



VEEAM

Exam Questions VMCE_v12

Veeam Certified Engineer v12

NEW QUESTION 1

What feature is only available with the Veeam Agent for Linux?

- A. File-level backup
- B. Application-aware processing of
- C. Backup from native snapshots
- D. Volume backup

Answer: C

Explanation:

The feature that is unique to Veeam Agent for Linux and not available in other Veeam Agent configurations is C: Backup from native snapshots. Veeam Agent for Linux includes the ability to leverage native snapshot capabilities of the Linux kernel, such as LVM (Logical Volume Manager) snapshots or Btrfs subvolume snapshots, to create consistent point-in-time copies of data. This capability allows for application-consistent backups even in complex Linux environments, ensuring that data is captured in a consistent state without the need for custom scripting or downtime. Native snapshot support in Veeam Agent for Linux enhances the flexibility and reliability of backups, particularly in environments where Linux-based applications and databases are critical to business operations.

NEW QUESTION 2

In Veeam Enterprise Manager, what granular restoration permission can be set for the Restore Operator role?

- A. SharePoint items
- B. Active Directory items
- C. OneDrive items
- D. File items

Answer: D

Explanation:

In Veeam Enterprise Manager, granular restoration permissions can be set for the Restore Operator role, including the restoration of file items. This role allows designated users to perform file-level restores, ensuring that they can recover individual files without having full administrative privileges over the backup infrastructure. References:

? Veeam Enterprise Manager Guide: Roles and Permissions

? Veeam Help Center: Configuring Granular Restore Permissions

NEW QUESTION 3

It is required that some replicated VMs start on a time delay during a failover. How can this be accomplished?

- A. Create a failover plan.
- B. Adjust boot delay in application group.
- C. Use a pre-freeze/post-thaw script.
- D. Modify the failover template file.

Answer: A

Explanation:

A failover plan in Veeam Backup & Replication allows you to define the order in which VMs are started during a failover operation. It can include startup delays for certain VMs, ensuring that some VMs can be configured to start after a defined time delay. This meets the requirement of having some replicated VMs start on a time delay. References:

? Veeam Backup & Replication User Guide: Failover Plans

? Veeam Help Center: Creating Failover Plans

NEW QUESTION 4

A company wants to ensure that, during a replication failover, database servers boot before the application servers. How can this be accomplished?

- A. Create a failover plan
- B. Create a disaster recovery template
- C. Create a replica mapping
- D. Create a planned failover

Answer: A

Explanation:

Veeam Backup & Replication allows you to set up failover plans to control the order in which VMs are started during a failover. By creating a failover plan, you can specify that the database servers boot before the application servers. This is achieved by setting up VM dependencies within the failover plan, ensuring that the database servers (which are critical for the application servers to function) are operational before the application servers start. Creating a disaster recovery template (B), creating a replica mapping (C), or initiating a planned failover (D) does not directly address the boot order of the servers during failover.

NEW QUESTION 5

A Windows Server using the ReFS filesystem has been used as a standalone Veeam repository for several years and is due for replacement. A new Windows server using the ReFS filesystem has been created to replace the old server, with twice the capacity. Backup files need to be transferred to the new server with no disruptions to the existing backup chains.

The Veeam engineer has begun to move backup files to the new repository but is now getting alerts that it is running out of space.

How could the engineer have avoided this issue?

- A. Use a Backup Copy Job
- B. Use the "Copy backup..." function
- C. Use the "Move backup..." function

D. Use Robocopy with the /compress switch


Answer: C

Explanation:

To avoid running out of space when moving backups to a new repository, the "Move backup..." function in Veeam Backup & Replication should be used. This function allows you to relocate backup files to a new repository without duplicating data, which can save space. Unlike a simple copy action, the move function ensures that the backup chain remains intact and does not require additional space for a copy of the backups during the transfer. When the move is initiated, Veeam will also automatically update the configuration to point to the new backup location, thus preventing any disruptions in the backup chain.

NEW QUESTION 6

Hourly backup jobs are configured to local repositories. Daily backups must be sent from Site B to the repository on Site A. The connection between the two sites does not allow a direct backup to complete within the backup window. Impact on the source VMs at Site B should be minimized. Which type of job and Veeam components should be used to provide the necessary protection?



- A. A backup copy job and components 6 ? 4 ? 3 ? 2
- B. A backup job and components 5 ? 4 ? 3 ? 2
- C. A backup job and components 5 ? 2
- D. A backup copy job and components 6?5?4?3?1 ? 2

Answer: A

Explanation:

For the given scenario where daily backups need to be sent from Site B to a repository at Site A and the connection between the two sites is limited, thus impacting the ability to complete direct backups within the backup window, a Backup Copy Job would be appropriate. A Backup Copy Job is a feature in Veeam Backup & Replication that allows you to create several instances of the same backup files across different locations (repositories).

