Veeam Backup & Replication

What's New in V12?

Veeam® Backup & Replication™ v12, the foundational product component of Veeam Data Platform, advances enterprise-grade recovery capabilities that ensure confidence in the face of disaster or cyber-attacks across the hybrid cloud. The following is a list of the major new features and enhancements added in V12.

Direct backup to object storage

Take full advantage of the unlimited scalability, built-in reliability and resiliency of on-premises and cloud object storage without having to sacrifice backup and restore performance. With V12, you can leverage object storage both as a regular backup repository and as ANY Scale-out Backup Repository™ (SOBR) tier. Object storage repositories can also be a target for both primary (i.e., backup) and secondary (i.e., backup copy) jobs.

The unique benefits of object storage integration in V12 include:

Direct-to-Object — Whenever possible, protected data will be transferred from backup proxies and agents directly to object storage, avoiding intermediary hops that could potentially affect performance and reliability. Plus, in scenarios when a direct network connection to object storage is not available, you can route the traffic through a redundant gateway servers pool.

Direct-to-Cloud — Reduce the overall cost and complexity of managing backups at the edge, such as in Remote Office Branch Office (ROBO) environments, with direct backups to cloud object storage. Please note that we continue recommending local backups as the first step for data center environments due to the performance implications and extra cost of achieving the 3-2-1 Rule when backing up directly to cloud object storage.

Immutable backups — Keep your backups safe in object storage with support for native immutability technologies provided by on-premises and cloud object storage vendors. Malicious actors will not be able to tamper with, encrypt or delete your recent backups for the specified number of days following their creation. GFS backups are automatically protected for the entire duration of the retention policy to satisfy compliance requirements.

Spaceless full backups — The spaceless full backup technology of ReFS and XFS-based Veeam backup repositories is also available when you backup directly to object storage. Remain flexible with your retention policies with the knowledge that your GFS full backups will not consume additional object storage space.

Improved storage format — Using our learnings from years of observing our customers storing their real-world workloads on object storage, we were able to significantly simplify our storage format without impacting performance or increasing object storage costs. As a result of removing the local index and the need to synchronize it between local and object storage, this new format brings significantly improved reliability.

Health check light — Managed-by-Agent backup jobs that use an object storage repository as a target will perform light storage-level data corruption guard scans that will only check if all required objects exist. This provides some balance considering the performance implications and cloud egress costs of full content verification. All backup jobs managed by a backup server avoid these issues with the help of automatically provisioned, in-cloud helper appliances that allow users to perform full health checks.

Wasabi integration — As one of the most popular object storage options among Veeam users, Wasabi now joins other directly integrated cloud object storage partners with a new, dedicated UI for registering a Wasabi Hot Cloud Storage repository.

Smart Object Storage API — This newly introduced Smart Object Storage API (SOSAPI) allows object storage vendors to integrate with Veeam Backup & Replication more deeply and improve performance and user experience. Its first version enables vendors to report remaining physical disk space in a bucket and gives them machine-level control over Veeam backup streams. For example, they can instruct Veeam to write backup data of a particular machine directly to the specific cluster node for faster performance. The launch partners for SOSAPI integrations are Scality (software defined storage) and Object First (hardware appliance), with other object storage vendors planning to release their integrations soon.

Trusted immutability

Ensure that your backups can always be restored after a cyberattack with comprehensive and enterprise-grade immutability options from trusted vendors, including on-premises objects, block and file storage, hardened repositories, deduplicating storage appliances, cloud object storage and tape.

Immutability for more backup types — In addition to existing functionalities for image-level backups, V12 adds immutability options for NAS backups, standalone agent backups, backups of AWS and Microsoft Azure-hosted workloads, transaction log backups and enterprise application plug-in backups.

Hardened repository improvements — Upgrading hardened repository components post V12 will no longer require you to enable the SSH Server for the duration of the upgrade process. This dramatically simplifies repository management without compromising security. Furthermore, we made hardened repositories more prominent in the UI and added a dedicated wizard to ensure its secure configuration. The UI now has a dedicated repository type and existing hardened repositories get converted to one automatically during the upgrade.

Microsoft Azure Blob Storage immutability support — We also added support for immutable backups on Blob storage repositories. Note that, due to existing Azure object-level immutability API limitations, this functionality is not currently supported for Managed-by-Agent backup jobs and for Veeam Cloud Connect repositories when an object storage repository is configured for direct transfer mode (without gateway servers).

HPE StoreOnce immutability support — V12 supports the creation of immutable backups on Catalyst Stores with the ISV Controlled Data Immutability feature enabled.

Cyber resiliency

Harden your backup infrastructure and stop cyber criminals at the door with the help of the following features and improvements:

Multi-factor authentication — Secure access to the backup console with optional two-factor authentication (2FA) that's based on Time-Based One-Time Passwords (TOTP) as per RFC 6238. You can enable 2FA for individual accounts in the Users and Roles settings of your backup server and enroll in an authenticator application of your choice to receive these one-time codes.

Kerberos-only authentication — V12 can be deployed in environments with NTLM authentication disabled for enhanced security. This includes all backup infrastructure components, backup agents, enterprise application plug-ins and proxy appliances. Kerberos-only authentication is supported by V12 right out of the box as long as managed servers and protected machines are registered with the backup server through valid, resolvable DNS names (IP addresses are not supported by Kerberos). NFS workloads require additional NFS Server and Client configurations, please refer to the User Guide for more information.

Note: If you have already been using our existing capability that allows application-aware guest processing in a network with NTLM disabled, please refer to the KB4393 before performing the upgrade.

IPv6 support — V12 brings support for IPv6 communication in both IPv6-only and dual-stack networks, giving preference to IPv6 over IPv4 addresses when both are available. This flexibility allows you to set your backup infrastructure up in an IPv6-only network right away or gradually migrate it to IPv6 by running it alongside your IPv4 infrastructure. IPv6 communication can be enabled using the Preferred Networks dialog in the backup server settings.

gMSA accounts for Windows — Perform application-aware processing of Microsoft Windows guests through password-less Group Managed Service Accounts (gMSA) without having to store full credentials, including passwords in the backup server configuration. In Microsoft's own words, "Group Managed Service Accounts are the most secure type of service account for on-premises needs. If you can move to one, you should!"

Single-use credentials for Linux — Perform Linux guest processing, recover files and manage Linux backup agents through a preinstalled management agent and certificate-based authentication without having to store powerful guest credentials in your backup server configuration. For further hardening, you can even disable the SSH Server completely after a management agent has been installed.

Improved audit logs and alerts — Over 90 additional events were added to the Windows Event Log and built-in audit logs based on customer feedback, including various tasks performed by backup administrators. In addition, whenever a backup server fails to create an event log item, it will now send the corresponding SNMP trap alert to notify users of this situation.

Automatic console lockouts — A configurable console lockout/timeout has been added to automatically close idle backup console sessions. Worry not if picking up that coffee is taking longer than expected!

