

Free Questions for VMCE_v12

Shared by Levine on 24-05-2024

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)

Question 1

Question Type: MultipleChoice

A Microsoft SQL server is running in a VMware VM. The VM is very sensitive to snapshots, and can only be snapshotted once per day at 3 AM. However, the RPO for the databases running inside the VM is 1 hour.

Which two Application-Aware Processing Settings should be selected in the backup job so that the 1 hour RPO for the databases is met? (Choose two.)

Options:

- A- SQL Tab: Choose how this job should process Microsoft SQL Server Transaction Logs: Back up logs periodically
- B- SQL Tab: Choose how this job should process Microsoft SQL Server Transaction Logs: Do not truncate logs
- C- General Tab: VSS Settings: Perform copy only
- D- SQL Tab: Choose how this job should process Microsoft SQL Server Transaction Logs: Truncate Logs
- E- General Tab: VSS Settings: Process transaction logs with this job

Answer:

A, E

Explanation:

To meet the 1-hour RPO for the databases within the VM, while only being able to snapshot the VM once per day, the best strategy is to enable regular transaction log backups (option A) and process transaction logs with the job (option E). This setup allows the system to capture changes in the SQL Server transaction logs at intervals less than the RPO requirement, without needing to snapshot the entire VM more frequently. Reference: Veeam Backup & Replication Documentation, Veeam Backup & Replication Best Practices Guide

Question 2

Question Type: MultipleChoice

To be able to increase backup retention, the company has bought a Data Domain deduplication appliance.

After setting up the jobs to use it, the backup administrator observes an increase of resource consumption on the backup server. The proxy configuration has not been modified.

What is causing the issue?

Options:

- A- The backup appliance does not have enough resources to process the data.
- B- The SSL certificate on the backup appliance needs to be updated.
- C- The backup server has to fulfill the gateway role.
- D- The backup server needs additional resources to use a deduplication appliance.

Answer:

C

Explanation:

When integrating a Data Domain deduplication appliance with Veeam Backup & Replication, it is typically used as a backup repository. The backup server may need to take on the gateway role, especially if the Data Domain is integrated over NFS or CIFS. This means that the backup server will be responsible for processing the data flow between the Veeam proxies and the deduplication appliance. If the gateway server (backup server) is not well-resourced, this additional workload can cause an increase in resource consumption on the backup server. The appliance's resources and the SSL certificate are not related to increased resource consumption on the backup server, and simply needing additional resources for deduplication (D) is not specific enough without indicating the gateway role.

Question 3

Question Type: MultipleChoice

A number of VMs are running as interdependent applications. They need to fail over, one by one, as a group. What method should be used to do this?

Options:

- A- Replica failover
- B- Replication plan
- C- Planned failover

D- Failover plan

Answer:

D

Explanation:

To ensure VMs running interdependent applications fail over one by one, as a group, the method to use is D: Failover plan. In Veeam Backup & Replication, a failover plan allows for the orchestration of a group of replicas to fail over in a predefined sequence. This includes the capability to set up delays between starting each VM, which is crucial for interdependent applications that must be started in a specific order to function correctly. The failover plan ensures that dependencies among the group are respected and that the startup sequence follows the correct order, enabling a smooth and organized transition to the failover state.

Question 4

Question Type: MultipleChoice

A Veeam administrator is creating a protection group that needs to back up all servers for the accounting department using Veeam Agent for Microsoft Windows. New accounting servers must be automatically added to the protection group. How is this accomplished?

Options:

A- Create an Active Directory based protection group targeted at an Active Directory security group.

B- Install a configuration file on each of the accounting servers that adds it to the protection group.

C- Create a protection group and set it to automatically scan every three hours for new servers.

D- Create a protection group that reads a CSV file listing the accounting servers.

Answer:

A

Explanation:

To ensure that all servers for the accounting department are backed up and any new servers

added to the department are automatically included in the protection group, the best approach is to use an Active Directory (AD) based protection group. Therefore, the correct answer is A: Create an Active Directory based protection group targeted at an Active Directory security group.

