Veeam Backup & Replication v13.1

What's New

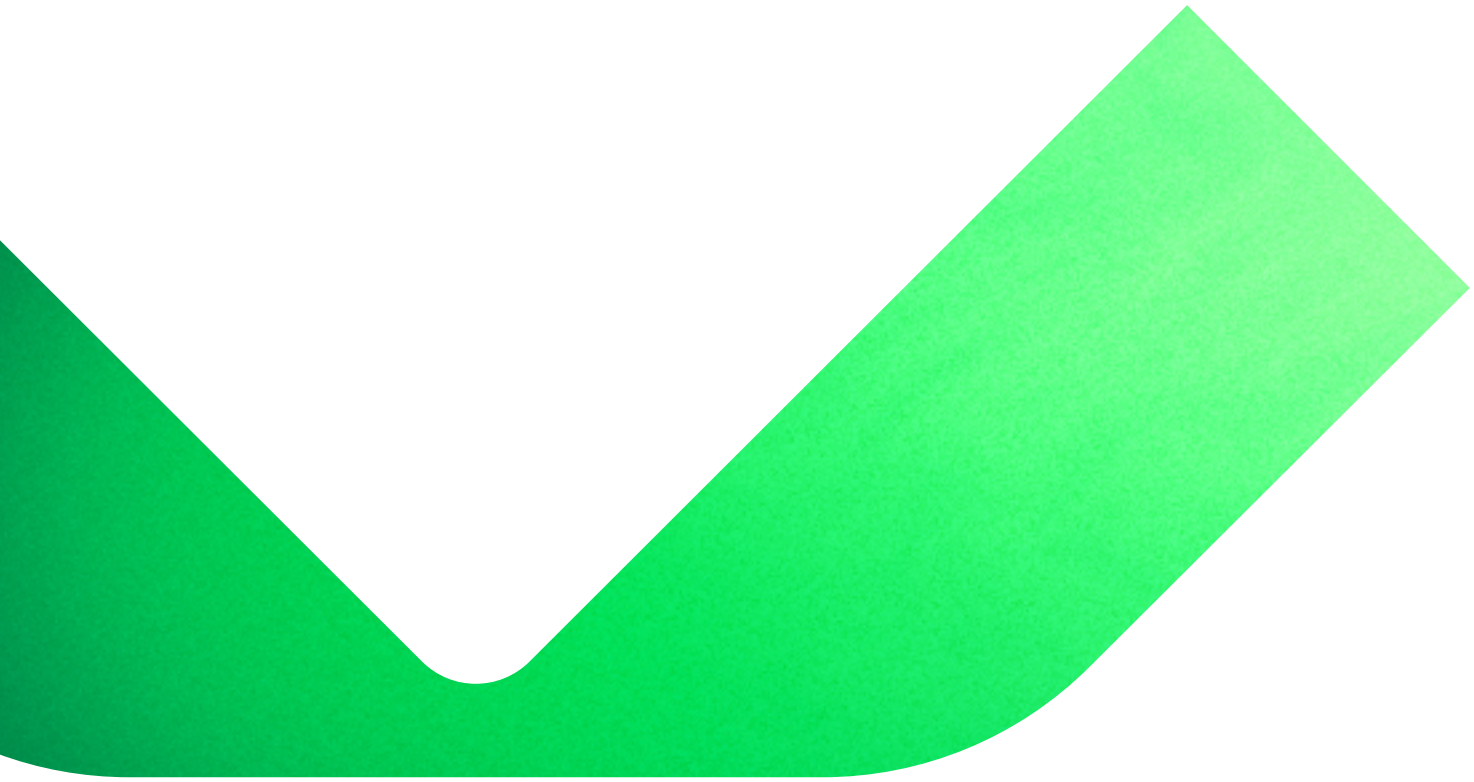

Contents

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 data protection landscape is undergoing its most significant transformation in a decade. Customers are navigating hypervisor migrations at pace, facing an identity-centric threat environment where a single compromised Identity System or Provider can halt an entire organization, and managing sprawling hybrid estates that span physical servers, cloud workloads, and applications at enterprise scale. Veeam Backup & Replication v13.1 addresses all three with precision.

Built for a multi-hypervisor future

The market is moving away from widely established hypervisor dependency faster than most organizations anticipated, and the question every IT leader is asking is not whether to migrate, but where to land and how to get there without disruption. Veeam Data Platform v13.1 makes that answer straightforward. Native support for Citrix XenServer, XCP-ng, and Sangfor aSV joins an already broad hypervisor portfolio. The new Universal Hypervisor API (UHAPI) opens a Veeam certified, vendor-driven path for additional platforms to join the ecosystem; with Platform9 and VergeOS integrated in this release. A standardized KubeVirt backend accelerates the next wave of KubeVirt-based integrations. With cross-platform VM recovery, disk export, and seamless recovery of guest files and application data, Veeam is far more than a backup solution for your destination hypervisor - it's the foundation that makes your workloads portable, wherever you're headed.

Identity is the perimeter - Veeam protects it

Compromised credentials are the entry point for the majority of cyberattacks today, and Active Directory/ Entra ID remains the single most targeted piece of infrastructure in the enterprise. When the forest falls, every authentication-dependent service falls with it. Veeam v13.1 delivers automated Active Directory Forest Recovery - turning what was once a multi-day expert-led crisis response requiring over 40 manual Microsoft-documented steps into a structured, wizard-driven process any trained administrator can execute, covering single-domain, multi-domain, and multi-tree topologies. Microsoft Entra ID protection gains persistent object relationship reconstruction, expanded coverage for organization contacts, device objects, and BitLocker recovery keys, and context-aware restore that handles encrypted properties inline. Identity is at the heart of your resilience strategy.

Enterprise-scale protection without enterprise-scale complexity

From ultra-large database protection at 100+ TB scale via the new Veeam Snap Scale engine, to Application Backup Repository turning the Veeam Hardened Repository into a managed, immutable NFS target for Oracle RMAN Incremental Merge and any application (including IoT devices) that can export or dump backups, v13.1 extends Veeam's reach into workloads that once required separate, specialized tooling or appliances. Backup targets gain more flexibility, with expanded archive-class storage support and scale-out options. On security, additional enhanced malware detection for more workloads, Hybrid FIPS + Post-Quantum cryptography processing, firewall friendly single-port transport, and MFA across sensitive configuration operations harden the platform against today's threats. And for administrators who want it, Veeam Intelligence becomes an optional companion that investigates and helps resolve issues directly; saving time without taking control away.

