




# Veeam Backup & Replication

## v13

What's New



---

# Contents

<b>Introduction</b>	<b>4</b>
<b>Veeam Software Appliance</b>	<b>4</b>
Pre-built	4
Pre-hardened	5
Predictable	5
All-in-one Appliance	5
Licensing	6
<b>Veeam Infrastructure Appliance</b>	<b>6</b>
<b>Web UI</b>	<b>7</b>
<b>High Availability</b>	<b>8</b>
<b>Enhanced Role-Based Access Control (RBAC)</b>	<b>9</b>
<b>Single Sign On (SSO)</b>	<b>10</b>
<b>Veeam Intelligence</b>	<b>11</b>
<b>Other Features</b>	<b>12</b>
Platform	12
Configuration Database	13
Backup Engine	13
Malware Detection	14
Application-Aware Processing	14
Host-based Backup	15
Continuous Data Protection (CDP)	15
Agent-based Backup	16
Backup Appliances	19
Recovery from Image-Level Backups	19
Unstructured Data Backup	20
Enterprise Applications	21



Backup Infrastructure	24
Primary Storage Integrations	25
Secondary Storage Integrations	26
Backup Console	27
API Enhancements	28

---

## Introduction

Veeam Backup & Replication, the foundational component of Veeam Data Platform, delivers enterprise grade data resilience for your entire hybrid estate, providing confidence in your protection, response, and recovery in the face of disaster and cyberthreats. The following is a list of the major new features and enhancements added in Veeam Backup & Replication v13. All capabilities here are transacted as the Veeam Data Platform, with certain features available only at the Advanced or Premium editions. Follow [this link](#) for a detailed edition comparison.

Note: All features and capabilities not available in the early release of Veeam Software Appliance (version 13.0.0) are labeled accordingly. This early release is intended for net new deployments, with upgrade and migration options coming in a future release. [Sign up](#) to be the first to know when this migration option is available!

---

## Veeam Software Appliance

In addition to the existing Windows-based installable software that will remain available, V13 now also offers a new software appliance deployment option for significantly reduced operational costs and security risks. By lowering deployment complexity and applying optimizations and DISA STIG hardening to the OS automatically, time to value is reduced while risky misconfigurations are avoided and the ongoing maintenance burden is decreased.

The unique benefits of Veeam Software Appliance (VSA) include:

### Pre-built

With VSA, the base OS aka “just enough” operating system (JeOS) and the backup software are packaged together to create a software appliance configured to run optimally on industry-standard server hardware or in a virtual machine. The entire VSA stack is fully maintained by Veeam, including JeOS, backup software and 3rd party components updates, so you could focus more on backup and recovery and less on backup infrastructure management.

VSA is delivered as bootable installation media (ISO) for quick deployment on any physical, virtual or cloud machine that supports boot from USB or from an ISO file. Installation of VSA requires just a few clicks with automated disk partitioning intelligently configuring volumes for optimal performance and OS security. In addition, for even faster deployment as a VMware vSphere VM, VSA is also offered as a Virtual Appliance (OVA).

Despite JeOS is based on Linux, no Linux OS expertise is required because we abstract all essential OS management tasks with a simple text-based user interface (TUI) and a web-based console covering common host OS management tasks such as network settings configuration. And should something ever go sideways, purpose-built bootable rescue media will help to get your VSA back up in no time by reinstalling the base OS while preserving its configuration and backups.

## Pre-hardened

VSA is secure-by-default thanks to its base OS preconfigured according to DISA STIG guidelines to minimize attack surface and block known attack vectors. No manual hardening, no guessing games — industry best practices are baked in! Plus, we go even further to make VSA harder for hackers to break into by restricting remote OS access over SSH out of the box and mandating multi-factor authentication (MFA) for all sensitive operations.

In addition, core services have been reworked to run under a low-privileged OS account, thereby vastly reducing privilege escalations opportunities in case of software vulnerabilities, ensuring an attacker is unable to easily take over the OS and extract sensitive information such as saved infrastructure credentials. We also sandboxed the execution of custom scripts used by various functions to ensure they cannot do any damage to VSA.

Secure-by-default from the start, the software appliance also stays hardened over time thanks to fully automated security patches and hardening updates. By making these updates mandatory, Veeam is taking ownership of security outcomes for our customers as we committed with our [CISA Secure by Design pledge](#).

## Predictable

VSA is designed for True Zero Trust operations with no base OS privileges available to backup server administrators or any other roles. This prevents backup admins from applying configuration changes with unpredictable results that may impact backup performance, reliability and security, helping businesses to maintain the baseline posture through the entire software appliance lifecycle.

Further, any legitimate but sensitive host management operations such as the creation of new admin accounts or the deployment of remote management agents must be approved by the dedicated user with the Security Officer role, who at the same time is not allowed to initiate any operations on their own. Designating a Security Officer from your company's Security Team is highly recommended to enable true Zero Trust operations but can be skipped for small IT environments with a single administrator responsible for everything.

Want even more predictability? New to V13 VSA is the new backup infrastructure lockdown mode which once enabled prevents unauthorized addition of backup infrastructure components, which can not only impact backup performance by changing the network traffic flow but also be used for data exfiltration if newly added infrastructure components are controlled by a malicious actor.

## All-in-one Appliance

Thanks to low deployment complexity and predictability, VSA particularly shines as an all-in-one backup appliance — especially thanks to offering the immutability option for backups stored in its built-in repository — making it an appealing solution for SMB and ROBO environments. Installed on an industry-standard storage-optimized server chassis, VSA offers unmatched performance and up to 1PB of immutable backup storage in a single box! And while the built-in repository does not provide the same level of protection against cyber-attacks as Veeam Hardened Repository due to the added attack surface of the backup management software, immutability still offers major benefits by protecting backups against accidental or malicious deletion. But for ultimate protection, we recommend separating management server and backup storage by deploying standalone hardened repositories — which V13 makes it a breeze to do with Veeam Infrastructure Appliance described next.

## Licensing

Veeam Software Appliance (VSA) requires a Veeam Data Platform Veeam Universal License (VUL) and is not available for legacy Socket-based licenses or as a Community Edition, however we continue offering Veeam Backup & Replication in the form of Windows-installable software for such licenses. Converting from Sockets to VUL is easier than ever at the price of your maintenance renewal, and it enables the flexibility of a growing list of hypervisor and cloud alternatives, unstructured data and more.

For Veeam Data Platform Essentials license holders, initially VSA will be supported only for deployment in a virtual machine on any hypervisor supported by Veeam. This limitation is introduced due to the high cost and complexity of support cases that involve troubleshooting hardware-specific issues on various non-Enterprise grade hardware our SMB customers tend to use, as well as the sheer size of our Veeam Data Platform Essentials customer base.

---

## Veeam Infrastructure Appliance

While many of our customers use all-in-one Veeam installations running on a single server, there may come a time when the number of protected workloads becomes large enough to require scaling the backup infrastructure out, for example by adding backup proxies to reduce your backup window. Veeam has always shined with its flexibility, allowing to expand only the needed resource (storage or compute) without forcing customers to purchase pre-built cluster nodes ending up with unneeded resources. However, until now, maintaining a fleet of remote backup infrastructure components has been time consuming.

