DUMPSBOSS.

Veeam Certified Engineer v12

Veeam VMCE_v12

Version Demo

Total Demo Questions: 21

Total Premium Questions: 213

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Architecture	35
Topic 2, Configuration	48
Topic 3, Protection	45
Topic 4, Recovery	62
Topic 5, Verification	16
Topic 6, Troubleshooting	7
Total	213

QUESTION NO: 1

An engineer needs to be able to perform all functions without needing to access the backup server locally or over remote desktop.

What should the engineer do?

A. Use SSH to connect to the Veeam Backup & Replication server.

B. Install the required Veeam Explorers locally

C. Install the Veeam console on the engineer's desktop.

D. Use Veeam Backup Enterprise Manager to connect to the Veeam Backup & Replication server.

ANSWER: C

Explanation:

Installing the Veeam Backup & Replication console on the engineer's desktop provides the required full remote administration capability. The console is the primary management interface for Veeam Backup & Replication and can connect to a backup server across the network. After connecting, an authorized engineer can perform the same administrative activities available through the console installed on the backup server itself, including configuring backup and replication jobs, managing repositories and proxies, monitoring sessions, performing restores, and maintaining the backup infrastructure. This approach avoids the need to interactively sign in to the backup server at its physical console or through Remote Desktop Protocol, while retaining centralized control through the Veeam management interface. The desktop must meet the console system requirements and be able to communicate with the backup server; the engineer must also use an account with appropriate Veeam Backup & Replication permissions. Veeam documents that the console may be installed on a separate machine and used to connect remotely to the backup server. See the Veeam Backup & Replication Console documentation and the Veeam installation documentation.

QUESTION NO: 2

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

A. FreeBSD

B. Ubuntu

C. IBM iSeries

D. iOS

E. Windows Server

ANSWER: B E

Explanation:

Veeam Agents can protect both Ubuntu and Windows Server workloads. Veeam Agent for Linux supports backup of supported Linux distributions, including Ubuntu. The agent performs image-level backup of the computer and can protect the full system, selected volumes, or individual files and folders, depending on the configured backup mode. This makes it suitable for physical servers, cloud instances, and other Linux machines that need centrally managed or standalone protection.

Veeam Agent for Microsoft Windows supports Windows Server operating systems. It provides image-level backups for server workloads and supports recovery scenarios such as full-machine recovery, volume-level recovery, file-level recovery, and application-aware processing for supported Microsoft applications. Agent backups may be managed through Veeam Backup & Replication or used independently, according to the deployed edition and configuration. Veeam’s platform documentation identifies separate agents for Microsoft Windows and Linux and lists the supported operating systems and requirements. See the Veeam Agent for Microsoft Windows system requirements and the Veeam Agent for Linux system requirements.

QUESTION NO: 3

Company security policy states that backups must be copied to a remote location within 8 hours.What step can a backup administrator take to receive an alert automatically if the backup copy job violates the company's policy?

A. Set a Data Protection alert in Enterprise Manager

B. Set a backup copy window alert on the repository

C. Set up post-thaw scripts to send an email

D. Enable the appropriate RPO monitor

ANSWER: D

Explanation:

Enable the appropriate RPO monitor. In Veeam Backup & Replication, an RPO monitor evaluates whether protected workloads meet a defined recovery point objective by checking the age of their most recent valid restore point. For this requirement, the administrator can configure the monitor with an 8-hour RPO and apply it to the backup copy job or its protected workloads. If the copy process does not produce an off-site restore point within that interval, the RPO monitor identifies the policy breach and generates an alarm/notification through the configured Veeam notification mechanisms. This provides an automated, policy-focused control rather than requiring someone to manually inspect job sessions or calculate copy timing from reports. RPO monitoring is designed to detect when the available recovery point is older than the organization’s permitted threshold, which directly matches a requirement that backups be present at a remote location within eight hours. The monitor should be configured using the intended backup-copy scope and notification settings so the responsible team receives the alert promptly. See the Veeam documentation on RPO monitoring and notification configuration.

QUESTION NO: 4

Which actions can be used to locate the exact extent of a Scale-Out Backup Repository containing a particular VM's backup data? (Choose two.)