New Features

V13.1 delivers hundreds of enhancements and improvements in response to real-world customer feedback as well as additional refinements from ongoing R&D. All capabilities here are transacted as the Veeam Data Platform, with certain features available only at the Advanced or Premium editions. The most significant are highlighted below.

Workloads

Identity Protection

Microsoft Active Directory

Active Directory is the identity backbone of the enterprise. When a cyberattack or catastrophic failure compromises the Active Directory Forest, every authentication-dependent service goes with it, and the complexity of manual forest recovery has historically made this one of the most feared and time-consuming incidents an IT team can face. Microsoft lists over 40 manual steps in the process. Veeam eliminates that complexity entirely.

Automated Active Directory Forest Recovery - Automated recovery of an entire Active Directory forest is now available through a guided recovery wizard that walks administrators through each step of the restoration process, including selecting the recovery point, defining domain controllers for every domain, choosing source backups, and configuring target settings. By combining Active Directory–specific recovery intelligence with a familiar VM restore workflow, the solution minimizes complexity and helps reduce operational risk during critical recovery scenarios.

During backup operations, a guest component automatically collects and preserves Active Directory forest metadata, ensuring that all information required for a successful forest rebuild is readily available when needed. The recovery process then orchestrates VM restoration and all required post-recovery activities automatically, bringing identity services back online with minimal administrative effort and no manual recovery steps.

This release supports recovery to both VMware vSphere and Microsoft Hyper-V environments and covers single-domain, multi-domain, and multi-tree Active Directory forests. Available with paid Veeam Universal License editions (Essentials, Advanced, and Premium), this capability transforms forest-level recovery from a complex, expert-driven operation into a repeatable and guided process that can be executed with confidence by trained administrators.

Microsoft Entra ID

Restoring Microsoft Entra ID has never been more complete - this release preserves and reconstructs complex object relationships automatically, eliminates manual reconfiguration after restore, and extends protection to previously uncovered object types including Organization

Contacts, device objects, and BitLocker recovery keys. Administrators gain precise, context-aware recovery across the full identity directory, with JSON export providing a reliable audit and recovery path for every protected object regardless of platform restore limitations.

Persistent original ID - Restoring Entra ID objects could result in new identifiers being assigned, breaking relationships between complex objects. With this release, such relationships are preserved and reconstructed during restore operations: the system now identifies object references and automatically re-links them - either directly when the original ID is retained, or by matching against previously stored IDs when objects are recreated. This capability eliminates the need for manual reconfiguration, reducing the risk of incomplete restores, preventing security gaps, and accelerating recovery of complex identity environments.

Export to JSON - A new restore option is now introduced allowing administrators to export any protected directory object, including users, groups, and other objects, to a JSON file. Each export captures all object properties, relationships, and the exact schema version as it existed at the time of backup, with no conversion of older restore points.

Organization Contacts - Entra ID now protects Organization Contacts, a directory object type managed by administrators and synchronized from on-premises directories or Exchange Online. Although these contacts are read-only in Microsoft Graph, they are now fully covered by backup, including delta backups to capture incremental changes efficiently and all standard comparison methods for the object so administrators can review and analyze changes across restore points.

Add public key during application restore - Some object properties, such as certificates or license keys, are stored encrypted in the backup or cannot be protected at all due to platform limitations, which previously prevented them from being fully restored. Entra ID now supports context-aware restore that detects when an object requires supplementary information and prompts for it through the restore wizard, applying the supplied data precisely where needed as the object is recreated. For application restores, this means the wizard can pre-fill the public key, allowing administrators to provide the certificate or key inline and complete the restore without manual post-restore steps.

Devices and BitLocker Keys - with v13.1 users now can protect device objects alongside their associated BitLocker recovery keys. Because device objects cannot be restored back into Entra ID, this scenario centers on review and export: administrators can view device metadata, most importantly the Intune policy a device is assigned to and access the underlying data through JSON export. BitLocker keys can be viewed and copied directly, and device metadata can be exported to JSON, giving administrators the visibility and recovery information they need without a full object restore.

RBAC Support - For Entra ID now applies a role-based access control (RBAC) permission model to Entra ID backup and restore operations. Administrators can grant granular, role-scoped permissions that govern who can configure and run backup jobs and perform restores, ensuring that each operation is carried out only by appropriately authorized users and aligning Entra ID data protection with least-privilege access practices.

Hypervisor Protection

The data center landscape is shifting. Workload migrations are accelerating, and Veeam is built to move with them. This release significantly expands native hypervisor coverage by adding the following hypervisors to the already broad portfolio of supported platforms.

- Sangfor aSV
- Citrix XenServer
- XCP-ng
- Platform9
- VergeOS

Sangfor aSV

Sangfor aSV is now a fully supported hypervisor in Veeam Backup & Replication, with a comprehensive set of backup and restore capabilities aligned with Veeam's cross-platform protection standards.

Backup jobs - Backup jobs use Changed Block Tracking (CBT) for efficient incremental processing. Snapshot quiescence is supported for Windows virtual machines.

Backup Scan - Backup health verification, Threat Hunter and YARA-based content scanning.

Entire VM Restore - Restore virtual machines to Sangfor aSV from any supported hypervisor, cloud VM (AWS, Azure, or GCP), or physical server backup (Veeam Agent for Windows or Linux).Equally, Sangfor aSV VM backups can be restored to any other supported hypervisor or cloud target, including AWS, Azure, and GCP.

Instant Recovery to VMware vSphere, Microsoft Hyper-V, Nutanix AHV and Proxmox VE - is supported from Sangfor aSV backups. Proxmox VE Instant VM Recovery is under Experimental Support.

Disk Mount - VM disk from backups can be mounted to a target guest host in read-only mode. It is supported for both Windows and Linux guest systems.

File Level Restore - Guest File Level Restore is available without requiring a full VM recovery.

Export of Backups - Backup data can be exported to VHD, VHDX, or VMDK formats for use across platforms.

Veeam Explorer based application with Veeam heuristics detection - Veeam's application-aware restore processing can restore databases (point of backup) or application items from Oracle databases, Active Directory, SharePoint, Microsoft Exchange, and Microsoft SQL Server by detecting the applications within the VM backups at restore.