This is where V13 changes the game with its flexible backup infrastructure appliance, which can be deployed on any physical, virtual or cloud machine to assume any backup infrastructure role. Based on the same JeOS as VSA, the Veeam Infrastructure Appliance (VIA) shares the benefits of being pre-built, pre-hardened and predictable to significantly reduce deployment complexity and costs. And with centralized patching and updates for the entire fleet of such appliances, Day 2 management becomes a breeze — making VIA an ideal foundation for secure, efficient, and scalable backup environments.

To reduce the attack surface to the absolute possible minimum, VIA offers various installation profiles tailored to the intended use case: a general-purpose appliance that can take any backup infrastructure role, an appliance with added iSCSI and NVMe over TCP/IP modules for when direct SAN connectivity is required, and a Veeam Hardened Repository appliance. And if you want even smaller footprint, you can remove components not required in your environment (e.g. SMB Client) using the backup console.

To dramatically simplify onboarding and improve security, we also made the management connection to VIA passwordless. Once deployed, it will simply wait for a connection from a backup server, with certificate thumbprints available to validate authenticity of both parties. Following the initial connection, certificate-based authentication will be used for subsequent connections.

---

## Web UI

V13 brings a preview of the next-gen, web-based user interface that enables easy, scalable access to backup management operations. By eliminating software and platform dependencies, the web UI removes barriers and unlocks accessibility for all IT teams. The new web UI also features WCAG (Web Content Accessibility Guidelines) ready design aligned with global accessibility standards, supporting compliance with EAA (European Accessibility Act) and ADA (Americans with Disabilities Act) regulations.

When designing the new web UI, our major focus has been in ensuring that it remains familiar for existing Veeam Backup & Replication (VBR) users and requires zero learning curve to adopt. At the same time, we used this opportunity to make the overall experience more convenient, for example by making our wizards use the entire browser screen estate to make it more convenient to browse large lists of objects, as well by adding advanced filtering and search tools to streamline managing data protection in larger environments.

In response to popular feedback from VDP Foundation customers, the new web UI also features a built-in monitoring dashboard that provides a convenient overview of most critical metrics of your backup infrastructure. While VDP Advanced and Premium users will find advanced monitoring, reporting and analytics right at their fingertips, conveniently integrated directly into the VBR web UI.

The web UI preview covers the two most used workloads (VMware vSphere and Microsoft Hyper-V) and allows for managing primary backup and Backup Copy jobs; scale-out backup repositories and Veeam Cloud Vault; as well as perform most popular recovery types including Instant VM Recovery®, entire VM restores, and guest file restores.

We will iterate quickly and bring more and more features over to the new web UI with every minor release, prioritizing workloads and features based on their actual usage. Meanwhile, a familiar Windows-based backup console — now with refreshed UI and the dark mode — will remain available with full functionality for when you need to manage settings not yet available in the new web UI.

While most enterprise platforms either bury critical actions in clunky, outdated UIs or do not include them in UI at all, forcing customers to trigger REST API directly to control many settings, Veeam Data Platform provides a choice of intuitive interfaces: the new, quickly expanding browser-based web UI and a refreshed, fully-featured and familiar Windows-based client.

---

## High Availability

Despite Veeam backups being portable and therefore easily importable into a new VBR installation in case of a backup server or entire data center loss, in large environments the import process can take significant time impacting your RTOs. The same stands for product configuration — while we make it super easy to restore your configuration backup into a fresh VBR install, your backup jobs cannot continue until the process completes, therefore impacting your RPOs.

To address this, V13 Veeam Software Appliance can be deployed in an active/passive backup server cluster configuration with the second node standing by and ready to take over. This ensures backup and recovery remain available — even during outages or entire data center disasters.

Under the hood, the HA cluster is implemented by continuously replicating the configuration database between cluster nodes, including necessary transformations to ensure all secrets encrypted with machine-specific key on the primary node are also usable by the secondary node. The HA cluster automatically reprotects itself upon failover, reversing replication direction to go from the currently active node. The replication engine fully supports high-latency network links (such as U.S. East to West coast) providing the complete freedom of cluster nodes placement, and is designed to reliably withstand long uplink or secondary node downtimes by caching the changes locally to be sent when the secondary node comes back online.

Users can perform one-click failover and failback for both unplanned (primary backup server is already down) and planned outages (when primary backup server is still active but a natural disaster is coming). While users of VDP Premium with Veeam ONE deployed can in addition enjoy automated failovers with Veeam ONE serving as the cluster witness.

High Availability functionality requires VSA 13.0.1 and VDP Premium edition.

---

## Enhanced Role-Based Access Control (RBAC)

While Veeam Backup & Replication has been offering RBAC since close to inception over 15 years ago, our larger customers felt constrained by pre-defined roles and asked us for the ability to define custom roles to reduce risk and simplify compliance. The new, enhanced RBAC framework in v13 introduces fine-grained access control for backup and restore operations across your environment.

Designed to enforce least-privilege principle and improve operational security, enhanced RBAC empowers administrators to delegate responsibilities with precision ensuring users are able to only access workloads and perform operations relevant to their roles. With support for custom roles with scoped permissions, organizations can align backup and restore management with internal policies, compliance mandates, and operational workflows.

The new Custom Role wizard provides an interface for defining roles and their associated permissions. For backup operators, you can define the inventory scope to limit production workloads that can be protected, and which backup repositories the role can use. And for restore operators, you can specify accessible backups, allowed production infrastructure scope restores can be performed to, and available restore types. The newly created roles can then be assigned to individual users or groups, ensuring their access level aligns with organizational policies and operational needs.

In this initial release, the enhanced RBAC framework can already be applied to over 90% of protected workloads including host-based backup for VMware vSphere and Microsoft Hyper-V, agent-based backups, and unstructured data backups (file share and object storage). The future releases will expand advanced RBAC coverage to all protected workloads.

---

## Single Sign On (SSO)

SSO is an industry-standard mechanism for outsourcing the authentication to a third-party identity provider, such as Microsoft Entra ID or Okta, enabling IT and Security teams to be able to effectively manage user accounts across different software. By centralizing user access and authentication, password loss risk and IT overhead are reduced. For example, when an employee leaves the company, the IT team can immediately disable their access to all applications, rather than logging into many different user management portals. SSO also improves experience of end users by preventing application-specific Time-based One-Time Password (TOTP) sprawl.

V13 delivers federated authentication enabling you to leverage external Identity Providers that support Security Assertion Markup Language (SAML) 2.0 with shared OAuth authorization service for accessing both the new web UI and the Windows-based backup console.

SSO integration is particularly important for frictionless backup console access when using Veeam Software Appliance, by allowing to specify external accounts in Users & Roles settings. This lets you avoid the hassle of local account authentication with its long and complex passwords in compliance with DISA STIG requirements. Further, as a security best practice, using VSA administrator account for day-to-day backup management activities is not recommended in favor of assigning a lower privileged role to you own external user account, which is now possible thanks to the SSO integration.

