operational restore window from the performance tier to the capacity tier. This directly satisfies the requirement for automatic relocation after a specified period, while using a different storage medium and avoiding the need for a separate backup copy job.

Veeam manages the offload process as part of the Scale-out Backup Repository policy. After data is moved, Veeam retains the backup's usability and catalog information so that restore operations can still access the required restore points, retrieving data from object storage when necessary. The operational restore window is the key setting used to determine how long backups remain on the local performance tier before they are moved to the off-site capacity tier. Veeam documents capacity-tier offload and the Move backups policy in its Capacity Tier documentation and its Move Backups to Capacity Tier documentation.

QUESTION NO: 3

An Availability engineer has configured Veeam Backup & Replication to replicate a mission critical VM to a DR site on an hourly basis and with 24 restore points. The production site has suffered a storage failure and needs to bring up the mission critical machine as soon as possible. Which step should the Availability engineer do next?

A. Create a failover plan

B. Failover to the replica VM

C. Failback of the replica VM

D. Planned failover to me replica VM

ANSWER: B

Explanation:

Failover to the replica VM is the appropriate immediate recovery action after a production storage failure. Veeam Backup & Replication maintains the replicated virtual machine at the disaster-recovery site, including the configured replica restore points. During an unplanned outage, failover starts the selected replica restore point at the DR site so that the critical workload can resume operation without waiting for the failed production storage to be restored. The engineer can select the latest available restore point to minimize data loss, or select an earlier point if operational circumstances require it. Veeam records the failover state and provides subsequent recovery workflow actions, such as permanent failover or failback, after the production environment becomes available again. This workflow is specifically intended to rapidly restore service when the source VM or its production infrastructure is unavailable. Veeam documents failover as the operation used to switch from the original VM to its replica in the event of a disaster, and notes that the replica is powered on from the chosen restore point at the target site. See the Veeam User Guide: Failover and Veeam User Guide: Replication.

QUESTION NO: 4

Which of the following are valid export options for Microsoft SQL Server items when using the Veeam Explorer for Microsoft SQL? (Choose two.)

A. .mdb file

B. .bak file

C. .mdf file

D. .sql file

E. .xls file

ANSWER: B C

Explanation:

Veeam Explorer for Microsoft SQL Server supports exporting recovered SQL Server database content in formats that can be used for subsequent SQL Server recovery or attachment operations. A **.bak file** is a SQL Server backup file. Exporting to this format produces a database backup that can be restored through SQL Server tools, making it appropriate when the goal is to recover the database into a SQL Server instance using a conventional backup-and-restore workflow. An **.mdf file** is the primary data-file format used by SQL Server databases. Veeam Explorer can export a database as its files, enabling recovery scenarios in which the database is attached to an SQL Server instance. These export capabilities complement Veeam Explorer's item-level browsing and restore workflow by providing portable database-level recovery outputs. The applicable export format depends on the recovery objective: a backup file supports restore operations, while a database data file supports attachment-based recovery where the required database files are available. Veeam documents the available export methods and the steps for exporting SQL Server data in its Veeam Explorer for Microsoft SQL Server User Guide: Exporting SQL Server Data. Additional product documentation is available from the Veeam Help Center.

QUESTION NO: 5

What is Veeam vPower NFS Service used for?

A. It is used to send data to the original location or to a new location specified by the user when restoring It.

B. It allows Veeam Backup & Replication to present VM files in a native format to an ESXi host, from a compressed and deduplicated backup file.

C. It is used to mount a compressed and deduplicated backup file directly to the backup server and perform file and item level restore.

D. It continuously checks the configuration settings of jobs and starts them in accordance with their schedules.

ANSWER: B

Explanation:

Veeam vPower NFS Service allows Veeam Backup & Replication to publish virtual-machine disks from backup files as an NFS datastore that an ESXi host can access. This makes the VM files available in their native VMware format even though the source backup is stored in Veeam's compressed and deduplicated format. The service is a key component of Veeam's vPower technology and enables recovery workflows that use backup data without first fully restoring that data to production storage. For example, it supports Instant Recovery by letting an ESXi host run a VM directly from the published backup datastore, while Veeam later migrates the VM to appropriate production storage. It also underpins operations such as SureBackup and Virtual Lab processing, where VMs need to be started from backup data in an isolated environment. The vPower NFS Service runs on the backup server by default, though the vPower NFS role can be assigned to a suitable server when required by the deployment design. Veeam describes this mechanism in its vPower NFS Service documentation and explains its use with instant recovery in the Instant Recovery guide.

