

Veeam Backup & Replication 13.1 Release Notes

This document provides last-minute information about Veeam Backup & Replication 13.1, including system requirements, installation and upgrade procedures, as well as relevant information on technical support, documentation, online resources, and so on.

The release version of Veeam Backup & Replication 13.1 is available for download at [veeam.com/backup-replication-download.html](https://www.veeam.com/backup-replication-download.html) starting from July 29, 2026.

If you are upgrading to Veeam Backup & Replication 13.1 from previous versions, please review the upgrade checklist closely before performing the upgrade.

See next:

- [System Requirements](#)
- [Known Issues and Limitations](#)
- [Installing Software](#)
- [Technical Documentation References](#)
- [Technical Support](#)
- [Contacting Veeam Software](#)

System Requirements

For the most up-to-date system requirements and compatibility information, see the Help Center articles linked below.

These pages are maintained continuously and reflect the latest supported workloads, backup infrastructure requirements, and platform compatibility. This section includes:

- [Workloads](#): Describes the operating systems, hypervisors, applications, and workloads that can be protected with Veeam Backup & Replication.
- [Backup Infrastructure](#): Lists the hardware and software requirements for Veeam Backup & Replication components to ensure smooth deployment and reliable operation.

Review them before installing or upgrading Veeam Backup & Replication.

Known Issues and Limitations

VMware

- Backup jobs of Linux machines may fail after upgrading from Veeam Backup & Replication 12 to version 13 when the system contains mount points with encoded spaces in the fstab configuration.
- RHEL 10 cannot be used as a VMware proxy out of the box due to the missing libcrypt.so.1 library.
- Replication of encrypted VMs with datasets is not supported for vSphere v8 or later.
- Replica failback operation forces digest recalculation for both source and target VMs. The quick rollback option will be ignored.
- Debian and Ubuntu-based Linux backup proxies require that DNS names of vCenter Server and ESXi hosts are resolvable from the proxy server. Otherwise, jobs will fail with the “NFC storage connection is unavailable” error.
- Linux-based backup proxies do not support the processing of VMs with virtual disks without dbd.uuid unique IDs in the hot add mode. Normally, such disks may only be created by specific P2V/V2V conversion tools.
- DirectSAN backup mode: Multipathing (MPIO) for Linux-based backup proxies leverages only path failovers and not load balancing. These are limitations of the VMware VDDK, and the distributions supported for MPIO in DirectSAN are listed in the Virtual Disk Development Kit release notes corresponding to your vSphere version.
- Processing virtual disks with the @ symbol in the disk name is not supported by proxies in VMware VDDK-based transport modes.
- A virtual backup proxy server cannot be used to back up, replicate, or copy itself in the virtual appliance (hot add) mode. Jobs configured to do this will automatically failover to the Network processing mode. CBT will be disabled for proxy VMs.
- Virtual Windows-based backup proxy must have VMware Tools installed; otherwise, it will be considered as not running and will never be assigned any tasks.
- VMware vStorage API for Data Protection has limitations preventing the hot add process depending on VM configuration. For a complete list of hot add limitations, refer to [this Veeam KB article](#). With the default proxy settings, should the hot add operation fail, the job will switch to the network mode for a specific virtual disk.
- Hard Disk restore may fail with the “*Restore job failed Error: A specified parameter was not correct: unitNumber*” error when restoring the disk to an SCSI controller slot above 15. To work around this, add a paravirtual SCSI controller to the target VM by editing VM virtual hardware settings with vSphere Client.
- Disk mapping functionality is not supported for IDE, SATA, and NVMe disks in the Hard Disk Restore wizard.
- Restoring and replicating VMs between different ESXi versions requires the VMs’ virtual hardware version to be compatible with the target host.
- Restoring VM with non-standard virtual disk layout (such as converted from VMware Workstation or VMware Server) as thin may fail. To work around this issue, restore these disks as thick.
- Instant recovery of non-VMware Linux machines to VMware is not supported for backups of machines with `mkinitrd` missing or with mount points outside of /
- Virtual disk placement and type cannot be customized during full VM restore when restoring backups produced by version earlier than 6.1.
- Replication jobs may fail if the source or target datastore has special symbols in its name.
- Networkless interaction with Microsoft Windows guests having UAC enabled requires that the Local Administrator (`MACHINE\Administrator`) or Domain Administrator (`DOMAIN\Administrator`) account is provided on the Guest Processing step.
- Guest customization settings are not backed up and restored for Cloud Director VMs.
- The state of the Standalone VM option of Cloud Director is not preserved with the full VM restore.
- Virtual appliance (hot add) processing mode does not support IDE disks. This is by the design of the VMware hot add functionality, which requires SCSI or SATA adapters (SATA hot add support requires vSphere 6 or later).
- Direct NFS Access is not supported for VMs with existing snapshots when VMware quiescence is enabled.
- RDM disks in virtual compatibility mode are skipped during Backup from Storage Snapshot.
- Pre-freeze and post-thaw scripts for Linux do not perform elevation to root (sudo) when networkless processing (VIX) is used.
- The entire VM restore to the original location does not preserve disk IOPs limits, erroneously removing the associated records from the VMX file. However, restoring to another location will keep these parameters.
- Encryption storage policy is not supported for the instant recovery to First Class Disk (FCD).
- Quick Migration will not function if the HotAdd proxy and the Mount server are located on the same host. As a workaround, use a different proxy type (NBD or SAN), or configure the Mount server on another host.

VMware Cloud on AWS

- Only the hot add transport mode is supported due to API limitations.
- All vPower NFS-based functionality is not supported due to platform limitations.
- Networkless guest interaction is not supported due to API limitations.
- Re-IP addressing and file-level recovery for replicas are not supported.
- Only the cold switch is supported for Quick Migration.
- Replica Failback option is not supported.

CDP

- In case of any environmental issues, deploying the I/O filter will fail with the “*Operation is not allowed in the current state*” error returned by the VMware vSphere VIB deployment framework. In most cases, the reason is a DNS issue. Other possible reasons include infrastructure issues such as an expired vCenter EAM certificate.
- CDP policy cannot be created or started with the “*veecdp@REPLICATION was not found*” error in case some storage providers are offline. To fix the issue, refer to VMware [KB76633](#).
- I/O filter cannot be attached to VMs with snapshots. Remove all snapshots from VMs to be protected before creating a CDP policy.
- Setup programs of specific primary storage integration plug-ins leave the CDP service stopped after the installation. Please verify the service state following the installation and start it manually if needed.
- During package installation via UI, the proper variables are not written to ESXi servers in Advanced System Settings. ESXi server reboot is required to resolve the issue.

Universal CDP

- After upgrading Veeam Backup & Replication from version 13.0.1 or 13.0.2 to version 13.1, a rescan of a protection group that contains machines protected by Universal CDP may fail during CDP component installation on the source machine. To work around this issue, reboot the affected source machine and run the protection group rescan again.
- Compared to traditional CDP, Universal CDP currently does not support the following features:
 - Planned failover
 - Failback to the original location
 - Replica seeding when replicating to a Cloud Connect provider

Hyper-V

- The Hyper-V 2022 and higher hosts force the VM configuration version to be updated from 5.0 to the latest supported by the host.
- Restoring VMs, which were backed up from Hyper-V 2016 (or later) hosts in the crash-consistent state, to the Hyper-V 2022 host fails with the “*Writer ‘Microsoft Hyper-V VSS Writer’ is failed at ‘VSS_WS_FAILED_AT_POST_RESTORE*” error due to a bug in Hyper-V.
- Application-aware processing of VMs with Windows guest OS other than Windows Server 2016 and Windows 10 fails with the “*Failed to take in-guest VSS snapshot COM error: Code: 0x80042308*” error. This known Hyper-V 2016 compatibility issue is fixed by updating Hyper-V integration components on the affected guests with [KB3063109](#).
- Application-aware processing of Active Directory domain controllers running on a guest OS other than Windows Server 2016 fails with the “*Failed to create VM recovery checkpoint*” error (32770). To resolve this issue, ensure the latest Windows Updates are installed for the guest OS on the affected VMs.
- Backing up VMs from the Hyper-V cluster in the rolling upgrade is supported. However, RCT will not be leveraged until the upgrade is completed for all nodes and the cluster functional level is upgraded to Windows Server 2016 or later. Note that the VM virtual hardware version must be upgraded to version 8.0 before RCT can be leveraged on the VM.
- Virtual machines with VMPmemController virtual hardware are skipped from processing due to a Hyper-V limitation around checkpointing of such VMs. Additionally, such machines may cause restore operations to the same Hyper-V host to hang on Hyper-V 2016 versions earlier than the 1803 SAC release due to a Microsoft Hyper-V VSS Writer bug.
- VMs with pass-through virtual disks cannot be processed due to limitations in Hyper-V 2016 and later checkpoints.
- Virtual disks consisting of multiple files (such as from virtual machines originally created on Virtual Server 2005) are not supported for processing.
- CPU Type SCVMM parameter is not backed up and restored on Hyper-V VMs.
- SCVMM remote components may fail to update after upgrade to v13 in SCVMM HA configurations. Refer to [this Veeam KB article](#) for details and workaround.
- Off-host backup from deduplicated volume fails if the Data Deduplication feature is disabled on the backup proxy server.
- Restoring a VM into the root folder of SMB share fails. To work around the issue, restore the VM into a subfolder instead.
- Off-host backup of a VM with Shared Virtual Hard Disks is not supported.
- Restoring an Arc-enabled virtual machine will cause inconsistent behavior from Arc Management, if the original machine is still online.
- Azure Local (formerly Azure Stack HCI) environments with Windows Defender Application Control (WDAC) require a supplemental WDAC policy. For detailed instructions, refer to [this Veeam KB article](#). Ensure that the latest version of the supplemental WDAC policy is deployed on Azure Local hosts before upgrading Veeam Backup & Replication components on these hosts.
- PowerShell Direct fails for x86 (32-bit) Windows guest operating systems, causing the Test Credentials operation to hang on these systems.
- Additional CPU resources may be required on SCVMM servers due to the introduction of new communication components, which may increase resource consumption by system processes.
- Hyper-V replication does not preserve the original VHDX schema. As a result, certain Linux distributions based on RHEL, such as Rocky Linux 9.4, may fail to start after replication.

Nutanix AHV

- Under certain circumstances the VirtIO driver injection processing may re-occur for restored foreign hypervisor VMs that already contain these drivers.
- AHV backup jobs acting on backup chains copied to a new repository will maintain chain continuity rather than creating new chains.
- Web UI multi-select for bulk storage container selection in the disk mapping step of disk restore wizards is currently not enabled.
- During instant disk restore all disks from backup are published even if not all of them are selected for restore.
- For custom AHV RBAC file, disk instant restore and disk level restore operations, the RBAC role must contain the entire cluster scope. Otherwise, available storage containers and networks will not be visible.
- The custom AHV RBAC role “Perform Restore To Original” permission with no specific restore targets does not allow disk restore & instant disk recovery to original VM.
- In certain cases ineligible “empty” job containers are visible in the file-level restore, publish disks and export disks wizards in the full console, however out of scope job VMs are not displayed.

- For Web UI RBAC custom restore roles in Prism Central, only cluster and VMs scopes may be selected for backup source. Categories selection combinations are only available in the full console.
- For Web UI RBAC scopes duplicate name in the "Source" column rather than displaying the underlying IP or hostname.
- Migrate and stop instant recovery buttons are not available for custom roles specifying instant disk and VM recovery in the Web UI.

Proxmox VE

- The certificate change on the Backup Server requires a manual restart of the veeam-platform-service-pve service. Refer to this user guide page for instructions.
- The cluster credentials update session shows as failed despite the credentials update being successful.
- The full VM restore process may change VM disk names without changing the bus type and bus number.
- Workers with concurrent tasks set to 1 may turn off and on sporadically if there are 2 or more VMs in a backup job. We do not recommend setting the number of concurrent tasks to less than 4 (default).
- The Proxmox cluster is not visible in the Web UI. Only individual hosts are listed.
- The Web UI shows no warning for corrupted restore points.
- A replication job may fail if the source VM has disks whose size is not a multiple of the block size used by the target storage.
- Per-VM storage mapping is not applied to TPM device placement in replication jobs.
- Red Hat-based Linux VMs of version 10 or later, when restored to Proxmox VE from another hypervisor, may fail to boot. Veeam Plug-in for Proxmox VE uses the x86-64-v2-AES processor type by default, while distributions such as RHEL 10, CentOS 10, and Rocky Linux 10 require x86-64-v3. To work around this issue, change the processor type to x86-64-v3 in the Proxmox VE VM settings before booting the VM.

Scale Computing HyperCore

- When restoring a foreign platform backup to HyperCore the plug-in attempts to inject compatible drivers into the restored VM. As part of this process all restored VM disks are simultaneously attached to the worker VM processing the restore. HyperCore supports a maximum of 30 disks per VM (one disk of course occupied by the worker VM itself). If the restored VM disk count exceeds the number of available worker VM disk slots or if the worker is processing parallel restore tasks with disk attachments exceeding the maximum, driver injection may fail.

oVirt KVM

This section applies to the Red Hat Virtualization (RHV) and Oracle Linux Virtualization Manager (OLVM) platforms.

- When upgrading the Veeam Software Appliance to v13.1, the Veeam Plug-in for oVirt KVM is not deployed automatically. [Deploy the plug-in manually](#) using the [bundle](#) after the upgrade is complete.

