Veeam Backup & Replication v12.1

What's New

Contents

Introduction

Veeam Backup & Replication, the workhorse of Veeam Data Platform, delivers enterprise-grade resiliency capabilities that ensure confidence in your protection, response, and recovery in the face of both classic disaster and modern cyberattacks across the hybrid cloud.

The following is a list of the major new features and enhancements added in Veeam Backup & Replication v12.1. All capabilities here are transacted as the Veeam Data Platform with certain features available only at the Advanced or Premium editions.

Major new features

Detect and identify cyberthreats

Minimize data loss from a cyberattack by detecting known threats and backup anomalies inline during backup or in real-time leveraging your existing Endpoint Detection and Response tools (EDR), submitting them into a SIEM system of your choice, and initiating proactive threat hunting to reduce further risk to your data.

The unique approach to malware detection in V12.1 includes:

Inline malware detection — performs inline, low-impact entropy analysis of a data stream to immediately detect previously unencrypted data becoming encrypted by ransomware using a specially trained Machine Learning (ML) model. Further, the same engine detects other signs of a malware or cyberattack such as Onion links, directly in a backup stream.

Because the analysis is performed by a backup proxy or a backup agent on a raw data stream that passes through them, there are no prerequisites or additional requirements. Just be sure your backup proxy CPU is not already overloaded, as in this case this additional data processing may slow down the backup speed.

Suspicious file system activity detection — searches guest file system indexes for files with known malware file extensions, ransom notes, and similar flags of malware presence. Analyzes file system activity by comparing guest indexes to detect suspicious changes like bulk deletes or renames of known "good" files (such as documents and images), many new files with previously unknown extensions appearing, and other activities which can be a sign of ransomware, a hacker or malicious insider activity.

This functionality works independently from the inline malware detection and requires guest file system indexing enabled in the backup job settings. Keeping the list of malware definitions automatically updated requires Internet connectivity to the Veeam update servers. Alternatively, you can always download the latest malware definition file manually, please refer to KB4514.

Early threat detection — Veeam Incident API makes it easy for external EDR tools (including NDR/MDR/XDR) to notify the backup server of infections at earlier attack stages, ensuring all restore points created after the corresponding moment in time for the given machine are marked as infected. Further, you can enable Veeam to instantly create an out-of-band restore point of the affected machine, before malware has a chance to do much damage.

Respond and recover faster from malware

Empower your team with threat visibility directly in the restore wizards, while providing them efficient tools to perform backup analytics, triage detected threats to slash incident response time and avoid reinfections.

Scan backups with YARA — once the malware strain affecting your environment has been identified, Veeam makes it easy to efficiently pinpoint them in your backups with YARA rules. This enables you to quickly find clean restore points and thus prevent reintroduction of malware into your production environment. Content scans can be performed ad-hoc using the restore point selection algorithm most suitable for the situation — sequential search from most recent backup for recent infections, or binary search to reduce the number of scans required when dealing with sleeping malware. Alternatively, you can schedule content scans to be performed periodically with SureBackup, which is helpful in the post-recovery phase to confirm clean recovery and no malware reintroduction. This functionality is available in Veeam Data Platform Advanced and Premium editions only.

Avoid reinfection with threat tracking — Veeam uses a robust, event-based approach to tracking potential threats and automatically marks all restore points affected by the corresponding event as Suspicious or Infected to prevent accidental restores of infected machines back to your production environment. And if a certain event is confirmed to be a false positive, all affected restore points receive a clean state automatically, without having to click through every one of them.

Event forwarding — regardless of what type of threat has been identified, Veeam provides a convenient way to forward the corresponding event to a SIEM system of your choice with the new support for Syslog, so your operations team can immediately react, perform triage and reduce further risk to your data.

Ensure security and compliance

Four-eyes authorization — prevent accidental or erroneous deletion of backups or entire backup repositories, changes to users, roles and other access settings using a backup console by requiring an approval from a second Veeam Backup Administrator before requested changes to these sensitive backup server settings can be applied.

Key Management Server (KMS) integration — eliminate the need for manual backup encryption key lifecycle management by integrating with the dedicated KMS server. Veeam takes the benefits of using KMS to a whole new level by leveraging asymmetric encryption keys, so intercepting an encryption key does not enable the attacker to decrypt backups, as is the case with competing offerings which leverage symmetric encryption. This also gives our customers more options to address backups exfiltration during a cyberattack, backups lost during transportation, etc. by effectively enabling them to remotely "destroy" all backups through deleting the private encryption key from the KMS server and thus making backups completely unrecoverable.

The integration is built on the KMIP (Key Management Interoperability Protocol) specification version 1.4. The latest versions of the following solutions were validated as supported by Veeam Quality Assurance (QA) as a part of the feature acceptance testing: Fortanix Data Security Manager, IBM Security Guardium Key Lifecycle Manager and Thales CipherTrust Manager. This functionality is available in Veeam Data Platform Advanced and Premium editions only.

Security and compliance monitoring — avoid accidental or temporary infrastructure changes exposing your infrastructure to hackers and affecting recovery success with periodic, automated scans by the new Security & Compliance Analyzer wizard. This ensures all backup infrastructure hardening and data protection best practices remain implemented and consistent with the baseline you yourself establish. Further, V12.1 checks dozens more best practices in addition to the existing ones from the V12 Best Practices Analyzer, or 3x more.

Veeam Threat Center — highlight detected threats, identify risks, and see the overall data protection score of your backup environment based on the level of security and compliance best practices implementation. Powered by Veeam ONE, available in Veeam Data Platform Advanced and Premium editions, this new dashboard is integrated directly into the backup console as a part of the new Analytics tab that also includes a few other dashboards that are most popular with our Veeam ONE customers.

Veeam App for ServiceNow

This new ServiceNow plug-in delivers an open integration leveraging ServiceNow application workflows and the backup server REST API to provide access to backup session data directly in ServiceNow. Your backup servers can be added directly to ServiceNow to provide visibility into its operation. The initial release of the app focuses on VMware vSphere protection and provides the following capabilities:

Backup visibility — you can now see your backup jobs, backup sessions and restore points in ServiceNow. This data can be used to create custom reports and dashboards directly in ServiceNow or can be used in other ServiceNow components.