Even though SSO is considered an advanced security feature, we're making it available as a part of the core offering across all Veeam Data Platform edition tiers, to incentivize its use and encourage best security practices. SSO is a rare case where usability meets enhanced security, and we highly recommend all customers to implement this integration as soon as possible.

---

## Veeam Intelligence

V13 expands Veeam Intelligence capabilities while giving you more control over AI usage with three simple options: disable it completely, the enhanced Basic mode that provide answers using the official documentation and other publicly available sources, and the new Advanced mode that further enhances the quality and depth of answers using the backup infrastructure monitoring data.

With the enhanced Basic mode, Veeam Intelligence now uses an agent-based retrieval to query the Veeam Help Center, Support KB articles, R&D Forums, and other \*.veeam.com content in real time. Comparing to leveraging a static vector database populated by the User Guide only, this delivers results that are more accurate and better sourced. As a result operators spend less time searching for answers and more time focused on what matters; protecting your critical assets.

The Veeam Intelligence Advanced mode is powered by Veeam ONE, which provides Veeam Intelligence access to your backup infrastructure details and real-time monitoring data. This additional input unlocks the usage of new AI agents that help IT teams work more efficiently, use existing datasets more effectively, and get timely, custom expert data without having to create custom reports manually. This new Veeam Intelligence mode requires version 13.0.1 or later.

The new Malware Threat Analysis Agent helps you detect and respond to threats with the help of Veeam Intelligence. Get real-time insights, anomaly details, blast radius, and next best actions to guide you through the recovery steps, including triggering signature-based malware scans directly from Veeam Intelligence. This agent requires Veeam ONE 13.0.1 or later.

The new Deep Data Analysis Agent enables you to request and generate backup infrastructure reports by simply asking Veeam Intelligence what you need to see in the report, reducing time spent searching for and manually running or creating reports. This agent requires Veeam ONE 13.0.1 or later.

## Other Features

In addition to the major new features, Veeam Backup & Replication V13 includes hundreds other enhancements that are a response to customer feedback and ongoing R&D findings, the most significant of which are listed below:

### Platform

**64-bit architecture** — V13 concludes Veeam Backup & Replication transition to fully 64-bit CPU architecture, with all backup infrastructure components now running natively on 64-bit OS. This upgrade dramatically improves performance, scalability, and memory utilization for the remaining components that still leveraged 32-bit binaries, delivering a more robust and truly enterprise-scale backup infrastructure.

**Linux-only backup infrastructure** — ALL remaining backup infrastructure components that still did not support being deployed on Linux (like Mount Server, Gateway Server or Guest Interaction Proxy to name only a few) now support this, enabling you to deploy V13 in a fully Linux-based backup infrastructure with zero Windows machines required. Note: Windows remains supported for all backup infrastructure components.

**Reduced network port requirements** — Communication between backup infrastructure components now leverages reverse proxy technology to reduce port requirements and simplifying firewall configuration. Check the much shorter Required Ports section of the System Requirements to see the results of implementing this technology and expect it to be reduced even further in the close future.

**Secure communication with OpenSSL 3.0** — V13 leverages OpenSSL 3.0 for general- cryptography operations and secure communication. This upgrade strengthens platform security and ensures compliance with modern encryption standards delivering improved resilience against vulnerabilities and future-proofing your data protection infrastructure.

**Microsoft RPC and Microsoft WMI discontinuation** — V13 eliminates the usage of these protocols for communication between backup infrastructure components and to protected workloads in favor of cross-platform gRPC protocol. In addition to improving performance and reliability, this change reduces the number and the range of ports required for Veeam Backup & Replication to function, thereby reducing your network exposure to cyberattacks and simplifying deployment with reduced number of ports to open in firewalls. This does mean that you will need to work with your networking team to adjust firewall settings when moving to V13!

**NTLM Authentication deprecation** — NTLM usage for connection between backup infrastructure components and protected workloads is discontinued in V13 Software Appliance and is deprecated in V13 Windows installable software in favor of Kerberos authentication. This change improves security by eliminating legacy authentication protocols, reducing your exposure to known NTLM vulnerabilities and aligning with modern standards.

**Veeam Deployment Kit** — V13 introduces a lightweight package to enable secure, certificate-based authentication with the backup server for environments where Kerberos is not available. The kit can be installed manually or deployed via automation tools to provide means of secure authentication with managed Microsoft Windows servers, Hyper-V hosts, and backup agents protecting machines running Windows OS.

## Configuration Database

**PostgreSQL 17.6 configuration database** — V13 ships with the latest and greatest PostgreSQL version which provides noticeable database performance boost compared to the previously used version.

**Automatic upgrades** — As a part of the fully managed Veeam Software Appliance experience, PostgreSQL updates will be delivered along with JeOS and Veeam Backup & Replication updates, so you no longer have to worry about patching it manually.

**Continuous configuration autotuning** — PostgreSQL instance configuration parameters are now automatically adjusted at every Veeam Backup service restart according to the detected CPU core count and the amount of RAM, ensuring that PostgreSQL remains optimally configured as compute resources are added or removed on the backup server.

**Scalability enhancements** — V13 further optimizes interaction with local PostgreSQL database, enabling all-in-one backup appliance configuration to protect up to 10000 machines, simplifying deployment and reducing Veeam footprint in many environments down to a single server.

**PostgreSQL logs collection** — Local database instance logs can now be easily collected with the Export Log wizard, when required, for troubleshooting by our Customer Support team.

**Configuration restore improvements** — Unattended configuration restore process now allow overriding answer file parameters via the command line, enabling you to avoid storing credentials in the answer file and instead only pass them in the command line.

## Backup Engine

**Improved job manager scalability** — V13 delivers dramatic scalability gains, catering the need well beyond those of the largest existing Veeam customers. Depending on the protected platform and job type, the number of parallel jobs has been increased into hundreds, or up to a few times compared to V12. At the same time, the improved engine significantly reduces CPU and RAM consumption on the backup server — which means you won't have to upgrade your backup server hardware for much longer now.

**Fast and secure data hashing algorithm** — The usage of modern BLAKE3 hashing algorithm reduces backup proxy and backup agent CPU usage up to 30%, which directly translates into increased backup performance in scenarios where CPU is the bottleneck. As a bonus, the new algorithm offers a high level of resistance against collision and pre-image attacks which the previously used MD5 algorithm was susceptible to.

## Malware Detection

**Proactive backup scans** — Accelerate incident response by having signature-based scans performed automatically whenever suspicious activity is detected during backup, so you have more information at hand when you get to investigating the malware event. In addition, with an option to automatically resolve the malware events based on scan results, you can significantly reduce false positives at the cost of potentially missing malware that deletes itself after doing the damage.

**Malware detection for Linux machines** — By popular demand, all malware detection capabilities are now available for both host-based and agent-based Linux machine backups as well. This included suspicious file system activity analysis during backup, Veeam Threat Hunter and YARA scans after backups, and Veeam Incident API integration.