KubeVirt Proxy for virtualization plug-ins

Veeam is releasing a KubeVirt Proxy built in conjunction with our Kasten team, which provides a management interface for the growing ecosystem of KubeVirt-based Hypervisors. The first platform plug-in built on this integration layer ships shortly. The plug-in will be distributed when ready in the Veeam Virtualization Plug-in Download Center.

Universal Hypervisor API (UHAPI)

The Universal Hypervisor API gives hypervisor vendors a defined, industry-standard path to Veeam compatibility. Vendors implement the standard API, and Veeam reviews and certifies the platform. In this release, Platform9 and VergeOS are the first vendors integrated directly through UHAPI, with additional hypervisors to follow as vendors complete certification. The updated UHAPI plug-in is downloadable here.

Nutanix AHV

Extended Web UI coverage - With this release, AHV backup jobs, restore operations, and infrastructure management are fully available in the Web UI, aligned with existing VMware vSphere and Hyper-V support. This capability provides consistent management experience across hypervisors, enabling administrators to configure, monitor, and operate AHV backups from a single interface.

Advanced RBAC support - Administrators can define custom roles with scoped access to specific VMs, jobs, categories, or infrastructure objects to enable granular role-based access control, improving security and ensuring users only have access to resources relevant to their responsibilities.

VirtIO driver injection - VirtIO drivers are now automatically injected during restore, enabling seamless recovery of non Nutanix AHV backups to Nutanix AHV targets.

Instant Disk Recovery - Backed up disks can be instantly recovered for immediate read/write access without a full restore. Changes can be committed back to production or discarded after use. This capability enables faster data access for recovery, testing, and troubleshooting scenarios, reducing recovery time and improving operational flexibility.

Prism Central category AND logic - Policy scope filters support the logical AND operator for category (tag) combinations, enabling administrators to define multi-criteria selection policies.

Improved restore job statistics - Restore job progress and statistics reporting is now available in the Web UI to provide a clear insight into restore activity, helping administrators better track progress, identify issues, and manage recovery operations more effectively.

Proxmox VE

VM Replication - VM-level replication jobs come to Proxmox VE, covering the full range of deployment realities: direct host-to-host replication for lean branch office setups without shared storage, and cross-data center replication for broader DR strategies. It is possible to replicate cross storage types. The result is native disaster recovery that improves availability and reduces recovery time where it matters most.

Instant VM Recovery - With this release, Instant VM Recovery enables administrators to run Proxmox VE VMs directly from any other supported platform backups and then migrate them to production storage. This capability minimizes downtime and helps organizations meet recovery time objectives (RTO) by making workloads available immediately during recovery operations. This feature is officially supported under Veeam experimental support status conditions in 13.1.

VLAN tagging - VLAN tags now can be assigned to both worker VMs and restored VMs in Proxmox VE environments. This is crucial for migration scenarios in existing VLAN-based networks that haven't fully finalized adopting the new SDN (vNet) approach.

VirtIO driver injection - VirtIO drivers are now automatically injected during restore, enabling seamless recovery of non Proxmox VE backups to Proxmox VE targets.

Extended Web UI coverage - With this release, Proxmox VE backup jobs, restore operations, and infrastructure management are fully available in the Web UI, aligned with existing VMware vSphere and Hyper-V support. This capability provides consistent management experience across hypervisors, enabling administrators to configure, monitor, and operate Proxmox VE backups from a single interface.

Advanced RBAC support - Administrators can define custom roles with scoped access to specific VMs, jobs, categories, or infrastructure objects to enable granular role-based access control, improving security and ensuring users only have access to resources relevant to their responsibilities. This feature has experimental support status in 13.1, please refer to Veeam Experimental Support Statement for details.

Improved restore job statistics - Restore job progress and statistics reporting is now available in the Web UI to provide a clear insight into restore activity, helping administrators better track progress, identify issues, and manage recovery operations more effectively.

HPE Morpheus VM Essentials

Integrated deployment - The HPE VM Essentials plug-in is now included by default in Windows-based backup server and Veeam Software Appliance deployments, eliminating the need for separate installation and simplifying initial setup.

Application-aware processing - Veeam's application-aware image processing produces transactionally consistent backups and enable granular, point-in-time recovery across the applications your infrastructure depends on: Microsoft SQL Server, Oracle Database, PostgreSQL, Microsoft Exchange, Active Directory, and SharePoint. For SQL Server, Oracle, and PostgreSQL it ships transaction and archived log backups to the repository, enabling recovery to any point between jobs, and truncates logs automatically - no extra tooling or manual DBA work required.

VirtIO driver injection - VirtIO drivers are now automatically injected during restore, enabling seamless recovery of non-HPE backups to HPE VM Essentials targets.

Scale Computing HyperCore

VirtIO driver injection - VirtIO drivers are now automatically injected during restore, enabling seamless recovery of non-HyperCore backups to HyperCore targets.

Red Hat Virtualization (oVirt KVM)

Integrated deployment - Veeam Plug-in for oVirt KVM is now included by default in Veeam Software Appliance (VSA) deployments.

Storage optimization settings - With this release, administrators can customize data block size for backup jobs. This capability enables more efficient data transfer and storage usage by aligning backup behavior with workload and storage characteristics.

Oracle Linux Virtualization Manager OLVM (oVirt KVM)

Integrated deployment - Veeam Plug-in for oVirt KVM is now included by default in Veeam Software Appliance (VSA) deployments.

Storage optimization settings - With this release, administrators can customize data block size for backup jobs. This capability enables more efficient data transfer and storage usage by aligning backup behavior with workload and storage characteristics.

VMware

VMware VDDK 9 support - Full support for VMware Virtual Disk Development Kit (VDDK) version 9, required for compatibility with the latest vSphere releases.

VMware Cloud Foundation (VCF) 9.1 support - Enabling latest VMware vSphere based hypervisor and related technologies for data protection, VM replication and CDP processing.

Cluster-wide vPowerNFS datastore mount for Instant Recovery - During Instant Recovery, the vPowerNFS datastore can now be mounted cluster-wide, making the recovered disks accessible from any host in the vSphere cluster which enables HA, DRS and manual host maintenance scenarios for systems with Instant Recovery VMs.