Automatic incident creation — failed backups can automatically create ServiceNow incidents and will bring its failure messages into Veeam ONE automatically.

Backup dashboard — the app ships with a pre-built dashboard displaying recent backup history, recent incidents, and includes drill down capability for when you need more details.

Backup job templates — you can select any backup job to be used as a template for automatic creation of backup jobs during the VM provisioning process. An example catalog request workflow is provided, however the process can be integrated into your existing workflows as well.

Veeam App for ServiceNow is available in Veeam Data Platform Advanced and Premium editions only.

Backup of object storage

Protect production data residing in both on-premises and cloud object storage with the powerful new backup and recovery functionality. Built on a storage-agnostic architecture, Veeam helps you to achieve your object storage recovery objectives without additional hardware investments. And when it comes to recovery, we ensure your recovery time objectives (RTO) are met by providing flexible restore options tailored to typical disaster scenarios including: entire bucket recoveries to address complete hardware loss or object storage migration scenarios, point-in-time bucket state rollback for recoveries from ransomware and malicious insider attacks, and object-level recovery for day-to-day operational restores.

The unique benefits of Veeam's object storage backup engine include:

Scalable, storage-agnostic architecture — based on a proprietary distributed file system specifically built for the protection of billions of objects of PBs in size to a storage target of your choice. You can backup directly to another object storage or use a scale-out backup repository based on commodity server hardware with internal or directly attached storage to easily scale both performance and capacity of the distributed file system. Or you can even back up your object storage directly to tape!

Extensive data source support — including S3-compatible object storage on-premises or in a public cloud, Amazon S3 and Microsoft Azure Blob storage.

Flexible protection scoping — protect entire storage accounts, or only specific buckets/containers, objects with specific tags or prefixes, or individual objects. The same flexibility is supported for exclusions as well.

Efficient forever-incremental backup engine — that does not require periodic active full backups to be performed — making it possible to protect PB-size buckets with significantly reduced RPOs.

Storage-agnostic changed object tracking — our innovative approach provides industry-leading incremental backup performance, allowing you to achieve low RPOs without requiring native changed object tracking capabilities from object storage.

And it goes without saying that object storage backup jobs also include most of the standard Veeam capabilities like backup encryption, backup copy, pre- and post-job scripting for automation, and multiple notification options.

Veeam CDP enhancements

V12.1 is a major release for Veeam Continuous Data Protection (CDP) with significant improvements for performance and reliability. Thanks to 4x increase of the number of VMs and virtual disks supported under protection and 2x decrease in compute requirements, the improved CDP engine can confidently protect T1 workloads in some of the largest environments that our competition cannot even touch and is simply walking away from based on scalability considerations — the real story from a 12.1 beta tester!

Further, updates to the CDP engine now make it possible to run item-level recovery and recovery verification (test failover) on CDP replicas without stopping the actual replication. This and other enhancements enabled us to deliver the following new features and enhancements:

File-level recovery — use CDP replicas to perform file-level recovery from any point-in-time state, for example, right before a ransomware attack.

Application item-level recovery — restore application items from CDP replicas using Veeam Explorers. Unlike file-level recovery, this is supported from long-term, application-consistent restore points only.

Automated recoverability testing — based on popular demand, we added support for full recoverability testing (a.k.a. SureReplica) to CDP replicas, so you can be confident you can recover when the time comes.

I/O Anomaly Visualizer — perform point-in-time recovery to the moment just prior to malware infection or cyberattack with the new I/O Anomaly Visualizer helping to select the best time to failover to (right before the I/O activity burst), ensuring the lowest possible data loss.

Planned failover — zero data loss failover is now also supported for CDP replicas. For example, this allows you to easily switch over from your main data center in light of an incoming hurricane or similar disasters.

Failback to cluster — you can now select a vSphere cluster as a failback destination, and we will pick the right ESXi host for each machine automatically.

More control — change the target VM name during failback operation and set the desired CDP replica VMs disk type for VMs added to the CDP policy after its initial configuration.

vSphere DataSets support — CDP policies will now protect extended VM information stored in DataSets, replicating their current state to the replica VM and storing its previous versions as a part of long-term restore points.

Veeam AI Assistant

Get product usage advice and help from a GPT assistant built directly into your backup console! The AI assistant uses OpenAI model with all of our technical documentation as a context to provide quality answers about how to best use Veeam products.

The AI assistant logic runs in Microsoft Azure in a private, dedicated instance managed by Veeam so no conversation data is ever shared with Microsoft or with any other 3rd party. We also do not provide the AI Assistant any customer-specific information, neither from your production environment nor from your customer records.

Should you wish to disable this capability completely, create the *AIAssistantDisabled* (DWORD, 1) registry value under *HKLM\SOFTWARE\Veeam\Veeam Backup and Replication key* on the backup server.

Other Features and Enhancements

In addition to the above major new features, V12.1 includes hundreds of other enhancements that are a response to customer feedback and ongoing R&D findings, the most significant of which are listed below:

Backup engine

TLS 1.3 readiness — all backup infrastructure components have been verified and updated not to use hardcoded TLS versions to ensure future compatibility with TLS version 1.3.

Stronger backup encryption — more secure SHA256 hash is now used in the PBKDF2 algorithm for encryption keys generation. This change should dramatically elongate brute-force attacks on encrypted backup files. In addition, per FIPS 140-2 guidelines for 2024, PKCS #1 padding was replaced with more secure OAEP padding to better protect against other attacks on the encryption.

Sensitive data protection — saved credentials stored in the configuration database in an encrypted form using Microsoft Data Protect API (DPAPI) are now additionally protected with a randomly generated entropy value providing additional key derivation. This entropy is stored in the backup server registry key only accessible to the Local Administrators group, providing protection against attacks by unprivileged users and processes.