Red Hat Virtualization 4.4 SP1

- You cannot back up hosted-engine VMs. You can use engine-backup to create an oVirt configuration backup. Related RHV bug:
 - https://bugzilla.redhat.com/show_bug.cgi?id=1989121
- An RHV host cannot be switched to maintenance mode due to active data transfer being performed. However, no backup operations are actually being processed. To resolve the issue, contact Red Hat technical support. Related RHV bug:
 - https://bugzilla.redhat.com/show_bug.cgi?id=2037057
- For an RHV cluster configured with 2 network adapters, manual configuration of network routing may be required. Additionally, if a SAN network is reachable from the ovirtmgmt network through a gateway, RHV will use the ovirtmgmt network instead of the selected SAN network. To avoid this issue, configure the network so that the SAN network cannot be reached from the ovirtmgmt network. Related RHV bug:
 - https://bugzilla.redhat.com/show_bug.cgi?id=2084186
- By default, Veeam Plug-in for oVirt KVM applies the following deduplication and compression settings to backed-up data:
 - Deduplication: Enabled
 - Data compression level: Optimal
 - Storage optimization: 1MBDue to technical limitations, you cannot change these settings while configuring backup jobs.
- A repeated Health Check of corrupted data returns a green session (the affected disk is skipped from processing).
- Backup Copy exclusions do not accept oVirt KVM jobs or objects.
- If you want to back up a VM configured with an [oVirt KVM Cloud-Init custom script](#), first remove the script from the VM, as it may contain sensitive data (such as credentials and authorised keys) that would appear in Veeam Plug-in for oVirt KVM backup logs.
- Only file-level restores are supported for StoreOnce Cloud Bank.
- When restoring an oVirt VM with the QXL GPU device to the Oracle Linux KVM environment, Veeam Plug-in for oVirt KVM changes the GPU device to VGA.
- When restoring an oVirt VM with the CIRRUS GPU device to the Red Hat Virtualization environment, Veeam Plug-in for oVirt KVM changes the GPU device to VGA.
- If you perform a revert-to-snapshot operation, Veeam Plug-in for oVirt KVM runs a full scan of the VM disks during the next session and creates an incremental backup.

HPE Morpheus VM Essentials

- Backing up VMs discovered by Morpheus VM Essentials (for example, VMs deployed natively through libvirt) is not supported.
- Backup of the HPE Morpheus VM Essentials Manager VM is not supported.
- For VM restores to file-based storage (GFS2, NFS, Ceph RBD, and so on), a worker is required in the same cluster as the restored VM.
- A backup job may fail for a VM with an active backing chain of disks. This issue may occur after a native clone/migration operation in VM Essentials leaving references to different backing devices. Contact HPE Support to resolve.
- In the Backup Job wizard's disk Exclusions window, only the bus type and unit number can be specified. If a VM has disks that share the same unit number on different bus controllers (e.g. SCSI 0:3, SCSI 1:3, SCSI 2:3), they cannot be addressed individually - a selection by unit number applies to all disks with that unit number across controllers, so they are always processed together.
- AnyToHpe restore of a UEFI VM with an IDE disk fails to boot. When restoring a UEFI VM to HPE VME from another platform, a source disk on the IDE bus causes the VM to be provisioned with legacy BIOS instead of UEFI, so it fails to boot ("no boot disk").
- On heavily loaded environments, running multiple parallel restores with VM volumes on HPE Alletra may occasionally attempt to attach a disk to a worker using a device name that already exists (e.g. "target sdX already exists"), which can cause the restore to fail. File-based datastores (e.g. NFS) are not affected.
- If "AUTOMATICALLY POWER ON VMS" in HPE Morpheus cluster settings is disabled for a worker's target cluster, the first worker test will fail. Consecutive worker preparations are fine, affects only the very first test when worker is deployed/redeployed (after upgrade)
- In rare cases, during concurrent restores, attaching a disk from the VM being created by the restore to the worker can fail with "Volume is attached to another server", causing the restore to fail.

Citrix XenServer and XCP-ng

- If the pool coordinator node is migrated from the original node specified in Veeam Backup & Replication managed servers, any further cluster operations are suspended. To mitigate this issue update the pool coordinator in the Veeam console and accept the new coordinator certificate. Potential workarounds include 1) installing a self-signed certificate with all pool host IPs set in Alternative Names to all cluster nodes 2) pre-accepting all pool nodes on VBR in advance. Note improved methods are being researched.
- Storage repositories hosted on SMAPiv3 protocol are not supported for Hotadd due the platform limitation of snapshots only existing in read-only mode. NBD is supported.
- NBD backup transport requires OpenSSL version 3.0.9 (part of XCP-ng version 8.3 LTS from March 2026). Note that NBD transport is not enabled by default and is highly dependent on customer's network configuration. To avoid any potential NBD-related issues it is generally advisable to create a dedicated worker on each host.
- When a backup snapshot is invoked we immediately relabel it using our own naming convention. Under some circumstances the API returns success but the rename operation fails resulting in an orphaned snapshot.
- Xen RAW type disks do not support snapshot operations which make them ineligible for backup. To backup VMs with this type of disk(s) simply exclude any RAW disks from processing using backup job disk exclusions.
- On restore to original location we are unable to preserve the original VM disk ID(s) which has the net result of the next incremental backup effectively creating another full-sized backup point. Note that if disk exclusions were set on the original VM these will need to be reconfigured due to the new disk ID(s).
- Foreign hypervisor or agent backup restores to Xen will fail with "Default template not found" if there is no default VM template defined for the pool. Note to restore default templates run the following command on the pool coordinator - /usr/bin/create-guest-templates
- Due to a known Xen issue attempting to restore a qcow2 format disk to an SMB storage repository will fail but leave disk on datastore which breaks all further storage repository rescan operations. This may be corrected by deleting the disk from the storage repository by directly accessing pool coordinator through ssh.
- Xen Orchestra (XOA) ancillary metadata is NOT restored during normal restore operations
- Xen supports very long VM names which will be preserved on restore-to-original-location operations. Restore-to-new-location will truncate longer VM names to 80 characters.
- If a host NIC is added to a worker which exists in a disconnected stated, worker testing worker will fail with a NO_HOSTS_AVAILABLE error. Remove the disconnected NIC from the worker configuration.
- Hot-Add backup will fail on GFS2 storage even if the worker and target VM reside on same host

Sangfor aSV

- Veeam Plug-in for Sangfor aSV is not installed automatically during Veeam Backup & Replication installation. You need to download it from veeam.com and install it manually.
- Backup of encrypted VMs is not supported.
- Sangfor Cloud Platform groups are not supported as a backup VM scope.
- Linked VMs cannot be restored to the original location, and they are no longer linked after being restored to a new location (they become flattened).
- Restore to the original location is possible only if the original machine still exists.
- Restore to the original location fails for VMs that have network adapters and are in compatibility mode.
- Restore fails for VMs in compatibility mode.
- Windows VMs with 5 or more disks (including CD-ROMs) cannot be restored.
- A VM may be restored to a storage other than the one specified in the restore wizard.
- The maximum number of disks during restore is limited to 63.
- Restore from an Archive-tier SOBR does not allow selecting a target network for the VM.
- Deploying more than one worker at a time is not recommended, as concurrent image upload might fail.
- Worker deployment targeting a tenant is not possible when a Communication Domain has not been created.

Platform9

- Due to the storage format, changed block tracking (CBT) does not work. Incremental backups perform a full scan and compare to calculate changes.
- Restore to the original location is not supported for Platform9.
- VM restore to Platform9 always powers on the restored VM.
- Worker deployment to Platform9 always fails when the worker is configured with a static IP address. Static IP address and DNS server settings are not supported.
- Entire VM restore of Windows VMs to Platform9 from other platforms may encounter issues with VirtIO drivers and fail to boot properly.
- Target storage limitations in the platform restrict restores to the RAW format.

VergeOS

- Spaces in VM names are not supported. During restore, each space is replaced with '.'.

Unstructured Data

NAS Backup

- Symbolic links located on Windows-hosted NFS shares cannot be backed up. This limitation is caused by the Windows NFS server returning input/output errors when processing symbolic link requests.
- Repositories containing unstructured data backups with immutable or GFS-protected restore points do not support operations that change the immutability state of the backup chain. This includes disabling repository immutability and copying backups with GFS restore points from a non-immutable repository to an immutable repository.
- When files are backed up directly from the Windows server, NTFS sparse files are handled as regular files, and thus are inflated during the restore.
- When files are backed up directly from Linux servers, the maximum full file path length is limited to 4096.
- Symlinks within Linux-based SMB shares are not supported and prevent the backup jobs from executing correctly.
- Force removal of SOBR extent (without backup evacuation) requires running the health check on impacted backups twice: after removing the extent and after running the backup job for the first time.

Object Storage Backup

- Restoring object storage backups to an object storage system other than the one they were taken from is only supported when the target system correctly implements all API calls from the S3 standard, as leveraging required S3 API calls that are incorrectly or incompletely implemented may lead to unexpected results including the data loss.
- For Backup from Object Storage, backup and restore to different object storage systems and software versions are only supported when all API calls from the AWS S3 standard are supported. Not following this requirement can lead to unexpected situations, including data loss.

Physical and Cloud Backup

Agent Management

- Active Directory integration and LDAPS connections over port 636 require a domain with a forest functional level of Windows Server 2012 R2 or above.
- File-level recovery in compare mode may request credentials even if the agent was connected using the Deployment Kit.
- Upgrading Debian with the Linux agent installed may cause backup jobs to fail.
- Universal and Domain local groups are not supported as containers for Microsoft Active Directory-based protection groups. Use Global groups instead.
- The processing rate for agent backup jobs may show incorrect values (much higher than actual).
- Password loss protection is not supported for agents backing up data to an object storage repository.
- File-level recovery via a Linux Mount server to a Cloud-native agent machine may fail with the error: "Unable to establish connection with the machine".
- Veeam Agent for Windows in the "Managed by Agent" job type may ignore the setting "Perform full backup when the required extent is offline" in some cases when the backup is targeted to a Veeam Cloud Connect repository backed by a Scale-out Backup Repository.
- The Get-VBRBackupServerDeployerCertificate PowerShell cmdlet has been deprecated and is no longer available. To generate a new deployment kit, use the Generate-VBRBackupServerDeployerKit cmdlet or the corresponding user interface controls.

Microsoft Windows Server Failover Cluster

- Workgroup clusters, multi-domain clusters, and mixed OS version clusters are not supported for Veeam agent backup.
- Only failover clusters with shared disks are supported. CSV (Cluster Shared Volume) based disks are not supported for Veeam agent backup.
- SQL Server AlwaysON Clusterless Availability Groups and Availability Groups based on multiple SQL Server Failover Clusters are not supported for Veeam agent backup.
- NetBIOS and DNS names for all failover cluster nodes must be resolvable from the backup server.
- Failover clusters with the same NetBIOS names are not supported even when joined in different domains.
- Adding a new node into the failover cluster will result in a full backup performed for all shared disks.
- Bare Metal Recovery restore is not supported for shared disks. Such disks will be filtered out and not displayed in the corresponding wizard. To restore the content of such disks, use volume-level recovery or disk export functionality.
- Instant Recovery automatically skips clustered volumes during recovery.

Veeam Agent for Windows

- Certain antivirus applications may cause various issues during the product setup or crash the operating system when file level recovery is initiated. To prevent this from happening, temporarily disable the antivirus for the duration of the product installation and exclude Veeam Agent for Microsoft Windows service and application folders from real-time antivirus monitoring.
- Microsoft Smart App Control feature must be disabled on the machine as otherwise the installation process will be blocked.
- Backup job may fail with the "Cannot create a shadow copy of the volumes containing writer's data. VSS asynchronous operation is not completed" error if file-level recovery is being performed at the same time.
- Backup job cannot wake a computer up from sleep or detect missed backup on mobile devices in Connected Standby and Modern Standby power saving mode.
- Backup job cannot wake a computer up from sleep to perform a retry cycle if "Sleep after" power saving scheme timeout is set to less than 10 minutes.
- Encrypted files and folders are skipped during file level backup with the "Access denied" error. To work around the issue, use volume-level backup.
- Backup job targeted at the Wasabi object storage in the eu-central-1 region and throttled to 15 MB/s may have occasional connection drops or reconnections.
- If a backup job is targeted at the Amazon S3 object storage and the selected bucket resides in US or Canada, make sure that the bucket name does not contain the dot character (.). Otherwise, the backup job will fail.
- "File is ready for archiving" and "Allow this file to have content indexed in addition to file properties" advanced file and folder attributes are enabled on files recovered with the file level recovery.
- Changed Block Tracking driver for file-level backup tracks files of 50 MB and larger. Smaller files are not tracked by the driver and are processed regularly without Changed Block Tracking driver.
- Changed Block Tracking driver is unable to track the volume changes made by non-Windows operating systems while the original Windows operating system with CBT driver was offline. Consider using MFT changed-block tracking with such setup or perform an active full after the changes were made for volumes.
- Size estimation may not work correctly when performing backup to a scale-out repository, deduplication repository or repository with ReFS configuration.
- Restore to the original location via Veeam Enterprise Manager may not work if the original computer and Enterprise Manager server are in different domains.
- During Remote Bare Metal Recovery, BitLocker-encrypted volumes cannot be manually unlocked when using the Custom Mapping option.
- The legacy Agent UI customization mechanism based on registry keys has been deprecated and is no longer supported starting with version 13.1. Please use the new Agent UI Branding and Customization functionality instead.
- The VeeamVolumeCT driver does not support Microsoft Windows BypassIO technology. As a result, volumes protected by change tracking continue to use the standard I/O path.