**Malware detection for cloud machines** — We added support for scanning Veeam Backup for Microsoft Azure, AWS or Google backups with Veeam Threat Hunter to enhance your cloud data safety.

**Linux mount server support** — Signature-based scans with Veeam Threat Hunter or Bring Your Own (BYO) antivirus are now supported when using Linux-based mount servers, with out-of-the-box support for the following Linux-based antiviruses: ClamAV, ESET and Sophos.

**Suspicious file system activity detection improvements** — To reduce false positives, the number of deleted files is now tracked per volume instead of for the entire backup.

**Encryption detection improvements** — Disk blocks belonging to RPM packages are now automatically excluded from inline entropy analysis as they were reported to be a common source of false positives.

## Application-Aware Processing

**Linux distribution support** — V13 added support for application-aware processing during host-based backup for the following Linux distributions: Alma Linux 8.10 and 9.4; Debian 12.9, 12.10 and 12.11; RHEL 9.5, 9.6 and 10.0; Rocky Linux 8.10, 9.4 and 9.6; SLES 15 SP7; Oracle Linux 9.6.

**Microsoft Exchange Server SE support** — V13 added support for application-aware processing with both host-based and agent-based backup, including application item-level recovery with Veeam Explorer for Microsoft Exchange. In addition, you can now enable the usage of secure LDAPS protocol instead of LDAP in the Veeam Explorer for Microsoft Exchange settings.

**Microsoft SharePoint SE 25 H1 support** — V13 added support for application-aware processing with both host-based and agent-based backup, including application item-level recovery with Veeam Explorer for Microsoft SharePoint.

## Host-based Backup

### Nutanix AHV

**Fully integrated** — The AHV backup appliance has been integrated directly into Veeam Backup & Replication, eliminating the need for managing a separate appliance with its own local web UI.

**Worker distribution improvements** — Worker images will now be uploaded to Prism Central-managed clusters on as-needed basis, as opposed to being deployed to all clusters right away, which led to timeouts in large Nutanix AHV deployments. In addition, workers now support cross-Prism Central operation.

**Persistent guest agent support** — Remove the need for using a high-privileged account and network port requirements for application-aware guest processing by deploying a persistent guest agent to each protected VM.

**vTPM support** — VMs with virtual Trusted Platform Modules are now fully supported for backup and restore.

**Job-level email notifications** — In addition to global email notifications settings, you can now customize email notifications on a per-job level.

### Proxmox VE

**Application-aware processing support** — Added Microsoft VSS integration for application-consistent backups; application-item recovery by Veeam Explorers for Microsoft Active Directory, Microsoft Exchange and Microsoft SharePoint; transaction log shipping and point-in-time database recoveries with Veeam Explorers for Microsoft SQL Server, Oracle, and PostgreSQL; support for custom pre-freeze/postthaw in-guest scripts. This capability requires version 13.0.1 or later.

### VMware vSphere

**Full vSphere 9.0 support** — As opposed to the compatibility-level support for vSphere 9.0 provided by version 12.3.1, V13 delivers full support for this vSphere version.

## Continuous Data Protection (CDP)

**Universal CDP** — V13 added support for agent-based CDP that can be used for continuous replication of any Windows machine (physical, virtual or cloud) to VMware Cloud Director, with support for additional targets coming later. This feature requires version 13.0.1 or later.

**Improved resilience to network issues** — Enhanced data stream validation detects and automatically recovers from network transmission issues causing severe packet loss or reordering, ensuring robust and reliable replication even over unreliable network links.

**Improved I/O journal scalability** — The new I/O journal format eliminates its previous size constraints which led to automatically shortening the configured short-term retention policy by truncating the journal after its maximum size was exceeded. With V13 however, you can retain as many short-term replica restore points as required.

## Agent-based Backup

### Agent Management

**Distribution server on Linux** — You can now deploy agent distribution servers on Linux OS to enable managed agent-based backup in fully Linux-based backup infrastructures.

**Job manager scalability** — The upgraded job engine now supports up to 5000 managed backup agents per backup server, which is double the amount compared to v12.

**Active Directory based protection groups** — Kerberos authentication with protected machines is now supported without requiring the backup server to be domain joined. In addition, LDAP/S support and TLS certificate validation improvements further enhance security and deployment flexibility.

**Computers with pre-installed agents protection groups** — The agent package export can now be performed to a desired remote location, as opposed to only locally; with the package featuring improved directory layout to reduce the size of the Veeam Agent for Linux distribution. In addition, backup jobs created for such protection groups are now always set to Managed by Agent type by default for better user's guidance.

**Individual computers protection groups** — Support certificate-based authentication is added, enabling secure communication with the backup agent with pre-installed Veeam Deployment Kit, streamlining deployment process and enhancing security posture.

**Block-level filtering for file-level backup** — Backup agents in file-level backup mode will now filter out disk blocks already present in the backup chain resulting in much smaller incremental backups.

**Removal of admin\$ share requirement** — After the Veeam Agent for Microsoft Windows has been deployed, backup server will no longer require the protected machine's admin\$ share to be available, allowing administrators to disable it, if desired, to reduce security risks.

**Consistent backup agents versioning** — Starting from V13, version numbers of all Veeam Agents will be aligned with their respective Veeam Backup & Replication release. This will make it much easier to determine if the agent version matches your backup server version, or if the agent is behind and needs to be updated.

### Veeam Agent for Microsoft Windows

**Significantly improved backup performance** — Thanks to multiple under the hood enhancements, V13 more than doubles the backup throughput on the same hardware vs. previous versions, so long as it does not hit any environmental performance bottlenecks.

**Better missed backup handling** — The agent will now create an appropriate type of missed backup (incremental, synthetic full or active full) upon system power on if a scheduled backup was missed while the system was powered off. This improvement is particularly important to ensure periodic full backups are not missed, as retention policy processing is dependent on them being present.

**Monthly synthetic fulls** — The ability to schedule synthetic full backup to happen monthly was added to enable users to reduce full backup frequency if desired. Keep in mind that monthly fulls delay the removal of oldest incremental chain by the retention policy.

**Reparse points restore** — Symlinks and junction points are now restores as original reparse points, thereby preventing recursive restores and excessive disk space consumption by duplicate data.

**Notifications about extra-long backup chains** — Users will now be alerted if their preferred GFS full backups schedule will result in extremely long incremental backup chains, helping to prevent inefficient backup storage usage and to ensure optimal backup performance.

**Upgrade compatibility validation** — The Agent upgrade process will now notify users if the backup job configuration references backup or cloud repositories incompatible with the new agent version. This eliminates the risk of upgraded agents being left unable to continue functioning.

### **Veeam Agent for Linux**

Added support for agent-based backup for the following latest versions of Linux distributions:

**Significantly improved backup performance** — Thanks to multiple under the hood enhancements, V13 more than doubles backup throughput on the same hardware vs. previous versions, so long as it does not hit any environmental performance bottlenecks.