QUESTION NO: 6

At 18.00, after a power outage, a VM is stuck in an infinite boot loop. The VM is configured as follows: -

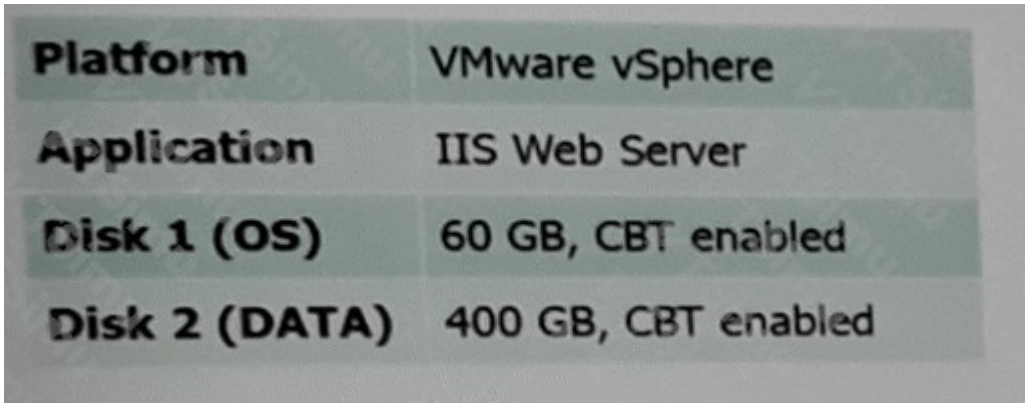

The last backup was created at 01:00. Choose the recovery approach that results in the least amount of data loss and lowest chance of data corruption.

A. Instant VM recovery with quick rollback

B. Instant VM Recovery

C. VM disk recovery with quick rollback

D. VM disk recovery

ANSWER: B

Explanation:

Instant VM Recovery is the appropriate approach because it starts a recovered VM directly from the selected 01:00 backup on the backup repository, rather than attempting to boot from or immediately overwrite the VM's potentially damaged production disks. The available restore point defines the recoverable data state, so this approach recovers the VM to the latest known good backup and limits unavoidable data loss to changes made after 01:00. It also provides a clean, isolated recovery path after the unplanned power event: the recovered VM can be validated and made available while the original failed VM and its storage remain untouched for investigation or later remediation.

Running the VM from the backup avoids depending on the corrupted boot state that caused the loop and avoids making irreversible changes to the original VM during the initial recovery. Once the VM is confirmed healthy, Veeam can migrate it back to production storage using its recovery-finalization workflow. Veeam documents Instant Recovery as a method for immediately running a VM from a backup without waiting for a full restore to production storage; see Veeam Help Center: Performing Instant Recovery and Veeam: Instant Recovery overview.

QUESTION NO: 7

When restoring Linux or Unix guest OS files, which components can be used to mount the volume containing files to be restored? (CHOSE 2) -

A. Enterprise Manager Server

B. Helper Host

C. NFS Mount

D. Guest interaction proxy

E. Helper appliance

ANSWER: B E

Explanation:

Helper Host and **Helper appliance** are the components used in Veeam's Linux/Unix file-level recovery workflow to provide access to the restored filesystem. During file-level recovery, Veeam makes the backed-up guest disks available through a recovery environment rather than requiring a full VM restore. A helper appliance is deployed and used to process and expose supported Linux or Unix filesystem volumes from image-level VM backups. This temporary appliance enables Veeam to mount the required virtual disks safely and browse their contents for copying individual files.

A helper host provides the corresponding recovery-host capability where Veeam needs a host system to mount and expose volumes for file-level restore operations. This architecture allows files and folders to be recovered directly from the backup while preserving the original backup and avoiding changes to the production guest. The selected component depends on the protected workload and recovery workflow, but both are recognized Veeam components for mounting volumes used in Linux/Unix file recovery. See Veeam's Linux File-Level Recovery documentation and restore documentation for the recovery process and infrastructure details.

QUESTION NO: 8

Which data is included in a Veeam Backup & Replication configuration backup? Choose two options.

A. Backup and replication job settings

B. Stored credentials

C. VM backup files

D. Replica virtual disks at the target site

E. VMware vSphere snapshots

ANSWER: A B

Explanation:

A configuration backup protects the Veeam Backup & Replication configuration database. This includes **backup and replication job settings**, along with the configuration of managed servers, repositories, proxies, and other backup infrastructure components. It also includes **stored credentials**, which are protected using encryption.

A configuration backup does not include the actual VM backup files stored in repositories. Those files must be protected separately through appropriate repository, Backup Copy, tape, or object-storage policies. See the Configuration Backup documentation.

QUESTION NO: 9

When you back up regular or standalone VMs in vCloud Director, which of the following is captured and stored in Veeam's Backup & Replication backup file? (Choose three.)

A. vApp metadata

B. VM disk content

C. Raw device mapping

D. In-guest iSCSI disk

E. VM metadata

F. Linked clone VM's prior to vCenter 6.5

ANSWER: A B E

Explanation:

Veeam Backup & Replication captures the VM disk content because the virtual disks contain the protected workload's operating system, applications, and data. This disk data is the primary recoverable payload in the backup and enables restores of the VM and its virtual disks.