Veeam Agent for Linux

- [For Debian and Ubuntu] When installing Veeam Agent version 13.0.2.2, you may encounter an error related to invalid (expired) GPG key that Veeam uses to sign its packages. To resolve this issue, download and install the latest version of the `veeam-release` package from [this Veeam webpage](#).
- User interface may appear unreadable in some terminal programs due to using incorrect symbols for pseudographics. To fix the issue in the Kitty terminal emulator, select the **Allow ACS line drawing in UTF** check box.
- If a directory contains file with a long name consisting of 2-byte characters, the content of this directory may not be displayed in the file system tree views of the user interface.
- Modifying layouts on disks under tracking may lead to problems with unloading the Veeam kernel modules.
- [For Veeam Agent for Linux on Power] In command line interface, you can target backup job at a Veeam backup repository only if an active product license is installed.
- Volume-level backup jobs that rely on a device name under the `/dev` directory require volumes to have their corresponding device names under the `/dev` directory staying persistent — for example: `/dev/md-127`, `/dev/dm-1`. Otherwise, the job will back up the wrong volume.
- Directories encrypted with eCryptfs can be backed up only in the following conditions:
 - eCryptfs is unmounted (files will be encrypted inside the backup)
 - snapshot-less mode is used (files will be decrypted inside the backup)
- Network shares (both CIFS and NFS) can be selected as a source only in the snapshot-less mode.
- LVM physical volumes cannot be backed up, as well as RAID members. Only LVM/RAID logical volumes containing file systems can be backed up. If a physical volume is selected, then all logical volumes residing on that physical volume will be backed up.
- Disks partitioned with `cfdisk` version older than 2.25 are not supported.
- File-level backup masks do not apply to directories.
- Backup job type cannot be switched between file and volume levels.
- Very high data change rate during backup may cause snapshot overflow. Should this happen, the job will attempt to retry processing up to 3 times with a larger snapshot.
- Sparse files are backed up as regular files.
- Configurations with degraded LVM mirrors cannot be backed up.
- [For systems with GPT disk layout scheme] To guarantee successful data restore, the backed-up disk must not contain more than 128 partitions.
- File-level restore cannot be performed on kernels versions 4.0.x–4.1.33 due to the bug in FUSE kernel module. To work around the issue, either upgrade the kernel, or perform file-level restore from the Recovery Media, or on another machine.
- File-level restore for files residing on eCryptfs directories does not mount eCryptfs filesystem automatically.
- Concurrent file-level restore sessions from different restore points are not supported because only one restore point can be mounted at a time. To perform a file-level restore from another restore point, unmount currently mounted one first.
- Restore of encrypted devices is not supported.
- If root partition and bootloader reside on software RAID (mdadm), the following operations will have to be performed manually via usage of the tools provided in the Recovery Media:
 - bootloader installation
 - software RAID configuration
 - fstab and grub configuration adjustments
- It is required to switch the agent from "Managed by backup server" mode to "Standalone" mode in order to restore from Veeam Cloud Connect repository.
- Sparse files are restored as regular files.
- If you restore data from an encrypted backup and create a LUKS volume using the product graphical user interface, the product creates a LUKS2 volume. The LUKS2 volume can be decrypted only on Linux kernel version 4.19 or later.
- [For SLES-based systems with BTRFS file system and GNOME desktop environment version 3 or later] When Veeam Agent mounts a backup to the machine, GNOME detects the loop device and automatically creates its own mount point. After the restore process, you must manually remove the mount point created by GNOME, otherwise BTRFS will block the next attempt to mount the backup file.
- Encrypted backups created with versions older than v3.0 cannot be imported to Veeam Agent for Linux directly.
- Instant Restore feature is not available for backups created with versions older than 3.0.
- Databases residing on ASM volumes cannot be backed up. However, they will be prepared by AAIP engine and will be shown in Veeam Explorer for Oracle without possibility to be restored.
- Size estimation may not work correctly when performing backup to a scale-out repository, deduplication repository or repository with ReFS configuration.
- For GPT disks: bootloader will not be backed up if BTRFS root and BIOS boot partition reside on different physical disks.
- Configuration with a swap file residing on a BTRFS subvolume cannot be backed up.
- Bugs and issues in the kernel that affect BTRFS snapshots and send/receive functionality can cause backup jobs to fail.
- Backup progress for subvolumes cannot be properly shown on systems with `btrfsprogs` versions lower than v4.1.

- Large overhead can be observed in the **Transferred** statistics field for backup jobs that include BTRFS that has no subvolumes.
- If the root subvolume (id 5) is more than 50% full, restore of the subvolume may fail.
- If a subvolume is deleted during restore, it is required to remove the corresponding entry from fstab manually. Otherwise, the system will not boot.
- Pseudo file systems (such as `/proc`, `/sys`, `tmpfs`, and such) cannot be backed up in the snapshot-less mode.
- User interface may not show backups that contain localized symbols in their names.
- Using Veeam Recovery Media, you may experience certain rare issues. To learn more, see [this Veeam KB article](#).
- The launch of the recovery UI may take more time if there are DNS or Internet access issues in the network.
- You cannot boot the system from a custom Veeam Recovery Media for Ubuntu 16.04 because this Linux distribution contains outdated drivers for the OverlayFS filesystem.
- Statistics for file-level backup jobs managed by the backup server may be inaccurate.
- A backup job that uses a network folder as a target cannot be created on systems with Kaspersky Endpoint Security 10.
- If an LVM physical volume is selected for backup, all volume groups that fully reside in this physical volume will be backed up.
- [For Debian 13] Veeam Agent installs the latest version of linux-headers-amd64 with the matching kernel and its headers. This can cause the kernel to be upgraded, which might affect drivers or software that rely on a specific kernel version. After kernel upgrade, you may need to rebuild or reinstall any custom kernel modules that you have on the system.
- The **Dedicated snapshot storage** feature currently has [experimental support](#) status and is not supported in Entire Machine Backup mode.
- During bare-metal recovery, if an invalid recovery token is entered before a valid one, a blocking dialog incorrectly reports the certificate as successfully connected and cannot be dismissed. To work around this, reboot and enter the valid token on the first attempt.
- During ISO patching, the product does not verify that sufficient temporary disk space is available before starting the operation. As a result, patching may fail if the temporary directory runs out of free space.

Veeam Agent for Mac

- macOS saves CIFS/SMB credentials in cache. As a result, if credentials have been changed, Veeam Agent for Mac uses obsolete credentials and fails to connect to shared folders. To refresh cached credentials, reboot the Veeam Agent computer.
- Veeam Agent for Mac does not restore **Finder** comments.
- Size estimation may not work correctly when performing backup to a scale-out repository, deduplication repository or repository with ReFS configuration.
- To perform file-level restore, AD network user must be added to the `veeam` group.

Veeam Agent for IBM AIX

- While a backup is mounted for file-level restore within a WPAR, it is not possible to start a backup job on the WPAR's host server.
- [For Veeam Agent for IBM AIX installed in WPAR] If the network connection between Veeam backup repository and Veeam Agent is interrupted during backup or restore operation, you must restart the operation because Veeam Agent for IBM AIX cannot automatically reconnect to Veeam backup server.
- The "Transferred" value in Text-based User Interface File Level Restore sessions is always displayed as "0 B", regardless of the actual value.
- Restoring multiple large files to a target with insufficient disk space may end with a warning, "Failed to restore some files" instead of an error. Incomplete files may appear in the target directory; verify file sizes after restore.
- [For bare metal recovery with Recovery wizard] If you restore from a backup file residing in a repository managed by Veeam backup server, regardless of the user role, you must explicitly allow them access to that repository in Veeam backup console. For more information, see [Setting Up User Permissions on Backup Repositories](#).
- Disks are not automatically assigned to WPARs after bare-metal recovery.
- Paths longer than 1006 symbols cannot be restored during bare metal recovery and must be restored manually afterwards.
- After bare metal recovery, you may need to reset, or close and open the HMC terminal to resolve any possible issues with its behavior.
- If you performed bare metal recovery from a backup located in an SMB (CIFS) network shared folder, some volume groups may become unavailable after reboot. For details, see [this Veeam KB article](#).
- During bare-metal recovery, access (atime), modified (mtime) and changed (ctime) timestamps are reset for all non-empty files or directories.
- [For bare metal recovery of IBM cloud machines] After bare metal recovery, non-rootvg volume groups may not be automatically imported after reboot. You must import such volume groups manually.
- [For flexible mapping] In certain cases, when you restore rootvg and also restore a logical volume to an existing volume group on the current system, the target volume group can get from the recovery environment a major device number that differs from its original major device number on the current system. Mismatching of the major device numbers can cause an infinite boot loop on the recovered system.
- [For flexible mapping] Volume groups marked for destruction on the current system are not listed at the Recovery Summary step of the Recovery wizard.
- Veeam Agent for IBM AIX can perform file system indexing only with the **mlocate** utility version 0.26-1 that comes with the product installation package.
- Veeam Agent for IBM AIX does not back up extended attributes of backed-up files and directories.
- Veeam Agent for IBM AIX does not support JFS2 snapshots.
- Veeam Agent for IBM AIX does not support sparse files.
- If the overall source data size significantly increases in comparison to the size of the data backed up during the previous backup job session, Veeam Agent for IBM AIX will create an active full backup.
- Restoring only from file-level backups created by Veeam Agent version 4.0 or later is supported.
- Bare-metal recovery to dissimilar hardware is not supported.
- Bare-metal recovery from backups created from within WPARs is not supported.
- Bare-metal recovery requires a backup that includes the full contents of the root directory (the `--includedirs` option must be set to `/`).
- [For flexible mapping] To avoid possible OS inconsistency after restore, rootvg must be restored as a whole volume group.
- [For flexible mapping] Logical volumes can be restored only to their local copy on the current system.
- [For flexible mapping] During restore, Veeam Agent ignores mirrored copies of logical volumes.
- GFS retention flags are not applied to backups stored in object storage repositories. Restore points that do not have stored GFS flags expire according to the regular retention policy, regardless of the configured GFS schedule. To ensure that GFS retention is applied to backups in object storage repositories, contact Veeam support for assistance.
- If Veeam Agent for Unix package is already installed, `veeamtransport` cannot be installed via certificate-based connection. To work around the issue, uninstall Veeam Agent for Unix, it will be reinstalled automatically when deployed from the Veeam backup server.
- Bare-metal recovery of machines running IBM AIX 7.3.4 in IBM Cloud is currently not possible.

Veeam Agent for Oracle Solaris

- Veeam Agent for Oracle Solaris does not back up extended attributes of backed-up files and directories.
- Veeam Agent for Oracle Solaris backs up and restores sparse files as standard data files.
- Veeam Agent for Oracle Solaris backs up and restores file system snapshots as standard data files.
- If the overall source data size significantly increases in comparison to the size of the data backed up during the previous backup job session, Veeam Agent for Oracle Solaris will create an active full backup.
- Bare metal recovery to dissimilar hardware is not supported.
- Bare metal recovery of SVM configurations is not supported.
- Bare metal recovery from backups created from within non-global zones is not supported.
- Bare metal recovery only from file-level backups created by Veeam Agent version 4.0 or later is supported.
- Bare metal recovery requires a backup that includes the full contents of the root directory (the `--includedirs` option must be set to `/`).
- During bare metal recovery, content of ZFS volumes is not recovered.
- [For Oracle Solaris on x86 architecture] After booting from Veeam Recovery Media, the system that has only a VMXNET3 network adapter may receive only an IPv6 address and not an IPv4 address.
- For file system whose encryption has a keysource unsupported by Veeam Agent, Veeam Agent will offer to specify a passphrase to keep the encryption after restore.
- [For flexible mapping] You cannot map a ZFS pool that resides on a disk slice.
- [For flexible mapping] You cannot map the system **rpool** to multiple disks.
- [For flexible mapping] Any special ZFS pool configurations, such as RAIDZ, mirroring, cache devices, and so on, are ignored. All the disks assigned to the source ZFS pool are simply added together to form the target ZFS pool.
- If Veeam Agent for Oracle Solaris is installed in a global zone on Oracle Solaris 10, it is not installed automatically into non-global zones.
- If Veeam Agent for Oracle Solaris is installed in a global zone, store path in non-global zones will not be excluded automatically.
- Veeam Agent for Oracle Solaris is unable to perform file system indexing if multiple **mlocate** tools are present in the system.
- [For Oracle Solaris 10] If the network connection between Veeam Agent for Oracle Solaris and Veeam backup repository is interrupted during a backup or restore operation, the operation must be restarted because Veeam Agent for Oracle Solaris cannot automatically reconnect to Veeam backup server.
- Paths longer than 1006 symbols cannot be restored during bare metal recovery and must be restored manually afterwards.
- During bare metal recovery, access (atime), modified (mtime) and changed (ctime) timestamps are reset for all non-empty files or directories.
- After bare metal recovery, you may need to reset, or close and open the OVMM terminal to resolve any possible issues with its behavior.
- [For flexible mapping] ZFS pools marked for destruction on the current system are not listed at the Recovery Summary step of the Recovery wizard.
- GFS retention flags are not applied to backups stored in object storage repositories. Restore points that do not have stored GFS flags expire according to the regular retention policy, regardless of the configured GFS schedule. To ensure that GFS retention is applied to backups in object storage repositories, contact Veeam support for assistance.

Direct Restore to Amazon EC2

- Direct restore of disks larger than 5TB requires using a proxy appliance; otherwise, restore will fail with the *"Object is too large"* error.
- Restore of AWS-native VMs (originating from AWS) with boot type UEFI can have the wrong boot type and may not boot properly. To restore a machine, use the AmazonDefaultWindowsImage registry key with specifying AMI, which will support the EFI boot type.