**Linux distribution support** — Support was added for x86\_64 distributions of AlmaLinux 10.0, Debian 12.11 and 13, RHEL 10.0, Rocky Linux 10.0, Oracle Linux 10.0, SLES 15 SP7; and IBM Power distributions of RHEL 10.0 and SLES 15 SP7.

**Time-based retention policy** — As a part of platform-wide changes, restore point-based retention policy has been replaced with a time-based retention. Not only does this approach match how business requirements are typically set, but it's also the essential enabler for a number of key features such as immutability.

**Synthetic full backups** — Avoid resource-intensive active full backups by creating synthetic full backups using data from previous restore points, reducing load on production storage and improving backup performance. And if your backup repository supports block cloning, it will make your synthetic full backups lightning fast and spaceless, enabling you to use GFS retention policy with no additional storage requirements.

**Dedicated snapshot storage** — The agent can now be configured to use the dedicated block device as the snapshot storage. This enables storing snapshot data directly on a raw (unformatted) disk, removing performance overhead of a file system. And by using the dedicated disk, you can be sure that snapshot storage will never impact available production storage capacity.

**Faster file-level recovery** — Managed agents should enjoy significantly improved file-level restore performance thanks to the usage of fuse3 for file system mount.

**Enhanced error logging** — The agent will now log more informative errors when the kernel module fails to load due to missing kernel headers or Secure Boot issues. This helps users troubleshoot these common issues faster and without help from our Customer Support.

**Simplified setup on UEFI Secure Boot systems** — To streamline the installation process, blksnap-ueficert and veeamsnap-ueficert packages are now merged into the single veeam-ueficert package, allowing for more seamless agent installation to UEFI Secure Boot systems.

**OCSP certificate validation** — Online Certificate Status Protocol support is now enabled by default and works alongside CRL (Certificate Revocation List) to provide secure and reliable certificate validation. Additionally, users are now able to configure an Internet proxy to use for establishing a connection to the OCSP responder in environments with restricted Internet access.

## **Veeam Agent for Mac**

**Time-based retention policy** — As a part of platform-wide changes, the restore point-based retention policy has been replaced with a time-based retention. Not only does this approach match how business requirements are typically set, but it's also the essential enabler for a number of key features such as immutability.

**Synthetic full backups** — Avoid resource-intensive active full backups by creating synthetic full backups using data from previous restore points, reducing load on production storage and improving backup performance. And if your backup repository supports block cloning, it will make your synthetic full backups lightning fast and spaceless, enabling you to use GFS retention policy with no additional storage requirements.

**OCSP certificate validation** — Online Certificate Status Protocol support is now enabled by default and works alongside CRL (Certificate Revocation List) to provide secure and reliable certificate validation. Additionally, users are now able to configure an Internet proxy to use for establishing a connection to the OCSP responder in environments with restricted Internet access.

## **Veeam Agents for IBM AIX and Oracle Solaris**

**Backup to S3-compatible object storage** — V13 brings the long-awaited support for agents to back up directly to S3-compatible object storage in both Standalone (CLI only) and Managed by Agent modes. Specifically validated object storage providers include Amazon S3, IBM Cloud Object Storage, Wasabi, MinIO, and 11:11.

**Time-based retention policy** — As a part of platform-wide changes, restore point-based retention policy has been replaced with a time-based retention. Not only this approach matches how business requirements are typically set, but it's also the essential enabler for a number of key features such as immutability.

**Synthetic full backups** — Avoid resource-intensive active full backups by creating synthetic full backups using data from previous restore points, reducing load on production storage and improving backup performance. And if your backup repository supports block cloning, it will make your synthetic full backups lightning fast and spaceless, enabling you to use GFS retention policy with no additional storage requirements.

**Disk exclusion** — We added the ability to exclude specific disks from backup. This is especially useful for skipping SAN LUNs and multipath devices, helping to avoid conflicts and redundant backups.

**ACL backup and restore** — Backups will now include file and folder access control lists (ACL) that will be restored along during bare metal recovery. This requires the usage of NFSv4 protocol version, when unavailable due to the system configuration the agent will failover to using NFSv3 which does not support ACL backup and restore.

**Additional backup server connections for DR** — You can now add a backup server in "read-only mode" in both the Standalone and the Managed by Agent modes to enable the agents to access its backups. This provides the ability to perform restore operations without consuming a license, offering more flexibility for recovery scenarios in environments with multiple backup servers.

**IBM AIX 7.3 TL3 support** — Veeam Agent for IBM AIX now supports AIX 7.3 TL3, expanding compatibility with the latest enterprise-grade UNIX environments.

## Backup Appliances

### General

**Mount server selection for external repositories** — You can now choose the preferred mount server to access backups stored on external backup repositories. This option gives administrators greater control over the data flow path, improving performance during restore and backup operations while reducing cloud egress costs.

### Veeam Backup for Microsoft Azure

**Veeam Data Cloud Vault support** — Veeam Backup for Microsoft Azure can now use Veeam Vault as the backup target, providing seamless user experience while enhancing security and resilience with immutable cloud storage that is managed by Veeam.

**Resource group selection for VM snapshot policies** — You can now override the default resource group selection when creating VM snapshots. This allows backup admins to apply locks on production VM resource groups, which enhances the security posture of production infrastructure.

**Generation 2 VMs support for backup appliance** — Veeam Backup & Replication will now use Gen2 VM type when deploying new backup appliances in Microsoft Azure. The usage of Gen2 VMs improves backup performance and strengthens infrastructure security.

### Veeam Backup for AWS

**Private deployment of external repository** — External repositories can now be added from AWS backup appliances deployed in private networks. This enhancement eliminates network connectivity errors issue enables secure integration in isolated cloud environments.

**Custom retention for manual EC2 snapshots** — You can now define a retention period for how long manual EC2 snapshots created with Veeam Backup for AWS must be retained. This enhancement prevents lingering snapshots after performing tasks like testing an application or OS upgrade.

**Email reporting for retention processing** — This new report delivers insight into retention tasks performed that have been performed by Veeam Backup for AWS, providing backup admins with better visibility on retention policy activities.

## Recovery from Image-Level Backups

### General

#### Instant VM Recovery

**Instant recovery engine improvements** — Improved I/O performance for VMs running directly from backups up to 50% compared to v12 significantly accelerating instant recovery and all capabilities that rely on this engine.

#### Entire VM Restore

**Mass VM restore improvements** — Improvements to the entire VM restore engine significantly reduces CPU and RAM usage on the backup server. For example, V13 now supports up to 1000 concurrent entire VM restore sessions on VMware vSphere.

## Virtual Disk Restore

**Support for additional controller types** — Support for IDE, SATA, and NVMe virtual disks have been added to the Virtual Disk Restore wizard.

## File-level Restore (FLR)

**Mount server on Linux** — You can now deploy mount servers on Linux OS to enable file-level recovery in fully Linux-based backup infrastructures.

**Unified FLR** — The ability to configure both Windows and Linux-based mount servers for each backup repository has been added, with file-level restore picking the suitable mount served depending on the type of backup. While Linux-based support can be sufficient to basic file-level restores due to its wide file system support, designating a Windows-based server unlocks support for advanced Windows file system configurations such as dynamic disks, ReFS volumes, volumes with deduplication enabled etc.