In the image provided, the components involved in the Backup Copy Job from Site B to Site A would be:

- ? 6 (Repository at Site B): This is the source repository where the hourly backup jobs are stored.
 - ? 4 and 3 (WAN Accelerators at both sites): These components optimize data transfer over the WAN.
 - ? 2 (Repository at Site A): This is the target repository where the backup copies will be stored.
- This setup would minimize impact on the source VMs at Site B since the Backup Copy Job works with backup data rather than directly with the production VMs, thus reducing the load on those VMs during the process.

NEW QUESTION 7

What is the primary benefit of configuring replica seeding?

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

Answer: D

Explanation:

The primary benefit of configuring replica seeding in Veeam Backup & Replication is reduced WAN traffic. Replica seeding allows for the initial replica to be created using a backup copy that is transported to the DR site, which significantly reduces the amount of data that needs to be transferred over the WAN during the initial replication process. References: Veeam Backup & Replication User Guide, Veeam Replica Seeding Guide

NEW QUESTION 8

An engineer is using Veeam Backup and Replication v12. The only backup repository is a Microsoft Windows server with direct attached Fibre Channel storage array. The engineer realizes that none of their backups are immutable. A second copy of the backup on a different site and a different media is required. Which option should be used to provide immutable backups on a secondary site with a different media?

- A. Create a Scale Out Backup Repository with the existing Microsoft Windows Server as the performance tier and an HPe StoreOnce Catalyst share with immutability enabled as the capacity tier.
- B. Create a Scale Out Backup Repository with the existing Microsoft Windows Server as the performance tier and AWS S3 bucket with immutability enabled as the capacity tier.
- C. Create a new hardened repository on a new Microsoft Windows Server, mark it as immutable and create a backup copy job on it.
- D. Create a Scale Out Backup Repository with the existing Microsoft Windows Server as the performance tier and Google Cloud Object Storage with immutability enabled as the capacity tier.

Answer: B

Explanation:

To provide immutable backups on a secondary site with a different media, the best option given the context is B: Create a Scale Out Backup Repository (SOBR) with the existing Microsoft Windows Server as the performance tier and an AWS S3 bucket with immutability enabled as the capacity tier. This approach involves leveraging the existing backup infrastructure (Microsoft Windows Server with direct-attached storage) as the performance tier of the SOBR, where the most recent backups are stored for fast access. For long-term storage and immutability, backups can be offloaded to an AWS S3 bucket configured with Object Lock. The Object Lock feature in AWS S3 provides an additional layer of data protection by making the backup data immutable, meaning it cannot be deleted or modified for a specified duration. This setup ensures that backup data is protected against accidental deletion, ransomware, and other malicious activities. By implementing this configuration, the engineer can achieve the desired level of data protection and immutability, utilizing cloud storage as a secure and scalable secondary backup location, distinct from the primary on-premises storage media.

NEW QUESTION 9

Which Veeam Backup & Replication functionality achieves the lowest RPO?

- A. Backup Copy Job
- B. Continuous Data Protection
- C. Snapshot-based Replication
- D. File to Tape

Answer: B

Explanation:

Within Veeam Backup & Replication, the functionality that achieves the lowest Recovery Point Objective (RPO) is B: Continuous Data Protection (CDP). CDP is designed to protect critical workloads by continuously capturing changes and replicating them to a target site, allowing for very low RPOs, often measured in seconds. This is achieved through the use of VMware's vSphere APIs for I/O Filtering (VAIO), which intercepts and replicates I/O streams almost in real-time. CDP is particularly valuable for applications that require high levels of availability and cannot tolerate significant data loss. By providing near-continuous replication, CDP ensures that in the event of a failure or disaster, data loss can be minimized to a very narrow window, significantly reducing the potential impact on business operations.

NEW QUESTION 10

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.)

- 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: AE

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. References: Veeam Backup & Replication Documentation, Veeam Backup & Replication Best Practices Guide

NEW QUESTION 10

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

- 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.

NEW QUESTION 13

A Veeam engineer creates a Scale-Out Backup Repository (SOBR) that uses AWS S3 as the Performance Tier. The backup job is configured to "Keep monthly full backups for: 12 months". The engineer wants the backups to move to Glacier after 90 days.

What should the engineer do first to achieve this goal?