Veeam also stores VM metadata. This includes the configuration information required to reconstruct the virtual machine appropriately during recovery, rather than restoring only its disk files. Preserving this information supports recovery workflows that recreate the protected VM with its relevant virtual hardware and configuration characteristics.

For workloads managed through VMware Cloud Director, Veeam additionally preserves vApp metadata. vApp-level information is important because Cloud Director uses vApps as logical containers and applies associated settings and relationships at that level. Capturing this metadata allows a restore to retain the Cloud Director context of the protected VM, alongside its disks and individual VM configuration.

These three data categories together provide a recoverable representation of both the VM's contents and its Cloud Director management context. Veeam documents VMware Cloud Director integration and its backup behavior in the Veeam Backup & Replication User Guide; see also Veeam's VMware backup overview.

QUESTION NO: 10

A daily backup job for seven Hyper-V VMs has been configured at the main site, keeping 14 days worth of backup files. They want to get a copy of the WM backups to a repository at the disaster recovery site. They want to keep six months' worth of the backup files at the disaster recovery site. They also need to be able to restore the VMs to any given day within two months and any given week within the six months. How should a backup copy job be configured to meet these requirements?

A. Use periodic copy (pruning) mode, keeping 180 days of retention.

B. Use immediate copy (mirroring) mode, keeping 62 days of retention and 26 weekly GFS restore points

C. Use immediate copy (mirroring) mode, keeping 180 days of retention.

D. Use periodic copy (pruning) mode, keeping 62 days of retention and six monthly GFS restore points.

ANSWER: B

Explanation:

Use immediate copy (mirroring) mode, keeping 62 days of retention and 26 weekly GFS restore points. Immediate copy mode is appropriate because it copies new source backup restore points as they become available, maintaining a current off-site copy at the disaster-recovery repository. The 62-day operational retention setting provides daily restore-point availability for approximately two months, satisfying the requirement to restore to any individual day during that period. Enabling the weekly Grandfather-Father-Son (GFS) policy and retaining 26 weekly restore points preserves one selected weekly restore point for each week over approximately six months. This combines short-term, day-level recovery granularity with longer-term, week-level recovery granularity without requiring every daily point to be retained for the entire six-month period. Veeam Backup Copy Jobs support both immediate copying and GFS archival retention, allowing the copied backup chain to meet different operational and long-term recovery objectives at the DR site. See the Veeam Backup Copy Job modes documentation and Veeam GFS retention policy documentation.

QUESTION NO: 11

For general data protection regulation (GDPR) compliance, Veeam can add a location tag to which component?

A. Scale-out Backup Repositories

B. WAN accelerator

C. File copy job

D. Proxy Server

ANSWER: A

Explanation:

Scale-out Backup Repositories is correct. Veeam Backup & Replication supports assigning a location tag to a scale-out backup repository to identify the geographic location in which its backup data is stored. This enables administrators to document and control data-placement considerations that are important for GDPR compliance, particularly where an organization must understand the storage location of protected personal data and demonstrate that data residency requirements are being observed.

A scale-out backup repository combines one or more storage extents into a single logical backup target and can include performance, capacity, and archive tiers. Applying a location tag gives the logical repository a meaningful location designation that can be used when planning and operating backup infrastructure according to organizational data-governance policies. The tag is metadata used for location awareness; it does not itself encrypt, move, or otherwise alter backup data. Administrators should still ensure the underlying storage configuration, data-transfer processes, access controls, retention settings, and encryption choices meet the organization's applicable GDPR obligations.

For Veeam's documentation on scale-out backup repository configuration and behavior, see Veeam Help Center: Scale-out Backup Repository. For GDPR's principles concerning processing and protection of personal data, see GDPR Article 5 overview.

QUESTION NO: 12

A backup administrator needs to restore two data drives to a VM where the operating system remains intact.

The target machine must remain powered on. What recovery option is most appropriate to perform this operation?

A. Instant VM Recovery

B. Instant disk recovery

C. Restore entire VM

D. Restore VM files

ANSWER: B

Explanation:

Instant disk recovery is the appropriate recovery option because it restores selected virtual disks from a VM backup to an existing virtual machine without requiring recovery of the entire VM. This directly fits a situation in which only two data drives must be recovered while the operating system disk remains intact. The recovery process lets the administrator select the necessary disks from the restore point and attach them to the target VM, preserving the existing OS and its configuration.

Veeam supports restoring the disks to the original VM or to another compatible target VM. The target VM can remain powered on during this operation, which minimizes disruption to running services. Once the disks are attached, their contents become available to the guest operating system subject to normal OS-level disk detection and any required mounting or application steps. This makes instant disk recovery especially useful when data disks are lost, corrupted, or need to be quickly made available without the downtime and scope associated with a complete virtual-machine recovery.