Microsoft Hyper-V

SCVMM with Windows Server High Availability support - Backup server now natively supports SCVMM servers deployed on a Windows Server Failover Cluster, enabling backup management in highly available SCVMM configurations.

Citrix XenServer

Citrix XenServer is now a fully supported hypervisor in Veeam Backup & Replication, with a comprehensive set of backup and restore capabilities aligned with Veeam's cross-platform protection standards.

Backup jobs - Backup jobs use Changed Block Tracking (CBT) for efficient incremental processing.

Backup Scan - Backup health verification, Threat Hunter and YARA-based content scanning.

Entire VM Restore - from any supported hypervisor, cloud VM (AWS, Azure, or GCP), or physical server backup (Veeam Agent for Windows or Linux). Equally, Citrix XenServer VM backups can be restored to any other supported hypervisor or cloud target, including AWS, Azure, and GCP.

Instant Recovery to VMware vSphere, Microsoft Hyper-V, Nutanix AHV and Proxmox VE - is supported from Citrix XenServer backups. Proxmox VE Instant VM Recovery is under Experimental Support.

Disk Mount - VM disk from backups can be mounted to a target guest host in read-only mode. It is supported for both Windows and Linux guest systems.

File Level Restore - Direct Guest File Level Restore is available without requiring a full VM recovery.

Export of Backups - Backup data can be exported to VHD, VHDX, or VMDK formats for use across platforms.

Veeam Explorer based application with Veeam heuristics detection - Veeam's application-aware restore processing can restore databases (point of backup) or application items from Oracle databases, Active Directory, SharePoint, Microsoft Exchange, and Microsoft SQL Server by detecting the applications within the VM backups at restore.

XCP-ng

XCP-ng is now a fully supported hypervisor in Veeam Backup & Replication, with a comprehensive set of backup and restore capabilities aligned with Veeam's cross-platform protection standards.

Backup jobs - Backup jobs use Changed Block Tracking (CBT) for efficient incremental processing.

Backup Scan - Backup health verification, Threat Hunter and YARA-based content scanning.

Entire VM Restore - Restore virtual machines to XCP-ng from any supported hypervisor, cloud VM (AWS, Azure, or GCP), or physical server backup (Veeam Agent for Windows or Linux). Equally, XCP-ng VM backups can be restored to any other supported hypervisor or cloud target, including AWS, Azure, and GCP.

Instant Recovery to VMware vSphere, Microsoft Hyper-V, Nutanix AHV and Proxmox VE - is supported from XCP-ng backups. Proxmox VE Instant VM Recovery is under Experimental Support.

Disk Mount - VM disk from backups can be mounted to a target guest host in read-only mode. It is supported for both Windows and Linux guest systems.

File Level Restore - Direct Guest File Level Restore is available without requiring a full VM recovery.

Export of Backups - Backup data can be exported to VHD, VHDX, or VMDK formats for use across platforms.

Veeam Explorer based application with Veeam heuristics detection - Veeam's application-aware restore processing can restore databases (point of backup) or application items from Oracle databases, Active Directory, SharePoint, Microsoft Exchange, and Microsoft SQL Server by detecting the applications within the VM backups at restore.

Platform9

Platform9 has integrated their hypervisor into the Veeam Universal Hypervisor API and is now a fully supported hypervisor in Veeam Backup & Replication, with a comprehensive set of backup and restore capabilities aligned with Veeam's cross-platform protection standards.

Backup jobs - VM backup jobs allow protection of your workloads.

Backup scan - Backup health verification, Threat Hunter and YARA-based content scanning.

Entire VM Restore - Restore virtual machines to Platform9 from any supported hypervisor, cloud VM (AWS, Azure, or GCP), or physical server backup (Veeam Agent for Windows or Linux). Equally, Platform9 VM backups can be restored to any other supported hypervisor or cloud target, including AWS, Azure, and GCP.

Instant Recovery to VMware vSphere, Microsoft Hyper-V, Nutanix AHV and Proxmox VE - is supported from Platform9 backups. Proxmox VE Instant VM Recovery is under Experimental Support.

Disk Mount - VM disk from backups can be mounted to a target guest host in read-only mode. It is supported for both Windows and Linux guest systems.

File Level Restore - Direct Guest File Level Restore is available without requiring a full VM recovery.

Export of Backups - Backup data can be exported to VHD, VHDX, or VMDK formats for use across platforms.

Veeam Explorer based application with Veeam heuristics detection - Veeam's application-aware restore processing can restore databases (point of backup) or application items from Oracle databases, Active Directory, SharePoint, Microsoft Exchange, and Microsoft SQL Server by detecting the applications within the VM backups at restore.

VergeOS

VergeOS has been integrated into the Veeam Universal Hypervisor API and is now a fully supported hypervisor in Veeam Backup & Replication, with a comprehensive set of backup and restore capabilities aligned with Veeam's cross-platform protection standards.

Backup jobs - Backup jobs use Changed Block Tracking (CBT) for efficient incremental processing.

Backup Scan - Backup health verification, Threat Hunter and YARA-based content scanning.

Entire VM Restore - Restore virtual machines to VergeOS from any supported hypervisor, cloud VM (AWS, Azure, or GCP), or physical server backup (Veeam Agent for Windows or Linux). Equally, VergeOS VM backups can be restored to any other supported hypervisor or cloud target, including AWS, Azure, and GCP.

Instant Recovery to VMware vSphere, Microsoft Hyper-V, Nutanix AHV and Proxmox VE - is supported from VergeOS backups. Proxmox VE Instant VM Recovery is under Experimental Support.

Disk Mount - VM disk from backups can be mounted to a target guest host in read-only mode. It is supported for both Windows and Linux guest systems.

File Level Restore - Direct Guest File Level Restore is available without requiring a full VM recovery.

Export of Backups - Backup data can be exported to VHD, VHDX, or VMDK formats for use across platforms.

Veeam Explorer based application with Veeam heuristics detection - Veeam's application-aware restore processing can restore databases (point of backup) or application items from Oracle databases, Active Directory, SharePoint, Microsoft Exchange, and Microsoft SQL Server by detecting the applications within the VM backups at restore.

Platform-wide Hypervisor Improvements

Universal worker image - All hypervisor workers use a single, unified image that is FIPS and DISA STIG compliant. This capability simplifies maintenance and patching across environments while strengthening the overall security posture.