**Improved junction points and symlinks handling** — Junction points and symlinks are now restored as the corresponding objects rather than by dereferencing their contents, thereby preventing recursive restores and excessive disk space consumption by duplicate data.

## Restore to Public Cloud

**Instant Recovery to Microsoft Azure** — Instantly recovery ANY image-level backup residing as an Azure VM in under 5 min (time to OS logon screen) by running it directly from backups residing in Veeam Vault or Microsoft Azure Blob Storage. Available for both Windows and Linux machines, this functionality enables seamlessly leveraging public cloud for DR, backup testing, dev/test labs and so on.

**Direct Restore to AWS EC2 enhancements** — You can now assign a static private IP address to restored EC2 instance to simplify migration of servers with static IPs from on-premises environments, preserving network topology and minimizing connectivity or application issues. We have also improved the logic image generation selection when restoring backups of native EC2 instances and the helper appliance default instance type selection.

## Unstructured Data Backup

**Backup proxy on Linux** — You can now deploy general-purpose backup proxies on Linux OS (including Veeam JeOS) to enable unstructured data backup in fully Linux-based backup infrastructures. This includes support for backing up SMB and NFS file shares, as well as NetApp filers including SnapDiff v3 integration and FlexGroup support.

**Backup engine improvements** — Thanks to multiple under-the-hood enhancements, V13 achieves an average of 25% performance improvement over already industry-leading unstructured data backup performance of V12.

**NFS backup performance** — NFS3 shares backup performance has been improved through leveraging multiple TCP connections. While the specific improvements depends heavily on production storage and backup infrastructure, we have observed up to 3x improvements in some QA labs.

**Folder-level restore from archive** — By popular request, we added the ability to restore entire folders from archive repositories instead, eliminate the need to manually select large numbers of individual files during restore operations.

**Orphaned backups retention** — Unstructured data backups and archives are now also subject to orphaned retention identical to one offered for image-level backups. This ensures that the last known retention policy is still processed for backups that no longer have a job associated with them, improving backup storage consumption hygiene and removing the need for manual processes.

**Improved symbolic links handling** — The content of symbolic links in SMB shares and Windows-based file servers now can be processed and included in backups to ensure no unintended data loss when backing up linked data structures.

## Enterprise Applications

### General

**Backup to object storage** — All application plug-ins now support backup to object storage, enabling scalable and cost-effective protection for critical database workloads. Available in both standalone and managed modes, plug-ins can now leverage any object storage supported by Veeam as a primary backup target, or as a secondary backup target for Backup Copy jobs. Plug-ins also honor the backup immutability settings set on the chosen repository to guarantee that backups cannot be deleted or tampered with, ensuring you can always recover from a cyberattack. And as plug-in backups are not image-level, we're able to use much larger default object size (8MB) as to keep API calls costs low without impacting backup storage consumption.

**Backup encryption support** — All application plug-ins now offer backup encryption at source to protect backups from unauthorized access and data exfiltration. For plug-ins running in managed mode, encryption settings are available directly in the backup policy settings, while those in standalone mode can encrypt backups by enabling encryption in the target repository settings.

**Object storage offload engine** — We improved efficiency of scale-out backup repository offload engine, significantly reducing backup server CPU and RAM usage when offloading or copying plug-in backups to object storage.

**Single port communication** — All application plug-ins now use a single port for local communication between internal components running on the database server. This helps prevent port exhaustion issues when many backup streams are running in parallel.

**Seamless backup import** — You can now import plug-in backups by simply rescanning the backup repository and without the need of backup metadata, eliminating the existing manual process of recreating it. This feature is designed to enable effortless recovery from immutable repositories following a cyberattack or a disaster that led to a complete loss of backup infrastructure.

**Restore action log** — You can now easily monitor the progress and statistics of restore sessions initiated on the protected application side directly in the backup console. The restore session action log uses the same familiar format of backup session logs and provides detailed database-level information.

### Veeam Plug-in for IBM Db2

**SLES 15 SP6 support** — The application plug-in for IBM Db2 is now fully compatible with the latest version of SUSE Linux Enterprise Server 15.

## Veeam Plug-in for Microsoft SQL Server

**Centralized management** — V13 brings the ability to centrally manage Veeam plug-ins for Microsoft SQL Server, similar to the centralized management of Oracle RMAN and SAP plug-ins, enabling the following familiar capabilities:

Fast and easy roll-out through Protection Groups, which have been expanded with an option to control the installation and upgrade of SQL Server plug-in, are now enabled. This capability includes continuous automated analysis of SQL Server instance topology during protection group rescans in order to detect the appearance or changes to failover clusters and Availability Groups configuration. Policy-driven database protection directly from the backup console provides centralized real-time monitoring and reporting for database and transaction log backups. This eliminates the overhead of configuring plug-ins and maintaining SQL Server Agent jobs and backup scripts separately on each protected database server. And with each policy targeting its own scope of databases, you maintain full flexibility of granular, database-level protection settings including backup schedule and repository used.

Now, customers have a choice of letting the database administrator (DBA) maintain full control of backups or letting the backup administrator implement a centralized, policy-based protection model with little to no DBA involvement. Or use a combination of this approach thanks to backup administrator's ability to create recovery tokens — time-limited access keys to access Microsoft SQL Server backups, which can be shared with DBAs to enable them to perform database restores without having to assign them any role.

Centralized managed for Microsoft SQL Server plug-in requires version 13.0.1 or later.

**Veeam Explorer for Microsoft SQL Server support** — You can now perform recovery from backups created by Veeam Plug-in for Microsoft SQL Server using the Veeam Explorer for Microsoft SQL, which provides the dedicated experience when detecting plug-in backup which maintains a close connection to the restore wizard of the plug-in while keeping the experience familiar to Veeam Explorer users.

**Incremental database recovery** — Instead of restoring the entire database, you can now perform recovery much faster by applying only transaction log backups or differential backups to an existing database. This significantly reduces RTO and streamlines migration and database replication scenarios by allowing you to restore changes.

**Retention policy improvements** — The new set-force-delete command in the configuration tool allows for deleting backup files older than the specific number of days. This functionality removes backups of databases that are no longer processed and therefore are skipped by the retention engine; as well as files that exist on the repository file system without an associated record in the configuration database.

**Microsoft SQL Server 2025 support** — The application plug-in for Microsoft SQL Server is now fully compatible with the latest version of Microsoft SQL Server.

**Microsoft SQL Server Management Studio 21.0 support** — The application plug-in for Microsoft SQL Server is now fully compatible with the latest version of SQL Server Management Studio, allowing you to run Configuration, Backup, and Restore wizards directly from it.

## Veeam Plug-in for Oracle RMAN

**Backup performance improvements** — Data retrieval from RMAN is now up to twice as fast, thanks to the use of shared memory buffer by the Veeam Data Mover running on the same database server with RMAN.