Best practices analyzer — Most security breaches, data losses and failed recoveries can be avoided by following just a few simple best practices. V12 introduces a built-in Best Practices Analyzer that checks the backup server and product configurations and suggests important changes that can improve security and increase the chance of successful recovery. We're starting with just a few checks in V12, but already have many more in the pipeline!

Hybrid cloud optimized

Achieve even greater efficiency and security across the hybrid cloud with immutable backups for AWS and Azure-hosted workloads as well as industry-first and cloud-integrated, application-aware backup agents that offer the following unique benefits:

Network-less discovery and deployment — The discovery of new workloads and automatic installations of backup agents are all done through a native cloud API without a direct network connection to those protected machines. Forget the network management and security concerns of accessing your cloud environment from an on-premises backup infrastructure.

Dynamic protection scope — Much like you can protect multiple on-premises virtual machines (VMs) by specifying the containers and tags you want to be protected, you can now create Protection Groups for cloud VMs via identical public cloud management constructs.

In-cloud data flow — Back up your cloud VMs directly to object storage in the same cloud provider without backup traffic traversing the internet and causing additional costs from egress charges. Make sure you copy the created backups back to your on-premises data center or to another cloud periodically to meet the 3-2-1 backup rule.

Full portability — Resulting cloud VM backups can be restored to any hyper-scaler or back to an on-premises hypervisor VM, therefore not leaving you locked into a specific public cloud.

Other features

In addition to the aforementioned major new features, V12 includes over 500 other enhancements that are a response to customer feedback and ongoing R&D findings, the most significant of which are listed below:

Platform

PostgreSQL support for a configuration database — Avoid the 10GB database size cap and artificial performance limitations of Microsoft SQL Server Express Edition without breaking the bank with support for free, multi-platform PostgreSQL as a configuration database engine. Microsoft SQL Server continues to be supported and, for the time being, is recommended for environments larger than a few thousand workloads. However, SQL Server is no longer included with the product and must be installed manually.

Backup engine

New backup chain metadata format — New, per-machine metadata allows for much greater scalability and more granular protection management. For example, V12 now allows users to perform an Active Full or a Retry operation on individual machines as opposed to an entire job, temporarily disabling the processing of some machines while they undergo maintenance, move machines between jobs, etc. In addition, dependent job operations, like applying a retention policy or transforming a backup chain, no longer need to wait for the entire job to finish first. In order to take advantage of these new capabilities for existing backup chains, you must upgrade them.

Policy-like jobs — V12 combines the flexibility of job-based management with the convenience of policy-based management, bringing the best of both worlds for best-in-class protection management. In addition to the ability to create very large jobs with a few thousand machines, you can now easily move backups between jobs — just like changing a protection policy — without triggering a new full backup or abandoning the existing backup chain. Just choose a new job for your selected backup, and we will take care of transferring the existing backup files to its repository, including adding the machine to the protection scope of the new job and excluding it from the old one.

Background retention — Stop worrying about repository space being consumed by expired backups. In addition to existing background retention processing for GFS backups, background retention in V12 will now process ALL backups with a time-based retention policy, including backups linked to a disabled job as well as orphaned backups by using their last known retention policy. This process starts each day at midnight, however, you can always trigger it manually by right clicking the Backups node. Want some backups to be kept indefinitely instead? We've got you covered, read on!

Background health check — By popular demand, instead of being attached to a backup job, storage-level corruption guard scans can now be scheduled to happen at a specific time, independently of the backup job. This keeps your backup windows in check and removes extra load from backup repositories while other backup jobs are still running.

New compression algorithm — The new algorithm behind High and Extreme compression levels improves data reduction ratios by up to 20%, lowers CPU usage by up to 3x and increases restore performance by up to 2x over the previously used one. Compared to the default Optimal compression algorithm, which remains unchanged in V12, High compression now provides up to 20% (OS disks) and 60% (database disks) more data reduction at the cost of 2x CPU usage and 2x slower restores. Extreme compression provides up to 25% (OS disk) and 80% (database disk) better data reduction at the cost of 10x CPU usage and 2x slower restores.

Unified symbolic link backup and restore behavior — Based on multiple requests, symbolic links are now processed and restored as symbolic links without expanding them across all backup and restore activities.

Block-cloning performance — This update also brings improved performance of full backup transformation operations with asynchronous metadata writes. These related engine changes resulted in the removal of *UseUnbufferedAccess* and *DisableHtAsyncIo* registry values, which were replaced with a single *DataMoverLegacyIOMode* (DWORD, 1) value.

NFS 4.1 client performance — V12 also has improved read and write performance from/to NFS 4.1 file shares used as a data source or as a backup repository.

Backup data management

VeeaMover — V12 delivers a native backup movement engine that can move backups between any backup repository type. While such tasks were relatively easy to perform with block and file storage through regular file management operations, the proliferation of object storage and deduplication appliances called for a built-in capability for backup management. Best of all, the source repository type does not need to match the target; VeeaMover can move backups between any type!

Block clone awareness — The highly sought after capability of saving space from ReFS/XFS/ object storage block cloning during backup migrations is finally here! Not only does VeeaMover preserve space savings while moving backup files between block clone aware repositories like between ReFS and XFS, but it also *creates* space savings even if a source repository does not support block cloning, like when moving backup files from SMB to XFS. How cool is that?

Streamlined repository change processes — Upgrading your backup storage has never been easier! Just choose a new repository in the backup job's settings and you will have the option to transfer all existing backups to that location automatically with VeeaMover. Done!

Move backups between jobs — Just choose the backup you want to move and we will take care of the rest automatically. This includes moving existing backups to the repository that was used by the target job with VeeaMover and automatically updating the inclusion and exclusion lists in the source and target backup job so you don't have to worry about doing that manually!

Copy backups between repositories — Want to put an entire backup chain to the side so that it isn't affected by the original retention policy? Now this process only takes five clicks because we do all the work for you. You can even specify a custom retention policy for the copied backups if you'd like them to be deleted automatically after a specific period of time.

Backup infrastructure

Backup repository

Multiple gateway server support — In addition to the automatic gateway selection mode, you can now specify a pool of gateway servers you want to be used for reading and writing data to/from a given backup repository. Gateways within a pool are prioritized based on network connectivity and existing task load, with backup proxies that are involved in a current job being top priority. Using multiple gateway servers is particularly beneficial for Fiber Channel (FC) backup storage configurations in order to ensure redundancy while addressing the performance issues of a single gateway server becoming a bottleneck.

Rotated drives cleanup — Repositories backed by rotated media now have the ability to automatically clean up newly inserted disks from existing backups. Available options include continuing an existing backup chain (i.e., existing behavior) and two new options which are: delete backups that belong to the current backup job only, or delete all the backups found on the drive, both of which start the new backup chain.

Scale-out Backup Repository (SOBR)

Object storage as performance extent — In addition to block and file storage-based extents, Performance Tier can now use object storage extents as well. For performance considerations, we recommend using on-premises object storage as a performance extent.