Backup mapping - Additional hypervisor backup jobs support backup mapping, allowing existing backup chains to be continued after repository changes or job reconfiguration to reduce the need for additional full backups, saving storage and time while simplifying backup management workflows.

Configurable Backup Window - Additional hypervisor backup jobs support now Backup Window Configuration, allowing administrators to restrict job execution to defined operational time windows. This capability helps prevent interference with critical workloads, ensuring backups run only during approved time periods while improving scheduling flexibility.

Trackable events - Backup and restore session events are generated to enable integration with Veeam ONE alarms, SIEM systems (syslog), and Incident API workflows, improving monitoring, automation, and incident response.

Cloud Workload Protection

For organizations that require full control over their cloud backup policies, data sovereignty, and infrastructure decisions, Veeam Backup for Azure, AWS, and Google Cloud delivers customer-managed protection for cloud workloads; complementing the fully managed Backup-as-a-Service coverage available through Veeam Data Cloud for those who prefer to offload that responsibility entirely.

Amazon AWS

Veeam Data Cloud Vault support - V13.1 lets you now connect Veeam Data Cloud Vault directly into Veeam Backup & Replication through the Veeam AWS appliances that manage it. Connect a Vault-backed appliance to the Veeam Backup & Replication console, or provision the repository on the appliance from within Veeam Backup & Replication itself. Either way, the repository registers in Veeam Backup & Replication as an external repository, no shared keys required, and becomes immediately available for restores, backup copy jobs, and other secondary workflows. Access is governed by federated, role-based authentication rather than long-lived credentials, keeping the integration both operationally clean and secure.

AWS China Region Support - The backup appliance can now be deployed in AWS China regions, deployed and managed from Veeam Backup & Replication; which must be running in the China region, either on-premises or on an EC2 instance in an AWS China region.

S3 Standard-Infrequent Access for Backup Repositories - Backups can now be stored in S3 Standard-Infrequent Access as the storage class for backup repositories within the backup appliance.

Microsoft Azure

Veeam Data Cloud Vault support - V13.1 lets you now connect Veeam Data Cloud Vault directly into Veeam Backup & Replication through the Veeam Microsoft Azure appliances that manage it. Connect a Vault-backed appliance to the Veeam Backup & Replication console or provision the repository on the appliance from within Veeam Backup & Replication itself. Either way, the repository registers in Veeam Backup & Replication as an external repository, no shared keys required, and becomes immediately available for restores, backup copy jobs, and other secondary workflows. Access is governed by federated, role-based authentication rather than long-lived credentials, keeping the integration both operationally clean and secure.

Microsoft Entra ID Authentication Support for Azure SQL - Backup policies protecting Azure SQL databases can now authenticate using a Microsoft Entra ID (formerly Azure AD) application instead of SQL credentials. This aligns Azure SQL backup with modern identity standards and enables customers to enforce Entra ID-based access controls consistently across their Azure SQL environments - eliminating the need to manage SQL logins for backup service accounts.

Malware Detection for Azure VM Backups - Azure VM backups support malware detection using guest indexing analytics and entropy analysis in addition to already available signature-based analysis and YARA rule scanning. This capability enables organizations to detect suspicious activity directly within Azure VM backups, improving threat visibility and strengthening ransomware protection for cloud workloads.

Disk Exclusion in VM Backup Policies - Exclude specific VM disks from backup policies using either a disk ID or Azure resource tags.

Auto-Approval of Private Endpoints for Azure SQL - In private network deployment configurations, the backup appliance can now automatically approve private endpoint connection requests for Azure SQL workloads by enabling the feature for the specific service account.

Direct Restore to Microsoft Azure

Network Security Groups - A new "Do not assign (use subnet settings)" option lets restored workloads inherit Network Security Groups pre-attached to the chosen subnet, so security admins can enforce their own hardened inbound rules instead of the loose default 0.0.0.0 source. "Empty" option has also been changed to "Create New" as it better reflects the underlying action. Applies to the Azure restore proxy along with VMs processed in scope of Instant and Direct Restore to Azure.

NVMe Disks - VMs processed in scope of Direct Restore to Azure and Azure Restore Proxies now support setting v6 and v7 VM sizes with NVMe disks.

Instant Restore to Microsoft Azure

Scalability - For a single Instant Recovery to Azure session, the maximum recommended number of processed VMs is now set to 250. This value is provided as guidance to help ensure stable and predictable performance during large restore operations. To reduce the possibility of Azure-side throttling and related slowdowns, VMs are now processed in parallel in batches of 80.

Google Cloud

Immutable repositories support - Veeam Backup for Google Cloud can now use immutable Google storage as a backup repository. This prevents backup files from being deleted, modified, or encrypted by malicious actors or accidental admin actions for a predefined retention period. This feature requires the Veeam Plug-in for Google Cloud v8 release.

Agent-based Protection

With this release, Veeam expands capabilities for scalable agent deployment, stronger security, improved data protection, and more flexible recovery options. You will find updates that simplify large-scale management, enhance ransomware resilience, improve backup reliability on modern endpoints, and extend platform support, along with new features that enable centralized control and consistent backup operations across diverse environments.

Agent Management

Direct Agent Backups to Veeam Data Cloud Vault and Microsoft Azure Blob Storage - Unmanaged Veeam Agents and managed by agent jobs can back up directly to Veeam Data Cloud Vault and Microsoft Azure Blob Storage. This capability provides ransomware-resilient, always-immutable storage for direct-to-object agent backups, enabling organizations to achieve the same level of data protection already available across the rest of the platform.

Access to historical backups - Previously, backups created by agent-managed machines were associated with the original device, making historical restore points inaccessible to end users after a machine was replaced, decommissioned, or reassigned. With Veeam Backup & Replication 13.1, administrators can grant managed agents access to historical backups, preserving self-service file-level recovery capabilities. As a result, users can recover data directly from the agent UI using backups created by previously used machines, without relying on the original device. For large-scale hardware refresh and migration projects, backup administrators can automate backup-to-device matching using PowerShell and the REST API.

Increased retention for agent backups - Agent-managed backup jobs were limited by a maximum retention value enforced in the UI. With this release, this limitation has been removed, allowing administrators to define retention periods without predefined limits. This enables organizations to meet compliance and data governance requirements while ensuring long-term data availability without workarounds.