See the Veeam User Guide section on Instant Disk Recovery and the VM Restore overview.

QUESTION NO: 13

To suppress Veeam ONE alarms under specific conditions or when a specific event occurs, which of these options can be selected? (Choose three.)

A. High CPU usage

B. Veeam backup activity

C. Snapshot creation

D. Storage snapshot deletion

E. Storage snapshot creation

F. Snapshot deletion

ANSWER: B C F

Explanation:

Veeam ONE alarm suppression is designed to prevent expected, temporary infrastructure changes from generating unnecessary alarms. The selectable conditions include **Veeam backup activity**, **Snapshot creation**, and **Snapshot deletion**. Backup processing can temporarily affect VM performance and create expected changes in the environment, so suppressing alarms during Veeam backup activity helps administrators focus on conditions that require attention rather than planned protection operations. Snapshot creation and deletion can likewise produce short-lived changes to VM state, datastore activity, and infrastructure performance. Configuring suppression for these events avoids alert noise while those expected snapshot operations are in progress.

These settings are applied as part of Veeam ONE alarm-management configuration and allow monitoring to remain meaningful without disabling the alarm permanently. After the relevant activity or snapshot event ends, Veeam ONE resumes normal alarm evaluation, preserving visibility into genuine issues. Veeam documents alarm suppression and its available conditions in the Veeam ONE User Guide: Veeam ONE alarm suppression. Additional Veeam ONE monitoring and alarm-management guidance is available in the Veeam ONE alarms documentation.

QUESTION NO: 14

A company wants to leverage Veeam DataLabs to automate the creation of a test environment. As testing one large, multi-VM application could require more than a single host, a dedicated VMware ESXi cluster has been identified. The sandbox is to be refreshed weekly or on demand. Select the benefits of using replication jobs in this environment. (Choose two.)

A. There are additional tests available in DataLabs when replicas are the source

B. Replication jobs do not stop the sandbox, unlike backup jobs

C. The VMs are running directly on a high performance datastore, uncompressed

D. There is a bigger number of VMs supported for multi-VM application tests when replicas are used

E. Replicas can take advantage of advanced multi-host virtual labs

ANSWER: C E

Explanation:

The VMs are running directly on a high performance datastore, uncompressed is a benefit because replicas are maintained as native VMware virtual machines at the target site. A replica-based virtual lab can therefore power on the replica VMs from the replication target datastore instead of first mounting backup data and presenting it through the backup infrastructure. This is especially useful for a regularly refreshed test environment where predictable VM startup and datastore performance are important. Veeam describes replicas as VM copies stored in their native format and ready for failover or testing.

Replicas can take advantage of advanced multi-host virtual labs is also correct. Advanced multi-host virtual labs are designed for SureReplica testing and allow the isolated multi-VM application to use resources across multiple ESXi hosts in the designated cluster. This removes the single-host placement limitation that can otherwise constrain larger application test environments. The virtual lab still provides isolated networking, enabling application tests to run without affecting the production network, while the replica infrastructure supplies the VMs used in the test. Veeam documents advanced multi-host virtual labs and their use with SureReplica jobs at Veeam Help Center. General replication behavior is described in the Veeam Replication Guide.

QUESTION NO: 15

Which of the following are valid backup locations for a licensed version of the Veeam Agent forMacOSt (Choose 2)?

A. Microsoft OneDrive

B. Veeam backup repository

C. Google drive

D. Amazon S3

E. Firewire drive

ANSWER: B D E

Explanation:

Veeam Agent for Mac can protect a Mac to a Veeam Backup & Replication backup repository, allowing the agent-managed computer to send backups to repository storage managed by Veeam Backup & Replication. This is a supported destination and can include compatible deduplication-appliance repositories. A Firewire drive is also a valid target because the agent supports direct-attached storage, explicitly including USB, eSATA, and FireWire external drives. Such storage can be selected as a local/directly attached backup destination, although relying only on storage attached to the protected computer is generally less resilient than using a separate repository. Amazon S3 is valid as a cloud object-storage destination: the agent supports backup to on-premises or cloud-based object storage, and Amazon S3 is a supported object-storage platform. These destination types are documented in the Veeam Agent for Mac system requirements and backup-target guidance. See the Veeam Agent for Mac System Requirements and the Veeam Backup & Replication Backup Repository documentation.

QUESTION NO: 16

A scale-out backup repository is configured with data locality policy. How is an extent selected during an active full backup? (Choose two.)

A. Always the same extent as the previous full backup

B. An extent determined based on load control settings and free space

C. A different extent from the previous chain to avoid single point of failure, if possible on free space

D. It tries the same location as the previous full if using a deduplication appliance

ANSWER: B D

Explanation:

With the data locality placement policy, Veeam keeps files that belong to an individual backup chain together on one performance-tier extent whenever possible. An active full creates a new full backup file and therefore requires Veeam to select an extent for that new chain. The selection is made dynamically using the Scale-out Backup Repository placement logic, which considers the extent's available free capacity and its load-control status. This distributes new full-backup activity to an extent that has sufficient capacity and is suitable for processing the workload, while preserving locality for subsequent files in the chain.