SHA256 fingerprints — SSH connections to Linux hosts will now use more secure SHA256-based fingerprints for host authentication.

Improved dedupe-friendly compression — the algorithm behind this compression level was optimized to reduce CPU usage and improve performance by 30%.

Backup data management

Retention options customization — after years of requests to add more options and/or to hide some existing options in the "Delete this backup in" functionality of VeeamZIP, Export Backup and Copy Backup features, we decided to simply enable users to customize this list to their liking via a flexible JSON configuration file that supports any number of custom retention values. For more information on how to create your perfect list of retention options to display to backup console operators, please refer to the User Guide.

Disable background retention processing — you shared with us that there are situations when you want to completely freeze backup retention processing, for example immediately following a cyberattack — to ensure your last good backups don't get deleted by the retention policy. Since disabling your backup jobs is no longer enough due to the presence of the background retention processing, we added the ability to disable this functionality as well by right-clicking the Backups node.

Move backup improvements — VeeaMover will now detect if both source and target backup repositories reside on the same storage volume, in which case it will leverage native file system functionality for near-instant data movement.

Health check email notification — storage-level corruption guard functionality will now send email notifications with the session results. This functionality was lost in V12 due to the separation of health checks into a separate activity, instead of being an integration part of backup jobs and thus its email notifications.

Backup infrastructure

Configuration database

PosgreSQL database scalability — following multiple optimizations based on support cases from early V12 adopters, V12.1 doubles the recommended scalability limit of using PostgreSQL as a configuration database, and based on extensive V12.1 performance testing is now supported for backup servers protecting up to 10,000 machines with up to 1,000 concurrent tasks.

Database load reduction — optimized backup service interaction with the database, noticeably reducing background configuration database load regardless of the database engine used.

Backup proxies

Improved multipath support — Linux-based VMware backup proxies should now leverage multipathing more reliably when performing backup from storage snapshots, by using larger timeouts to allow for the path selection process to complete. The new logic is automatically enabled whenever we detect the *multipathd* daemon running on the backup proxy server.

Backup repositories

Accelerated backup import — increased backup repository rescan performance, including during the backup import process — up to 3x faster for image-level backups and up to 50x faster for NAS backups.

Hardened backup repositories

Time change protection for hardened repositories — our hardened repositories will now automatically cease removing immutable flags from backup files following a significant system time jump, which can be a sign of attempted cyberattack. Should this happen, you will be required to reset time protection by removing the */etc/veeam/immureposvc/retainLock* lock file as root.

To protect against all time-based attack vectors, the logic ensures that the time progresses forward continuously and monotonously, aside from short downtimes caused by reboots and maintenance. For this reason, environments where long repository downtimes are expected (more than 24 hours per year, SLA 99.73%) may consider changing the default thresholds. Please refer to the User Guide for additional information on how to customize the tolerated time deviance values.

Reduced number of active ports — to further reduce attack surface, hardened repositories will now continuously listen to the single Transport port (TCP 6162) only, instead of two ports. The Deployer port (TCP 6160) will be automatically opened and closed only for the duration of the hardened repository components upgrade.

Object storage repositories

Object storage improvements — multiple under-the-hood optimizations were implemented to significantly improve the performance of object storage repository operations like backup rescan, import, upgrade and removal while reducing object storage load at the same time.

Governance immutability mode — S3 Object Lock command used to make backups immutable on S3-compatible object storage can now be set to the Governance mode. This mode allows users with specific permissions to delete objects before their retention period ends, which can be useful for example to Managed Service Providers (MSPs). This setting is global, affects all object storage repositories and cannot be enabled on existing backup servers where S3 immutability is already in use. Enabling this setting will also trigger a warning in the Security & Compliance Analyzer which can be suppressed. To enable this mode, create the *S3GovernanceImmutabilityMode* (DWORD, 1) registry value under *HKLM\SOFTWARE\Veeam\Veeam Backup and Replication key* on the backup server.

Scale-out backup repositories

Export backups enhancements — exporting restore points residing on Capacity Tier can now be performed without downloading data to Performance Tier first, by leveraging native copy commands of object storage whenever object storage supports such commands.

Backup consistency check — you can now schedule periodic verification of backups stored on Capacity Tier. Known as "health check lite" and first introduced in V12 specifically for object storage repositories, this will verify consistency of the latest restore point by ensuring all required objects are still present on object storage, but without downloading actual data to keep the costs of this operation in check.

Restore from Capacity Tier improvements — initiating a restore from backups stored in the Capacity Tier (using the Backups > Object Storage node) should no longer read matching data blocks from the Performance Tier when restoring backups other than latest. This capability can be useful for performing periodic non-disruptive restore tests directly from the cloud without putting your Performance Tier into maintenance mode.

Decryption performance improvements — increased encrypted backups import performance from Capacity Tier by accelerating decryption by up to 5x compared to V12.

Immutability enforcement — to prevent accidental configuration mistakes, scale-out repositories will enforce the same immutability settings for all extents of a particular tier and will not allow a mix of immutable and non-immutable extents for all supported extent types.

Configuration backup

Immutable configuration backups — configuration backups will now be made immutable when using object storage backup repositories with immutability enabled.

Unattended restore — to facilitate unattended or mass deployment scenarios, you can now perform automatic configuration restore silently by using an answer file which can be obtained directly from a configuration restore wizard or generated from a command line.

Database verification — the product configuration analysis first introduced in the new V12 setup will now be performed during the configuration restore as well, making you aware of potential problems like unsupported or not recommended configurations which proceeding with the configuration restore may result in.

Image-level backup

Backup Copy

Improved scalability — following many optimizations, V12.1 quadruples the recommended scalability limit for backup copy jobs to a maximum of 2,000 machines per job.

SureBackup

SureBackup lite — the new backup verification mode enables you to check backup data consistency and perform a backup content scan without testing the backup for recoverability. This mode does not require configuring virtual labs, which significantly lowers the barrier of entry to the backup verification.