Support for multiple object storage extents — Both the Performance Tier and the Capacity Tier can now be configured to use multiple object storage extents. This capability can be useful if your object storage supports only a limited number of objects per bucket, in which case object storage vendors usually recommend creating multiple buckets. Whenever multiple object storage extents are used, SOBR will use a round-robin placement algorithm on a per-machine level.

Direct to archive — When using AWS S3 or Microsoft Azure Blob Storage as Performance Tier extents, you now have the option to skip configuring the Capacity Tier altogether. In such configurations, backups will be archived directly from the Performance Tier to the Archive Tier.

Glacier Instant Retrieval support — This update also brings added support for the new Glacier storage class for Archive Tier. With the same throughput and access latency as the S3 Standard and S3 Standard-IA storage classes, this more expensive archive storage class is best suited for scenarios when you expect that you'll need to urgently access your archived backups.

SOBR rebalance — You can now perform a rebalance of storage consumption across block and file storage-based Performance Tier extents to equalize the data distribution among them. Rebalancing is designed for when you want to add new extents to Performance Tier and is not considered a maintenance operation that you are expected to perform periodically. Note that a rebalance operation puts all Performance Tier extents into maintenance mode for the duration of this activity.

VeeaMover integration — Extent evacuation and SOBR rebalance operations leverage the new VeeaMover engine to preserve space savings from block cloning while backups are moved between extents.

Strict data placement policy enforcement — By default, SOBRs violate your chosen placement policy whenever data placement policies cannot be met to ensure that a backup can still be created. V12 introduces a new option that can make backup jobs fail in these circumstances instead.

Network traffic management

Support for multiple internet rules — By popular demand, we've added the possibility to add several "Internet" rules to address the different internet traffic throttling needs of multi-site and multi-network environments.

Time-based throttling level — Our customers have asked for more flexibility around the control of throttling levels depending on the time of day to enable them to reduce it outside of business hours without having to disable throttling completely. V12 makes this possible with a new Unthrottle setting that allows for a different bandwidth throttling threshold during specified off-hours.

Restore traffic throttling options — Network traffic throttling rules were previously not applied to network traffic from any restore task to avoid having an impact on recovery performance. However, with some of our restores becoming way too fast due to continuous optimization, V12 adds an option to apply throttling to restore operations as well.

Email notifications

OAuth 2.0 support for email notifications — In addition to basic SMTP authentication, V12 now supports secure authorization and access-token-based authentication for Google Gmail and Microsoft 365 through the modern OAuth 2.0 protocol.

Active instant recoveries — New daily summary reports will remind you of all your active Instant Recovery sessions that are still waiting to be finalized. This report can be suppressed by creating the *IRSuppressDailyReport* (DWORD, 1) registry value under the *HKLM\SOFTWARE\Veeam\Veeam Backup and Replication* key on the backup server.

Full product version display — All email reports now include a full server version, including the cumulative patch level displayed in the footer.

VMware vSphere

Linux backup proxy enhancements — V12 brings added support for Direct storage access transport mode, which enables Linux proxies to perform direct backup from storage snapshots of NFS storage.

Hardened repository as a backup proxy — You can now use your hardened repository as a VMware backup proxy in Network (NBD) transport mode. Advanced transport modes are not available since they require backup proxy components to run as root, which dramatically expands the attack surface.

Improved replication performance — VMs with SeSparse snapshots (i.e., the default snapshot format for VMFS-6) should see faster replication performance thanks to asynchronous write-to-replica VM disks.

Backup I/O control latency minimums — In light of All-Flash storage proliferations, the minimum allowed value for both latency controls was reduced to one millisecond.

Microsoft Hyper-V

Infrastructure caching — Similar to existing vSphere infrastructure caching, the backup server will now also cache Hyper-V host information to reduce the number and frequency of WMI queries. This reduces the load on Hyper-V hosts, improves job performance and makes the user interface much more responsive in large infrastructures.

Compatibility checker enhancements — You can now restore and replicate VMs to Hyper-V hosts of lower versions compared to the original host as long as the VM hardware version is supported by the target host.

Image-level backup

Application-aware processing

PostgreSQL support — V12 added full application-aware processing, including transaction log backups for point-in-time instance recoveries for PostgreSQL on Linux, like the existing functionality for Microsoft SQL Server and Oracle databases.

Persistent guest agent enhancements — In the case of no direct network connectivity, guest interaction proxies will now communicate with a guest agent over network-less, hypervisor-based protocols like VIX and PSDirect.

Backup scheduling

More synthetic full scheduling options — By popular demand, in addition to weekly synthetic full backups, synthetic full backups can now be scheduled to happen monthly. The usage of

monthly synthetic full backups is not recommended for legacy backup repositories and should only be considered for modern block cloning aware backup storage.

More monthly GFS scheduling options — To satisfy the compliance requirements that some of our customers are facing, monthly GFS backups can now be scheduled for creation on the first, second, third, fourth and final week of the month. This option may also be useful for spreading the load of monthly full backup creation by multiple jobs pointed to a legacy backup repository without block cloning support.

Increased GFS restore point limits — You can now create weekly GFS retention policies of up to 9999 weeks long. Surprisingly, the previous UI limit of 999 weeks was not enough for some customers who face certain compliance requirements!

Backup copy

Backup Copy jobs were made multi-platform to avoid the need to create individual platform-specific Backup Copy jobs, which allows multiple workload types to be added into a single backup copy job. New Backup Copy jobs come with the following changes and improvements:

Per-machine backup chains — New Backup Copy jobs will always create per-machine backup chains through the new per-machine metadata format to enable the compatibility of backup copies with many of the new V12 backup management features. Existing Backup Copy jobs will continue to operate as is until upgraded.

Expanded platform support — You can now copy backups that were created by agent-based backup jobs in Managed-by-Agent mode, backups created in Veeam repositories by standalone agents, and backups created by Veeam Backup *for Nutanix AHV* and Veeam Backup *for Red Hat Virtualization*. This includes transaction log backups created by these backup jobs (in Immediate copy mode only).

Ability to change backup copy mode — All newly created Backup Copy jobs will support changing modes from Immediate to Periodic and visa-versa at any time.

Active full for GFS restore points — By popular demand, we've added an option to create GFS restore points through the active full approach — by reading the entire restore point content from the source backup — to the Immediate copy mode as well. This is in addition to the current synthetic full backup approach, which synthesizes GFS restore points from existing data in the target repository.

Restore point selection option — With enhanced Periodic Backup Copy jobs, you can now specify whether you want to wait for the running source job to finalize the restore point that's currently being created or immediately pick the latest available restore point instead.

Daily email reporting option — For those who want to reduce reporting noise from Backup Copy jobs, you can now choose to receive a daily summary backup copy email report instead of a separate report every time a backup copy job finishes processing.

SureBackup

Agent-based backup support — In addition to host-based backups of VMware and Hyper-V VMs, SureBackup® jobs can now process agent-based backups of cloud, physical and virtual Windows and Linux machines.