Direct Restore to Microsoft Azure

- Specific Linux computer configurations may require Azure VM configuration to be adjusted upon the restore to Microsoft Azure. If your VM fails to boot, contact Veeam Support for assistance.
- The VM name, VM group name, and VM size are not validated for compatibility with the Microsoft Azure naming policy and storage account type and may cause the restore to fail.
- After upgrading to Veeam Backup & Replication 13.1, existing Azure Restore Proxies must be removed and redeployed.

Instant Recovery to Microsoft Azure

- If you upgraded from a version earlier than 13.0 and previously added an Azure Compute Account using an existing Microsoft Entra ID application with limited permissions, update the application permissions as described in the product documentation to enable the new Instant Recovery to Microsoft Azure functionality.

Direct Restore to Google Cloud

- Restore from points located in the Archive Tier backed by AWS S3 Glacier is not supported.

Windows File Level Restore

- To perform file-level recovery from a backed-up ReFS volume Windows mount server is required.
- File-level restore may fail if a VM you are restoring from lacks free disk space at the time of backup.
- Storage Spaces volumes are not supported for file-level recovery. Consider using Instant VM Recovery to recover guest files from such VMs. Note that Microsoft does not support Storage Spaces within a VM.
- To restore files from deduplicated volumes, ensure the mount server and backup console are installed on Windows Server with the Data Deduplication feature enabled and the Windows Server version is the same or greater than the version of the VM you are restoring from. Otherwise, the deduplication driver incompatibility will cause file-level recovery errors with false data corruption reports.
- Windows file-level restore over PowerShell Direct fails if the source VM has any backup job with the Use Persistent Guest Agent option enabled.

Multi-OS File-Level Restore with Mount Host

- The mount host kernel must support the source file system from the backup. Otherwise, the mount operation may result in kernel panic in some cases.
- LVM snapshots are not supported.
- For restores from the ZFS pool, the mount server kernel must support ZFS and have the zfsutils-linux package installed.
- Restore from BTRFS volumes is only possible with the mount host that is not the original host because there cannot be multiple BTRFS volumes with the same UUID attached to the same host.
- Restoring to the original Windows host requires the Deployment Kit to be pre-installed on the target host, a configured Mount Server, and the target host must be part of the same domain with Kerberos authentication configured.
- XFS devices created in Rocky Linux 10 or RHEL 10 cannot be mounted to the software appliance backup server. This is due to the newer XFS version used in these operating systems, which includes features that are incompatible with the current XFS version on the backup server.
- UFS file systems cannot be mounted using FUSE.
- The file-level restore wizard does not filter out backups that have no restore points.
- The registry key that makes the restore Reason field mandatory does not work correctly in the file-level restore wizard.

Multi-OS File-Level Restore with Helper Appliance

- Legacy Logical Volume Manager version 1 (LVM1) volumes are not supported.
- Encrypted LVM volumes are not supported.
- The helper appliance fails to deploy if the backup contains mirrored Linux LVM partitions.
- Spanned, striped, mirrored, and RAID-5 Windows dynamic disks are not supported. To work around the issue, use Windows File Level Restore instead.
- Non-standard file system configuration support is limited (for example, configurations where the file system journal is located on another volume, separately from the actual file system, are only supported for the ext3 file system, but not for other file systems).
- For Cloud Director (vCD) backups, Linux file-level restore does not show the previously configured helper appliance settings in the customization window.

Replica Failover

- Starting a replicated VM using means other than the product user interface (including vSphere Client, Hyper-V Manager, SCVMM, and PowerShell) disables the advanced replication functionality such as Re-IP and failback.

Guest File System Indexing

- File ownership data is not collected for files on non-NTFS volumes and for guest files of Hyper-V VMs.
- Indexes cannot be published if there are locales or Unicode characters in the file system path.
- On Linux guests with an NSS file system, guest file system indexing fails when using the new veeamupdatedb binary.
- Guest file system indexing with the new veeamupdatedb binary might take longer than with the previous binary.

Malware Detection

- Granular exclusions for platform services in the Web UI are available as an Experimental Preview feature.
- Microsoft Windows Defender, ESET (Linux version) and Sophos do not support the Secure Restore option to stop antivirus scanning after the first virus has been found, so the entire volume will always be scanned.
- ClamAV returns a warning when attempting to scan archive files larger than 105,906,176 bytes.
- ESET (Linux version) does not provide proper scan output log.
- Sophos is unable to scan files when the resulting file path after mount operation exceeds 260 characters.
- Malware autoscan is not supported for guest operating systems with BitLocker enabled or other encrypted disk types which cannot be decrypted automatically on the mount host.

Veeam Threat Hunter

- The mount host that will perform backup scanning requires persistent connection to the signatures update server <http://avupdate.veeam.com/av64bit>. At the same time, backup server requires persistent connection to the endpoint <https://vbr.butler.veeam.com/v1/get-token> to authorize signatures update process. If connection is lost for more than 14 days, Veeam Threat Hunter functionality will stop working.
- First signatures update on the mount host will take more time in order to build initial threat database.
- Veeam Threat Hunter scanner might (and most likely will) be blocked by existing antivirus software on the mount host. Read <http://www.veeam.com/kb1999> and configure AV exclusions accordingly.
- Updates server endpoint certificate validation might fail in IPv6-only deployments, leading to a lack of feature functionality.

Role-based access control (RBAC)

- It is strongly recommended to reopen the Console if global permissions for a role have been modified. Minor changes may take up to 15 minutes to take effect.
- Replicas are visible but unavailable for use.
- Backups located on cloud repositories are not shown to users with custom roles but such repositories may remain visible on the tenant side.
- If an available backup contains a VM that is also included in another unavailable backup, all restore points for that VM may be accessible in the Instant Recovery and Entire Recovery wizards.
- Users with the Manage Restores permission cannot cancel an active Publish Disk session.
- In the restore wizards, the original location may be pre-filled with destination that is outside of the permitted restore scope for the user.
- Encryption passwords are not accessible to users with the Manage Restores permission.
- Custom operators may not retrieve all search results when using the search bar in the Backups node.
- Standalone agent backups may not appear in restore wizards for users with custom operator roles; however, these backups should be visible in the Backups node.
- Custom roles with Inventory scope limited to specific VMware objects (VM, pool, or folder) do not show those objects in the VMware hierarchy of the Web UI backup job wizard. Workaround: use search, a different dropdown view, the Folder view (pools and folders), or the Remote Console.
- Users with a custom role see only the sessions they started; sessions started by other users with the same role are not visible.
- The Web UI Custom Role wizard does not flag already-selected scope objects that were later removed or became unavailable.
- Backups on the Capacity Tier do not inherit the job's ACL, and their Permissions tab is not greyed out.

SureBackup

- The automatic virtual lab configuration is not supported for networks with non-private network addresses.
- The automatic virtual lab networking configuration process may fail with the *"Unable to resolve default network settings"* error. To work around the issue, return to the wizard and try again.
- The automatic virtual lab networking configuration may sometimes fail when DVS are present in the virtual environment. In such cases, use the advanced configuration mode to manually set up the virtual lab networking.
- SureBackup job fails on VM with unsupported or excluded virtual disks that were not explicitly set to be removed from configuration (as a part of disk exclusion settings in the backup job) because the test VM cannot find its disks and cannot start.
- Automatic physical mode RDM disk exclusion in the backup job may lead to a situation where the test VM can connect to the RDM disk and make irreversible changes on the disk. To avoid this, always explicitly exclude the physical RDM disk from the backup job, selecting the option to remove the excluded disks from the configuration.
- Some antivirus applications are known to cause BSOD on the backup repository server when the SureBackup job is started. To prevent this, exclude backup folders from monitoring.
- The SureBackup 'SQL Server' role is not supported on Linux-based Veeam Backup & Replication servers.

Backup Copy

- After upgrading to the latest version, a backup copy job may produce a larger-than-expected incremental restore point if the primary backup job has recently created an active full backup. This is a one-time size increase caused by improvements to the data hashing logic.
- Backup copy jobs no longer support selecting workloads from "from infrastructure."
- In the immediate copy mode, the backup copy job processes only the latest backup files chain. To make such jobs copy all existing backups, create BackupCopyMirrorAll (REG_MULTI_SZ) registry value under the HKLM\SOFTWARE\Veeam\Veeam Backup and Replication key on the backup server. This value should be populated with Backup Copy job names.
- A synthetic full backup is not created according to the backup copy job schedule for Cloud Director workloads if the primary backup job did not create any restore points for the same date.

Cloud Connect Backup

- Cloud Connect functionality is not supported in the Web UI, use the Remote Console instead.
- Microsoft SQL Server, Oracle, and PostgreSQL transaction log backup to a cloud repository is not supported. However, Backup Copy jobs in the immediate copy mode are supported for copying transaction logs backups to a cloud repository.

Cloud Connect Replication

- Cloud replica status is not refreshed in real-time but rather periodically. You can press F5 to retrieve the most current state.
- vApp-level networks are not supported for the network mapping functionality (only org-level networks).
- For non-Windows VMs, guest network settings cannot be detected automatically. Because of that, network mapping in the replication job wizard must be performed manually. Additionally, if you are replicating non-Windows VMs only, you need to manually specify the default gateways for each production network by using the Manage Default Gateways dialog on the Service providers node of the Backup infrastructure tab of the management tree.
- Network extension functionality requires static IP addresses assigned to the processed VMs. Dynamic IP addressing via DHCP is not supported, and DHCP must be disabled.
- The failover plan session cannot be viewed once the failover has been performed.
- Enabling VM auto-import in Cloud Director settings may impact replication functionality.
- Using the move functionality of Cloud Director to move cloud replicas may cause virtual disk loss.
- The Remote Console and Remote Desktop features are not supported for Cloud Director-based tenants.
- If a cloud provider renames the Virtual Switch in Hyper-V and updates the Hardware plan, network mapping will stop working for this Hardware plan. The replicas will be in the "Not connected" status, and failover will not work.
- During the execution of the CDP Test Failover Plan, configuration changes to the source VM, such as adding or removing a disk or similar, are not replicated to the target replica VM.

NFS/SMB Share Repository

- Immutability period extension may fail with 'Access is denied' on WORM repositories backed by Grau Data.

Object Storage Repository

- Not all Amazon and Microsoft Azure regions provide archive storage classes used by the Scale out Backup Repository Archive Tier, or the compute resources required for deploying archiver appliances. Consider this limitation when selecting a region.
- Using the immutability feature with the existing S3 bucket containing backups created with Veeam Backup & Replication 9.5 Update 4 requires that both Versioning and Object Lock are enabled on the bucket simultaneously before the immutability feature is enabled. Any other approach will lead to backup offload failures and the inability to interact correctly with backups in the bucket.
- Data in an object storage bucket/container must be managed solely by Veeam, including retention and data management. Enabling lifecycle rules is not supported and may result in backup and restore failures.
- Automatic bucket provisioning is not supported by S3-integrated object storage systems that implement load balancing at the smart entity level. Please consult with your storage vendor to determine whether they recommend enabling this option.

Scale-Out Backup Repository

- Rebalancing the Scale-Out Backup Repository with the Performance Tier comprising immutable or object storage repositories is not supported. The first run of SOBR rebalancing may skip Microsoft SQL log backups and the related image-level backups. These backups will be relocated on the next run.
- Health check for backups located in a SOBR consisting of Dell Data Domain or HPE StoreOnce extents is not supported.
- When a Veeam Data Cloud Vault Archive repository is used as the Archive Tier of a scale-out backup repository and multiple gateway servers are configured, archive processing may use only one gateway server selected from the configured gateway list. Other configured gateway servers are not used, which may prevent the expected throughput scaling. To avoid unpredictable gateway selection, configure a single preferred gateway server with sufficient network bandwidth and resources.

Application Backup Repository

- Application backup repository backup copy job may permanently fail with 'Source snapshot is already present in target dataset' if the network connection is interrupted during a snapshot transfer.

Dell Data Domain

- Storage copy between Dell Data Domain repositories may fail if the source repository stores backup chains in forever incremental mode.

HPE StoreOnce

- Target repositories running HPE StoreOnce version lower than 5.2 don't support the BLAKE3 hashing algorithm. Health check for the jobs pointing to such repositories is skipped.

IBM FlashSystem

- After upgrading to v13.1, snapshot creation may fail with 'CMMVC6035E The action failed as the object already exists' if the IBM FlashSystem Plug-in for Veeam Backup & Replication is not updated to the latest version (2.4.2).

Veeam Explorer for Storage Snapshots

- Multi-home restore is supported for snapshots created by jobs only, not the ones created during a rescan operation. For other storage snapshots, only disks residing on the same datastore with the VMX file will be available for all restore types.

Tape

- After upgrading, all database plug-in backups stored on repositories associated with tape jobs are automatically included in subsequent tape processing.
- When a tape job containing database plug-in backups runs for the first time on a day scheduled for a synthetic full, an additional synthetic full may be created on a subsequent run.
- The same backup on tape cannot be restored simultaneously to multiple targets. Concurrent restore sessions from the same backup may result in a restore operation deadlock.
- NAS Backup to Tape jobs may fail when the root directory of an NFS share contains file names that use unsupported character encodings.
- NAS Backup to Tape jobs may fail during inventory processing if the protected NAS share contains both a file and a folder with the same name in the root directory.
- Changes to tape libraries may not be displayed immediately in the Tape Infrastructure view. Refresh the view manually to display the current library list.
- The first incremental file to tape backup after upgrading to version 13 may take significantly longer than usual. This behavior is expected and occurs only for the initial incremental job following the upgrade.
- Source backup chain upgrade for VM backups will cause Backup to Tape to process the upgraded chain from scratch to ensure the data and metadata consistency on the tape media.
- File to Tape job fails to build a list of files to process if the catalog contains files with certain Unicode symbols.
- Backup to Tape job will perform a full backup during each run if the source forever forward incremental backup job or backup copy job in the immediate copy mode has a retention of less than 3 restore points or if the source backup copy job in the periodic copy mode has a retention of less than 4 restore points.
- SQL, PostgreSQL, and Oracle transaction log backup to tape are not supported.
- If you manage several tape libraries with the same Veeam backup server and use barcodes to identify tapes in these libraries, all barcodes must be unique across all tape libraries.
- File-to-tape operations from SMB source may fail during reconnects when backing up large amounts of data.
- The restore wizard in the UI may freeze or crash if there are millions of objects in the root of a volume, share, or bucket.
- Tape job progress may display incorrect information if retries occur within the job or if the initial enumeration takes an extended time, even though data is being processed correctly to tape.
- Occasionally, backup files may not appear in the tenant-to-tape backup properties.