Machine exclusions — by popular demand, we added an ability to exclude objects in linked backup jobs from processing by SureBackup jobs.

Random machines testing — you can now instruct SureBackup jobs to only test a certain number of random machines from the linked backup jobs. Note that the list of machines will always be truly random and does not reference any previous SureBackup job runs when picking a list of machines to process in the given run.

VMware NSX-T networking support — VMware VMs connected to NSX-T segments will now be handled properly inside virtual labs. Both production and isolated networks of a virtual lab itself can now also be set to an existing NSX-T segment.

Recovery from image-level backups

Instant recovery

Instant disk publish — you can now mount any disk from backup to any Windows or Linux machine directly from the backup console, enabling you to instantly access the content of any Veeam backup. This capability can be leveraged for special file-level recovery scenarios like bulk recovery with native tools, or — with the help of third-party applications and scripts — also for data mining, security analysis, data compliance checks, and other scenarios requiring scanning of data in backups.

File-level recovery

Mount to the original host — based on popular demand, the Linux file-level recovery wizard will now suggest mounting Linux backups back to the original machine by default, streamlining and accelerating the file restore experience.

Restore To performance improvements for Linux — restores from Linux agent-based backups to non-original Linux servers (using the Restore To option) are now over 10x faster for large directories.

Restore to the original location for AIX/Solaris — restoring files to the original location from AIX/Solaris agent-based backups can now be done entirely from the backup console, without requiring users to copy files to their original location manually.

Secure Restore

Additional antivirus integration — V12.1 adds experimental out-of-the-box integration with McAfee VirusScan Enterprise and SentinelOne Endpoint Security based on the input from Veeam community experts.

Veeam Explorer for PostgreSQL

Instant Recovery of PostgreSQL instances — an accidental update query in a database led to data loss or corruption? Not a problem because Instant Recovery leverages the already existing Publish engine and makes an instance available for production use in just a few seconds! All changes made to a published instance are preserved in cache while the backup itself stays intact. Actual data restore to the production storage happens in the background, with additional synchronization of the published and the target instances thanks to the native PostgreSQL streaming replication technology. To finalize the recovery, you need to switch the instance over to running from the production storage, which can be done with minimal downtime. This switchover can be done manually or scheduled to happen automatically or as soon as the synchronization catches up.

PostgreSQL database export — to simplify Dev/Test scenarios, you can now export single or multiple PostgreSQL databases in the native *pg_dump* format to any point in time, which can be restored later using the *pg_restore* tool. Export can be done to the local host or any Linux server either directly from a backup or from a published PostgreSQL instance while preserving the actual database state with all changes that happened while the instance was published.

Veeam Explorer *for Microsoft SQL Server*

Database restore without admin share — we enabled restoring of Microsoft SQL Server databases without leveraging admin$ share on the target SQL Server. The restore process will now first attempt to connect to the Veeam Installer Service which must be deployed manually or automatically (by adding a server to the Managed Servers in Veeam) to the target SQL Server. If this service is not available, we will failover to using the legacy behavior that leverages the *admin$* share.

Veeam Explorer *for Oracle*

Multiple Oracle homes support — configurations with several Oracle homes of different versions are now supported on Linux-based VMs thanks to a new Oracle proxy process that interacts with the database environment and circumvents the dependency of the OCI library on a specific Oracle version.

Veeam Explorers

Gmail and Microsoft 365 support email notifications — for Veeam Explorer email notifications, in addition to basic SMTP servers, V12.1 now supports Google Gmail and Microsoft 365 with their OAuth 2.0 protocol-based secure authorization and access-token-based authentication.

Agent Management

AIX & Solaris

Centralized agent management — deployment and management for AIX and Oracle Solaris backup agents now can be performed in the same convenient way as with Veeam Agents for other OSes, without having to manually manage agents directly on the protected system.

Linux

Deployer service — manage Linux servers without requiring SSH connectivity and without storing OS credentials on the backup server with the new Deployer service that allows secure connection to the protected system over a single port and uses certificate-based authentication.

Reduced sudo permissions — based on common feedback from enterprise security teams, we significantly reduced the number of special sudo permissions required for backup servers to manage Linux agent.

Nosnap official support — nosnap packages are now officially supported for protection groups with preinstalled agents, enabling faster Linux upgrades for customers who are ready to use LVM snapshots over using the Veeam kernel module. This includes nosnap packages for newly supported Linux on Power.

Protection groups with pre-installed agents

Reduced footprint — the Installer service no longer needs to be deployed to Windows computers with pre-installed agents, making the agent installation easier and smaller.

Managed by Agent jobs

Enhanced scale-out repository support — backup agents can now back up to scale-out repositories with Capacity and Archive tiers enabled, allowing you to copy your backups for redundancy and offload older restore points to a cheaper object storage class to reduce costs.

Background retention enhancements — orphaned backups created by Managed by Agent jobs pointed to an object storage repository are now also subject to the background retention processing.

Agents

Veeam Agent *for Microsoft Windows* 6.1

Microsoft Windows 11 23H2 support — added full support for the latest version (23H2) of Microsoft Windows 11 operating system. This includes support for recovery media creation and correct operating system version detection.

Malware detection — when backing up to a Veeam repository, standalone and self-managing agents include the suspicious file system activity detection functionality explained at the beginning of this document (Managed by Server agent-based backup jobs provide full malware detection capabilities).

Scale-out backup repositories from object storage support — starting from this version, standalone agents can store their backups in a scale-out repository with Performance Tier backed by a single object storage extent.

Backup mapping support for object storage repositories — standalone agents now support mapping of the existing backups stored in an object storage repository to a new backup job.

Veeam Agent *for Linux* 6.1

Linux on Power support — we've ported the agent to IBM Power architecture with support for SLES version 12 SP5, 15 SP3 and 15 SP4 (including SLES for SAP); as well as RHEL versions 8.4 and 8.6 (including RHEL for SAP).