Multi-platform jobs and application groups — Starting from V12 SureBackup jobs and Application Groups are no longer platform-specific and can include a mix of VMware, Hyper-V and agent backups. Processed machines will be converted to the data lab platform (i.e., VMware or Hyper-V) through an on-the-fly P2V/V2V process.

Disable Windows Firewall — Application Groups now have the option to automatically disable the Windows Firewall on processed machines before starting them up in a Virtual Lab. This option can be useful if some of your custom application tests cannot be performed due to a Windows Firewall blocking your external connection attempts.

Granular email notifications — Just like with backup jobs, you now can set up granular email notifications per SureBackup job which will take priority over global notification settings.

Recovery from image-level backups

Application item-level recovery

Veeam Explorer *for PostgreSQL* — A new addition to the Veeam Explorer™ family allows you to restore PostgreSQL instances with ease without the need for an extensive PostgreSQL administration background. In addition, you can publish a point-in-time state of any instance directly from backup to the selected database server for Dev/Test, and any changes made to the published database can either be exported or discarded. With the service-based architecture, you don't have to depend on the user interface running while restoring or publishing an instance. Any failed tasks caused by intermittent infrastructure issues can be easily retried without having to go through the Publish wizard again.

Veeam Explorer *for Microsoft SQL* bulk restore — You can now pick and choose multiple databases when performing a mass database export, restore or Instant Recovery. The Veeam Explorer will then automatically create and start multiple tasks, one for each processed database.

File-level recovery for Windows

Compare with production — Compare a selected restore point to the production machine with the new Backup Browser view that enables you to effortlessly identify all the file system objects that were changed or deleted since the backup was taken.

Restore changes only — This new restore mode allows you to quickly initiate the recovery of all the items that were changed or deleted since the selected backup was taken, which minimizes downtime after a ransomware attack or user error.

Compare attributes — Easily check differences in file system attributes for individual files and folders through a new dialog that shows their values in production and backup side by side.

Restore permissions only — Restore access control lists (ACL) of your files and folders without needing to overwrite the contents of the file. This restore mode can be useful in case of an administrative error during the file system permission management process.

Restore directly to another machine — The ability to select a different target machine to restore files to — previously available for Linux file-level recovery only — is now available for Windows machines as well.

Alternate Data Streams (ADS) restore — When restoring a file or a folder from a backed up NTFS volume to an NTFS volume, ADS content will be restored as well.

ReFS volume autodetection — V12 detects the presence of ReFS volumes in backups and automatically mounts the ones with a VHD mount API. This helps to prevent BSOD with certain combinations of a mount server and protected machine OS versions. This previously required the use of the *ForceVhdMount* registry value as a workaround.

File-level recovery for Linux

Restore from storage snapshots without a helper appliance — File restores from storage snapshots of Linux file systems can now be performed by mounting backups to ANY Linux machine, whether that's dedicated, targeted or the original one, which is always guaranteed to understand the file system you're trying to restore from.

FLR helper appliance improvements — V12 added the ability to specify a DNS Server to the helper appliance to properly resolve DNS names.

Instant VM Recovery

Instant VM Recovery® to Hyper-V without space pre-allocation — When performing an instant recovery of any backup to a Microsoft Hyper-V VM, you now have the option to not pre-allocate the entire disk space on the target host's storage. This reduces the time it takes to start a published VM and removes the need to have the required physical disk space available.

Secure Restore

Bitdefender support — V12 added out-of-the-box integrations with Bitdefender antivirus.

Export

Export backups to another location — Improved Export Backup functionality now allows you to select any destination for an exported restore point instead of having to save it into the same repository where the original backup resides.

Export as a virtual disk enhancement — Export as a virtual disk functionality, which was only available for agent-based backups, is now extended to host-based VM backups.

Continuous Data Protection (CDP)

CDP proxy on Linux — Lower your operational costs and save the expense of OS licenses by deploying CDP proxies on Linux servers.

Veeam Cloud Connect support — In addition to regular replication jobs, you can now also target CDP policies to a cloud host that's provided by the Veeam service provider of your choice.

VMware Cloud Director support — Eliminate downtime and minimize data loss for Tier-1 vApps by continuously replicating them within and across your Cloud Director instances. Perform immediate recovery to the latest state or desired point in time for individual VMs and entire vApps.

Enhanced security — Communication between all CDP components is now secured with the TLS protocol.

Support for native VVOL snapshots — Native VVOL snapshot awareness should lead to a reduced number of objects that are stored on VVOL-datastores, which improves CDP reliability on storage devices that have low VVOL scalability limits.

Sparse block detection — Sparse disk blocks are now outright skipped from processing which should result in improved initial replication performance.

Improved transaction log storage formats — Transaction logs are now stored as virtual disks instead of as regular files, which significantly increases log write speed.

Proxy memory overflow protection — A long-standing issue with the replication of a particular virtual disk should no longer impact the processing of other disks by the same proxy.

Agents

Agent Management

Recovery tokens — V12 introduces a simplified way to provide users who are performing a Bare Metal Recovery with access to a particular backup. Backup administrators are now able to generate time-limited access keys, or recovery tokens, that can be shared with users and enable them to connect to a Veeam repository when performing Bare Metal Recovery.

Changed-block tracking for Windows workstations — In addition to Windows servers, the Protection Group wizard now includes the option to install a Veeam changed block tracking (CBT) driver on workstations that run on Microsoft Windows 10 or later for when a faster incremental backup is desired.

Multiple cluster volume recovery — Multiple volumes are now restored to a cluster disk concurrently in order to reduce recovery time.

The following agents are included in the Veeam Backup & Replication V12 redistributable:

Veeam Agent *for Microsoft Windows*

New V6 functionalities include backing up directly to object storage, a new SQLite-based configuration database, CBT-based file-level backups for servers and workstations, OAuth 2.0 support for email notifications, support for the latest Microsoft Windows versions and more. For a complete list of new features, please refer to the corresponding What's New document.

Veeam Agent *for Linux*

New V6 functionalities receive advanced data fetcher technology and includes backup directly to object storage, GFS retention policy, support for the latest distribution versions and more. For a complete list of new features, please refer to the corresponding What's New document.

Veeam Agent *for Mac*

New V2 functionality receives a graphical user interface and includes backup directly to object storage, the ability to resume backup functionality, native support for M1 and M2 chips, FIPS compliance, support for the latest MacOS versions and more. For a complete list of new features, please refer to the corresponding What's New document.

Veeam Agent for AIX

V4 brings Bare Metal Recovery with fully automated disk mapping on the new system.

Veeam Agent for Solaris

V4 brings Bare Metal Recovery with fully automated disk mapping on the new system.

Application plug-ins

Plug-in management

V12 adds the ability to centrally manage enterprise application plug-ins in a manner that's similar to our backup agents. This includes the following features and functionalities:

Centralized application plug-in management — The Protection Group wizard has been expanded with additional options to control the installation and upgrade of application plug-ins on included servers. This includes the collection of application topology and the detection of Oracle RAC and SAP HANA scale-out systems during protection group rescans.