Applications

- MongoDB backup fails if there are hidden nodes in the replica set. To work around this issue, alter the Hidden node member priority to 0, or uncheck the node in the list of nodes for backup components installation in the Protection Group wizard.
- The Mount server will be used as a gateway server if application plug-in backups are pointed to an object storage repository and the direct connection mode is selected in the repository settings.
- The plug-in backup job will fail if the plug-in is running a version lower than 13.0.2 and the job targets an encrypted repository, or if encryption is enabled in the application backup policy settings.
- After performing a switchover or failover in a VBR High Availability cluster, the Entra ID Tenant backup job may fail with the following error: "The number of restore points in the configuration database and in the Microsoft Entra ID backup repository is different. Rescan the repository to proceed." To resolve this issue, rescan the Entra ID repository.
- Database recovery performed by Veeam Explorer can be retried up to 20 minutes after the failover in the VBR High Availability cluster is completed.
- The application backup policy will fail after a failover in a VBR High Availability cluster if a Protection Group rescan is not performed.
- The Windows temporary log path for PostgreSQL cannot be set in the web UI. In the guest processing settings of the backup job, only the Linux temporary log path can be configured.
- The "Remove deleted items after option" is disabled and removed from the Advanced settings of Microsoft SQL Server backup policies to prevent potential data corruption.
- When restoring an InterSystems IRIS instance to its original location, the instance on the target ODB server must be stopped manually before the restore is started.
- Removing a storage system from Veeam Backup & Replication makes the InterSystems IRIS restore points based on its storage snapshots unavailable for restore. Volumes that have already been restored may be deleted from the storage system, and the target ODB server will lose access to the restored volumes.

Veeam Kasten Integration

- Manually exported Veeam Kasten backups fail to import to Veeam Backup & Replication with the "*Object reference not set to an instance of an object*" error.
- Exporting backups fail to a backup server fails when domain credentials are used for the location profile.

Veeam Explorer for Microsoft Exchange

- Restoring public folder items from the system "In-Place Hold Items" folder to the original location restores them to the newly created folder with the same name instead of the actual system folder.
- When launching an Explorer outside of the console (using a separate application), MFA code is requested only once for saved users.

Veeam Explorer for Oracle

- Instant recovery of the entire Oracle Data Guard is not supported. However, you can restore individual databases.
- The database restore for Oracle Data Guard with tnsnames.ora and listener.ora located in non-default paths is not supported.
- The database export functionality may fail for databases larger than 1TB in size due to the SSH command limit. If you have such databases in your environment, contact Veeam Support to enable the workaround.
- The database restore may fail if the backed-up Oracle server version and the target server version have different patch levels.
- Oracle Real Application Clusters (RAC) and Oracle Data Guard deployments with snapshot standby option enabled are not supported with OCI-based integration. To work around this issue, use Veeam Plug-in for Oracle RMAN.
- Oracle XE on Linux is not supported.
- Configuration with multiple Oracle versions on the same machine is not supported.
- ASM-based Oracle deployments running in virtual machines with Open-VM-Tools installed are not supported.
- For restore operations using a private key as credentials, the following algorithms are supported:
 - ssh-keygen -t ed25519 (works with passphrase)
 - ssh-keygen -t rsa -b 4096
 - ssh-keygen -t dsa
 - ssh-keygen -t ecdsa -b 521

Veeam Explorer for Microsoft SharePoint

- Modified By field of restored documents is updated with the account performing the restore.
- Restored Issue list items are assigned new Issue IDs.
- Restore of Time Card list is not supported.
- Versioning settings of SharePoint lists are not preserved on restore.
- Restoring Generic List and Pages Library may fail with the "*No content type 'XXX' found in web YYYY*" error.
- Importing Picture Library export may result in IDs being changed for some items.
- Importing Project Tasks list export does not preserve column order.
- Importing SharePoint list export does not preserve Validation Settings.
- Some Rating Settings of Discussion list values are not restored.

Veeam Explorer for Microsoft SQL Server

- Instant database publishing to the SQL Server cluster requires a free drive letter on all cluster nodes according to the number of clustered disks in the backup. Instant database recovery requires twice the number of free drive letters.

- A point-in-time restore with fine-tuning requires all nodes of the same AlwaysOn availability group to be in the same time zone.
- Transaction log backup requires that at least one image-level backup of the SQL Server machine is performed. That means that transaction log backup will not function after the full SQL Server machine restore or for newly appearing databases until the next image-level backup is performed.
- Transaction log backups are not supported for Windows Server 2008 or earlier VMs on Hyper-V 2012 R2.
- SQL Server 2017 and later Graph Tables are not currently supported.

Veeam Explorer for PostgreSQL

- Publish and export operation on community edition is not possible over UI, use PowerShell as workaround.

Backup infrastructure

- Veeam Agent for Linux cannot be installed on Linux backup infrastructure hosts.
- Antivirus software with features like SSL/TLS Filtering may block connectivity for backup infrastructure components (e.g., remote PowerShell or Remote Backup Console). Ensure backup infrastructure activity is whitelisted.
- All registered server names must be resolvable into IPv4 or IPv6 addresses.
- The roles of Mount Server, CDP Proxy, and WAN Accelerator cannot be assigned to Linux servers where Veeam Agent for Linux (VAL) is installed. For more information, please refer to [this Veeam KB article](#).
- SSH private keys generated with DSA are no longer supported. This applies to any feature utilizing SSH private key authentication.
- Storage infrastructure is rescanned every 10 minutes. Perform the “Rescan Storage” operation manually after storage infrastructure changes. Otherwise, Veeam may not “see” newly added volumes immediately.
- All sensitive information, such as user credentials or encryption keys, is stored in the configuration database encrypted with a machine-specific private key of the backup server. Accordingly, a newly installed backup server cannot decrypt such information if attached to the existing database, so any encrypted information must be supplied manually. To work around this, use the configuration backup and restore functionality for backup server migrations.
- The length of VM virtual disk file and configuration file names must not exceed 128 symbols.
- A backup proxy cannot transfer a backup to an SMB-based repository when the share is located on the same server as the proxy. To work around this, create a Windows-based backup repository on the proxy server instead.
- Backup jobs containing VMs with non-Latin characters in their names and using Quantum DXi, Fsas ETERNUS CS800, or Infinidat InfiniGuard as a backup repository fail with an SSH error.
- Repositories backed up by Microsoft Windows Server 2025 and Windows 11 24H2 ReFS volumes with deduplication and encryption enabled are not supported.
- When adding an Veeam Infrastructure Appliance version 13.0 to a 13.0.2 Backup Server, you may see the warning “Failed to synchronize update settings. Error: Failed to invoke SetSettings method.” This does not affect functionality — the appliance remains fully operational.
- After changing/updating the VBR certificate, it is required to restart certain plugin services related to various platforms on the VBR server such as HPE Morpheus VM Essentials and other hypervisors.

Veeam Software Appliance

- Copy Backup, Export to VBK, VeeamZip and other activities may fail to use local or SMB paths as a destination, use backup repository option as a workaround.
- The user name for Veeam Software Appliance is limited to 32 characters long.
- When attempting to join a Windows domain, Veeam Software Appliance may fail to join if the domain forest is not at a Windows Server 2016 level or lacks the required encryption types. Veeam Software Appliance supports the following encryption types: aes256-cts-hmac-sha384-192, aes256-cts-hmac-sha1-96, aes128-cts-hmac-sha256-128, and aes128-cts-hmac-sha1-96.
- Veeam Infrastructure Appliance virtual machines may fail to start on AMD-based hosts and become stuck at the “Bootimg ‘Veeam Infrastructure Appliance’” stage. In some cases, the boot process may terminate due to a kernel panic related to FIPS-compliant entropy generation.
- CHAP authentication for external iSCSI storage is supported only when the storage system uses SHA-256 or SHA3-256 hash algorithms.

High Availability

- Whenever a backup server is promoted to the Veeam High Availability cluster, it is mandatory to reapply the backup configuration to all machines in protection groups with pre-installed agents.
- The console may take longer to update the cluster state after a successful failover whenever the old primary node comes back online. To resolve this, reconnect the console to refresh the cluster state.
- Snapshot operations may fail during application-aware processing if the old primary node comes back online after a successful failover. To resolve this, restart the Veeam Backup Service on the secondary node.
- The backup console may incorrectly prompt for a failover immediately after a successful switchover. To resolve this, wait a few minutes before reconnecting to the cluster IP.
- Automatic update of the secondary node may not retry if interrupted by a power-off, restart, or network disconnect. To resolve this, restart the update on the secondary node using the Host Management Console.
- Upgrading a cluster node may fail at the “Upgrading the backup server database” step. To resolve this, reboot the cluster node.
- The Entra ID tenant backup job may fail after a failover or switchover. To resolve this, perform an Entra ID repository rescan.
- Private fixes installed on a backup server before the HA cluster is assembled are not synchronized to the secondary node. If the primary server has any private fixes installed, install the same private fixes on the secondary node before building the cluster.
- After you change the cluster IP address by editing an HA cluster that was upgraded from version 13.0, the previous cluster IP address stays permanently assigned to the cluster alongside the new one. To resolve this, remove the obsolete cluster IP address manually.
- After a private fix is installed on the primary node, switchover is not blocked until private fix synchronization to the secondary node starts. To perform a switchover, wait until synchronization is completed.
- Synchronized users are not removed from the secondary node after an HA cluster is disassembled. If you need to use the former secondary server as a standalone backup server, sign in with the credentials from the primary node.
- After failover and recovery of a cross-subnet, IPv6-only HA cluster, the console does not reconnect automatically. To resolve this, restart the console.
- After modifying an IPv6 cross-subnet cluster configuration, the PostgreSQL configuration is not updated correctly. As a result, cluster switchover may fail and failover operations may not function as expected.

Setup and Update

- Installation of Veeam Backup & Replication on Microsoft Windows may fail on systems configured with Turkish or other non-ASCII locales due to a limitation in PostgreSQL 17 locale handling.
- Upgrade may fail if the remote PostgreSQL server is not configured with the recommended settings, as advised by Veeam Compliance Analyzer.
- During the installation of the Veeam Software Appliance or Veeam Infrastructure Appliance the Anaconda setup may occasionally freeze at the “Configuring addons” stage mainly on systems with HDDs. Restarting the installation resolves this issue.
- Veeam Software Appliance fails to boot when deployed to a Hyper-V VM with Secure Boot enabled if the VM is not using the “Microsoft UEFI Certificate Authority” Secure Boot template.
- Auto-updates from public repository over IPv6 are not supported, as a workaround use IPv4 connection or mirror update repository to a custom one and distribute updates locally over IPv6 afterwards.
- Updates and hot-fix installation will automatically perform service restart; please ensure you do not have any scheduled jobs within the planned maintenance window.
- Deployment to Microsoft Windows may fail at the PostgreSQL step if the system locale name contains non-ASCII characters. To resolve this, switch to an ASCII-named locale or preinstall PostgreSQL with an ASCII locale.
- Upgrade may fail if the remote PostgreSQL server is not configured with the recommended settings: max_connections = 3000 and max_wal_senders = 0, as advised by Veeam Compliance Analyzer. For details on the configuration file, see <https://www.postgresql.org/docs/16/config-setting.html#CONFIG-SETTING-CONFIGURATION-FILE>.

Veeam Updater

- In the Software Updates Summary email report, a Veeam Infrastructure Appliance (VIA) host appears multiple times if it had several update-install sessions in the last 24 hours.
- After updating Veeam Backup & Replication from 13.0 to 13.1, JeOS (VIA) hosts with the WAN Accelerator role remain in the “Out of Date” state, failing with “*Failed to perform appliance OS upgrade: one or more updates could not be installed*”. Re-run the Host Upgrade wizard for the affected host.

Web UI

- This release of the Web UI provides core management features but does not yet include the full functionality set available in the Remote Console. For the complete experience, the Remote Console can be downloaded directly from the Web UI.
- Backups configured with GFS settings cannot be deleted.
- Oracle guest processing retention option cannot be changed via the Web UI.
- Web UI on Firefox web browser does not refresh access token on expiry.
- During Microsoft Entra ID and Active Directory Restore only the first 100 child objects of a container are displayed. As a result, some objects may appear to be missing even though they exist. Use search or the classic Veeam Explorer to browse all child objects.
- During Microsoft Entra ID and Active Directory Restore DNS records are not displayed correctly. SRV records cannot be viewed, and data under the _tcp and _udp nodes may not be displayed. Use the classic Veeam Explorer to browse DNS records.