Nosnap package — nosnap packages are now officially supported, enabling faster Linux upgrades for customers preferring LVM snapshots over using the Veeam kernel module.

Latest Linux distributions versions support — this agent version adds full support for Debian 12.1 and 12.2; Fedora 39; RHEL 8.9, 9.2 and 9.3; Ubuntu 23.10.

Veeam Agent *for Mac* 2.1

macOS 14 (Sonoma) support — added full support for the latest version of macOS operating system.

Enhanced backup scheduler — Veeam Agent *for Mac* now features the same extensive job scheduling and retry options as other backup agents.

GFS retention support — macOS machines protected with a server license can now leverage GFS retention for the created backups.

Veeam Agents *for AIX* and *Solaris* 4.1

Since our Unix agents share much of the codebase, most new functionality applies to both agents.

Backup engine improvements — reduced RAM consumption by agents by over 5x while improving the backup speed when backing up a large number of files.

Recovery tokens — a simplified way to provide users who are performing a Bare Metal Recovery with access to a particular backup that was first introduced in V12 has now made its way to our AIX and Solaris agents as well. 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.

File-level recovery TUI — added a new text-based user interface (TUI) that dramatically simplifies day-to-day file restore activities by allowing you to search for and restore individual files from backups.

GFS retention — you can now specify GFS retention for the created full backups.

Recovery media as OVA — the recovery media is now available in an OVA format to allow performing restores in production infrastructures and systems that do not support ISO mount (AIX only).

Hardware accelerated CRC — the new agent version uses hardware-assisted checksum calculation to boost backup speed (AIX only).

Restore performance improvements — restore performance has been improved up to 2x for global zones and up to 8x for non-global zones (Solaris only).

ZFS compression awareness — the agent is now aware of ZFS compression and will consider it when providing estimations on the backup size (Solaris-only).

Application plug-ins

General

Scalability improvements — to improve performance and reliability in large environments, application plug-ins in V12.1 will reuse the same target backup file instead of creating multiple backup files, the approach first piloted in the Veeam Plug-in for Microsoft SQL Server released in V12. The target backup file will remain open for up to 24 hours by default and consolidate data produced by multiple backup sessions. This approach dramatically reduces the number of backup files, relieving backup storage and configuration database from additional load while accelerating repository rescan.

IBM Db2

Veeam Plug-in for IBM Db2 — the new addition to our application plug-ins family allows you to stream native Db2 backups directly to Veeam repositories. The plug-in provides support for a large scope of deployment types including standalone database servers, HADR systems and failover clusters managed by IBM PowerHA, TSA or Pacemaker. And from a platform support perspective, the plug-in is supported on both Linux x86_64 and AIX platforms.

SAP HANA

Veeam Explorer for SAP HANA — wizard-driven SAP HANA restore experience is now available for application plug-in backups. This allows backup admins to restore SAP HANA databases with ease even without having extensive SAP administration background. Restore is possible to the original or a different location, to any point in time or to the specific data backup.

Veeam Explorer for SAP HANA combines a simple and intuitive Veeam restore experience while preserving close connection to the native recovery tools like SAP HANA studio. And thanks to the service-based architecture, you don't need to keep the backup console open while performing the restore.

The explorer does not leverage SSH and instead connects to the SAP system via web service using HTTP(S). To ensure secure communication with SAP HANA services with TLS enabled, SAP Common Crypto library must be installed on the backup server.

IBM Power support — looking for an SAP-certified backup solution to protect your SAP HANA systems running on Linux on IBM Power? Just use the updated Veeam plug-in for SAP HANA which now provides its full functionality set also on Power. The plug-in works in standalone mode and can handle various SAP HANA deployment types: standalone HANA servers, scale-out systems and includes support for HANA system replication.

Backup appliances

General

Appliance instance type selection — tune backup appliance compute resources according to your needs to improve scalability or reduce costs by changing the default cloud machine instance type.

Veeam Backup *for AWS*

Amazon S3 protection — backup of Amazon S3 buckets and objects to any Veeam supported repository, including to another Amazon S3 repository, Azure Blob and more. Flexible restore options include entire buckets to a new bucket for DR or migrations, point-in-time rollbacks, and granular object-level recovery.

Amazon DynamoDB protection — manage your DynamoDB database protection alongside other AWS workloads you're protecting with Veeam. Use flexible policy-based scheduling and lightning-fast database recovery to a point-in-time to the original or a new location. Make backup copies across AWS regions to protect yourself against many types of possible outages and data loss.

Amazon RDS for PostgreSQL protection — add an additional layer to your PostgreSQL database protection on top of native snapshots by enabling archiving of backups to Amazon S3, S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive, including optional immutability through Amazon S3 Object Lock.

Granular per-workload permissions — you can now grant only required permissions for IAM roles to work with specific workloads for backup and/or restore operations. This allows maximum flexibility in infrastructures with tight security regulations.

Private network deployment — the Veeam Backup *for AWS* appliance and its workers can now be deployed in networks with no public access to endpoints or object storage to meet the internal security requirements and regulations. The updated Getting Started section of the user interface will guide you through the required configuration steps to make the backup appliance operate while being accessible over the private IP address only.

Worker tagging — increase visibility and control over automatically provisioned worker appliances in your production infrastructure by assigning predefined or custom tags to workers. This will allow you to meet the requirements for newly created cloud resources that are often enforced by cloud management and security teams, enabling them to easily track their usage and resource consumption.

Backup performance improvements — increased EC2 instances backup speed up to 5x depending on scenario, through leveraging multiple copies of the EBS volumes with round-robin read logic.

Standard Accelerated Retrieval support — you can use this new retrieval mode when you need to retrieve your data from Amazon S3 Glacier Flexible Retrieval storage even faster.

Appliance autoconfiguration — appliances on EC2 instances with ephemeral disks will now be automatically configured with swap files.

Worker OS update — operating systems for workers has been upgraded to the latest Amazon Linux v2023.

Repository ownership monitoring — the same backup repository can no longer be added to different backup appliances to increase reliability and protect against potential data loss.