Certificate-based TLS authentication — Protection Groups use certificate-based authentication when establishing a network connection to plug-ins.

Application backup policy — Orchestrate backups of Oracle RMAN, SAP HANA and SAP on Oracle in a policy-driven way directly from a backup console, which eliminates the overhead that comes with configuring plug-ins and maintaining backup scripts manually on each database server. This allows customers to choose between keeping backups in full control of the database administrator (DBA) and using a backup-administrator-centric, policy-driven protection with minimal DBA involvement.

Granular database-level protection settings — To achieve greater flexibility, application backup policies support different settings for individual databases. Each policy can target its own scope of databases (including entire database servers and clusters) and have its own backup schedule and repository.

Centralized backup monitoring — Application backup policies tap into native backup tool outputs on each server and provide real-time monitoring, statistics and reporting for database and redo log backups.

Recovery tokens — Backup administrators are now able to generate time-limited access keys, or recovery tokens, to access plug-in backups of other servers. These tokens can be shared with DBAs to enable them to connect to a Veeam repository and perform database restores with native tools without having to assign them to any backup infrastructure role.

Plug-in for Microsoft SQL Server

This new application plug-in provides deep integration through a built-in Microsoft SQL Server backup process (VDI) that allows you to perform native database backups directly to Veeam repositories.

VDI-based plug-ins leverage native capabilities to ensure backup consistency and, unlike snapshot-based backups, does not rely on Microsoft VSS and enables you to back up advanced SQL Server configurations like Windows Server Failover Clusters with shared volumes. This is great from a scalability perspective too — the plug-in can handle protecting up to 2,000 databases per SQL server and up to 10,000 databases per backup server.

This plug-in was designed from the ground up to follow the best practices of protecting other popular SQL Server deployment configurations like SQL Always On, where the backup jobs will automatically read data from your preferred replica server.

To simplify configuration, backup and recovery tasks, this plug-in has a user interface (UI) that is integrated directly into Microsoft SQL Server Management Studio where it can be launched right from the toolbar! However, DBAs are still allowed to use any preferred external tools or favorite scripts in order to schedule and execute full, differential or transaction log backups into Veeam repositories. The plug-in UI even includes the ability to create SQL Server Agent jobs.

Achieving the 3-2-1 Rule is as easy as ever with Veeam. Backup Copy jobs support copying native SQL backups to secondary repositories, and if you want to avoid managing these additional jobs, then you just need to use a SOBR with a Capacity Tier in the Copy mode. Our proprietary backup

chain analysis enables you to offload inactive backup chains with the Move policy to free up disk space in a Performance Tier, and helps to ensure that the immutability period is correctly extended for still-active backup chains.

Plug-ins for Oracle RMAN and SAP HANA

Performance improvements — V12 improves plug-in backup and restore speeds by up to three times! Most of this performance boost was achieved by teaching data movers to fetch data directly from an application, bypassing the plug-in manager process in between.

Reduced impact on production environment — We will no longer start a dedicated data mover for each backup channel in order to reduce the number of data movers that run on a database server. The schedule will ensure that the maximum number of data movers never exceed 2x the available CPU cores. To further reduce the CPU usage, a plug-in configuration option, *calculateDigestsOnRepository*, is available to move hash calculations to the repository. With all these optimizations, the CPU load on the application server can be reduced by up to two times.

Linux ACL support for configuration files — You can now leverage Linux groups to restrict changes to the plug-in configuration by backup operators.

Backup proxy appliances

Veeam Backup *for AWS*

New V6 functionalities available in March 2023 include immutable EC2 instance backups in AWS S3 storage, brings scalability improvements for large-scale AWS environments, GP3 disk support for appliances, multi-tenant container Oracle database support, OAuth 2.0 support for email notifications, integration with Veeam Service Provider Console and support for Veeam Universal License (VUL). For a complete list of new features, please refer to the corresponding What's New document.

Veeam Backup *for Microsoft Azure*

New V5 functionalities available in March 2023 include immutable Azure VM and Azure SQL backups in Azure Blob storage, scalability improvements for large-scale Azure environments, load control for backup repositories, OAuth 2.0 support for email notifications, integration with Veeam Service Provider Console and support for Veeam Universal License (VUL). For a complete list of new features, please refer to the corresponding What's New document.

Veeam Backup *for Google Cloud*

New V4 functionalities include granular protection for PostgreSQL databases and the ability to store instance snapshots in a single region. This version also lowers management overhead with support for cross-project service accounts, the ability to use organization folders as a policy source and support for Veeam Universal License (VUL). For a complete list of new features, please refer to the corresponding What's New document.

Veeam Backup *for Nutanix AHV*

New V4 functionalities include a new modern UI, backup directly to object storage, immutable backups, synthetic full backups, GFS retention, storage-level corruption guards and REST API to manage backups and restores. We've also increased restore performance and added support for Instant VM Recovery from Veeam Cloud Connect repositories. For a complete list of new features, please refer to the corresponding latest Release Notes.

Veeam Backup *for Red Hat Virtualization* (RHV)

New V3 functionalities include backup directly to object storage, immutable backups, synthetic full backups, GFS retention, storage-level corruption guard and multiple UI improvements. For a complete list of new features, please refer to the corresponding Release Notes.

Container backups

Kasten K10 Backup for Kubernetes integration — K10 instances can now be registered with a backup server which allows users to view all Kubernetes backup policies, sessions and backups directly in the backup console, regardless of whether those backups are stored in Veeam backup repositories or in another location. Invoking backup policy editing and restore operations will redirect users to the contextually appropriate K10 workflow and allow for operation completion through the K10 web UI.

NAS backup

General

SMB over QUIC support — File share backup jobs are now able to retrieve data from SMB files securely and more efficiently with the QUIC protocol whenever a file share supports it.

Recovery throttling support — Network traffic throttling functionalities now support bandwidth management for NAS restore activities as well.

Improved search performance — Optimized file search functionalities deliver search results up to 25% faster in both the backup console and in the Veeam Backup Enterprise Manager web UI.

Storage-level corruption guard enhancements — Health check processes are now able to detect and notify users about possible issues with the archived metadata as well. The repair process is not initiated automatically and therefore must be manually triggered by a user after addressing the root cause.

NFS fs_locations support — File share backup jobs now understand the concept of NFS referrals, which enables you to more easily configure the protection of complex, distributed NFS file systems.

Backup to disk

Copy mode — In addition to archiving older file versions that are no longer under the backup retention policy, you now have the option to immediately archive current versions as well. In this case, the archive will contain an entire copy of your backup, which will allow you to perform full share restores to the latest state directly from an archive repository in the case of a complete loss of the backup repository. Or, you can heal your backups in case a partial loss of data in a backup repository by downloading the missing data from the archive.