When the repository uses a deduplication appliance, Veeam preferentially attempts to place an active full on the extent holding the preceding full backup. Keeping the full backups together can allow the appliance to identify and reuse already stored blocks, maximizing the appliance's deduplication efficiency. This preference is applied in addition to the normal suitability checks, including available capacity and extent availability. Veeam documents data locality as a policy intended to keep a backup chain on the same extent, and documents placement considerations for active full backups on deduplicating storage. See the Veeam Scale-out Backup Repository documentation and the Veeam SOBR configuration documentation.

QUESTION NO: 17

A backup administrator is creating a SureBackup environment up using Hyper-V as the hypervisor. Which virtual labs are available for Hyper-V? (Choose two.)

A. Advanced multi-host (automatic configuration)

B. Advanced single-host (automatic configuration)

C. Basic single-host (automatic configuration)

D. Advanced multi-host (manual configuration)

E. Advanced single-host (manual configuration)

ANSWER: C E

Explanation:

For a Hyper-V SureBackup environment, Veeam supports a basic single-host virtual lab with automatic network configuration and an advanced single-host virtual lab with manual network configuration. The basic single-host deployment is intended for a straightforward setup on one Hyper-V host; Veeam automatically creates and configures the isolated virtual network required for testing recovered workloads. This makes it suitable where the available Hyper-V networking design permits Veeam-managed configuration.

The advanced single-host deployment is also supported on Hyper-V and is used when the administrator must control the isolated-network configuration. It allows the virtual lab to be mapped to a manually prepared Hyper-V virtual switch and associated network settings, accommodating environments with specific networking, routing, or addressing requirements. In both cases, SureBackup can start VMs from backup data in an isolated environment and perform verification without affecting production workloads. Veeam documents the Hyper-V virtual-lab deployment choices and configuration workflow in its User Guide: Veeam Backup & Replication User Guide — Virtual Lab for Hyper-V. Additional SureBackup workflow details are available at Veeam Backup & Replication User Guide — SureBackup.

QUESTION NO: 18

Which actions can be performed after a VMware vSphere Instant VM Recovery has been started? Choose two options.

A. Quick Migration

B. Storage vMotion

C. Replica failback

D. Backup chain transformation

E. Storage snapshot consolidation

ANSWER: A B

Explanation:

During Instant VM Recovery, Veeam starts the recovered VM directly from the backup repository through the vPower NFS service. This provides rapid availability while the VM disks are still running from backup storage. After the VM has been validated, an administrator can migrate it to production storage using **Quick Migration** or **Storage vMotion**.

Quick Migration is a Veeam-assisted migration method that transfers the recovered VM to the selected production datastore. Storage vMotion can also be used when supported and available in the VMware vSphere environment. Both methods allow the VM to be moved from the temporary instant-recovery location to permanent production storage. See the Instant VM Recovery documentation.

QUESTION NO: 19

Veeam ONE's _____ provides a single pane of glass view for all Veeam backup servers, proxies, and repositories in a backup infrastructure.

A. Business View

B. Data Protection View

C. Infrastructure View

D. System View

ANSWER: B

Explanation:

Veeam ONE's **Data Protection View** is the centralized monitoring view intended for Veeam data-protection infrastructure. It provides an at-a-glance, single-pane-of-glass representation of the Veeam Backup & Replication environment, including backup servers, backup proxies, and backup repositories. From this view, an administrator can assess the overall operational state of the backup environment and quickly navigate to relevant dashboards, alarms, job-status information, performance charts, and events. This consolidated visibility is especially useful when an organization operates multiple backup components or servers and needs to identify capacity, performance, availability, or job-related conditions without examining each component separately. Veeam documentation describes Data Protection View as the area for monitoring the Veeam backup infrastructure and its protected workloads, including related data-protection activity. It is therefore the feature that directly matches the requirement for a unified view of backup servers, proxies, and repositories. For additional details, see the Veeam ONE Data Protection View documentation and the Veeam ONE monitoring overview.

QUESTION NO: 20

A financial company is looking for a new backup application. One of their main requirements is the ability to have quick access to their backup data. What features in Veeam Backup & Replication support the rapid availability of data, applications and VMs? (Choose three.)