Remote Console

- Veeam Backup & Replication does not prompt for a repeat MFA code when re-logging into the remote console or web UI during the cookie grace period of 1 minute.
- If a NETBIOS domain name differs from a fully qualified domain name, the AD browser dialog will resolve the NETBIOS domain name incorrectly when new credentials are added. To work around the issue, fix credentials manually.
- Remote console auto-update functionality is supported starting from version 12.3.2 Patch 1 (build 12.3.2.4165).
- Job filter functionality includes unmanaged Linux agent jobs under both Server and Workstation workload types.
- For group-based authentication, the SAML token issued by the identity provider (IdP) must contain the group name in a claim with the name <http://schemas.xmlsoap.org/claims/Group>. If the IdP does not support custom claim names for group attributes, group-based authentication will not function.

Enterprise Manager

- Adding a Veeam Backup & Replication v12 backup server to Veeam Enterprise Manager v13.1 running on the Veeam Software Appliance is not supported.
- NETBIOS names of backup servers must be resolvable on the Enterprise Manager server.
- Reverse DNS lookup on the Enterprise Manager server must be functional for setting up the self-service recovery delegation scope.
- Instant Recovery to another location will fail with the “*Failed to prepare change storage*” error if the Redirect cache option is used.

PowerShell

- Restores from imported backups residing on an SMB (CIFS) share are not supported through PowerShell.
- The Veeam PowerShell module does not load automatically on Linux systems. To load the module, enter “Import-Module /opt/veeam/powershell/Veeam.Backup.PowerShell.dll” in a PowerShell session.

Configuration Backup and Restore

- Under certain circumstances, some encrypted backups may get disconnected from the corresponding job and appear as Imported. To work around the issue, use the backup mapping functionality to reconnect the job to the backup files.
- Immediately after the configuration restore, Enterprise Manager may show duplicate jobs, and some jobs may be missing. The issue will go away by itself after some time.
- Configuration restore fails to download the .bco file from a network share.
- A configuration backup (.bco) created on v13.0.1 cannot be restored on 13.1 if an Entra ID backup exists.
- Migration from Windows-based Veeam Backup & Replication to the Veeam Software Appliance (VSA) is not supported in 13.1.
- A configuration backup (.bco) created by a tenant does not become immutable when the configuration backup job targets a Veeam Cloud Connect repository backed by a Veeam Hardened Repository.

Veeam Plug-in for AWS

Infrastructure

- If you change the backup appliance instance type to a one with less memory as recommended as a minimum, PostgreSQL used for the local database might not be able to start. To resolve the issue, contact the Veeam Customer Support Team.
- In private deployment mode, initial creation of repositories, which involves operations such as listing buckets and validation, is performed using public endpoints. Once the repository is created, Veeam Backup for AWS switches to using private endpoints exclusively.
- When using AWS Organizations, the number of logs has increased and it is recommended to extend the data volume on the appliance. In case of need of assistance, contact the Veeam Customer Support Team.
- After a backup appliance is migrated from one Veeam Backup & Replication server to another and an existing vault is reassigned to the new installation, backup jobs targeting vaults created under the previous VBR installation fail with an sts:AssumeRole authorization error, because the appliance continues to use the vault IAM role issued by the original installation. As a workaround, remove the affected vaults from the appliance and register them again under the new installation.
- Private network deployment mode cannot be enabled on a backup appliance if it has Veeam Data Cloud Vault storage configured as a backup repository. Error: "Private network deployment cannot be enabled while Veeam Data Cloud Vault storage is added to the backup appliance. To enable it, remove all storage vaults from the appliance."
- Veeam Data Cloud Vault storage cannot be added on a backup appliance that has private network deployment enabled. This is validated on the first step of the Add Vault Repository wizard. Error: "Adding Veeam Data Cloud Vault storage is not supported while private network deployment is enabled on the backup appliance. To proceed with the wizard, disable it in the backup appliance Web UI."
- If a configuration backup being restored to another backup appliance contains storage vaults from the original appliance, the new appliance will not be able to access these vaults due to an impersonation IAM role mismatch. As a result, backup policies, restore operations, and retention tasks that require access to these vaults will fail. To resolve this, remove the affected storage vaults from the backup appliance and add them again using the Veeam Backup & Replication console.

Backup

- Backup of EC2 instances with arm64 architecture containing [product codes](#) (deployed using any product AMIs) is not supported.
- RDS backup: replication of RDS databases with custom parameters (such as TDE, Timezone and so on) within an option group is not supported.
- When you enable image-level backups creating in the RDS backup policy settings, cost estimation of creating and storing in backup repositories is not calculated.

Restore

- In AWS China Regions, inbound access on the default File-Level Recovery browser port (443) is available only for AWS accounts with an ICP (Internet Content Provider) license. If your AWS account does not have an ICP license, you can configure a custom port for the FLR browser instead of 443; after changing the port, you must also update the security group specified in worker instance settings to allow inbound traffic on the new port. For step-by-step instructions, see [Configuring Custom Port for File-Level Recovery Browser in AWS China Regions](#) in the Veeam Backup & Replication User Guide.

REST APIs

- To review the detailed change log and breaking changes in the latest REST API version, see the [Veeam Backup for AWS REST API Reference](#).

Veeam Backup & Replication Integration

- It is not recommended that you add several backup appliances from the same AWS account. This may cause unexpected behavior.
- If access keys of an IAM user are not specified for an S3 standard backup repository, encrypted backups stored in the repository will be displayed as non-encrypted ones (there will be no key on a backup icon).
- If access keys of an IAM user are not specified for an S3 standard backup repository, you will only be able to perform entire EC2 instance restore. All other features described in the [Performing Restore](#) section will be disabled.
- From the Veeam Backup & Replication console, you can only create and delete archive repositories from the Veeam Backup & Replication infrastructure. To perform other management operations, use the backup appliance Web UI.
- From archive backups, you can only restore entire EC2 instances to Amazon EC2.
- If you change the repository encryption settings, for example, change password encryption to KMS key encryption, or vice versa, you need to update the repository settings before the next run of the policy where this repository is specified. To do that, complete either the [Edit Repository wizard](#) or the [Edit Appliance wizard](#) in the Veeam Backup & Replication console.
- Enterprise Manager does not support management of policies created in Veeam Backup for AWS.
- Volume-level restore to the original location is not supported from the Veeam Backup & Replication console.
- File-level restore to the original location is not supported from the Veeam Backup & Replication console.
- File-level restore using cloud-native snapshots is not supported from the Veeam Backup & Replication console.
- You cannot remove backups and snapshots created by Veeam Backup for AWS using the Veeam Backup & Replication console.
- During the restore of a database backup, the Veeam Backup & Replication UI will show 0% in the session log until the restore operation is complete.
- If you reconfigure the backup appliance Web UI SSL certificate, licensing and collection of the appliance data may fail. To resolve this, use the Edit Appliance wizard to accept the new certificate.

Veeam Plug-in for Azure

General

- Additional permissions may be required to perform backup and restore operations. For more information, see the [Permissions Changelog](#) section in the Veeam Backup for Microsoft Azure User Guide.
- During backup policy creation, the **Cost Estimation** step of the wizard may display an incomplete list of protected Azure VMs. To resolve the issue, click **Rescan** at the **Sources** step of the backup policy wizard.
- Cost estimation calculation can be slow when a large number of snapshots is defined.
- The cost calculator may return N/A values for Azure SQL backup policies when using a staging server to process databases. This issue resolves itself after the first policy run.
- The cost calculator may return zero values for backup policies when using containers as a source for backups. This happens when the initial discovery process has not been completed yet. To resolve this issue, wait for 5-10 minutes and re-launch the policy wizard to review the cost estimation data.
- On the Protected Data page, filtering by a specific Virtual Network type may not work as expected.

Infrastructure

- If you enable boot diagnostics with managed storage account on worker instances, Microsoft Azure creates a dedicated storage account. This account will not be automatically removed by the backup appliance after the worker is deallocated.
- In private deployment mode, enabling the DNS zone management feature will create new private endpoints. To work around this issue, either disable private deployment mode first and then re-enable it alongside DNS zone management, or remove all previously created private endpoints before enabling the feature.
- Worker tags specified using parameters as tag values are not assigned to deployed worker VMs if an Azure Policy enforcing tagging rules is enabled in Microsoft Azure.
- Virtual network backup copy is not supported on backup appliances with a configured web proxy.
- The Add Archive Repository wizard in Veeam Backup & Replication may display a Mount Servers step, even though archive repositories are not fully registered in Veeam Backup & Replication and therefore do not use a mount server.

Upgrade

- The **Backup Size** column on the **Protected Data** tab will not display any values until a backup policy protecting the corresponding resource runs successfully at least once.
- After completing the upgrade operation, if you navigate to the Web UI using the link in the Veeam Backup & Replication console, the previous login session might still be active. To resolve the issue, refresh the page.
- When updating a Veeam Backup for Microsoft Azure appliance to a newer build within the same major version (v9 to a higher v9 build), the File-Level Recovery (FLR) component version may not be updated, even if a newer FLR version is available. FLR is updated correctly only when the appliance update involves a major version change.
- When updating a Veeam Backup for Microsoft Azure appliance from v8 to v9, the update completes successfully but may finish with a warning in the Apply step of the update wizard, without a clear description of what the warning refers to. This warning does not indicate an update failure.

Backup

- A backup appliance cannot detect Azure VM snapshots created by other backup appliances. This limitation does not apply when restoring a configuration backup to a new appliance.
- Due to Microsoft limitations, backup policies protecting Windows-based Azure VMs with the enabled guest processing option notify on success even if the PowerShell script fails. To work around the issue, add necessary exceptions to the script.
- It may take up to 10 minutes for Veeam Backup for Microsoft Azure to connect to a staging server when processing Azure SQL Managed instances.
- When backing up Azure SQL Managed instances using a staging server, both the Azure SQL backup and Azure SQL restore roles must be assigned to the service account selected in the policy settings.
- An archive session of an SLA-based backup policy may create a temporary backup of a VM, even if a backup session is scheduled for the same day. This temporary backup will be removed during the daily retention session.
- The virtual network backup policy schedule resets daily at 12 AM, causing backup intervals to restart from that point.
- Obsolete snapshot retention may remove snapshots older than the configured Obsolete Snapshots Retention period even if these snapshots are still within the retention period set by the policy protecting them.

Restore

- When restoring an Azure VM that contains a cloud-init script, the script becomes active during the first boot.
- Veeam Backup for Microsoft Azure ignores Microsoft Azure locks on existing databases when restoring databases to the original location.
- If a service account used by a backup policy for creating VM snapshots is removed, the entire VM and disk restore from these snapshots might fail. To resolve the issue, either manually rescan the region where the snapshots are located or wait for the automatic infrastructure rescan performed every 24 hours.
- Under certain circumstances, the restore speed of large Cosmos DB for MongoDB databases may decrease during the operation.
- After restoring a SQL database to its original location, the backup policy protecting the source database fails to detect it. As a workaround, remove the database from the backup policy, perform a rescan operation, and then re-add the database back to the policy scope.
- When performing a Virtual Machine restore to the original location for specific regions (Australia Central, UK South and all Germany regions), an error may occur in the restore session "Azure cloud exception occurred". To resolve the issue, verify that the Virtual Machine is restored in the region and in running state or contact the Veeam Customer Support Team for assistance.
- If a backup was created with the OS disk excluded from the snapshot, drive letters shown in the File-Level Recovery browser for that restore point may not be user-friendly. This happens because disk-to-drive-letter mapping information is stored on the OS disk, which was excluded from the snapshot, so the original mapping cannot be resolved.
- Azure Files shares in storage accounts with restricted network access may not be synced when creating a backup policy. Restoring Azure Files data to a location with restricted network access may fail even when auto-creation of a private endpoint is enabled.

Configuration Restore

- Connection to the backup appliance using SSH fails if the original user is no longer present in the OS list of users after the configuration restore.
- Under certain circumstances, the repository check performed after the configuration restore may fail when some Azure SQL resources were deleted but are still registered in the configuration backup. To resolve the issue, click **Recheck** at the **Configuration Check** step of the restore wizard.
- When restoring a configuration backup containing numerous virtual network restore points, there might not be enough space on the data disk to start the restore. To resolve this, increase the data disk size in Microsoft Azure and expand the ext4 partition on the data disk.
- During the configuration restore of a backup appliance with private deployment mode enabled, the private deployment settings check correctly shows 'Success' but also displays a misleading warning: 'Private deployment settings have not been changed'.
- When performing a configuration restore of a backup created for a private network deployment, make sure that the target Veeam Backup appliance for Microsoft Azure is configured using private network deployment mode. Under certain circumstances, the configuration restore may fail if the appliance is not configured using private network deployment mode.
- After restoring the configuration of a Veeam Backup for Microsoft Azure appliance onto a second appliance, the restored appliance may be temporarily listed twice in the Veeam Backup & Replication console under Managed Servers. The duplicate entry disappears after refreshing or restarting the console.
- When restoring a configuration backup taken on Veeam Backup for Microsoft Azure v8 onto a v9 appliance, the Veeam Backup & Replication console may display a non-blocking *"Failed to find log record with id"* error message once the restore process completes. The restore itself completes successfully (with warnings) despite this message.

REST APIs

- To review the detailed change log and breaking changes in the REST API v8, see the [Veeam Backup for Microsoft Azure REST API Reference](#).

Veeam Backup & Replication Integration

- When you deploy Veeam Backup for Microsoft Azure from the Veeam Backup & Replication console, you can choose to automatically create the default service account on the backup appliance. If you do, the created account will be associated with the same Entra application as the Azure compute account used for deployment and will share the same permissions. For more information, see the [Veeam Backup for Microsoft Azure User Guide](#).
- [Job and job session reports](#) are not supported for backup policies created in Veeam Backup for Microsoft Azure.