Expanded exclusion rules — Avoid protecting unnecessary data with the ability to exclude specific shares or folders by adding a path to excluded objects like *\share\folder\path* to the exclusion rules. Use the new wildcard-based exclusion rules like **\folder* to instruct the job to never back up folders with these names. (This is supported for first-level folders only). For NAS filers, hidden file shares *\ipc$ \admin$* and *\c$* are now excluded from processing by default, but you can always remove these corresponding rules if desired.

Advanced backup mapping — You can now avoid performing a full backup after renaming or moving a source file share through the Update-VBRNasBackupPath PowerShell cmdlet to point your existing backup to the new file share location.

Backup to tape

Scalable file backup engine — No more struggling with backing up a large number of files directly to tape with a tape-out engine that's designed to export a small number of large image-level backup files! The new, redesigned file backup engine in V12 was built to scale to billions of files in petabytes of size which allows you to back up enterprise-scale NAS deployments directly to tape with confidence. This new capability supports all the same data sources and is licensed in the same way as the existing NAS backup-to-disk functionality.

NAS backup to tape — By popular demand, you can now use file share backups as a source for Backup to Tape jobs, thus enabling the classic Disk to Disk to Tape (D2D2T) backup approach for your enterprise NAS filers and file shares. To ensure recoverability at any point in time, this functionality exports files and folders to tape in the native format as opposed to copying blobs of the proprietary NOSQL database behind NAS backups. Since it's a secondary backup, this functionality does not consume a license.

Instant Recovery

In addition to the existing capability to publish file share content directly from backups as a read-only share, you can now also perform instant file share recovery. This allows your end users to get back to work in just seconds after the catastrophic loss of production NAS. This functionality is enabled through these new capabilities:

Writable SMB file shares — Emulated SMB file shares that were published as a part of instant file share recovery are now writable, which enables end users to continue working with an emulated file share normally, including the process of modifying existing files and creating new ones. NAS backups are never modified and all changes are cached separately.

Migration to production — Once your new production NAS is back online, you can initiate the process of restoring the most recent file share state from backups. This operation is done in the background without impacting the end user's ability to interact with the published file share.

Switchover — Choose one of three available switchover types to finalize your instant recovery. The Automatic switchover mode will be performed as soon as the background restore process completes and the Scheduled mode allows you to delay automatic switchover until off hours. You can also choose the Manual switchover mode to interactively perform this process yourself. Keep in mind that the published file share will become unavailable to end users for the duration of the switchover in order to synchronize the final published file share state. This includes any content change to the production NAS.

Publish

NFS shares publishing — Protected NFS shares can now be published directly from a backup as emulated SMB shares. One of the use cases includes making the content of your backups available to data analytics software that support SMB file shares as a data source.

Filer integrations

Nutanix Files integration — Register entire Nutanix Files filers as data sources and perform file backups without having to obtain access permissions to each protected file share. In this kind of configuration, backups will be performed from native storage snapshots out of the box, which allows you to avoid issues with locked files without requiring complex setup and scripts to manage snapshots. In addition, incremental backups are dramatically accelerated thanks to the use of Nutanix Files Changed File Tracking (CFT) APIs to instantly query a list of changed or removed files without needing to scan a file share.

Dell PowerScale (Isilon) — V12 adds support for OneFS versions 9.3 and 9.4 for a native filer backup integration.

Backup storage integrations

Rotated drive support — File share backup jobs now support backup repositories that are backed by rotated drives.

AWS Snowball Edge support — In addition to existing support for Microsoft Azure Data Box, V12 adds support for initial full-backup seeding to AWS Snowball Edge object storage for further import into a corresponding public cloud.

Backup console

Granular operations — You can now retry processing or perform Active Full backups of individual machines without triggering this operation for all the machines in the job by right clicking the corresponding machine in the job session.

Global exclusions list — Manage permanent or temporary exclusions easily and centrally by specifying the master list of machines that you never want to be processed, even when added to a job explicitly by mistake. The Global Exclusions dialog can be accessed from the main menu. In addition, these kinds of machines will have the "Disable processing" option selected in the Inventory tab.

Detach backups from a job — You can now disconnect existing backups from the job, which will make the job start a new backup chain by performing an active full backup in the next run. Detached backups will appear in the Orphaned node under Backups and are still subject to their last known time-based retention policy. If you want to avoid this, use the Export Backup functionality or the new Copy Backup feature instead of detaching.

Dashboard redesign — Workload and storage registration dashboards have been redesigned to simplify navigation within an ever-growing armada of supported platforms, storage vendors and storage models, which Veeam Backup & Replication natively integrates with.

Last backup time — By popular demand, we will now display the date and time at which the last backup was taken in the list of machines on the Inventory tab.

Exported backup retention — If you choose to have your exported, copied or VeeamZIP backups be deleted automatically, the calculated backup removal date will be logged in the session log.

Repository membership and role — Newly added columns in the repository view will show whether the given repository will serve as an extent member of SOBR as well as the extent type (i.e., Performance, Capacity or Archive).

Instant recovery preferences — The state of checkboxes on the last page of the Instant Recovery wizard will now be saved on a per-user basis to help make commonly performed recoveries a little faster.

Session initiator logging — All Restore and System sessions will now display the account of the user who initiated the operation.

History tab improvements — All still-active sessions (i.e., those with no End Time) will now be shown at the top of the list when sorting it by End Time.

Update notification improvements — The update notifications engine has been improved to include notifications about cumulative patches.

Enterprise Manager

General

Certificate-based authentication — Backup administrator credentials for each managed backup server no longer need to be stored in the configuration database. Once the backup server has been registered with the Enterprise Manager server, the following communication will use certificate-based authentication with automatically generated certificates.

Backup server preferences — We will now remember your choice of backup servers on a dashboard page and will display previously selected servers by default.

Improved email reports — Multiple improvements in the email report content and layout.

Export support logs — Collect logs requested by Veeam Support more easily with the new Log Export functionality.

Backup

CDP for VMware Cloud Director — Manage existing VCD CDP policies and their replicas in the Enterprise Manager.

Catalyst Copy jobs — Manage HPE StoreOnce Catalyst Copy jobs in the Enterprise Manager.

Quick backup — Trigger Quick Backup operations in the Enterprise Manager for Managed-by-Server agent-based backup jobs.

High priority jobs — Change job priorities in the web UI and recognize high priority jobs in the job lists.

Restore

Instant VM Recovery — Perform instant recovery for VMware vSphere, VMware Cloud Director, Microsoft Hyper-V VMs and file shares directly from backups and storage snapshots through the web UI. The option to finalize recovery with the Migrate to Production operation is also available.

Restore to another location — In addition to in-place restores, you can now perform a full VM restore to another host or storage and specify all your other related placement settings.

PostgreSQL database restore — Perform point-in-time instance restores from application-aware backups of PostgreSQL servers yourself or delegate them to database administrators and help desk operators.

VM templates restore — You can now search for and restore VMware VM templates.

Exchange item-level restore improvements — You no longer have to store an AD account that's a member of Domain Administrator or Organization Management groups in Veeam Backup Enterprise Manager database permanently. Instead, you can interactively provide an account with limited rights to a reduced scope of mailboxes when starting a restore.