A. Check Scale-Out Backup Repository settings

B. Check the specific VM in the job statistics

C. Check the inventory view for the specific VM

D. Open the extent repositories manually

E. Check backup proxy settings

ANSWER: A B

Explanation:

Check Scale-Out Backup Repository settings and **Check the specific VM in the job statistics** are the appropriate actions. The Scale-Out Backup Repository configuration provides visibility into the repository’s performance-tier extents, including the repositories that participate in the scale-out repository. This establishes the available extent locations that Veeam Backup & Replication can use for backup chains and related files.

For a particular protected workload, the job or task-session statistics provide the operational details required to correlate that workload’s backup processing with the repository storage used during the session. Reviewing the individual VM’s task information is especially useful when a job protects several VMs, because it narrows the investigation to the selected VM rather than the job’s overall target configuration. Used together, repository configuration and VM-level job statistics let an administrator identify the relevant physical performance-tier extent for the VM’s backup data.

Veeam documents Scale-Out Backup Repository administration and performance-tier extent configuration in the Scale-Out Backup Repository management guide. For information on reviewing backup and replication job sessions and task details, see the Veeam job session monitoring documentation.

QUESTION NO: 5

Which of the following is used in order to avoid unwanted changes to VM replicas when performing verification or recovery?

A. .vlb file

B. Virtual Lab Manager

C. VM snapshot

D. .ctk file

ANSWER: C

Explanation:

VM snapshot is used to preserve the replica's original state when it is powered on for operations such as replica verification or failover/recovery testing. Before a replica is started, Veeam Backup & Replication relies on the virtualization platform's snapshot mechanism to capture its current disk state. Any writes generated while the replica runs—including operating-system activity, application updates, or test changes—are redirected to the snapshot's delta files instead of permanently modifying the replica disks. This makes it possible to safely test that the replica can boot and operate, or to run it temporarily during a recovery workflow, while retaining the ability to discard those temporary changes and return the replica to its pre-operation condition. If the recovery is confirmed as the desired final state, the appropriate failover workflow can instead commit the changes. This snapshot-based protection is fundamental to keeping a replica recoverable and avoiding unintended divergence from the replicated restore point. Veeam documents that replica failover uses snapshots to preserve and manage the replica state during recovery operations; see the Veeam failover documentation and the SureReplica verification documentation.

QUESTION NO: 6

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?

A. Performance Tier: Windows REFS, immutability set for 11 days

B. Performance Tier: Windows REFS, immutability set for 11 days

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

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

E. Performance Tier: Linux Hardened Repository, XFS, immutability set for 21 days; Capacity Tier: Amazon S3 with immutability set for 60 days; use copy and move mode

ANSWER: E

Explanation:

Performance Tier: Linux Hardened Repository, XFS, immutability set for 21 days; Capacity Tier: Amazon S3 with immutability set for 60 days; use copy and move mode is the appropriate design. The Linux hardened repository provides on-premises backup storage with Veeam's immutability protection, helping prevent alteration or deletion of recent restore points by compromised credentials or malware. Setting its immutability period to 21 days aligns with the available local capacity and ensures the backups retained locally receive protection throughout their local residence.

The Amazon S3 capacity tier provides the required off-site copy and long-term retention. Enabling immutability in the capacity tier uses Amazon S3 Object Lock, so the cloud backup objects remain protected for the required 60-day retention period. The copy policy ensures backups are copied to AWS promptly, while the move policy can offload backups from the performance tier after the 21-day operational restore window, avoiding exhaustion of the on-premises repository while

preserving the full retention period in AWS. Veeam documents capacity-tier offload and copy policies at Veeam Help Center and describes hardened repository immutability at Veeam Help Center.

QUESTION NO: 7

For which workload can Veeam Data Platform achieve image-level backups?

A. AS/400

B. Solaris

C. IOS

D. IHP-UX

ANSWER: B

Explanation:

Solaris is correct when the Solaris workload is hosted as a supported virtual machine, such as on VMware vSphere. Veeam Data Platform performs image-level backup by protecting the virtual machine at the hypervisor level. This captures the VM's virtual disks and configuration as a recoverable image, rather than requiring an in-guest, Solaris-specific backup agent. As a result, the guest operating system does not need to be a Veeam Agent-supported operating system for Veeam to protect the VM image. A Solaris VM can therefore be backed up and restored as a complete virtual-machine workload, including its operating system, installed applications, and stored data. Veeam's VMware backup architecture uses VMware snapshots and APIs to obtain a consistent point-in-time image of the virtual machine, then transfers and stores the VM data in the backup repository. This approach is a core capability of Veeam Backup & Replication within Veeam Data Platform. For details on VMware image-level processing, see the Veeam Backup & Replication VMware vSphere overview and the Veeam VMware backup process documentation.