- When removing the backup appliance from the Veeam Backup & Replication infrastructure, you also remove all Blob Storage backup repositories for which credentials are not set. To remove the connected resources from Microsoft Azure, see the [Removing Appliances](#) section in the Veeam Backup for Microsoft Azure User Guide.
- Backups and snapshots of Azure VMs, Azure SQL, Azure Files, Azure Cosmos DB accounts and Azure virtual network configurations cannot be removed using the Veeam Backup & Replication console.
- Azure compute account must be re-added to the backup console if it was registered as an existing one in Veeam Backup & Replication v11 and earlier.
- Backup and snapshot size for restore points created by backup appliances is not displayed in the Veeam Backup & Replication console.
- If you change the secret of the Entra application associated with the Azure compute account, Veeam Backup & Replication fails to connect to the application even after updating the secret on the backup server. To resolve this issue, restart the Veeam Azure Service.
- If you add a backup appliance to multiple Veeam Backup & Replication servers, backup policies may fail with the error: 'Failed to assign a license key to the resource.' To resolve the issue, restart the Veeam Azure Service.

Installing Software

Veeam Software Appliance

Installing Veeam Software Appliance

To install the Veeam Software Appliance with Veeam Backup & Replication v13 or Veeam Backup Enterprise Manager v13:

1. Download the [latest ISO version](#) since it will have the latest available cumulative patch built in.
2. Mount the ISO to a VM or physical machine (or use virtual machine template option) and load the GRUB installer. Select the **Veeam Backup & Replication** or **Veeam Backup Enterprise Manager** option and press Enter.
3. Once the installation is complete, proceed with the host configuration in the initialization wizard.
4. Review and accept the terms of the EULA, licensing policy, required software, and third-party license agreements.
5. Specify the host configuration settings, such as host name, network, and time server. These settings can be changed later in the web UI.
6. Specify the host administrator account, set a password for this account, and configure MFA (multi-factor authentication). Optional: configure a security officer account later in the web UI. You will be able to add new users in the web UI at any time. Finalize the host initialization process.
7. Once the installation and initialization are complete, a link to the Host Management web UI (accessible over port 10443) and a link to the application web UI (accessible over port 443) will be displayed on the splash screen.

Migrating to Veeam Software Appliance

Migration from a Microsoft Windows installation of Veeam Backup & Replication to Veeam Software Appliance is not supported in version 13.1. Customers interested in migrating should remain on the latest Veeam Backup & Replication 13.0 patch release and contact Veeam Support for migration assistance. The migration process is currently available only for Veeam Backup & Replication 13.0.

Upgrading Veeam Software Appliance

The Veeam Software Appliance includes a built-in automated update engine, simplifying the installation of patches, updates, and major versions. Update settings are centrally managed from the backup server and are accessible in the Backup Infrastructure view within the Veeam Backup & Replication web UI, enabling monitoring and control of updates across your backup environment. Although the update engine is preconfigured by default, it is strongly recommended to review and adjust maintenance window settings to meet your organization's operational requirements. Additionally, update configuration and status can be viewed or modified individually on each server using the Host Management interface.

Upgrading Veeam Hardened Repository

The Veeam Infrastructure Appliance allows you to upgrade a Hardened Repository v2 to the latest version using the corresponding option in the GRUB menu. Keep in mind that the upgrade process may take some time.

NOTE

If the Hardened Repository v2 also has the VMware proxy role, you must: Rerun the Linux Host Wizard, click Upgrade on the proxy component and perform a rescan operation.

Installable Software on Microsoft Windows

Installing Veeam Backup & Replication

To install the Veeam Backup & Replication 13 server and management console to a Microsoft Windows machine:

1. Download the [latest ISO version](#) since it will have the latest available cumulative patch built in.
2. Mount the ISO and use autorun or the `Setup.exe` file. Click **Install** in the **Veeam Backup & Replication 13** splash screen.
3. Select the **Install Veeam Backup & Replication** option.
4. Accept the terms of Veeam Backup & Replication and 3rd party components license agreements to install the product. You can find a copy at [veeam.com/eula.html](https://www.veeam.com/eula.html)
5. Provide the setup program with your license file.
6. Review the default installation settings. To change the defaults, click **Customize Settings** and go through additional wizard steps.
7. Click **Install** to start the installation and follow the setup wizard steps.
8. Launch the backup console by clicking the **Veeam Backup & Replication** product icon on your desktop and specify localhost to connect to the local backup server.

Installing Veeam Backup Enterprise Manager

If you want to manage one or more Veeam Backup servers with a centralized management web UI, install Veeam Backup Enterprise Manager. You only need one Enterprise Manager installation per environment.

To install Veeam Backup Enterprise Manager:

1. Mount the ISO and use autorun or the `Setup.exe` file. Click **Install** in the **Veeam Backup & Replication 13** splash screen.
2. Select the **Install Veeam Backup Enterprise Manager** option.
3. Accept the terms of the License Agreement to install the product.
4. Provide the setup program with your license file.
5. Review the default installation settings. To change the defaults, click **Customize Settings** and go through additional wizard steps.
6. Click Install to start the installation and follow the setup wizard steps.
7. Once the installation is complete, access the Veeam Backup Enterprise Manager web UI by clicking the **Veeam Backup Enterprise Manager** product icon on your desktop.

Uninstalling Veeam Backup & Replication or Veeam Backup Enterprise Manager

1. From the Start menu, select **Control Panel > Add or Remove Programs**.
2. In the programs list, select **Veeam Backup & Replication** or **Veeam Backup Enterprise Manager** and click the **Remove** button.
3. In the programs list, select and remove any additional remaining Veeam components.

Upgrading Veeam Backup & Replication

Veeam Backup & Replication 13 supports an automated in-place upgrade from version 12.3.1 (build 12.3.1.1139) or later, which preserves all product settings and configuration. To upgrade from earlier versions, contact [Veeam Customer Support](#).

Use the checklist below to ensure your infrastructure is ready for the Veeam Backup & Replication upgrade. The built-in configuration check mechanism of the Veeam Backup & Replication Upgrade wizard performs some of the checks. Still, you can control them manually before starting the upgrade procedure.

Licensing

1. Veeam Backup & Replication 13 supports the same license file format, so you can use your existing license file to install version 13.
2. Your support contract must be active as of the date when the product build you are installing was built. This is determined by the Support Expiration date in your installed license. If required, you can install a new license during the upgrade procedure.

System Requirements

Important: Customers planning to migrate from a Microsoft Windows installation to Veeam Software Appliance should not upgrade to version 13.1. Stay on the latest 13.0 patch release and contact Veeam Support for migration assistance.

1. Verify that the backup server to be upgraded is running a supported operating system as specified in [System Requirements](#). If not, perform a configuration backup, install Veeam Backup & Replication 13 on a supported OS, and restore the configuration backup created earlier. To learn how to perform the migration, see the [Migrating Veeam Backup & Replication to Another Backup Server](#) section in the Veeam Backup & Replication User Guide.
2. Ensure that the backup server has sufficient disk space. The minimum disk space is calculated dynamically during the system configuration check in the upgrade procedure. It is based on the list of required packages to be installed on the machine and usually is about 9 GB. We recommend allocating at least the minimum disk space calculated during the system configuration check, but not less than 55.5 GB: 3 x ISO size (50.5 GB) in the selected installation path (for example, D:) plus 5 GB for the database operations on the system volume (for example, C:).
3. Make sure that backup infrastructure components meet the system requirements listed in the [System Requirements](#) section of this document. In particular, ensure all backup infrastructure servers are based on 64-bit operating systems.
4. Ensure that the environment you will protect with Veeam Backup & Replication meets the requirements listed in the [System Requirements](#) section of this document.
5. Make sure that all necessary ports are open.

NOTE

Port 443 is now required for Web UI operations. The upgrade process checks its availability and may block the upgrade if the port is in use.

6. Make sure that all necessary permissions are granted.
7. Do you have **legacy per-machine backup format with single metadata file**? Detach or upgrade such backups to a new format to proceed with the upgrade.
8. Do you have **legacy backup copy jobs**? Switch your existing backup copy jobs to a supported mode before upgrading.
9. Do you have **nested repository paths**? Configurations where one repository is located inside another are no longer supported.
10. Do you use **UNC path in export vbk wizard**? Configurations with such settings are no longer supported.
11. Do you have **file copy jobs**? The backup server can no longer be used as a destination for File Copy jobs. Update the configuration for these jobs before proceeding with the upgrade.
12. Do you have **Veeam Hardened Repository v2** added to backup infrastructure? Consider upgrading Veeam Hardened Repository to the latest version using Veeam Infrastructure Appliance.
13. Do you have **FIPS mode** option enabled? Persistent components will be required for guest processing.
14. Do you have **trusted hosts hardening** set to Manual? After the upgrade, open the backup console and make sure all your backup infrastructure servers are trusted.
15. Do you have any jobs using the **transform previous backup chains into rollbacks** option? This option has been removed from the product, and such jobs are no longer supported.

16. The following **features are no longer available** in the product.

Deprecated features:

- Restore point–based retention is no longer available for newly created jobs.
- Reversed incremental backup mode is no longer available for newly created jobs.
- Single-storage backup format option is no longer available in the repository settings.
- Active Directory–based authentication is no longer available for new Veeam Cloud Connect tenants.

Discontinued features:

- Universal Application-Item Recovery wizard (U-AIR).
- Initiating restore by double-clicking VBK/VBM files in Windows Explorer.
- Burning Recovery Media to a CD/DVD/BD media.
- Veeam Cloud Connect Portal

17. Are you upgrading an environment with **backup jobs that run periodic Active Full backups**? If so, please review the following information.

To improve backup stability and portability, compression processing has been reworked. As a result, after the upgrade, environments that perform Active Full backups on primary (source) backup jobs may observe a one-time full-sized data processing impact on secondary operations that rely on those backups.

A one-time Active Full–sized operation may be observed for the following components:

- Backup Copy Jobs. When an Active Full backup is performed at the source, the next incremental run of the backup copy job processes data equivalent to a full backup.
- Backup to Tape from Backup Copy Jobs. When an Active Full backup is performed at the source, the subsequent incremental backup copy run becomes full-sized. As a result, incremental tape backups are written at full-backup size.
- Capacity Tier (Object Storage) – Primary Backup Jobs. When an Active Full backup is performed at the source, the offload to the capacity tier transfers data equivalent to the size of the active full backup.
- Capacity Tier (Object Storage) – Backup Copy Jobs. When an Active Full backup is performed at the source, the subsequent incremental backup copy run is full-sized. Consequently, offloads to the capacity tier are also full-sized.
- Archive Tier (Object Storage) - When GFS restore points are offloaded after an Active Full backup is performed at the source, deduplication efficiency may be reduced, as backup blocks are re-shared during processing. The actual impact depends on the data set and retention configuration.

18. Are you using installations of Veeam Backup & Replication and Veeam Backup for Microsoft 365 on the same machine? First upgrade Veeam Backup for Microsoft 365, second upgrade Veeam Backup & Replication.

19. Are you using Server 2019-based ReFS backup repositories? If yes, avoid upgrading them to Server 2022 or mounting ReFS volumes from Server 2019 to new Server 2022 installations until you read [this thread](#) on Veeam R&D forums. Microsoft has addressed the known regression in the ReFS format upgrade code, and the fix is now publicly available.

20. Are you using Scale-Out Backup Repositories with immutable performance tier extents? Make sure that all extents have the same immutability settings.

21. Are you using a customized AntivirusInfos.xml file? During the upgrade, Veeam Backup & Replication will replace it with the default file. Make sure that you save your customized file at another path and make necessary changes to the default file after the upgrade.

22. Azure compute accounts based on Azure AD user credentials (created with the Use the existing account option) are obsolete. Replace these accounts with new ones to restore workloads to Microsoft Azure, use the Microsoft Azure archive storage or Microsoft Azure Plug-in for Veeam Backup & Replication appliance.

23. Are you using integration with Veeam Backup for Microsoft Azure? If yes, after upgrading to Veeam Backup & Replication 13 and replacing the obsolete accounts from p.21, select the existing Microsoft Azure compute account in the Manage Cloud Credentials, click Edit, and go through the Microsoft Azure Compute Account wizard to update account permissions. Otherwise, you can face problems when adding an external repository with backups created by Veeam Backup for Microsoft Azure.

Integration with Veeam Management and Monitoring Products

1. Are you using **Veeam ONE** to monitor your backup infrastructure? If yes, upgrade it first. Veeam ONE v13 supports monitoring of backup servers version 12 or later.
2. Are you using **Veeam Backup Enterprise Manager**? If yes, start the upgrade procedure with this component. Veeam Backup & Replication should be upgraded after that. If you have a backup server installed on the same machine, upgrade it immediately after completing the upgrade of the Veeam Backup Enterprise Manager server. Otherwise, the Veeam Configuration Database Connection Utility (DBConfig) utility will not work correctly for Veeam Backup & Replication.
3. Are you using the **Veeam Backup Enterprise Manager** server added to **Veeam ONE**? If yes, first upgrade Veeam ONE, second upgrade Veeam Backup Enterprise Manager, and third upgrade Veeam Backup & Replication.
4. Are you using **Cloud Connect**? If yes, consider the following:
 - Check with your Cloud Connect service provider if they have already upgraded their systems to at least the version you are upgrading to.
 - Ensure your Cloud Connect tenants use the supported Veeam product versions. The minimal supported tenant versions are: Veeam Backup & Replication 12.3.2.3617, Veeam Agent for Windows 6.3.2, Veeam Agent for Linux 6.3.2, Veeam Agent for Mac 2.3.1. Tenants running Veeam Software Appliance 13.0.0.4967 are not supported and must be upgraded to 13.0.1 or later.
 - The Veeam Cloud Connect Portal functionality is deprecated and cannot be used anymore.