Search result sorting — Performing file-level recovery is now easier with files sorted alphabetically by name in the search results.

Setup

This new, completely reworked setup experience is aimed at one goal — making the installation and upgrade process as smooth as possible. Key highlights include:

New setup experience — We like to keep things simple, so we've preserved all the familiar functionalities of the old setup wizard while getting rid of all the annoying stuff and wrapping it up in a new and shiny modern UI. You no longer have to worry about prerequisites, select multiple checkboxes or customize anything — unless you really want to!

Streamlined upgrade — We've revised our upgrade pre-flight check functionalities and will now present all your detected configuration issues in a much easier to use format. See a full report with all the details, copy and share it with your colleagues, address the issues and upgrade with confidence!

Licensing

The best change is no changes, and we're happy to say this again! V12 still uses the same license file format that was introduced in V10. These license files are no longer tied to a particular software version which allows you to continue using your existing license file for V12 as well as long as your maintenance contract is still active.

Veeam Universal License (VUL)

Doubled license buffer — Users with the License auto update functionality enabled will now enjoy a two times larger license buffer that lets them exceed their VUL usage by up to 20% or 20 licenses (whichever is greater). Note that an increased buffer requires the license update check to happen successfully at least once per month, so make sure this connection is not blocked by your firewall.

Community Edition

Now with even more features — The updated Veeam Backup & Replication *Community Edition* V12 benefits from many of the new features and enhancements introduced in V12. For more information, refer to the edition comparison document.

Cloud-native backups — Veeam Backup *for Microsoft Azure*, *for AWS* and *for Google Cloud* are now included in Community Edition as part of the 10-license allowance.

Support Community Edition by upgrading to Veeam Data Platform Essentials if possible! This subscription costs less than a dinner in a fancy restaurant, while giving you access to ALL the features and 24.7.365 customer support! These conversions help us continue to offer free enterprise-class data protection to those who can't afford it.

Storage integrations

Primary storage

General

Storage discovery optimizations — Automatic storage rescan logic has been revised to avoid unnecessary scans that impact storage performance. The number of heavy VMFS rescan operations has been reduced significantly and they should only happen when collecting or updating volumes and snapshots information is actually required.

Universal Storage API v2

The new version of USAPI enables storage vendors to take their Veeam integration to the next level with a set of new interfaces that enable the following functionalities:

Snapshot replication orchestration — Allows backup jobs to leverage existing storage replication relationships to create storage-based snapshot replicas as additional restore points. This also enables you to perform backups from secondary storage arrays to avoid any load from backup activities on your primary storage array.

Snapshot archiving orchestration — This allows backup jobs to manage the offload of volume snapshots to a different type of storage device that's supported by a storage vendor and track them as additional available restore points.

Synchronous replication support — This allows backup jobs to create coordinated storage snapshots and track them as additional available restore points on dual storage system configurations that serve as the same production volume from multiple storage management backends. This enables integration with both Active/Active and Active/HotStandby configurations.

The development and launch partner for USAPI V2 is Pure Storage with their new plug-in for its FlashArray product line being available immediately.

Note: Existing USAPI V1 plug-ins remain fully supported with Veeam Backup & Replication v12.

Cisco HyperFlex

Multi-vCenter configuration support — V12 leverages the new HyperFlex API for VMware vCenter server to VMware ESXi host matching to enable full support for environments that have multiple HyperFlex clusters under the same vCenter Server and multiple HyperFlex clusters with a separate vCenter Server for each cluster.

IBM Spectrum Virtualize

Volume replication orchestration — Added support for both the snapshot replication (i.e., Global Mirror) and the synchronous replication (i.e., Metro Mirror and HyperSwap) technologies for IBM Spectrum Virtualize based storage. This enables secondary storage array support for backup from storage snapshots and snapshot-only jobs.

Compressed snapshots support — For IBM FlashSystem volumes with support for hardware compression that's based on IBM FlashCore modules, snapshot compression will be used automatically.

HPE Nimble and HPE Alletra 5000/6000

Synchronous replication support — V12 adds support for synchronous replication and coordinated snapshots, including deployment configurations with Peer Persistence for application-transparent failover. This enables secondary storage array support for backup from storage snapshots and for snapshot-only jobs. This functionality is supported for HPE Nimble and HPE Alletra 5000/6000 storage array with Nimble OS version 5.1.4 or later.

NetApp All SAN Array (ASA)

NetApp ASA support — V12 adds support for NetApp ASA for storage snapshot integration.

Secondary storage

ExaGrid

Fast cloning support — V12 integrates with native block cloning functionality that was introduced in ExaGrid version 6.2 for a dramatic increase in synthetic full backup performance.

Dell Data Domain

Longer incremental backup chains — The maximum number of increments before the next periodic full backup has been increased to 120 restore points in order to align with increased Data Domain scalability limits.

DD OS and DD Boost support — V12 has added support for DD OS versions up to 7.10 and updated DD Boost SDK to version 7.7.1 LTS.

HPE StoreOnce

HPE StoreOnce immutability support — V12 supports immutable backups on Catalyst Stores with the ISV Controlled Data Immutability feature enabled. This functionality requires Catalyst Stores in the compliance mode, which in turn forces Dual Authorization that effectively restricts the modification and deletion of immutable files by StoreOnce administrators without Security Officer role approval. This functionality requires StoreOnce firmware version 4.3.2 or later.

Fixed block size chunking — V12 enables the Enforce fixed block chunking setting on a newly created Catalyst store which improves incremental backup and synthetic full performance up to 4x-5x times per stream (based on HPE's own tests) compared to variable block processing. This performance boost is achieved thanks to the Catalyst Client using fixed block chunking backup files that were pre-aligned by Veeam with the Align backup file data blocks backup repository setting. This functionality requires StoreOnce firmware version 4.3.2 or later.

HPE Cloud Bank Storage support — With V12, you can use HPE Cloud Bank Storage — an extension of HPE StoreOnce Catalyst — for leveraging external lower-cost object storage as underlying storage, as a target backup repository for Catalyst Copy jobs, which helps to reduce long-term retention costs. HPE StoreOnce firmware version 4.3.2 is required.

Catalyst Copy support for more platforms — Backups that were created by Veeam Agents *for Microsoft Windows* and *for Linux* in both managed-by-agent and standalone modes as well as backups created by Veeam Backup *for Nutanix AHV* and Veeam Backup *for Red Hat Virtualization* can now be used as the source for Catalyst Copy jobs.

Infinidat InfiniGuard

V12 integrates natively with Infinidat InfiniGuard, including a dedicated UI wizard, storage detection logic and support for native block cloning capabilities.

Fujitsu CS800

V12 integrates natively with Fujitsu CS800, including a dedicated UI wizard, storage detection logic and support for native block cloning capabilities.