- A. Create a Vault in S3 Glacier
- B. Add AWS S3 as a Capacity Tier before they can use Glacier
- C. Add an Archive Tier with the appropriate Bucket from S3
- D. Reconfigure the SOBR to use block storage as the Performance Tier

Answer: B

Explanation:

In Veeam Backup & Replication, to utilize AWS S3 Glacier for long-term storage, the AWS S3 must first be added as a Capacity Tier within the Scale-Out Backup Repository (SOBR). This step is necessary before backups can be moved to Glacier. Once S3 is established as the Capacity Tier, the policies for moving backups to Glacier can be configured based on the retention requirements (in this case, after 90 days). This approach ensures that the monthly full backups are initially stored in S3 and then offloaded to Glacier for cost-effective long-term retention. References:

- ? Veeam Documentation: Scale-Out Backup Repository
- ? Veeam Help Center: AWS S3 as Capacity Tier and Archive Tier Configuration

NEW QUESTION 16

Veeam Backup & Replication is currently configured to keep backups for 14 days on a NAS repository. An engineer needs to enable GFS with 12 monthly and 5 yearly full backups to extend the retention for archival purposes but does not have enough space on the repository to store them. After adding the repository to Scale-Out-Backup Repository (SOBR) as Performance Tier, which option should the engineer choose to solve this issue?

- A. Add an Object Storage as Capacity tier and set the policy to copy backups to object storage as soon as they are created
- B. Add another NAS repository as Capacity tier and set the policy to move backups older than 14 days
- C. Add an Object Storage as Archive tier and set the policy to move all GFS backups
- D. Add an Object Storage as Capacity tier and set the policy to move backups older than 14 days

Answer: D

Explanation:

To solve the space issue while enabling GFS retention, the engineer should add an Object Storage as a Capacity tier to the Scale-Out Backup Repository (SOBR) and set the policy to move backups that are older than 14 days. This approach allows for efficient use of local NAS storage for short-term retention and leverages object storage for long-term archival purposes. References: Veeam Backup & Replication User Guide, Veeam SOBR Configuration Guide

NEW QUESTION 18

In the war against ransomware, a company decided to implement tape backup. The infrastructure contains Hyper-V VMs. What is the correct approach to getting all servers onto tape?

- A. Create media pools, then create file to tape jobs.
- B. Create media sets, then create backup to tape jobs.
- C. Create media sets, then create file to tape jobs.
- D. Create media pools, then create backup to tape jobs.

Answer: D

Explanation:

The correct approach for backing up Hyper-V VMs to tape in a ransomware protection strategy is to create media pools first and then configure backup to tape jobs. This method ensures organized and efficient tape management while securing VM backups on tape media. References: Veeam Backup & Replication User Guide, Veeam Tape Backup Configuration Guide

NEW QUESTION 20

A company is running nightly backups to satisfy their 24-hour RPO. There are two critical applications that cannot be offline more than 4 hours with no more than an hour of data loss. How can this be accomplished?

- A. Linux hardened repository with immutability
- B. A local SOBR with AWS Object storage
- C. Backup copy jobs with hourly VM replication
- D. Enable Continuous Backup Replication

Answer: C

Explanation:

To meet the stringent Recovery Time Objective (RTO) of 4 hours and Recovery Point Objective (RPO) of 1 hour for the two critical applications, the most suitable approach is to use a combination of backup copy jobs and hourly VM replication, which corresponds to answer C. Backup copy jobs can ensure that backups are available offsite for disaster recovery purposes, while VM replication provides the ability to quickly failover to a replica VM in case of a primary VM failure, thereby minimizing downtime. Veeam's replication capabilities are designed to create exact copies of VMs at regular intervals, which can then be rapidly activated in case of a failure. This strategy ensures that the applications can be brought back online within the 4-hour RTO, with data loss not exceeding the 1-hour RPO, thus meeting the company's stringent data protection requirements for these critical applications.

NEW QUESTION 21

Which two environments can Veeam Agents back up? (Choose two.)

- A. FreeBSD
- B. Ubuntu
- C. IBM iSeries
- D. iOS
- E. Windows Server

Answer: BE


Explanation:

Veeam Agents are designed to provide backup solutions for physical and cloud environments. They support various operating systems, including Ubuntu (a Linux-based OS) and Windows Server. These agents ensure that data on servers running these operating systems can be effectively backed up and restored. However, Veeam Agents do not support FreeBSD, IBM iSeries, or iOS as these platforms require different backup solutions or are not typically used in environments where Veeam operates. References:

- ? Veeam Agents User Guide
- ? Veeam Help Center: Supported Environments

NEW QUESTION 25

A backup administrator must enable guest file system indexing for a backup job of a Microsoft Windows file server VM. All required credentials are added to the Veeam Backup & Replication server. VMware Tools (VIX) cannot be used due to security regulations. What should be done to make it possible in the following environment?