By leveraging an AD-based protection group, the Veeam administrator can specify an AD security group that encompasses all user accounts or computer accounts (in this case, servers) associated with the accounting department. As new servers are introduced and added to this AD security group, Veeam Backup & Replication will automatically detect these new members during its periodic rescan (which can be scheduled as per the administrator's preferences) and include them in the protection group. This method not only automates the inclusion of new servers into the backup routine but also simplifies the management of backup policies for dynamically changing environments such as the accounting department's server infrastructure.

Question 5

Question Type: MultipleChoice

A customer wants to set up a Scale-Out Backup Repository. Due to malware concerns, immutability is recommended. An on-premises server can be used to hold primary backups, but it can only hold about 21 days of backups. A copy of the backups should be stored in AWS. The retention for all backups is 60 days.

Which configuration of a Scale-out Backup repository meets these requirements?

Options:

A- Copy mode

Performance Tier: Windows REFS, immutability set for 11 days

Capacity Tier: Amazon S3, immutability set for 60 days

B- Copy and move mode

Performance Tier: Windows REFS, immutability set for 11 days

Capacity Tier: Amazon S3, immutability set for 60 days

C- Copy mode

Performance Tier: Linux Hardened Repository, XFS, immutability set for 11 days

Capacity Tier: Amazon S3, immutability set for 60 days

D- Copy and move mode

Performance Tier: Linux Hardened Repository, XFS, immutability set for 11 days

Capacity Tier: Amazon S3, immutability set for 60 days

Answer:

D

Explanation:

To meet the requirements of setting up a Scale-Out Backup Repository (SOBR) with immutability for malware protection and specific retention policies, the most fitting configuration is D: Copy and move mode with a Performance Tier configured on a Linux Hardened Repository using the XFS file system and immutability set for 21 days, and a Capacity Tier on Amazon S3 with immutability set for 60 days. This setup utilizes the copy and move mode to ensure that backups are first stored on the on-premises Linux Hardened Repository with an immutability setting that prevents modifications to backups, providing protection against malware for the most recent 21 days of backups. As this on-premises server has limited capacity, older backups beyond 21 days are moved to the Capacity Tier in Amazon S3, where they are also protected with immutability for the entire 60-day retention period. This configuration leverages the strengths of both on-premises and cloud storage while ensuring that all backups are protected from modification or deletion by immutability, aligning with the customer's malware protection concerns and retention requirements.

Question 6

Question Type: MultipleChoice

A 3-node Microsoft SQL Always On cluster is running in a VMware environment.

The RPOs are:

- * 1 day for the cluster nodes
- * 15 minutes for the log files

How should the cluster be backed up?

Options:

- A-** Veeam Backup and Replication - Create a VMware backup job, include the 3 nodes of the cluster. Use Application-Aware Processing to back up the transaction logs every 15 minutes.
- B-** Veeam Agent for Windows - Create a managed by server backup job for all 3 nodes. Use Application-Aware Processing to back up the database once a day.
- C-** Veeam Agent for Windows - Create a managed by agent backup policy for all 3 nodes. Use a log shipping server to back up the transaction logs every 15 minutes.
- D-** Veeam Enterprise Plug-in for Microsoft SQL - Install the Plug-in on each node in cluster. Use Microsoft SQL Server Management Studio to create the backups.

Answer:

A

Explanation:

For backing up a 3-node Microsoft SQL Always On cluster running in a VMware environment with the specified Recovery Point Objectives (RPOs), the most suitable option is A: Veeam Backup and Replication - Create a VMware backup job, include the 3 nodes of the cluster. Use Application-Aware Processing to back up the transaction logs every 15 minutes. This approach allows for the entire SQL cluster nodes to be backed up as part of a regular VM backup job, which satisfies the 1-day RPO for the cluster nodes. The Application-Aware Processing feature of Veeam Backup & Replication ensures that the SQL databases are in a consistent state during backup. Moreover, this feature allows for the transaction logs to be backed up separately at a more frequent interval, in this case, every 15 minutes, meeting the 15-minute RPO requirement for the log files. This dual approach ensures comprehensive protection for both the SQL cluster nodes and the critical transaction logs, aligning with the specified RPOs.

Question 7

Question Type: MultipleChoice

What is the primary benefit of configuring replica mapping?