A. Full VM recovery without quick rollback

B. Microsoft SQL database, and Oracle database publishing

C. Instant VM Recovery

D. Instant VM disk recovery

E. Bare metal recovery

F. Restore to Amazon EC2

ANSWER: B C D

Explanation:

Rapid availability in Veeam Backup & Replication is provided by recovery capabilities that make protected workloads or their data usable directly from backup storage, rather than requiring a full restore to complete first. **Instant VM Recovery** starts a VMware VM directly from a backup file, allowing the organization to resume service quickly while the VM is later migrated back to production storage. **Instant VM disk recovery** similarly restores access at the virtual-disk level, enabling a required disk to be presented from the backup and reducing the time to recover an affected workload or its data. Veeam documents both capabilities as mechanisms for minimizing recovery time objectives; see the Instant VM Recovery documentation and Instant Disk Recovery documentation.

Microsoft SQL database, and Oracle database publishing supports rapid application and data availability by making database content available from backups for use or validation without a conventional full restore workflow. Publishing databases is especially useful where business users or application teams need prompt access to a recovered database copy. Together, these capabilities address the stated need across VM-level recovery, virtual-disk recovery, and application/database availability.

QUESTION NO: 21

When configuring a SureBackup job, what are some of the criteria for adding servers to an application group? (Choose two.)

A. VMs cannot be encrypted

B. VM backups must reside on the same repository

C. VMs cannot be added to an application group twice

D. VMs must be running the same OS version

E. VMs must have at least one valid restore point

ANSWER: C E

Explanation:

"VMs cannot be added to an application group twice" is correct because an application group is intended to define one ordered set of machines that SureBackup starts and verifies together. Each VM is represented only once in that group, where it can be assigned an appropriate startup role, boot delay, and application-test configuration. This prevents conflicting startup definitions for the same recovered machine and keeps the SureBackup validation workflow predictable.

"VMs must have at least one valid restore point" is also correct because SureBackup validates recoverability by starting VMs from their backup data in the isolated virtual lab. Veeam therefore needs a usable restore point for every VM selected in the application group. During execution, Veeam publishes the required backup data through vPower NFS, starts the VM from that restore point, then performs the configured heartbeat, ping, and optional application tests. A valid restore point is consequently the essential recoverable source from which the SureBackup job can create and test the virtual machine.

See Veeam's documentation on SureBackup jobs and application groups.

QUESTION NO: 22

A Scale-out Backup Repository with one local extent has been configured as follows.

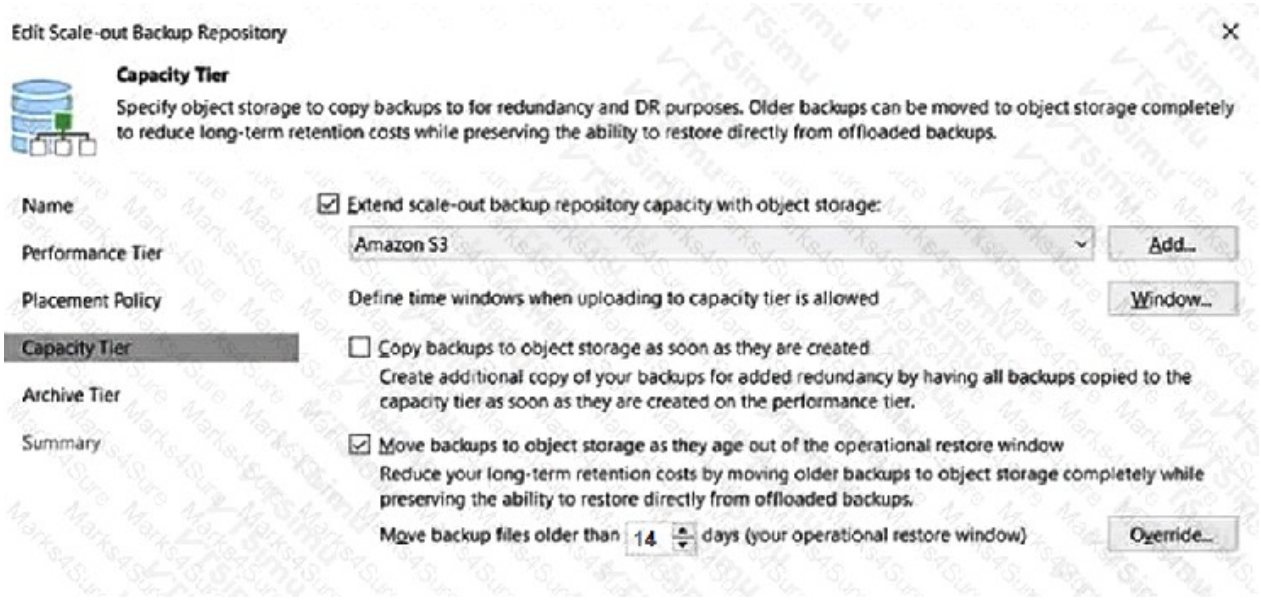

12 VMware VMs are being backed up daily by a forward incremental backup job, which has a 31-day retention, creating synthetic fulls on Saturdays. After running for more than 31 days, which of the following would be accurate?