**Non-SYSDBA processing support** — When the database authentication mode is used, Oracle databases can now be protected using only the SYSBACKUP role without requiring SYSDBA privileges. This enhances security by limiting the privileges required for backup operations. To enable this functionality, set the UseSysbackup plug-in configuration option to true.

**Linux distributions support** — The application plug-in for Oracle RMAN is now fully compatible with the following Linux distributions: RHEL 9.5, 9.6 and 10.0; Oracle Linux 9.6; SLES 15 SP 7.

## Veeam Plug-in for SAP HANA

**SAP HANA System Replication support** — SAP HANA deployment leveraging the System Replication feature are now fully supported. Ensure that the value of the plug-in configuration parameter customServerName is identical on all SAP HANA nodes. For instructions on using federated and independent modes, refer to KB4391.

**Linux distributions support** — The application plug-in for SAP HANA is now fully compatible with the following Linux distributions: RHEL 8.8 for SAP Solutions; SLES 15 SP5 for SAP Applications.

## Veeam Plug-in for SAP MaxDB

The newest addition to our application plug-ins family closes the final gap in supporting all databases that can serve as a backend for SAP Business Suite, by enabling native MaxDB backups to be streamed directly to Veeam repositories. The plug-in's architecture is modeled after our existing SAP plug-ins and leverages the Backint interface, providing enterprise-level scalability and high backup and restore performance. The plug-in supports both database clusters and standalone servers and offers a full range of platform features available to other plug-ins, including immutable backups in object storage or hardened repository, backup copy jobs, scale-out backup repository and more.

Veeam Plug-in for SAP MaxDB requires version 13.0.1 or later.

## Microsoft Entra ID

**Intune policies protection** — Seamlessly protect your Device Configuration profiles and Device Compliance policies along with other Entra ID tenant data, and restore entire policies or their individual properties with the help of convenient policy metadata comparison between production environment and selected restore point. This capability requires version 13.0.1 or later.

**Backup Copy support** — You can now leverage Backup Copy jobs for Entra ID backups, enabling automatic creation of secondary backups of your Entra ID data for added protection and compliance with 3-2-1 rule. Unlike primary backups of Entra ID, backup copies can be stored on immutable storage, safeguarding your data against accidental or malicious deletion or encryption by ransomware. This capability also enables you to keep a copy of your Entra ID backups offsite, helping to improve resiliency against natural disasters and meet key regulatory requirements.

**Granular backup scope and permissions** — The Entra ID tenant wizard now enables you to define a granular backup scope by selecting specific resource types for protection. Only the permissions required for these resources will be assigned to the backup application, improving security. Previously, all resources were backed up by default; now, you have full control over both backup content and permissions.

## MongoDB

**Replica set oplog backup** — The MongoDB backup policy now offers periodic oplog (aka operation log or transaction log) backup, enabling recovery to a specific point in time, to minimize potential data loss. For optimal performance and reliability, oplog backups follow the preferred node selection logic, similar to image-level backups. Point-in-time recovery for replica sets, databases, and collections is available in Veeam Explorer for MongoDB.

**Scale-out repository offload support** — MongoDB backups now support being offloaded or copied to Capacity Tier, including fully seamless and transparent restores directly from Capacity Tier.

**Instant database publication** — Veeam Explorer for MongoDB now allows for publishing a point-in-time state of a MongoDB instance to a selected server for dev/test purposes, by running the instance directly from backup. Additionally, you can recover databases and collections directly from the published state, including any changes that have been made to the database while it was published.

**X.509 authentication** — MongoDB protection groups and Veeam Explorer for MongoDB now provide authentication options that allow you to specify client and server certificates for secure, certificate-based connections to replica sets.

## Backup Infrastructure

### Hardened Repository

**VIA profile for hardened repository** — The direct descendant of managed Veeam Hardened Repository ISO, the dedicated Veeam Infrastructure Appliance profile allows for deploying hardened repositories with even more ease thanks to passwordless initial pairing and no requirement for Linux OS expertise for deployment and on-going management. Compared to other VIA profiles, it reduces hardened repository exposure by not having the host management web UI enabled by default and automatically disabling it after a timeout.

**Simplified target path selection** — The Hardened Repository registration wizard will no longer show system directories on the repository path selection step, reducing configuration errors and improving security.

**Simplified repairing** — Easily reconnect hardened repositories to another backup server using certificate-based authentication and time-limited PIN codes.

**Hardened repository visibility** — By popular demand, hardened repositories are now hidden from the Files node of the management tree and are also not available for selection in the File Copy jobs.

## Object Storage

**Immutability for Google Cloud Storage** — Backups stored in Google Cloud Storage can now also be made immutable thanks to Object Retention Lock functionality added by Google, helping you to safeguard your backups against accidental or malicious deletion or modification.

**Immutability architecture improvements** — The immutability model has been optimized to reduce storage overhead and the number of API calls by shifting from protecting the entire state of the backup chain to protecting individual restore points for the duration of their retention policy. With the new approach, object storage space requirements are reduced by 40%, the number of PutObjectRetention API call is reduced 8x, the number of ListBucket API call is reduced 3x, and the total number of API calls is reduced more than 2x.

At this time, the new immutability architecture applies to direct backup to object storage only, based on smaller footprint and on-prem object storage being much more impacted due to its strictly limited capacity. Once proven stable we will apply this to Capacity Tier of scale-out backup repository.

## Scale-out Backup Repository

**Expanded Archive Tier compatibility** — You can now use Archive Tier with any type of Performance or Capacity Tier in your Scale-Out Backup Repository, with all previous platform combination restrictions removed. This provides complete flexibility to design the backup archival solution irrespective of what storage type you are using for your other scale-out backup repository tiers.

## Primary Storage Integrations

**Universal Storage API v2.1** — This minor API framework update allows storage vendors to make their plug-ins compatible with Veeam Software Appliance (with the first Plug-in for IBM FlashSystem already available); delivers support for the NVMe-oF protocol family (NVMe-RDMA, NVMe-FC, NVMe-TCP); and enables vendors to expose native immutability capabilities, empowering users to create immutable snapshots that provide an extra layer of protection against accidental or malicious deletion.

**Time-based retention for storage snapshots** — You can now choose between restore point-based and time-based retention for snapshots created by snapshot-only backup jobs and for backup jobs with primary storage snapshot retention policy.

**vTPM support** — Backup from storage snapshots now supports VMs with virtual TPMs.

**HPE Alletra 9000/MP NVMe-oF support** — Enables lightning-fast connectivity between backup proxies and storage arrays over NVMe-RDMA, NVMe-FC, and NVMe-TCP protocols.

**Dell Unity REST API** — We've transitioned the existing integration from using uemcli to REST API, ensuring continuous support for Dell Unity XT and Unity storage systems, enabling compatibility with the Veeam Software Appliance and improving overall integration efficiency.

## Secondary Storage Integrations

### General

**Persistent Data Mover on deduplication appliances** — Veeam Data Mover is now installed permanently on ExaGrid, Quantum DXi, Infinidat InfiniGuard, and Fujitsu ETERNUS CS800 arrays upon initial connection. This provides V13 compatibility and eliminates the need for elevated privileges required to deploy runtime components for each connection, reducing potential security exposure. Note that this requires updated storage firmware.