Certificate-based authentication for Microsoft Active Directory protection groups - With this release, protection groups support certificate-based authentication, eliminating the need to store and manage privileged credentials. This capability improves security by reducing credential exposure and simplifies operations by removing the need for credential rotation and management.

Cluster Disk Transfer Optimization for Backup Copy Jobs - Although cluster shared disks are stored only once in backups, backup copy jobs previously performed consistency processing for each cluster node referencing the disk, generating additional transfer overhead. Backup copy jobs now recognize shared disks across cluster nodes and process them only once per job run, reducing redundant data transfer, improving efficiency, and accelerating copy operations for cluster-based environments.

Pre-built deployment script for pre-installed Windows Agents - A ready-to-use deployment script is included for Windows agents directly to a package export location. This capability simplifies large-scale agent deployment, reduces operational overhead, and ensures a consistent deployment approach across environments.

Independent Deployment Kits - Creating a new deployment kit no longer invalidates or deprecates existing deployment kits, enabling administrators to maintain separate deployment configurations for different environments, use cases, or rollout scenarios. This enhancement simplifies agent deployment management and provides greater flexibility when supporting multiple deployment strategies simultaneously.

Web UI support - Agent Management capabilities are expanded in the Web UI, including support for protection groups (both individual machines and Active Directory–based), Windows and Linux agent backup job management, and common recovery operations to simplify management of agent-based workloads, enabling administrators to configure, monitor, and recover systems directly from a single, browser-based interface.

Veeam Agent for Windows

Remote Bare Metal Recovery - Administrators can now initiate and manage Bare Metal Recovery (BMR) sessions for Windows machines directly from the Web UI. Recovery operations can be performed using pre-configured Recovery Media or a recovery partition available on the target machine. End users simply boot the affected machine into the pre-configured recovery environment, allowing administrators to perform the remainder of the recovery process remotely through the Web UI. Available in Advanced Edition and above, this capability reduces downtime, minimizes the need for hands-on IT intervention, and streamlines disaster recovery operations.

Virtual recovery partition - Backup server can create and store a recovery environment directly on the target machine to enable Remote Bare Metal Recovery without relying on PXE, providing greater deployment flexibility and ensuring recovery readiness across diverse environments.

Agent UI Branding and Customization - With this release, administrators can customize the branding of backup agents, including company logos, accent colors, and Support and Help Center links. This capability enables consistent, branded deployments that align with corporate identity standards across the organization. Available in Advanced Edition and above, these customization options can also be used by service providers to deliver white-labeled experiences tailored to their customers, helping create a more seamless and recognizable user experience.

Automated Bandwidth Throttling - Veeam Agent for Microsoft Windows running on workstation operating systems can now automatically adjust backup traffic based on current network conditions, eliminating the need for manually configured throttling rules. By dynamically adapting bandwidth consumption, backups can continue running transparently in the background while minimizing the impact on user activity and business applications. This capability helps maintain end-user productivity while ensuring backup operations complete efficiently, even on bandwidth-constrained connections.

Modern Standby handling - Previously, on devices using Modern Standby (S0 low-power idle), backup jobs could fail or be skipped because the operating system throttles background services and ignores wake timers, preventing agents from executing scheduled operations reliably. With this release, Veeam Agent for Windows detects Modern Standby mode and handles it explicitly by adapting backup behavior, including preventing unreliable job starts and avoiding interruptions when the system transitions into low-power states.

Job retry after system restart - Veeam Agent for Windows backup jobs interrupted by a system restart are now automatically retried, helping prevent missed backup points in environments with scheduled maintenance reboots.

Veeam Agent for Microsoft Windows on Microsoft Store - Veeam Agent for Microsoft Windows is now available through the Microsoft Store, providing a simplified and familiar installation experience for Windows users. By leveraging Microsoft Store distribution, organizations and individual users can more easily discover, deploy, and update the agent through trusted Windows software delivery mechanisms. This helps streamline agent deployment and maintenance while reducing the operational overhead associated with manual software distribution.

Veeam Agent for Linux

Centralized management for nosnap Veeam Agent for Linux - Introducing centralized deployment and management for the nosnap version of Veeam Agent for Linux, which leverages native snapshot capabilities of supported Linux file systems (LVM & Btrfs) and does not depend on Veeam's blksnap kernel module. Previously, this agent variant had to be manually installed on each machine. With v13.1, the nosnap Linux agent can now be deployed and managed centrally in the same way as the standard Linux agent. This functionality is available for x86_64 Linux systems.

Backup window support - Adding backup window enforcement for Veeam Agent for Linux when jobs are managed via Veeam Backup & Replication. Administrators can now define allowed and prohibited time intervals for agent backup jobs, preventing overlap with production hours and ensuring backup traffic does not impact business-critical operations.

Restore cross Backup & Replication Server management via CLI - V13.1 allows adding additional backup servers for Linux agents using CLI commands, enabling continued managed mode operation without resetting the agent state and avoiding backup chain disruption. The additional backup server is connected read-only and general backup processing can continue to run with the original configured backup server.

Active full backup scheduling enhancements - Active full backup scheduling in standalone mode now includes a monthly day-of-week option, consistent with synthetic full backup settings. You can configure active full backups to run on a specific day of the week on selected months or use the existing day-of-month schedule, providing greater flexibility for aligning backup schedules with operational calendars.

Linux distribution support - Added support for x86_64 distributions of Ubuntu 26.04 LTS; RHEL 9.8 and 10.2; Rocky Linux 9.8 and 10.2; AlmaLinux 9.8 and 10.2 and IBM Power distributions of RHEL 9.8 and 10.2. Support for the x86_64 distribution of RHEL 7.9 ELS has been added to restore compatibility with legacy RHEL environments.

Linux kernel support - Support was added for Linux kernel version 7.0.

Updated Veeam Recovery Media - V13.1 upgrades Veeam Recovery Media for Veeam Agent for Linux to Debian 13, ensuring alignment with the latest stable release and delivering broader hardware compatibility for reliable bare metal recovery operations.

Recovery Media patching on IBM Power (ppc64le) - Patching Recovery Media on the ppc64le architecture was added, enabling users to generate customized recovery ISO images that include the required drivers for IBM Power Systems.