A. Both the performance tier and capacity tier will contain all restore points.

B. The performance tier will contain all restore points. The capacity tier will contain zero restore points.

C. The performance tier will contain the most recent 15 to 22 days of restore points. The capacity tier will contain the oldest 14 days of restore points.

D. The performance tier will contain the most recent 14 to 20 days of restore points. The capacity tier will contain all restore points.

ANSWER: C

Explanation:

The performance tier will contain the most recent 15 to 22 days of restore points. The capacity tier will contain the oldest 14 days of restore points. This follows from the Scale-out Backup Repository capacity-tier move policy configured to offload backups once they are older than 14 days. Veeam does not split an active backup chain simply because an individual incremental restore point reaches the configured age. Instead, it retains the backup chain on the performance tier until that chain can be offloaded consistently, including the full backup file required for its incrementals.

Because the job creates a synthetic full every Saturday, the boundary at which an eligible chain can be moved varies according to the day on which the age threshold is reached. Consequently, the local performance tier can retain approximately 15 through 22 days of restore points rather than exactly 14 days. Restore points that age beyond the offload threshold are moved to the capacity tier, where the older retained restore points reside under the 31-day retention policy. Veeam documents that moving data to object storage is based on backup-chain handling and the configured age threshold, rather than a guarantee that precisely the same number of daily restore points remains in each tier. See Veeam Help Center: Capacity Tier and Veeam Help Center: Backup Chains.

QUESTION NO: 23

If the deduplication setting are changed for an existing job, when will the new setting be applied?

A. At the next run of the job

B. If any kind of full backup is created

C. If a new active full backup Is created

D. Only when you rescan the repository that stores the backup files

ANSWER: C

Explanation:

If a new active full backup Is created is correct. In Veeam Backup & Replication, the deduplication setting determines the block size used to identify and eliminate duplicate data within a backup job. Existing backup files were created with the block size that was in effect at the time, so changing this setting does not retroactively alter those files or immediately change the current backup chain. The new deduplication block size takes effect when the job creates an active full backup, because an active full reads all protected data from the source and generates a new full backup file. Subsequent incremental backup files in the chain then use the settings associated with that newly created active full backup. Administrators should therefore plan an active full backup after changing deduplication settings, accounting for the additional source, repository, and network load that a full backup may require. This behavior preserves chain consistency while allowing the revised data-reduction configuration to be used for newly created backup data. See Veeam's documentation on compression and deduplication settings and its active full backup description.

QUESTION NO: 24

The administrator of a VMware environment backed up by Veeam Backup &Replication has a critical server that has crashed and will not reboot. They were able to bring it back online quickly using Instant VM recovery so people continue to work. What else is required to complete the recovery?

A. Commit failback

B. Migrate to production

C. Merge delta file

D. Commit failover

ANSWER: B

Explanation:

Migrate to production is required to complete this recovery. Instant VM Recovery starts the recovered virtual machine directly from its backup files, using Veeam's vPower NFS service, so that the business can restore service quickly without waiting for the full virtual-machine disks to be restored to production storage. This is intended as an immediate operational recovery state rather than the final placement of the workload.

After users are working again, migration to production moves the VM's virtual disks from the backup repository to the selected production datastore while the VM remains available. Veeam uses VMware vMotion, Storage vMotion, or Quick Migration as appropriate for the environment. When the migration completes, Veeam finalizes the Instant VM Recovery session and the VM runs from normal production storage, rather than from the backup files. This provides the durable, fully restored production workload required after the emergency startup. Veeam documents this workflow as the migration-to-production stage of Instant VM Recovery. See Veeam Help Center: Instant VM Recovery and Veeam Help Center: Migrating Recovered VM to Production.

QUESTION NO: 25

Which of these repository types CANNOT be an extent of a scale-out backup respository? Choose two options.

A. A deduplicating storage appliance

B. A Windows server

C. A Cloud Connect cloud repository

D. A shared folder (CIFS/SMB share)

E. Another scale out repository

F. A Linux Server

ANSWER: C E

Explanation:

A Cloud Connect cloud repository cannot be used as a scale-out backup repository (SOBR) performance-tier extent. Cloud repositories represent storage resources made available through a Veeam Cloud Connect service provider and are used as an off-site backup target; they are not repositories that can be pooled into the local performance tier of a SOBR. Veeam treats the Cloud Connect repository as a distinct repository role and workflow rather than as a capacity-contributing extent.

Another scale out repository also cannot be selected as an extent. A SOBR is itself a logical repository that applies placement, performance-tier, capacity-tier, and optionally archive-tier policies across its underlying resources. Extents must be eligible underlying backup repositories, not another logical SOBR. Preventing a SOBR from being nested as an extent avoids recursive repository structures and ensures Veeam can manage backup placement and maintenance directly at the extent level. This design keeps the SOBR's policy engine and repository hierarchy unambiguous. Veeam's documentation for the performance tier identifies the eligible repository resources and the restrictions applicable when adding extents; see the Veeam Help Center: Performance Tier and the Veeam Help Center: Cloud Repositories.