Veeam Backup *for Microsoft Azure*

Azure Blob protection — backup of Azure Blob containers and objects to any Veeam supported repository, including to another Azure Blob repository, Amazon S3 and more. Flexible restore options include entire buckets to a new bucket for DR or migrations, point-in-time rollbacks, and granular object-level recovery.

Azure virtual network configuration protection — changes in virtual network settings may sometimes lead to unexpected results and production environment disruption. When protecting your network configuration settings with Veeam, you will always be able to easily identify what settings have changed in your production environment compared to the recent backups and easily restore only changed objects or properties. Additionally, you can re-create the entire production network configuration for migration or testing purposes by restoring all virtual network resources to a new location.

File-level recovery to the original location — restoring files back to the production environment takes just a few clicks now, helping you perform this most common and frequent recovery type much faster than before.

Granular, per-workload permissions — you can now grant only required permissions to service accounts used for backup and restore operations, limiting them to specific subscriptions (management groups) and to specific workloads. This allows maximum flexibility in infrastructures with tight security regulations.

Private network deployment — the Veeam Backup for Azure appliance and its workers can now be deployed in networks with no public access to endpoints or object storage to meet the internal security requirements and regulations. Further, we have switched to using Azure Queue Storage messaging service to reduce operating costs in the private network deployment mode. The updated Getting Started section of the user interface will guide you through the required configuration steps to make the backup appliance operate while being accessible over the private IP address only.

Worker tagging — increase visibility and control over automatically provisioned worker appliances in your production infrastructure by assigning predefined or custom tags to workers. This will allow you to meet the requirements for newly created cloud resources that are often enforced by cloud management and security teams, enabling them to easily track their usage and resource consumption.

Support for VM with SSD v2 and Ultra disks — you can start using this new VM disk type to improve the I/O performance without losing the ability to protect these VMs.

Veeam Backup *for Google Cloud*

Cloud Spanner protection — manage your Google Cloud Spanner protection alongside other Google Cloud workloads protected with Veeam. All Cloud Spanner configurations are supported, from a single database to a complex, multi-region Cloud Spanner deployment.

Cloud SQL backup policies enhancements — you can now configure user credentials on a per-instance basis for PostgreSQL and MySQL backup policies.

File-level recovery to the original location — restoring files back to the production environment now only takes a few clicks, helping you perform this most common and frequent recovery type much faster than before.

Shielded VMs support for workers — implement Google Cloud security best practices by enabling the Shielded VM option for automatically provisioned Veeam worker appliances. Once enabled via the appliance config file, Shielded VMs will be used for all Veeam workers. For instructions on editing the configuration file, please refer to the product documentation.

Private network deployment — the Veeam Backup for Google Cloud appliance can now be deployed in networks with no public access to endpoints and object storage to meet the internal security requirements and regulations. Private network deployment mode can be enabled via the appliance config file. Manual production network setting changes may be required, please refer to the product documentation. For instructions on editing the configuration file, please refer to the product documentation.

Worker tagging — increase visibility and control over automatically provisioned worker appliances in your production infrastructure by assigning predefined or custom tags to workers. This will allow you to meet the guidelines for newly created cloud resources that are often enforced by cloud management and security teams, enabling them to easily track their usage and resource consumption.

Veeam Backup *for Nutanix AHV* 5

V12.1 substantially extends the capabilities of the recently released Veeam Backup *for Nutanix AHV* v5 with newly introduced Veeam Data Platform features documented above, enabling you to:

- Test or scan your AHV backups with the new SureBackup-lite mode.
- Perform ad-hoc backup content scan with antivirus engines and YARA rules.
- Protect your AHV backups with “four-eyes” authorization and KMS-based encryption keys.
- Instantly publish disks to Windows or Linux hosts directly from AHV backups.

Veeam Backup *for Red Hat Virtualization* (RHV) 4.0

Fully integrated solution — RHV backup and restore wizards have been integrated directly into the backup console, just like the corresponding VMware and Hyper-V wizards.

Multiple workers support — accelerate backup and restore performance in large environments with parallel VM processing by using multiple workers, enabling Veeam to automatically distribute processed VMs among them.

BitLooker support — the existing Veeam Backup & Replication functionality that allows skipping blocks belonging to deleted files, as well as swap and hibernation files, is now available also for RHV VM backups, significantly decreasing their size.

VeeamZIP backups — you can now create ad-hoc VeeamZIP backups also for RHV VMs, for example for archiving VMs before they are decommissioned.

Expanded disk restore capabilities — you can now restore individual RHV VM disks, or instantly publish their content to any Windows or Linux machine.

Overload alerts — we will now warn you if the backup appliance CPU or RAM remains overloaded for an extended time, so that you could consider granting it more resources.

Update management — you can now control updates to be installed more granularly (i.e. Veeam component updates vs. OS component updates).

Containers backup

Restore sessions visibility — in addition to the backup session data, V12.1 adds support for displaying Kasten K10 restore sessions, providing more complete insight into K10 operations directly in the backup console.

Improved performance and scale — the Kasten K10 integration has been optimized and tested to perform well with up to 50 K10 instances with a total of over 1,000 K10 policies, allowing it to be used in some of the largest deployments.

NAS backup

Compare to production — compare a selected restore point to the production file share with the new Backup Browser view that enables you to effortlessly identify all files and folders 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.

Reworked inclusions and exclusion — include and exclude masks have been streamlined to provide a cleaner and simpler way to configure file share backup job inclusions and exclusions. And we've also added the ability to export and import previously created filters.

NetApp ONTAP FlexGroup support — added the ability to backup and recover files from/to file shares backed by FlexGroup volumes.

Enterprise load-balancers awareness — NAS backup job is now aware of NetApp ONTAP and Dell EMC Isilon load balancers and can perform backups from different storage nodes depending on their current load, instead of always using a single node.

Hardened repository optimizations — applying immutability to NAS backups and a few other operations should now be performed significantly faster thanks to the new caching service.