QUESTION NO: 8

Which capabilities are available in Veeam Explorer for Microsoft Exchange? (Choose three.)

A. Restore mailbox items to the original location

B. Restore mailbox items to another mailbox

C. Export mailbox items to a PST file

D. Restore Oracle archived redo logs

E. Convert an Exchange database to a VMware template

F. Create a Hyper-V replica of a mailbox

ANSWER: A B C

Explanation:

Restore mailbox items to the original location, **restore mailbox items to another mailbox**, and **export mailbox items to a PST file** are available capabilities in Veeam Explorer for Microsoft Exchange. The explorer provides granular access to Exchange data stored in Veeam backups, allowing administrators to locate messages, folders, contacts, calendar items, tasks, and other mailbox content without restoring an entire Exchange database.

Granular restore supports recovery to the original mailbox or to an alternate mailbox when the original destination is not appropriate. Exporting selected mailbox content to a PST file is useful for review, legal discovery, or providing recovered data without importing it immediately into Exchange. See the Veeam Explorer for Microsoft Exchange documentation.

QUESTION NO: 9

An engineer needs to test the impact of operating system updates on app functionality in an isolated non-production environment.

What Veeam Backup and Replication feature can be used to facilitate this?

A. Replication

B. Direct Storage Snapshots

C. Virtual Labs

D. Instant Recovery

ANSWER: C

Explanation:

Virtual Labs is the Veeam Backup & Replication feature designed for running isolated tests of applications and operating-system changes. A Virtual Lab uses VM backups to create an on-demand, fenced-off environment in which virtual machines can be started without affecting the production network. Its networking configuration, including the virtual lab appliance and isolated network settings, enables the recovered VMs to communicate with each other as needed for application validation while preventing conflicts with production systems. This makes it suitable for applying operating-system updates, rebooting systems, and confirming that dependent applications continue to function correctly before changes are made in production. Virtual Labs are used by SureBackup jobs, which automate recovery verification by starting VMs from backup files and performing configured tests such as heartbeat, ping, and application-level checks. The engineer can therefore validate both the update process and resulting application availability in a repeatable non-production environment. See the Veeam User Guide documentation for Virtual Labs and SureBackup jobs.

QUESTION NO: 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 for no VM snapshots during transit.

B. The proxies allow replication automatic restart after failure.

C. The proxies allow data compression during data transit across the WAN.

D. The proxies allow automatic WAN acceleration.

ANSWER: C

Explanation:

The proxies allow data compression during data transit across the WAN. In a replication job spanning two sites, Veeam backup proxies perform the data-movement work at the source and target locations. The source-side proxy reads and processes VM data, while the target-side proxy receives the transferred data and writes it to the replica infrastructure. When network traffic compression is enabled for the job, the source proxy compresses the transferred data before it crosses the WAN and the target proxy decompresses it after receipt. This reduces the volume of traffic sent over the intersite link, which can lower bandwidth consumption and improve replication efficiency, especially where WAN capacity is constrained. Locating a proxy in each site also keeps the primary data-processing operations close to the relevant production and disaster-recovery resources, avoiding unnecessary traversal of the WAN for local processing. Veeam documents that backup proxies are responsible for retrieving, processing, and transferring VM data, and that the traffic-compression setting controls compression of data transferred between source and target proxies. See Veeam Backup Proxies and Replica Job Network Settings.

QUESTION NO: 11

What are the three Veeam Backup & Replication technologies required to verify recoverability of the backup? (Choose three.)