Options:

- A- Compressed WAN traffic
- B- Deduplicated WAN traffic
- C- Reduced WAN traffic
- D- Encrypted WAN traffic

Answer:

C

Explanation:

The primary benefit of configuring replica mapping in Veeam Backup & Replication is C: Reduced WAN traffic. Replica mapping is a feature that allows an administrator to map an existing VM in the disaster recovery site to a replica job. This is particularly useful when the initial replication of

a large VM has been performed by other means (such as shipping a hard drive with the VM data to the DR site) or if a replica VM already exists and needs to be re-synced with the source VM. By using replica mapping, Veeam can avoid transferring the entire VM over the WAN again and instead only synchronize the differences between the source VM and the existing replica. This significantly reduces the amount of data that needs to be transferred over the WAN, conserving bandwidth and speeding up the replication process. It's an efficient way to manage replicas, especially for large VMs or in environments with limited bandwidth.

Question 8

Question Type: MultipleChoice

Management asks a backup administrator to deploy the Veeam Agent on a number of Amazon EC2 instances running Windows and Linux operating systems. A Veeam Protection Group is also required by management. The Veeam Distribution Server does not have network access to these instances.

What protection group type should be used to select these objects?

Options:

- A- Individual computers
- B- Microsoft Active Directory objects
- C- Computers listed in a CSV file
- D- Cloud machines

Answer:

D

Explanation:

For deploying the Veeam Agent on Amazon EC2 instances running Windows and Linux operating systems without direct network access from the Veeam Distribution Server, the appropriate type of Protection Group to use is D: Cloud machines. The 'Cloud machines' protection group type in Veeam Backup & Replication is specifically designed for protecting cloud-based workloads, including instances in public cloud environments like Amazon EC2. This protection group type allows the Veeam Agent to be deployed and managed remotely, even when the Veeam Distribution Server cannot directly access the instances over the network. It facilitates centralized management of backup tasks for cloud instances, ensuring that the EC2 instances are adequately protected as per management's request, despite the network accessibility constraints.

Question 9

Question Type: MultipleChoice

An administrator is asked to change a backup copy job from periodic mode to immediate mode. How can this be accomplished?

Options:

- A- Enable immediate in the backup copy job settings drop down.
- B- Right click on the job name and choose immediate.
- C- Copy the original backup copy job to a different folder and remap the job.
- D- Create a new backup copy job and delete the original job.

Answer:

D

Explanation:

To change a backup copy job from periodic mode to immediate mode, the most straightforward approach is D: Create a new backup copy job and delete the original job. Veeam Backup & Replication does not directly allow changing the mode of an existing backup copy job from periodic to immediate within the job settings. Therefore, the recommended practice is to set up a new backup copy job with the desired settings, in this case, immediate mode, which starts copying backups as soon as they are created by the primary backup job. After the new backup copy job is configured and tested to confirm it meets the requirements, the original periodic mode job can be safely deleted. This ensures a seamless transition to the immediate mode operation without risking data protection consistency or coverage.

Question 10

Question Type: MultipleChoice

Why is it recommended to have at least one backup proxy server in each site when defining a replica job?

Options:

- A- The proxies allow replication automatic restart after failure.
- B- The proxies allow for no VM snapshots during transit.
- C- It allows deduplication during data transit across the WAN
- D- The proxies allow automatic WAN acceleration.
- E- The proxies enable a stable connection for VM data transfer across sites.
- F- It allows for no VM snapshots

Answer:

E

Explanation:

Having at least one backup proxy server in each site when defining a replica job is recommended because the backup proxy servers are responsible for data processing and transfer. Having proxies in both sites enables a stable connection for VM data transfer across sites, as they handle the data compression, deduplication, and transfer processes. This setup ensures that data is efficiently moved from one site to another, thereby making replication more resilient and reliable. Proxies do not automatically restart replication after failure (A), remove the need for VM snapshots during transit (B and F), or enable automatic WAN acceleration (D), although they can work in conjunction with WAN accelerators if configured to do so. They also don't directly deduplicate data during transit across the WAN (C), although they do compress and optimize it for transfer.

To Get Premium Files for VMCE_v12 Visit

https://www.p2pexams.com/products/vmce_v12

For More Free Questions Visit

<https://www.p2pexams.com/veeam/pdf/vmce-v12>



20%
DISCOUNT

A red and white badge with a diamond pattern background. The text "20%" is in a large, bold, red font, and "DISCOUNT" is in a smaller, bold, red font below it.