Flat data sets optimizations — enumeration of large flat file share structures during backup and opening folders containing a very large number of files with Backup Browser was accelerated up to 2x.

Backup console

Dynamic data refresh rate — we will now dynamically change the UI update frequency depending on the environment size to avoid excessive load on the configuration database in large environments. The UI should now also update its content more reliably and consistently without having to refresh it manually by pressing F5.

Backup properties improvements — added the ability to search for an object by its name directly in the backup properties window and made other quality of life improvements to user experience as well. In addition, the content of this dialog should now load much faster in large environments.

Encryption password checks — the user interface will now check the supplied password length and warn the user that they need to be at least 12 characters long to ensure the backup encryption cannot be easily brute forced in case your backups get stolen. Further, the password hint will now be validated to ensure it does not contain your encryption password.

Files functionality access — due to the growing number of security-related concerns, we have restricted access to most operations on the Files node of the management tree to Backup Administrator role only.

Log export — accelerated support log bundle creation with a more efficient compression algorithm implementation and multi-threading.

Enterprise Manager

Managed by Agent policies support — such backup policies will now also be visible in the Enterprise Manager, in addition to the already visible Managed by Server jobs.

Job type filtering — added additional job filtering options to simplify searching for a job in large environments.

Bulk file restore — we made it easier to add the entire file search results into the Restore list with the new Add Everything option.

Microsoft Exchange restore improvements — you can now restore Microsoft Exchange application items to any domain. To enable this option, select the corresponding checkbox in the Active Directory account settings and the restore wizard will prompt you for the admin credentials of the desired domain.

Backup server version — Enterprise Manager will now display the full backup server build number including patch level.

Initial data collection — we will now start the first data collection session automatically for any new backup server added to the Enterprise Manager, so you no longer need to initiate it manually.

Licensing

Socket subscription license — added in-product support for a Socket-based license with the Subscription contract type. Socket-based licenses with the Perpetual contract type remain supported by the product as before.

Setup

Unattended install — to facilitate unattended or mass deployment scenarios, you can now install, upgrade, and remove Veeam Backup & Replication silently with a single command, by using an answer file to specify all the installation options. The new silent installation replaces the former unattended installation approach of automating installs of multiple MSI files separately.

Sign-in to license — retrieve and install your production license directly from the setup wizard! All you need is your Veeam Customer Portal account credentials. And if you don't have a license yet, you will be provided an option to get a trial license instead.

Disk space check — the setup wizard will now display free disk space recommendations to ensure successful installations and upgrades.

Private hotfixes check — the setup wizard will now detect private hotfixes present in your existing installation and refuse to proceed with the upgrade in case they are not included in the current version. This should help to prevent possible issues following the upgrade. If you get such a message, please contact our Customer Support for further assistance.

ISO size — we managed to reduce the ISO size by almost 10% despite adding many new components and capabilities!

Storage integrations

Object storage

Smart Object Storage API extension — the new archiving extension of SOSAPI supports S3 Glacier-compatible storage systems, enabling the usage of such storage as an Archive Tier of a scale-out backup repository.

AWS Glacier accelerated restores — added support for leveraging S3 Batch operations to increase data retrieval performance, speeding up restores from scale-out backup repository Archive Tier.

Microsoft Azure Blob Storage Cold Tier support — promising the same throughput and access latency as the Hot and Cool tiers, this new archive storage class makes the best candidate for Archive Tier storage when your use cases demand instant access to your archived backups.

Microsoft Azure Entra ID authentication support — you can now leverage service accounts (applications) to access Microsoft Azure Blob Storage resources, which is a more secure approach recommended by Microsoft over using shared keys.

Google Coldline storage class support — with slightly lower availability, 90-day minimum storage duration and higher costs of data retrieval, this new storage class may provide a better choice than Standard and Nearline classes for cost-effective long-term storage of infrequently accessed backups.

Primary storage arrays

Universal Storage API 2.0 enhancements — added automatic staging for restore operations from archived snapshots to avoid backup admins having to wait for the download to be completed before they can start a restore process.

HPE Alletra MP Fibre Channel (FC) support — added support for storage snapshot integration functionality to HPE Alletra MP on an FC fabric.

IBM Spectrum Virtualize — adopted new method of snapshot creation introduced in IBM Spectrum Virtualize version 8.5.1.

NetApp ONTAP FlexGroup support — added support for FlexGroup volumes for storage snapshot integration functionality.

Cisco HyperFlex spanned VM support — added support for native snapshot creation of VMs with virtual disks located on different datastores.

VMware vSAN Express Storage Architectures (ESA) support — in addition to the existing support for original vSAN architecture support, V12.1 adds support for modern vSAN ESA as well.

Secondary storage

HPE StoreOnce immutability for Catalyst Copy jobs — you can now leverage native StoreOnce immutability for backups copied between Catalyst stores with Catalyst-integrated backup copy jobs.

Dell Data Domain Retention Lock support — added support for creating immutable backups on Dell Data Domain by leveraging the native Retention Lock functionality.

Dell Data Domain DD OS and DD Boost version support — V12.1 adds support for DD OS versions up to 7.12 and updates the included DD Boost SDK to version 7.11.

Immutability period override — backups on HPE StoreOnce and Dell Data Domain are made immutable taking into account the minimum and maximum immutability period values in the storage configuration.

Import of backups — backup import functionality has been expanded to also support importing backups from HPE StoreOnce and Dell Data Domain.

Tape

IBM 3592 (Jaguar) tape support — leverage your existing investments into enterprise tape infrastructure for air-gapped, low-cost archive storage of your Veeam backups with added support for IBM 3592 tape libraries and media.

Tape content view is back — based on popular feedback, we've restored the ability to view tape content by going into the tape media's properties as before.

Distributed File System (DFS) as source — File to Tape jobs are now capable of crawling through the DFS, enabling backup of the entire DFS structure to tape without requiring workarounds.

File to Tape licensing details — when upgrading from previous versions, the setup wizard will notify you about V12 File to Tape licensing changes if the corresponding jobs are present.