A. Virtual lab

B. Application group

C. Backup proxy

D. SureBackup job

E. Gateway Server

F. Replication job

ANSWER: A B D

Explanation:

Veeam Backup & Replication verifies backup recoverability through SureBackup, which starts VMs directly from their backup files in an isolated environment and performs automated validation. A **Virtual lab** supplies that isolated environment. It uses networking components such as a proxy appliance to create a fenced network, allowing recovered VMs to start and communicate for testing without creating conflicts with the production network. An **Application group** defines the supporting VMs and their startup order when an application requires dependent infrastructure, such as domain, DNS, database, or other services, to be available during verification. A **SureBackup job** brings these components together: it specifies the virtual lab, selects the backup content to validate, optionally uses an application group, starts the required VMs, and runs boot and application-level tests. This automated workflow provides evidence that protected workloads can boot and that defined application services respond as expected, rather than merely confirming that backup data exists. Veeam documents the SureBackup job as the mechanism for testing recoverability in a virtual lab and documents application groups as the source of VMs required to support verification. See Veeam SureBackup documentation and Veeam Application Groups documentation.

QUESTION NO: 12

Which file extensions are associated with Veeam Backup & Replication? (Choose three.)

A. .vbk

B. .vbo

C. .vrb

D. .bco

E. .vlb

F. .vhb

G. .vib

ANSWER: A C G

Explanation:

The applicable backup-chain file extensions are **.vbk**, **.vib**, and **.vrb**. A **.vbk** file is a Veeam full backup file. It provides the baseline backup data for a backup chain and is produced for active full backups, synthetic full backups, and the initial full backup created by a job. A **.vib** file is an incremental backup file used by forward incremental backup chains; it contains the changes captured since the prior backup in the chain. A **.vrb** file is a reverse incremental backup file. In a reverse incremental chain, the most recent restore point is maintained in the full backup file, while reverse-increment data enables restoration to earlier points in time. These extensions identify the core Veeam Backup & Replication image-level backup file types and are the three requested by the question. Veeam documents these file types and their roles in its Backup Repository reference: Backup Repository File Types. The backup-chain behavior, including forward and reverse incremental processing, is also described in the Veeam Backup Methods documentation.

QUESTION NO: 13

Veeam Backup & Replication is currently configured to back up different workloads: Linux physical computers and Windows physical computers.A second copy of all the backups including transaction logs must be created.Which job configuration is correct?

A. Two "periodic copy" backup copy jobs

B. A single "immediate copy" backup copy job

C. Two "immediate copy" backup copy jobs

D. A single "Periodic copy" backup copy job

ANSWER: C

Explanation:

Two "immediate copy" backup copy jobs is correct because immediate-copy processing is designed to transfer newly created restore points to the target repository as soon as they become available. This behavior also allows the copy workflow to protect transaction-log backup data alongside the applicable backup chain, rather than waiting for a scheduled copy interval. The result is a secondary copy with a substantially smaller exposure window and a restore chain that includes the log data needed for application-aware recovery scenarios.

Separate backup copy jobs are required for the Linux physical-computer backups and the Windows physical-computer backups. Veeam treats these physical-machine backup sources according to their agent and workload type, so configuring an immediate-copy job for each applicable source set provides clear, supported processing boundaries and independently managed copy chains. Each job can be assigned its own target repository, retention policy, and copy interval settings while ensuring that new source restore points and their associated transaction logs are copied promptly. Veeam documents the behavior and configuration of backup copy jobs in the Veeam Backup & Replication User Guide: Backup Copy Jobs and describes application-aware transaction-log handling in the Application-Aware Processing documentation.

QUESTION NO: 14

Which recovery verification tests can be executed when verifying both backups and replicas? (Choose three.)

A. Heartbeat test

B. Ping test

C. Application test

D. Memory test

E. Operating System test

F. CRC check

ANSWER: A B C

Explanation:

“Heartbeat test,” “Ping test,” and “Application test” are the recovery-verification tests available in Veeam SureBackup and SureReplica workflows. These workflows start the VM from a backup in an isolated virtual lab or start a replica in an isolated environment, then validate that the recovered workload is usable rather than merely confirming that recovery data exists. The heartbeat test checks VMware Tools or Hyper-V Integration Services heartbeat status to establish that the guest operating system has started and is responsive. The ping test sends ICMP echo requests to the VM, validating network reachability in the virtual lab or isolated replica environment. The application test runs configurable scripts against the VM to validate application-specific functionality, such as a service response, database availability, or a web application check. Together, these tests provide progressively more meaningful evidence of recoverability: the guest is running, it can be reached over the network, and required application behavior can be verified. Veeam documents these verification

mechanisms for SureBackup jobs and identifies the same test types for SureReplica verification. See the Veeam SureBackup verification tests documentation and Veeam SureReplica documentation.