QUESTION NO: 26

Veeam ONE _____ can detect when a potential ransomware attack is happening and be configured to take mitigating actions.

A. Reporting

B. Alert Management

C. Alarm Management

D. Data Protection View

ANSWER: C

Explanation:

Veeam ONE **Alarm Management** is the component used to detect potential ransomware behavior through the built-in *Possible Ransomware Activity* alarm. This alarm evaluates indicators of unusual activity in the virtual infrastructure, such as an abnormally high number of changed or deleted files, which can signal that ransomware is encrypting or otherwise modifying data. When the alarm is triggered, administrators can use its configuration to generate a notification and initiate a defined response. Veeam ONE alarm actions can include sending email notifications, executing scripts, and opening a support case, allowing the organization to rapidly alert relevant teams and begin mitigation procedures. Alarm Management therefore provides both the monitoring logic that identifies the suspicious condition and the configurable actions needed to respond to it. For configuration details and the supported alarm actions, see the Veeam ONE Alarm Management documentation and Veeam's ransomware detection overview.

QUESTION NO: 27

What data sources can be connected to Veeam ONE? (Choose 2)

A. VMware vCenter server

B. Veeam Backup & Replication agents

C. Veeam Backup & Replication server

D. Microsoft Azure VMs

E. Nutanix Acropolis server

ANSWER: A C

Explanation:

Veeam ONE is designed to collect monitoring, reporting, and alerting data from both the VMware virtual infrastructure layer and the Veeam data-protection layer. A VMware vCenter server can be added as a managed server, allowing Veeam ONE to obtain inventory, performance, configuration, capacity-planning, and alarm data for the vSphere environment. This connection is the basis for Veeam ONE monitoring of vCenter-managed hosts, clusters, datastores, virtual machines, and related VMware components.

A Veeam Backup & Replication server can also be connected as a data source. Veeam ONE uses this connection to collect information about backup and replication jobs, sessions, repositories, proxies, protected workloads, and backup infrastructure health. That data supports Veeam ONE Backup & Replication alarms, dashboards, and reporting, including job status and protection-related operational visibility. In practice, connecting both systems gives administrators correlated visibility: the VMware environment supplies infrastructure and workload context, while Veeam Backup & Replication supplies protection-operation context. Veeam documents these supported managed-server connections in its procedures for adding a VMware vCenter Server and a Veeam Backup & Replication server to Veeam ONE.

QUESTION NO: 28

Which is the mount server component used?

A. When there are multiple backup servers

B. When backing up to tape at a remote site

C. When the remote backup repository is a CIPS share

D. When restoring files or application items with Veeam Explorers

ANSWER: D

Explanation:

When restoring files or application items with Veeam Explorers, Veeam Backup & Replication uses a mount server to make the required backup data available for browsing and recovery. During a file-level restore, the mount server publishes the contents of the backup so the operator can locate individual guest files and folders without performing a full VM restore. The same mechanism supports Veeam Explorer-based recovery workflows for supported application data, allowing Veeam to access the relevant backup contents and present the recovery interface for granular item restoration.

The mount server is therefore an infrastructure component selected for recovery operations that require backup files to be mounted and read. Its placement and connectivity should allow it to reach the backup repository containing the restore point and the systems involved in the recovery workflow. Veeam documents the mount server as the component that mounts backup files for file-level restore and application-item restore operations. See the Veeam Mount Server documentation and the Veeam File-Level Restore documentation.

QUESTION NO: 29

Which Veeam component deploys the runtime during the backup process when a Linux VM running on Microsoft Hyper V is protected using the file indexing option.

A. Backup & Replication server

B. Guest interaction proxy

C. Virtual lab

D. File proxy

ANSWER: B

Explanation:

The **Guest interaction proxy** deploys and uses the Veeam runtime components required to communicate with a Linux guest OS when guest processing features, including file-system indexing, are enabled. During processing, this role connects to the virtual machine through the configured guest credentials and supported network protocol, performs the required in-guest operations, and retrieves the file-system metadata that Veeam stores in the backup catalog. This catalog enables users to browse or search indexed guest files and restore selected files without needing to restore the complete VM first.

For Linux virtual machines, Veeam uses a runtime/deployer mechanism in the guest to perform these operations. The Guest interaction proxy is specifically intended to offload and manage this communication and runtime deployment activity; when no dedicated proxy is assigned, the backup server can fulfill that role by default. The relevant role is therefore determined by the guest OS interaction required for indexing, not by the Hyper-V data transport path. Veeam documents the Guest interaction proxy role and its use for guest processing in its Hyper-V Backup documentation. Additional indexing behavior is described in the file-system indexing documentation.