Tape

Backup

Expanded workload support — Backup-to-Tape jobs now support exporting any backup copies that were created by new multi-platform Backup Copy jobs in either Immediate or Periodic modes, regardless of workload type.

Virtual synthetic performance — Advanced data fetcher technology is now used for NFS-based backup repositories to significantly improve virtual synthetic-full export performance.

GFS monthly fulls with daily incrementals — Backups to a daily media set can now be configured with periodic monthly full backups in addition to the previously available weekly full backup option. This provides additional flexibility in tape GFS rotation.

Tape infrastructure

Tape server on Linux — In addition to Windows-based tape servers, you can now register tape libraries and tape drives that are connected to Linux servers.

LTO-9 support — All tape functionalities were made aware of the LTO-9 tape initialization process and will now correctly wait for the initialization to finish instead of timing out in case the initialization takes a long time.

Tape auto-eject — Tapes will now be automatically ejected from the tape library's drives upon the completion of Inventory and Catalog operations to prevent users from accidentally erasing them.

Waiting for tape notification enhancements — Tape jobs will now send a detailed email report whenever it's stuck in a situation that requires intervention from a backup administrator. This notification includes essential details like the tape library, media pool, media set, last written and oldest expired tape information to help you determine the most suitable tape to continue the job with.

Experimental features

The following experimental behavior mods can be used by creating the corresponding value under the HKLM\SOFTWARE\Veeam\Veeam Backup and Replication key on the backup server.

Disable extra barcode scans — Create the *TapeSuspendBarcodeValidation* (DWORD, 1) value to make tape media import operations avoid doing extra barcode scans. This can help accelerate the migration of tapes across libraries that are connected to the same backup server. Be sure you're only using the unique barcodes across all connected tape libraries before enabling this option.

Return expired tapes to free media pools — Create the *TapeMarkExpiredMediaFree* (DWORD, 1) value to immediately return any expired tapes into the Free media pool. With this mode, tapes will no longer be constantly assigned to the same media pool after it expires, which makes tape management less predictable. However, this also allows you to maximize tape capacity usage across all media pools.

Platform support

Microsoft Azure

Azure AD application support — Backup servers can now leverage service accounts (a.k.a. Azure AD Applications) to access Microsoft Azure resources like subscriptions, resource groups, storage accounts, etc.

Tags integration — You can now assign tags to restored Azure IaaS VMs to ensure they are still properly categorized according to your policies.

Microsoft Applications

Microsoft SQL Server 2022 — V12 adds support to SQL Server 2022 for application-aware processing, transaction log backups and in Veeam Explorer *for Microsoft SQL Server*. In addition, SQL Server 2022 is now supported for hosting a Veeam Backup & Replication configuration database.

Microsoft SharePoint Server Subscription Edition- Added support for the latest version of SharePoint Server, including application-aware processing and Veeam Explorer *for Microsoft SharePoint*.

System Center Virtual Machine Manager 2022 — Support for the latest version of SCVMM.

Microsoft Windows

Windows 10 22H2 and Windows 11 22H2 — This is supported as a guest OS for application-aware processing, for agent-based backup job protection and for the installation of Veeam Backup & Replication and its components.

Linux

V12 brings added support for the following version of supported Linux distributions as a guest OS for application-aware processing, for agent-based backup job protection and for the installation of Veeam Backup & Replication components. This includes:

- Ubuntu 22.04
- Red Hat Enterprise Linux (RHEL) 8.6, 8.7, 9.0 and 9.1
- SUSE Linux Enterprise Server (SLES) 15 SP4
- Oracle Linux (RHCK) 9.0 and 9.1

And in addition, for agent-based backup only:

- Oracle Linux (UEK) 9
- Fedora 36 and 37
- openSUSE Leap 15.4

VMware vSphere

vSphere 8.0 — In addition to basic compatibility provided by V11a, V12 delivers full vSphere 8 support, including new features like VMs with virtual hardware version 20 and vSphere DataSets backup and restore.

Cloud Connect

In addition to benefitting from core platform enhancements like IPv6 support, Veeam Cloud Connect V12 includes many new features and enhancements for service providers that offer BaaS, DRaaS and cloud repository as a service, the most significant of which are listed below:

CDP to Cloud Connect — In addition to regular replication jobs, CDP policies can now target a cloud host from a Veeam Cloud Connect service provider. On the service provider side, both VMware vSphere and VMware Cloud Director are supported as targets.

Instant Recovery as a Service — MSPs that provide BaaS can now instantly recover any tenant workload as VMware vSphere VMs by running them directly from unencrypted backups stored in cloud repositories.

Instant Recovery for Nutanix AHV backups — Tenants can now instantly recover AHV VMs from backups stored in cloud repositories.

Stronger NEA encryption algorithm — For increased security, we updated the encryption algorithms used by Network Extension Appliances (NEA).

Rental licensing enhancement — Incoming cloud-native backups that were created by Veeam Backup *for AWS* / *for Microsoft Azure* / *for Google Cloud* through a Rental license will no longer consume points on the Veeam Cloud Connect service provider side.

API Enhancements

In addition to adjusting our PowerShell SDK for compatibility with the aforementioned new features, here are just a few highlights of the most noteworthy additions to our APIs:

PowerShell

Backup engine

Upgrade to per-machine metadata — This new cmdlet for upgrading existing backup chain metadata may come in handy for an automatic upgrade of a large fleet of backup servers.

Active Full and Retry operations — New parameters in existing Start* cmdlets allow automation to trigger Active Full and Retry operations for individual machines in a backup job.

Detach backups — Use this new cmdlet to detach backups from a job. Detached backups will appear in the Orphaned node under Backups and are still subject to their last known time-based retention policy.

Apply retention policy — New cmdlet to granularly apply the retention policy to selected backups.

VeeaMover — New cmdlets to move machines between jobs, move and copy entire backups between repositories.

File-level recovery

Mount to backup console — A new cmdlet to mount backups to a server with the backup console.

Compare with production — New cmdlets and parameters that allow you to compare files in a selected restore point with the production machine and restore changed items only.

Stop recovery session — New cmdlet for stopping file-level and item-level recovery sessions and unmounting published backups.

Instant Recovery

Instant recovery to a Hyper-V VM — A new set of cmdlets for instant recovery of image-level backups from any platform to Microsoft Hyper-V.

Veeam Cloud Connect

Instant recovery for tenant backups — A new set of cmdlets that allow users to perform instant recovery from tenant backups on the service provider side.

REST API for backup server

Backup infrastructure

Backup server info — V12 has added a new endpoint with information on the backup server version, build and patch level.

SOBR management — Control extent mode, evacuate backups, rebalance and manage SOBR access permissions.

Backup

vSphere tags support — Specify tags in backup job settings, including application-aware processing.

Restore

Entire VM restore — Added support for entire VM restores, except for the Staged Restore mode.

Instant VM Recovery — Added full support for managing instant recovery, migration to production, obtaining a collection of all active IR mounts.

Recovery tokens — Added support for CRUD operations for agent recovery tokens, including expiration date prolongation.