Veeam Agents for Mac

Backup window support - Adding backup window enforcement for Veeam Agent for Mac when jobs are managed via Veeam Backup & Replication. Administrators can now define allowed and prohibited time intervals for agent backup jobs, preventing overlap with production hours and ensuring backup traffic does not impact business-critical operations.

Restore cross Backup & Replication Server management via CLI - V13.1 allows adding additional backup servers for Mac agents using CLI commands, enabling continued managed mode operation without resetting the agent state and avoiding backup chain disruption. The additional backup server is connected as read-only, and general backup processing can continue to run with the original configured backup server.

Network usage control - V13.1 adds network usage settings for Veeam Agent for Mac, allowing administrators to throttle backup job bandwidth and limit network consumption when creating protection groups with pre-installed agents.

Veeam Agents for Unix (AIX & Oracle Solaris)

Malware Detection for AIX and Solaris backups - With this release, malware detection capabilities such as including guest file analytics, Threat Hunter, YARA scans, and Veeam Incident API integration are fully supported for backups created with Veeam Agent for Unix (AIX and Solaris).

Certificate based authentication, Single-use credentials and Deployment Kit support for IBM AIX agents - V13.1 expands deployment and security options for Veeam Agent for IBM AIX in two key areas. First, single-use credentials are now supported when adding IBM AIX machines to individual protection groups, allowing SSH credentials to be used once for initial agent deployment only. As part of this, certificate-based authentication is now supported for IBM AIX agents, enabling more secure, passwordless communication between the agent and Veeam Backup & Replication after the initial deployment. Second, IBM AIX packages are now included in the Veeam Deployment Kit, enabling pre-installation of Veeam Deployer Service with a temporary certificate and bringing IBM AIX in line with the automated deployment capabilities already available for other platforms.

Managed by Server backup job mode for UNIX agents - Introducing support for the Managed by backup server job mode, allowing backup jobs for IBM AIX and Oracle Solaris to be configured, scheduled, and executed directly from Veeam Backup & Replication. Unlike agent-managed policies, server-managed jobs give backup administrators centralized control over job execution and scheduling, without relying on the agent to initiate sessions independently.

Backup window support - Adding backup window enforcement for Veeam Agents for IBM AIX and Oracle Solaris when jobs are managed via Veeam Backup & Replication. Administrators can now define allowed and prohibited time intervals for agent backup jobs, preventing overlap with production hours and ensuring backup traffic does not impact business-critical operations.

Intelligent data processing - File-level backup performance is improved through parallel data processing: data is transferred to the backup server using multiple parallel streams, with large and small files processed concurrently to better utilize available network bandwidth. The degree of parallelism is tuned dynamically based on observed performance from previous backups, improving progressively as statistics are collected.

Gateway mode for S3-compatible object storage - V13.1 adds support for Gateway mode when backing up Veeam Agents for IBM AIX and Oracle Solaris to S3-compatible object storage, complementing the existing Direct mode. In Gateway mode, all backup and restore traffic is routed through a gateway server assigned in the Veeam Backup & Replication console. This is better suited for environments with network restrictions for on-premises S3 storage or where direct internet access from endpoints to a cloud based S3 storage are not permitted.

Export skipped files as .csv file report - Users can now get a .csv-formatted list of files skipped during backup jobs for Veeam Agents for IBM AIX and Oracle Solaris directly from Veeam Backup & Replication when exporting backup job logs. This enables centralized review of skipped files from the backup server without inspecting logs on individual UNIX systems, improving compliance, and audit readiness. Alternatively, the .csv file can also be collected from each instance individually.

IBM AIX 6.1 support - Veeam Agent for IBM AIX reintroduces support for AIX version 6.1, restoring compatibility with legacy AIX environments in addition to existing AIX 7.x coverage.

Universal CDP

Universal Continuous Data Protection: Linux support - Following the introduction of agent-based Universal Continuous Data Protection (CDP) for Windows in V13, V13.1 extends near-zero RPO continuous replication to Linux workloads as well. You can now continuously replicate any Linux machine - physical, virtual, or cloud - to VMware vSphere or VMware Cloud Director (through Veeam Cloud Connect for Cloud Service Provider and Sovereign Cloud based DR) targets, protecting your entire mixed estate with a single CDP strategy instead of maintaining separate approaches for Windows and Linux.

The experience remains unchanged: Linux machines are organized through the same policy-driven Protection Groups and use the same building blocks as on Windows - the same lightweight I/O interception on the source and replication to your chosen target. As on Windows, the filter driver attaches to any block device without requiring a reboot, so replication can begin on a live production system immediately after installation, with full failover, undo failover, permanent failover, and failback to a new location all supported.

Improved long RPO support - CDP delivers the lowest possible RPO, but not every workload needs sub-second protection - and when a longer RPO (such as 30–60 minutes) is configured, you rightly expect it to reduce the load on your infrastructure, not increase it. V13.1 redesigns the CDP proxy cache and traffic queue to deliver exactly that. Overwrites occurring within the same RPO window are now coalesced before transfer, significantly reducing the amount of data sent from the source proxy to the target and shrinking the resulting replica writes. The cache also prioritizes already-confirmed replication traffic over not-yet-confirmed data and gracefully spills to disk during I/O bursts instead of slowing replication down. The net effect is intuitive at last: configuring a longer RPO now means less infrastructure load, exactly as it should.

Universal CDP: feature parity with standard CDP - When agent-based Universal CDP was introduced in V13, it focused on the core near-zero RPO replication workflow. V13.1 closes the gap with classic VMware-based CDP by bringing the day-2 operational capabilities customers rely on in real DR scenarios directly into Universal CDP policies. Veeam Explorers now enable application and item-level recovery from Universal CDP restore points; Re-IP rules automate IP address changes during failover; replica seeding accelerates initial replica creation while cutting the bandwidth required for the first sync; network mapping streamlines DR placement and configuration; and application-aware image processing delivers application-consistent long-term restore points. With the same near-zero RPO engine now backed by the operational flexibility expected from a production DR solution, Universal CDP is ready for real-world disaster recovery.

Unstructured Data Protection

Enhancements to NAS and unstructured data backup in this release are focused on simplifying long-term retention, reducing storage costs, and improving scalability and performance. Organizations can apply Grandfather-Father-Son (GFS) with weekly, monthly and yearly backups to long term retention directly in backup jobs, leverage archive storage in multiple roles (primary, secondary, and archival tier), and scale protection to multi-petabyte environments using object storage. Performance and data consistency are also improved through optimized processing of large file shares and enhanced tracking of file permissions.