QUESTION NO: 15

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

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 connects to a VMware vCenter server to collect inventory, performance, configuration, alarm, and capacity information for the managed VMware vSphere environment. vCenter Server is therefore a core Veeam ONE data source and provides the virtualization-layer data used by Veeam ONE Monitor, Business View, reporting, and alerting capabilities.

Veeam ONE also connects to a Veeam Backup & Replication server. This connection allows Veeam ONE to retrieve backup infrastructure details and job/session data, then monitor backup operations, identify protection issues, and produce reporting on backup and recovery activity. The Veeam Backup & Replication server is added as a managed server/data source in the Veeam ONE configuration, with Veeam ONE using the server's API to obtain this information.

These two connections complement each other: the VMware vCenter server supplies information about the virtual infrastructure, while the Veeam Backup & Replication server supplies information about the data-protection environment. Veeam documents the supported connection and monitoring configuration in its Veeam ONE deployment guidance and its VMware vSphere monitoring documentation. See Adding Veeam Backup & Replication Servers and VMware vSphere Servers.

QUESTION NO: 16

Which operations can application-aware processing perform during a VMware backup job? (Choose three.)

A. Create application-consistent backups

B. Process transaction logs

C. Run pre-freeze and post-thaw scripts

D. Convert the VM to a physical machine

E. Configure vSphere Distributed Switches

F. Create a replica at the target site

ANSWER: A B C

Explanation:

Create application-consistent backups, **process transaction logs**, and **run pre-freeze and post-thaw scripts** are valid application-aware processing operations. Veeam uses Microsoft Volume Shadow Copy Service (VSS) inside supported Windows guest operating systems to coordinate application quiescence before the VM snapshot is created. This produces an application-consistent restore point rather than relying only on crash-consistent disk data.

Application-aware processing can also perform transaction-log backup and truncation for supported applications, such as Microsoft SQL Server and Microsoft Exchange, according to the configured processing options. Administrators can also use pre-freeze and post-thaw scripts to perform workload-specific actions before and after VSS processing. See Application-Aware Processing in the Veeam Backup & Replication User Guide.

QUESTION NO: 17

Which VMware infrastructure components can be added to Veeam Backup & Replication as a virtual infrastructure server? (Choose two.)

A. VMware vCenter Server

B. Standalone ESXi host

C. VMware NSX Manager

D. VMware Aria Operations server

E. VMware Horizon Connection Server

ANSWER: A B

Explanation:

VMware vCenter Server and a **standalone ESXi host** can be added as VMware virtual infrastructure servers. Adding vCenter Server is the preferred method for environments managed through vCenter because it provides Veeam with centralized inventory and access to the vSphere hosts, clusters, datastores, and virtual machines managed by that vCenter instance.

A standalone ESXi host can also be added directly when a vCenter Server is not present. Veeam then connects to the individual host to inventory and protect the VMs it manages. See the Veeam documentation for adding VMware vSphere servers.

QUESTION NO: 18

In the event a VMware VM configuration file (VMX) is missing and a restore is needed, what Veeam restore process is the best choice for this situation?

A. VM file restore

B. Instant disk recovery

C. Instant VM recovery

D. Virtual disk restore

ANSWER: A

Explanation:

VM file restore is the appropriate process because a VMX file is a VMware virtual machine configuration file, not a virtual disk. It contains the VM's hardware and configuration settings, including references to its virtual disks, virtual hardware parameters, and other properties needed for VMware to register and run the machine correctly. Veeam's VM file restore workflow is designed to recover individual files that belong to a VMware virtual machine from a backup, including configuration files such as VMX files, NVRAM files, VMXF files, and log files. This enables the administrator to retrieve the missing configuration file without performing a full VM restore or replacing the VM's disks. During the workflow, Veeam mounts the backup and provides access to the selected VM files so the required VMX file can be restored to the appropriate datastore location. This is the least disruptive and most targeted recovery method when only the configuration file has been lost. Veeam documents this capability in its VM files restore documentation; additional VMware restore workflow details are available in the Veeam Restore Guide.