Reworked inclusions and exclusions — include and exclude masks have been streamlined to provide a cleaner and simpler way to configure File to Tape backup job inclusions and exclusions. And we've also added the ability to export and import previously created filters.

Import tape option — we've added the import option directly to the Tape node of the management tree to simplify your day-to-day tape management activities.

Action log optimizations — based on your feedback, we tuned the Backup to Tape action log to eliminate unnecessary log lines when processing backup jobs using updated V12 per-machine chains.

Scalability and performance improvements — through our support case learnings we've implemented multiple performance optimizations to our tape backup engine and accelerated many tasks such as image-level backup task building process, incremental backup job for NAS backups, file to tape backup and restores, loading the list of backups on tape in the UI, tape backup encryption and decryption and more.

Platform support

VMware vSphere

Full vSphere 8.0 Update 2 support — in addition to basic backup and restore compatibility with vSphere 8.0 U2 provided by V12, version 12.1 adds full support for this vSphere version, including support for VMs with virtual hardware version 21.

VMware Cloud Director

Full VMware Cloud Director 10.5 support — in addition to basic backup and restore compatibility provided by the corresponding V12 hotfix, V12.1 delivers full VMware Cloud Director 10.5 support including vSphere Web Client plug-in and all Veeam CDP functionality.

Cloud Director self-service portal customization — tailor self-service backup and restore portal actions available to tenants according to your needs with the ability to hide unnecessary tabs.

Amazon AWS

EC2 Helper Appliance IMDS update — we updated EC2 helper appliances' instance metadata service to the latest version for added security.

AWS SDK update — AWS SDK was updated to the latest version to enable support for new AWS regions (namely Israel).

Microsoft Azure

New instance types for helper appliances — the list of Azure VM types recommended for helper appliances has been updated with newly added instance types.

Cloud Connect

General

PostgreSQL database support — following multiple optimizations and QA testing, starting from v12.1 we're removing the previous recommendation not to use PostgreSQL as a configuration database for

Veeam Cloud Connect infrastructures.

Maintenance mode enhancements — running Cloud Connect Replication jobs and Cloud Connect CDP policies will now be stopped gracefully and prevented from starting automatically for the duration of the maintenance mode.

Tenant quota recalculation — actualize tenant resource consumption with the new *Sync-VBRCloudTenantResource* and *Sync-VBRCloudTenantReplicationResources* cmdlets.

Tenant logs export improvements — collect Cloud Connect logs only for the specific tenant requested by Veeam Support more easily with the new *-Tenant* parameter for *Export-VBRCloudTenantLogs* cmdlet.

Cloud Connect Backup

Backup copy job from external repository — backups created by Veeam Backup *for AWS*, Veeam Backup *for Microsoft Azure*, or Veeam Backup *for Google Cloud* can now be copied from external repositories to cloud repositories, enabling managed service providers (MSP) leveraging Cloud Connect to protect additional workloads for their clients.

Transaction log backup copy to object storage — backup copy job now supports processing of transaction log backups of Microsoft SQL Server, Oracle, and PostgreSQL to object storage backup repositories.

Backup copy optimizations — significant reduced load on the Cloud Connect infrastructure from backup copy jobs containing a large number of small machines.

New restore types for tenant backups — you can now publish backups content using *Publish-VBRBackupContent* and perform instant disk recovery with *Start-VBRViInstantVMDiskRecovery* cmdlets from unencrypted tenant backups.

Immutable backup chain management — revert immutable backup chains state to an earlier point in time for the particular tenant only with the enhanced *Sync-VBRObjectStorageRepositoryEntityState* and *Sync-VBRSOBREntityState* cmdlets now supporting *-TenantBackupId* as an input.

Cloud Connect Replication

Improved CDP visibility — Cloud Connect CDP policies and task sessions will now contain more detailed statistics and action results to help you with the initial troubleshooting.

Network Extension Appliances (NEA) isolation — vCD network extension appliances on the provider side are now deployed in a standalone vApp to prevent mutual influence between them and the replicas.

PowerShell SDK

We extended our PowerShell SDK to enable management of the aforementioned new features. Here are just a few highlights of the most noteworthy additions to our APIs which you are more likely to use in your automation efforts:

SureBackup — leverage new SureBackup lite mode, perform adhoc backup content scan, manage the new exclusions list.

Malware detection — manage malware detection exclusions, report on detection events and manage infected objects, configuring scanning settings and exclusions.

Object storage backup — full automation capabilities around object storage protection to disk and to tape, including adding and managing data sources, configuring jobs, reporting protection status and performing restores.

Syslog integration — add and configure syslog server parameters.

KMS integration — add and configure KMS servers, configure new and existing backup jobs and repositories to use KMS-based encryption keys.

Security & Compliance analyzer — trigger ad-hoc infrastructure check against the full list of best practices.

Global exclusions list management — manage the global list of machines that should be excluded from processing by any jobs.

Backup server info — retrieve backup server information like its build number and patch level.

REST API

Backup server REST API has been expanded with multiple new capabilities around existing and newly added features.

File-level recovery — perform individual file recovery from vSphere and agent-based backups, featuring a fully fetched backup content browser with custom search capabilities.

VMware Cloud Director backups — manage backup jobs and perform restores back to Cloud Director.

Replication jobs management — manage replication jobs for VMware vSphere VMs, perform failover and failbacks.

Malware detection — create your own malware detection events using the information from 3rd party cybersecurity systems and processes with new Veeam Incident API endpoints.

Syslog integration — add and configure syslog server parameters through REST API.

KMS integration — manage KMS servers with CRUD requests and configure new and existing backup jobs and repositories to use KMS-based encryption.

Google service credentials — added support for GCP service account credentials in REST API. Specify a JSON account key generated by Google Cloud and use it in your resources.

This document includes only features and enhancements first introduced in version 12.1. If you are looking for the list of new features for version 12.0, available since February 14, 2023, you can find them in the corresponding document following this link.