### Dell Data Domain

**Accelerated Direct Restore to Azure** — We adopted “sequential read, random write” approach for Direct Restore to Azure, significantly improving restore performance while reducing storage load.

Governance mode immutability support — The ability to skip Compliance mode check using the `DDBoostSkipComplianceModeCheck` configuration option has been added.

**Data Domain OS support** — Support for DD OS versions up to 8.3 by updating Data DD Boost SDK to version 8.4 is included.

### HPE StoreOnce

**Long backup chains support** — The maximum backup chain length increased up to 90 times by leveraging the Catalyst capability to open multiple objects within a single data session to overcome storage connection limits.

**Direct backup to HPE Cloud Bank** — Primary backup and backup copy jobs can now be pointed to Cloud Bank stores with HPE Alletra Storage MP X10000 as a backend. VMware and Hyper-V backups support is available immediately, and experimental support is available for other selected workloads.

**Accelerated Direct Restore to Azure** — We adopted the “sequential read, random write” approach for Direct Restore to Azure, significantly improving restore performance while reducing storage load.

**Catalyst Copy jobs for Proxmox backups** — Previously available with experimental support designation, this capability is now fully supported.

**StoreOnce OS support** — Support for StoreOnce OS version 5.1.0, updated HPE Catalyst SDK to version 4.3.9 is now included.

### Tape

**LTO10 support** — V13 provides full support for the latest LTO standard bringing cartridges with native storage capacity of 30 TB for most cost-efficient long-term data archival with true air-gap.

**Object storage backups export** — Backup to Tape jobs now support object storage backups as a source, enabling seamless long-term archival to tape.

**GFS retention for file backups** — File to Tape jobs now support GFS media pools with daily, weekly, and monthly media sets. This structured, policy-driven retention is ideal for organizations with strict archival and regulatory requirements, while intelligent retry capability ensures backup consistency by automatically rescheduling missed or failed jobs.

**Expanded virtual fulls availability** — Exporting virtual full backups is now supported for all backup modes, as opposed just forever forward incremental. This simplifies tape backup strategies and better positions customers for the deprecation of reverse incremental backup mode.

**Tape metadata processing enhancements** — The tape cataloging process has been redesigned to eliminate cross-process file access and reduce reliance on temporary directories. This improves performance and reliability, especially with resource-constrained Linux-based infrastructure.

**Managed servers as data sources** — All managed servers are now automatically added to the Inventory view and registered as available data sources for File to Tape jobs. This streamlines configuration and ensures that all eligible servers are readily available for tape backup operations without any additional steps.

## Backup Console

**Refreshed UI design** — We modernized the look and feel of the backup console, taking this opportunity to reduce clutter, ensure UI controls layout consistency and improve adherence to accessibility standards required by EAA (European Accessibility Act) and the ADA (Americans with Disabilities Act) regulations.

**Dark mode support** — The number one request for the backup console is finally here! This required us to redesign over 4000 icons into the vector graphics format, which is why it took us a while to deliver. We hope you will enjoy the result!

**Reporting built-in** — The new Reporting section in the Analytics node, powered by Veeam ONE, allows users to view reports using Veeam ONE reports directly from the backup console.

**Credential Manager improvements** — Scoping, sorting, and the ability to bulk delete credentials is added. This improves the credential management experience in large environments.

**Encryption password verification** — You can now verify stored backup encryption passwords to confirm your memory, or records of them, are still correct, helping to ensure you can recover your encrypted backups when needed. Thanks to built-in brute-force attack protection, this new capability cannot be used to perform attack on stored passwords.

**Security & Compliance Analyzer improvements** — In addition to dedicated VSA hardening checks, we added new stored password complexity and Data Domain immutability mode checks. In addition, we enabled the ability to generate reports on individual sessions' results from the History view to help streamline internal audits, and added new PowerShell cmdlets for automated test execution and results retrieval.

**SureBackup enhancements** — The SureBackup job wizard now provides "From Backup" selection option in Linked Jobs and Exclusions dialogs. We also added a "Last Run" column to the SureBackup jobs grid to provide visibility on the last job activity.

**Repository configuration validator** — The Add Backup Repository wizard now explicitly prohibits the creation of repositories with path nested in directories used by other repositories to prevent possible data losses caused by such configuration.

**File system access restrictions** — For added VSA security, backup console access to the local file system is limited to /var/lib/veeam directory, providing users with the designated place to upload files like YARA rules or scripts.

**Console download** — The backup console setup can now be downloaded directly from the backup server using the link at the bottom of the Web UI logon page.

## 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

**Standalone module for Linux** — You can now install the Veeam Backup & Replication PowerShell module to a RHEL 9 or Rocky Linux 9 based server with a dedicated PowerShell package and start using your PowerShell scripts with PowerShell 7.

**Update management** — New cmdlets to list and install available update packages to your Veeam Software Appliance or configure auto-update options are added.

**IBM Cloud object storage** — New cmdlets now allow you to register and configure your IBM object storage repositories.

**FLR for Windows** — Now specify a server to mount backups to when performing automated file-level recovery. You can use any Windows server with the Veeam Mount Service installed.

**Certificate-based authentication** — We added support for the new authentication type to cmdlets for Microsoft Hyper-V, Microsoft Windows or Linux servers registration.

**SOBR log collection** — Export-VBRLogs cmdlet now supports collecting logs from scale-out backup repositories with a dedicated parameter.

### REST API

Expanded backup server REST API coverage enabling:

**Agent Management** — Create the following protection group types: individual computers, computers from CSV file, Active Directory based and computers with pre-installed backup agent; configure backup jobs for managed Windows and Linux agents; perform file-level restores.

**Backup Copy jobs** — Now you can create backup copy jobs with backup jobs and backup repositories as supported sources.

**Microsoft Hyper-V protection** — Manage Hyper-V servers and backup proxies; create and run Hyper-V backup jobs; initiate entire VM restore, instant VM recovery and file-level restore.

**Certificate-based authentication** — Generate Veeam Deployment Kit and configure your Linux servers to allow certificate-based authentication.

**Datacenter credentials** — Configure your Group Managed Service Accounts (GMSA) to use for guest processing.

**Azure Blob Storage repository** — Configure immutability for your Azure Blob Storage repositories.



**Incident API** — Requests for creating malware detection events now supports querying machines by their object or restore point IDs and defining event severity as infected or suspicious. Additionally, the dedicated endpoint allows you to mark detected objects as clean.

**Instant Recovery to Microsoft Azure** — Manage the full instant recovery lifecycle from mounting backups to switching your machine over to production.

**PostgreSQL processing** — PostgreSQL application-aware processing options are now available for all backup and replication job types covered in backup server REST API.

**Veeam Data Cloud Vault** — Create Veeam Data Cloud Vault repositories and assign them to backup jobs now available.

**VMware vSphere global exclusions** — Add vSphere and Cloud Director machines to your global exclusions list.