QUESTION NO: 19

In Veeam Backup & Replication, encryption can be enabled at rest for the following: (Choose three.)

A. Backup job

B. Tapes in media pools

C. Backup Copy job

D. Backup to tape job

E. Replication job

ANSWER: A B C

Explanation:

Backup job, **Backup Copy job**, and **Tapes in media pools** are the applicable at-rest encryption locations in Veeam Backup & Replication. For a backup job, Veeam can encrypt the backup data written to the repository by enabling backup encryption in the job's storage settings and selecting a password. The encrypted backup chain is protected while stored, and that password is required when Veeam needs to access its contents. Backup copy jobs also support backup encryption independently, allowing copied restore points to be encrypted at their target repository. This is important because the copy can have its own retention and storage target and therefore needs its own protection configuration. For tape, encryption is configured at the media-pool level. Veeam then encrypts data written to tapes assigned to that pool, protecting the data while the physical media is stored or transported. Veeam uses password-based encryption management, so passwords should be securely documented and protected; loss of the required password can prevent recovery of encrypted backup data. See Veeam's Backup Encryption documentation and its tape media-pool encryption documentation.

QUESTION NO: 20

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?

A. Copy modePerformance Tier: Windows REFS, immutability set for 11 daysCapacity Tier: Amazon S3, immutability set for 60 days

B. Copy and move modePerformance Tier: Windows REFS, immutability set for 11 daysCapacity Tier: Amazon S3, immutability set for 60 days

C. Copy modePerformance Tier: Linux Hardened Repository, XFS, immutability set for 11 daysCapacity Tier: Amazon S3, immutability set for 60 days

D. Copy and move modePerformance Tier: Linux Hardened Repository, XFS, immutability set for 11 daysCapacity Tier: Amazon S3, immutability set for 60 days

ANSWER: D

Explanation:

Copy and move mode with a **Linux Hardened Repository, XFS, immutability set for 11 days** as the Performance Tier and **Amazon S3, immutability set for 60 days** as the Capacity Tier meets the stated design goals. The hardened Linux repository provides an on-premises primary backup location with Veeam immutability protection, helping prevent backup deletion or alteration during the configured immutability period. XFS is the supported filesystem for the Veeam Hardened Repository configuration used for this purpose.

Using copy and move mode ensures backup data is copied to the Amazon S3 Capacity Tier and that older backup chains can be moved from the limited on-premises Performance Tier when they are no longer required within the operational restore window. This allows the local repository to remain sized for approximately 21 days while AWS holds the longer-term

backup data required for the full 60-day retention period. Amazon S3 Object Lock immutability for 60 days protects the off-site backup copies for their entire required retention period, providing an immutable recovery source even if the primary environment is compromised. Veeam documents Capacity Tier offload behavior and immutability at Veeam Capacity Tier documentation and hardened repository requirements at Veeam Hardened Repository documentation.

QUESTION NO: 21

What is the advantage to using the Veeam Advanced Data Fetcher during a backup job?

A. Double I/O performance compared to the VMware VDDK

B. Double deduplication performance compared to the VMware VDDK

C. Fewer blocks of data have to be read during the backup job

D. Less impact on the VM during the snapshot consolidation

ANSWER: A

Explanation:

Veeam Advanced Data Fetcher is a proprietary Veeam data-retrieval technology for VMware backup jobs that is intended to improve how a backup proxy obtains virtual-machine disk data. Its key benefit is substantially higher I/O throughput than the standard VMware Virtual Disk Development Kit data path. Veeam documents that Advanced Data Fetcher can provide up to twice the I/O performance of VMware VDDK, helping backup jobs read and transfer VM data more efficiently. This improvement is especially useful where the data-fetching phase is a limiting factor in the backup window. The feature is enabled at the backup-proxy level and is used when the job and infrastructure meet its requirements; Veeam selects the applicable retrieval method during processing. Higher throughput can shorten backup processing time and allow the proxy infrastructure to make better use of its available network and storage resources. See Veeam's Advanced Data Fetcher documentation and the Veeam Backup & Replication v12 What's New guide for feature details and requirements.