- A. Use the proxy server as a guest interaction proxy
- B. Connect the ESXi host to the VM network
- C. Use the Veeam Backup & Replication server as a guest interaction proxy
- D. Install a file system indexing plugin on the Microsoft Windows file server

Answer: A

Explanation:

If VMware Tools (VIX) cannot be used due to security regulations, the backup administrator can use the proxy server as a guest interaction proxy to enable guest file system indexing for a backup job. The guest interaction proxy is responsible for the deployment and coordination of guest processing tasks, such as guest file system indexing,

which does not require VIX. References:

- ? Veeam Backup & Replication User Guide: Guest Interaction Proxy
- ? Veeam Knowledge Base: Configuring Guest Interaction Proxies

NEW QUESTION 26

What is the purpose of a Cache Repository when adding the file share?

- A. To walk the file shares and move data to the backup repository
- B. To control how fast the backup proxy can read data from the source file share
- C. To make sure backup retention is being managed correctly
- D. To store temporary metadata and track all objects that have changed

Answer: D

Explanation:

The purpose of a Cache Repository when adding a file share in Veeam Backup & Replication is D: To store temporary metadata and track all objects that have changed. When Veeam Backup & Replication backs up data from file shares, it uses the Cache Repository to store metadata related to the files and directories on the share. This metadata includes information about file versions, change logs, and the structure of the file share, which is essential for efficient backup operations, incremental backups, and restore operations. The Cache Repository plays a crucial role in tracking changes between backup jobs, enabling Veeam to perform quick incremental backups by only processing the data that has changed since the last backup, thereby optimizing backup performance and reducing network and storage load.

NEW QUESTION 30

A physical Linux file server needs to be migrated to VMware ESXi. The server has several volumes: /dev/sda (60GB), /dev/sdb (1TB) and /dev/sdc (4TB). What is the quickest option to migrate the server?

- A. Create an empty VM and perform bare metal recovery inside the VM.
- B. Perform Instant VM Recovery.
- C. Perform Instant Disk Recovery for each volume.
- D. Export all disks, create an empty VM and attach the disks.

Answer: B

Explanation:

The quickest option for migrating a physical Linux file server with multiple volumes to VMware ESXi is to perform Instant VM Recovery. This feature allows you to run the server directly from the backup file in a VMware environment, significantly reducing the time and complexity involved in migration. References: Veeam Backup & Replication User Guide, Veeam Instant VM Recovery for Physical Servers

NEW QUESTION 35

A physical Linux server protected by a centrally managed Veeam agent is physically damaged. A VMware vSphere infrastructure is available, and the physical server is eligible for virtualization. Which recovery step provides the lowest possible RTO?

- A. Use Instant VM Recovery to VMware
- B. Use Export Disk Content as Virtual Disk to create a new VM.
- C. Use Bare Metal Restore to VMware vSphere VM.
- D. Use Instant Disk Recovery to VMware vSphere.

Answer: A

Explanation:

Instant VM Recovery to VMware offers the lowest possible RTO for restoring a physically damaged Linux server to a VMware vSphere infrastructure. This feature rapidly restores service by running the server directly from the backup file in a VMware environment. References: Veeam Backup & Replication User Guide, Veeam Agent for Linux Guide

NEW QUESTION 39

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 VM backups to a repository at the disaster recovery site. They want to keep six months' worth of 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 62 days of retention and six monthly GFS restore points.
- B. Use immediate copy (mirroring) mode, keeping 180 days of retention.
- C. Use periodic copy (pruning) mode, keeping 180 days of retention.
- D. Use immediate copy (mirroring) mode, keeping 62 days of retention and 26 weekly GFS restore points.

Answer: A

Explanation:

To meet the specified requirements, configuring a backup copy job in periodic copy (pruning) mode with 62 days of retention and six monthly Grandfather-Father-Son (GFS) restore points is the best approach. This setup allows for daily backups for up to two months and weekly backups for up to six months, aligning with the desired restore capabilities. References: Veeam Backup & Replication User Guide, Veeam GFS Retention Policy Explanation

NEW QUESTION 43

Why is it recommended to install Veeam ONE before Veeam Backup & Replication?

- A. Veeam ONE stores the license file necessary for Veeam Backup & Replication.
- B. Veeam ONE can verify that the server meets the minimum requirements to install Veeam Backup & Replication.
- C. Veeam ONE is a prerequisite for Veeam Backup & Replication installation.
- D. Veeam ONE can potentially identify infrastructure issues prior to performing a backup.

Answer: D

Explanation:

It is recommended to install Veeam ONE before Veeam Backup & Replication because Veeam ONE can potentially identify infrastructure issues that might affect backup operations before they are performed. This proactive approach ensures that the environment is optimized for successful backup and replication tasks. References: Veeam Installation and Configuration Guide, Veeam ONE Deployment Best Practices

NEW QUESTION 45

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

VMCE_v12 Practice Exam Features:

- * VMCE_v12 Questions and Answers Updated Frequently
- * VMCE_v12 Practice Questions Verified by Expert Senior Certified Staff
- * VMCE_v12 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * VMCE_v12 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The VMCE_v12 Practice Test Here](#)