Integration with Veeam Backup Agents and Enterprise Plug-Ins

1. Are you using **Veeam Agents** managed by Veeam Backup & Replication?
 - If you use **Veeam Agent for Microsoft Windows** or **Veeam Agent for Linux** below 6.3.1, they will stop working after upgrading to Veeam Backup & Replication 13. In this case, we recommend immediately upgrading **Veeam Agent for Microsoft Windows** or **Veeam Agent for Linux** to 13. If you use **Veeam Agent for Microsoft Windows** or **Veeam Agent for Linux** 6.3.1 or later, they will continue working after upgrading to Veeam Backup & Replication 13, but new features implemented in Veeam Backup & Replication 13 will not be supported. In this case, you can upgrade **Veeam Agent for Microsoft Windows** or **Veeam Agent for Linux** to 13 later if the support of new features is not critical for you.
 - If you use **Veeam Agent for Mac** below 2.3.1, it will stop working after upgrading to Veeam Backup & Replication 13. In this case, we recommend immediately upgrading **Veeam Agent for Mac** to 13. If you use version 2.3.1 or later of the agent, it will continue working with Veeam Backup & Replication 13, but new features of Veeam Backup & Replication 13 will not be supported. We recommend upgrading the agent to version 13 at your earliest convenience.
 - If you use **Veeam Agent for IBM AIX** or **Veeam Agent for Oracle Solaris** below 4.6.1, it will stop working after upgrading to Veeam Backup & Replication 13. In this case, we recommend immediately upgrading **Veeam Agent for IBM AIX** or **Veeam Agent for Oracle Solaris** to 13. If you use **Veeam Agent for IBM AIX** or **Veeam Agent for Oracle Solaris** 4.6.1 or later, it will continue working after upgrading to Veeam Backup & Replication 13, but new features implemented in Veeam Backup & Replication 13 will not be supported. In this case, you can upgrade **Veeam Agent for IBM AIX** or **Veeam Agent for Oracle Solaris** to 13 later if the support of new features is not critical for you.

For details about Veeam Agent upgrade, see [Upgrading Veeam Agent](#).

2. Are you using Veeam Plug-Ins for Enterprise Applications (**Veeam Plug-In for Oracle RMAN**, **Veeam Plug-In for SAP HANA**, **Veeam Plug-In for SAP on Oracle**, **Veeam Plug-In for Microsoft SQL Server** or **Veeam Plug-In for IBM Db2**)?

If you use Veeam Plug-Ins for Enterprise Applications below 12.3.2.4165, they will stop working after upgrading to Veeam Backup & Replication 13. In this case, we recommend immediately upgrading Veeam Plug-Ins to 13. If you use Veeam Plug-Ins for Enterprise Applications 12.3.2.4165 or later, they will continue working after upgrading to Veeam Backup & Replication 13, but new features implemented in Veeam Backup & Replication 13 will not be supported. In this case, you can upgrade Veeam Plug-Ins to 13 later if the support of new features is not critical for you.

For details about upgrade of Veeam Plug-Ins for Enterprise Applications, see [Upgrading Veeam Plug-In](#).

3. Are you using **MongoDB Backup**?

If you use MongoDB Backup with Veeam Agents below 12.3.2.4165, they will stop working after upgrading to Veeam Backup & Replication 13. In this case, we recommend immediately upgrading Veeam Agents to 13. If you use MongoDB Backup with Veeam Agents 12.3.2.4165 or later, they will continue working after upgrading to Veeam Backup & Replication 13, but new features implemented in Veeam Backup & Replication 13 will not be supported. In this case, you can upgrade Veeam Agents to 13 later if the support of new features is not critical for you.

For details about upgrade of Veeam Agents, see [Upgrading Veeam Agent for MongoDB Backup](#).

Integration with Storage Systems

1. **IBM FlashSystem**: If you are using IBM FlashSystem Plug-In for Veeam Backup & Replication, upgrade it to version 2.3.77 or later. If the plug-in version is lower than the minimum required, storage integration will not function.
2. **Hitachi VSP**: If you are using Hitachi Plug-In for Veeam Backup & Replication, upgrade it to version 2.2.271 or later. If the plug-in version is lower than the minimum required, storage integration will not function. Note that the integration with Hitachi is available starting from Veeam Backup & Replication version 13.0.1.180.
3. **HPE XP**: If you are using HPE XP Plug-In for Veeam Backup & Replication, upgrade it to version 2.2.271 or later. If the plug-in version is lower than the minimum required, storage integration will not function. Note that the integration with HPE XP is available starting from Veeam Backup & Replication version 13.0.1.180.
4. **NEC Storage V Series**: If you are using NEC Storage V Series Plug-In for Veeam Backup & Replication, storage integration will stop working after the upgrade. Plug-in version with support for Veeam Backup & Replication will be released by the vendor later.
5. Make sure your storage systems work on a supported operating system:
 - **Cisco HyperFlex**: the minimum supported operating system version is v5.0 (2x). Upgrade to it or remove the storage arrays from the backup server configuration. Also note that support for Cisco HyperFlex in this release is [experimental](#).
 - **HPE Nimble**: the minimum supported operating system version is 5.2. Upgrade to it or remove the storage arrays from the backup server configuration.
 - **Dell Data Domain**: the supported operating system version is 7.9 to 8.6. Upgrade to it, or the backup jobs pointed to this repository will fail to start.
 - **ExaGrid**: the minimum supported operating system version is 7.2.0 P08. Upgrade the storage array and Veeam Backup & Replication, open the corresponding Linux host properties and complete the wizard to install persistent data mover components on the storage array. If the same array was added twice under different names in the previous version, complete the wizard for one of the hosts and re-point all backup repositories to it.
 - **Fsas CS800**: the minimum supported operating system version is 5.2.0. Upgrade the storage array and Veeam Backup & Replication, open the corresponding Linux host properties and complete the wizard to install persistent data mover components on the storage array. If the same array was added twice under different names in the previous version, complete the wizard for one of the hosts and re-point all backup repositories to it.
 - **Infinidat InfiniGuard**: the minimum supported operating system version is 3.12. Upgrade the storage array and Veeam Backup & Replication, open the corresponding Linux host properties and complete the wizard to install persistent data mover components on the storage array. If the same array was added twice under different names in the previous version, complete the wizard for one of the hosts and re-point all backup repositories to it.
 - **Quantum DXi**: the minimum supported operating system version is 5.2.0. Upgrade the storage array and Veeam Backup & Replication, open the corresponding Linux host properties and complete the wizard to install persistent data mover components on the storage array. If the same array was added twice under different names in the previous version, complete the wizard for one of the hosts and re-point all backup repositories to it.

Upgrade Process

1. Make sure the latest run for all existing jobs has been completed successfully. Rerun any failed jobs.
2. Ensure there are no running jobs, restore sessions, Instant Recovery sessions, or SureBackup jobs. We recommend that you do not stop running jobs and let them complete successfully.
3. Disable any periodic and backup copy jobs temporarily to prevent them from starting during the upgrade.
4. Disable CDP policies.
5. Ensure there are no active tasks from standalone (unmanaged) agents.
6. Ensure there are no active Veeam Disaster Recovery Orchestrator tasks.
7. Perform the configuration backup, as described in [Running Configuration Backups Manually](#).

Ensure you have configuration backup encryption enabled. Otherwise, stored credentials will not be included in it. For more information, see [Creating Encrypted Configuration Backups](#).

Upgrading Veeam Backup & Replication Server

1. Download the [latest ISO version](#) since it will have the latest available cumulative patch built in.
2. Ensure there are no active processes, such as any running jobs and restore sessions. We recommend that you do not stop running jobs and let them complete successfully instead. Disable any periodic and Backup Copy jobs so that they do not start during the upgrade.
3. Perform a backup of the corresponding SQL Server configuration databases used by the backup server so that you can easily go back to the previous version in case of issues with the upgrade. You can also use the built-in configuration backup functionality.
4. Mount the product ISO and use autorun or run the `Setup.exe` file. Click **Upgrade** in the **Veeam Backup & Replication 13** splash screen.
5. Select the **Upgrade Veeam Backup & Replication** option.
6. Follow the upgrade wizard steps as outlined in the installation procedure above. Be sure to select the same SQL database and instance used by the previous product version.
7. Open the Veeam Backup & Replication user interface. If necessary, the automated upgrade wizard will automatically appear, prompting you to upgrade product components running on remote servers. Follow the wizard to complete the upgrade process.
8. If some remote servers are unavailable at the time of the upgrade, you can run the Upgrade wizard at any time later from the main product menu. Note that jobs cannot use out-of-date product components until updated to the backup server version.
9. Enable any scheduled jobs that you have disabled before the upgrade.

Please note that immediately after the upgrade, backup server performance may be impacted due to the configuration database being optimized by the maintenance job. This can take up to an hour, depending on the database size.

Upgrading Veeam Backup Enterprise Manager

To upgrade Veeam Backup Enterprise Manager to version 13, you must run version 12.3.1 (build 12.3.1.1139) or later on the supported operating system (refer to the [System Requirements](#) section of this document). To upgrade from previous versions, contact [Veeam Customer Support](#).

1. Download the [latest ISO version](#) since it will have the latest available cumulative patch built in.

2. Perform a backup of the corresponding SQL Server configuration databases used by the Enterprise Manager server so that you can go back to the previous version in case of issues with the upgrade.
3. Mount the product ISO and use autorun or run the `Setup.exe` file. Click **Upgrade** in the **Veeam Backup & Replication 13** splash screen.
4. Select the **Upgrade Veeam Backup Enterprise Manager** option.
5. Follow the setup wizard steps as outlined in the installation procedure above. Be sure to select the same SQL database and instance used by the previous Veeam Backup Enterprise Manager version.
6. If you have the Veeam Backup & Replication server installed on the server, upgrade it immediately after completing the upgrade of the Veeam Backup Enterprise Manager server. Otherwise, this local backup server will not be able to run jobs.

Please note that immediately after the upgrade, Enterprise Manager performance may be impacted due to the configuration database being optimized by the maintenance job. This can take up to an hour, depending on the database size.

License

Veeam Backup & Replication can be licensed per protected workload with Veeam Universal License (VUL). For more information, see Veeam Licensing Policy at [veeam.com/licensing-policy.html](https://www.veeam.com/licensing-policy.html)

The trial license key is sent to you automatically after downloading the product. The trial license is valid for 30 days from that moment and includes basic technical support.

To obtain a full license key, please refer to [veeam.com/buy-veeam-products-pricing.html](https://www.veeam.com/buy-veeam-products-pricing.html)

Subscription VUL and Perpetual VUL licenses include a maintenance plan with Premium support. To renew or upgrade your maintenance plan, contact Veeam Renewals at [veeam.com/renewal.html](https://www.veeam.com/renewal.html)

Updating License

The Veeam Backup & Replication server license can be managed centrally by the Enterprise Manager server. If you are using Enterprise Manager, do not update the license on individual backup servers directly, as Enterprise Manager will force its license to all connected backup servers.

To install the new license file to a backup server connected to the Enterprise Manager server:

1. Open the **Configuration > Licensing** tab in Enterprise Manager UI, and click **Install License**.
2. Browse to the license file (`.lic`) sent to you after registration to install the license. To learn more, see the [License](#) section.
3. The provided license file will be automatically propagated and applied to all Veeam Backup servers connected to this Enterprise Manager server.

To install the new license file to a standalone backup server that is not managed by the Enterprise Manager server:

1. Select License from the main menu.
2. Click the Install license button to browse for the license file (`.lic`) sent to you after registration to install the license. To learn more, see the [License](#) section.

Technical Documentation References

If you have any questions about Veeam Software Appliance, Veeam Infrastructure Appliance, Veeam Backup & Replication or any other Veeam products, you may use the following resources:

- Product web page: [veeam.com/products/veeam-data-platform.html](https://www.veeam.com/products/veeam-data-platform.html)
- User guides: helpcenter.veeam.com/
- R&D forums: [veeam.com/forums](https://www.veeam.com/forums)

To view the product help, press the **F1** key or select **Help > Online Help** from the main menu.

Technical Support

We offer email and phone technical support for customers with active maintenance agreements and during the official evaluation period. For a better experience, please provide the following when contacting our technical support:

- Version information for the product and all infrastructure components.
- Error message and/or accurate description of the problem you are having.
- Full debug log bundle. To export the log files, navigate to Veeam Host Management Console select **Logs and Services > Logs** tab and click create support log bundle to export the full set of log files.

To submit your support ticket or obtain additional information, please visit [veeam.com/support.html](https://www.veeam.com/support.html).

TIP

Before contacting Veeam Customer Support, consider searching for a resolution on Veeam R&D forums at [veeam.com/forums](https://www.veeam.com/forums).

Contacting Veeam Software

At Veeam Software, we pay close attention to comments from our customers — and make it our mission to listen to your input and build our products with your suggestions in mind. We encourage all customers to join Veeam R&D Forums at [veeam.com/forums](https://www.veeam.com/forums) and share their feedback directly with the R&D team.

Should you have a technical or licensing issue or question, contact our Customer Support organization directly. We have qualified technical and customer support staff available 24 hours a day, 7 days a week, who will help you with any inquiry that you may have.

Customer Support

For the most up-to-date information about our support practices, business hours, and contact details, please visit [veeam.com/support.html](https://www.veeam.com/support.html). You can also use this page to submit a support ticket and download the support policy guide.

Company Contacts

For the most up-to-date information about company contacts and office locations, please visit [veeam.com/contacts.html](https://www.veeam.com/contacts.html).

Document updated 2026-07-30

Document content applies to build 13.1.0.411