Malware Detection - With this release, malware detection capabilities are extended to unstructured data workloads, introducing file-level analytics and Indicators of Compromise (IoC) scanning. File-level detection analyzes suspicious file system activity, including mass file renames and deletions, helping organizations identify potential ransomware attacks and other malicious activity earlier in the incident lifecycle. Combined with IoC-based inspection, this capability provides deeper visibility into the integrity of protected unstructured data. Available with the Advanced Edition of Veeam Universal License, it enhances cyber resilience by helping administrators detect threats before they impact recovery operations.

Threat Hunter support - Veeam Threat Hunter now supports file share and object storage backups, extending signature-based malware detection and YARA rule scanning to unstructured data workloads. Administrators can perform both automated post-backup scans and on-demand inspections to identify known threats and suspicious content within protected data. By combining signature-based analysis with customizable YARA rules, Threat Hunter provides deeper visibility into potential malware infections and helps organizations validate backup integrity before recovery. Available with the Advanced Edition of Veeam Universal License, this capability strengthens cyber resilience for NAS and object storage environments.

GFS retention - Grandfather-Father-Son (GFS) retention can be configured directly within NAS and object storage backup jobs, allowing administrators to define weekly, monthly, and yearly restore points without additional jobs. The implementation is storage-efficient, as GFS points reference existing data from previous backups instead of creating new full copies. This capability simplifies backup configuration, reduces administrative overhead, and enables cost-efficient long-term retention while maintaining compliance with data retention policies.

Archive storage as direct backup target - NAS environments that rely on internal snapshot replication for day-to-day recovery now have a cost-effective path to long-term disaster recovery by storing an additional copy outside of their storage to an archive. In this scenario, rather than landing backups on hot object storage and tiering them out later, Veeam can write unstructured data backups directly to archive-class storage as the primary repository. The workflow itself stays intact: the same source, cache repository, proxy, and processing components remain in place. Only the destination changes. No architectural redesign, no additional infrastructure. Because archive storage is optimized for rarely accessed data, restores are not immediate; data must first be retrieved from the archive tier before it can be restored. Supported archive targets include Veeam Data Cloud Vault Archive, Azure Archive, and AWS S3 Glacier.

Archive storage as a secondary target - With this release, backup jobs can create an additional copy directly in archive-class storage as a secondary target alongside the primary repository. Supported targets include Amazon S3 Glacier, Azure Archive, and Veeam Data Cloud Vault Archive. This capability enables cost-efficient long-term retention and redundancy, allowing organizations to meet compliance and resiliency requirements without maintaining a second copy on expensive hot storage.

Tiered archival - With this release, archive-class object storage can be configured as an archive target. Backup data can be automatically moved to archive storage as it ages out of the primary retention window, or recent restore points can be copied to archive immediately upon creation. Supported targets include Veeam Data Cloud Vault Archive, Azure Archive, and Amazon S3 Glacier. This capability enables automated, cost-efficient long-term retention without additional workflows, helping organizations meet compliance requirements while reducing storage costs.

SOBR object storage extent support - Scale-out backup repositories have long been available for NAS backups; V13.1 now adds object storage as a supported extent type for them. This enables horizontal scaling across large on-premises and cloud-based object storage deployments for Unstructured Data Backup. As your data grows, the SOBR grows with it: add capacity incrementally by extending the repository with additional object storage extents, without redesigning your backup architecture or interrupting running jobs.

Improved single-share backup performance - Source-side reader thread parallelism has been improved, enabling multiple threads to process data from a single share concurrently. This capability increases backup throughput for large file shares, reducing backup windows and improving efficiency when protecting high-capacity NAS workloads.

Incremental Processing of Large Files - Unstructured Data Backup now supports incremental processing of files larger than 32 MB. Instead of reprocessing the entire file when changes occur, Veeam can identify and process only the modified portions, significantly reducing backup windows, network utilization, and repository consumption for large files that change frequently.

ACL changes support - NAS backup tracks and processes now ACL changes, capturing them in incremental backups and reapplying them during restore operations.

Full Web UI coverage - With this release, unstructured data sources, backup jobs for file shares and object storage, along with monitoring and recovery operations, are fully available in the Web UI.

Enterprise Application and Database Protection

Application backup repository

The Veeam Infrastructure Appliance operating in Hardened Repository mode now introduces Application Backup Repository (ABR), a capability that transforms local storage into a managed pool of virtual volumes, each exposed as an individual NFS target. External applications and services write directly to their assigned NFS targets without requiring awareness of the underlying Veeam infrastructure, while Veeam automatically manages snapshots, backup copies, immutability, and access control. This provides a secure-by-design foundation for application-generated backup, export, and dump data.

Application Backup Repository supports a broad range of workloads that can write data to NFS. For structured applications such as Oracle RMAN Incremental Merge, it provides live-access NFS targets with fast snapshot reverts that integrate naturally into database protection workflows requiring frequent recovery points. For less structured environments, including network devices, IoT systems, and standalone databases, Application Backup Repository offers a simple and immutable destination for backup and configuration data that might otherwise reside on unprotected storage. This functionality is available when a Veeam Universal License is installed on the backup server. The following features are available with the initial release:

Full Application Backup Repository management in the Veeam Host Management - Application Backup Repository is managed directly through the Host Management Console with disk provisioning and pool configuration, keeping the configuration footprint within the existing Hardened Repository management experience. The feature is opt-in by design. Application Backup Repository must be explicitly enabled on newly installed Hardened Repositories and is NOT activated on existing deployments through updates. This ensures that no changes are made to production hardened repository infrastructure without deliberate administrative action. The Application Backup Repository wizard allows the definition of virtual volume, NFS export and schedules.

Instant restore - Snapshot-based restore points with configurable retention give Application Backup Repository the ability to revert an entire NFS share to any selected point in time instantly. External backup processes can pick up from an older restore point and continue building new backup chains from there, without manual intervention or data reconstruction.

Instant export - For scenarios requiring access without disruption to the most actual data stored, a temporary snapshot export to an alternative NFS path provides read-only access to the backed-up data. This allows inspection, validation, or cross-system restores to run in parallel against the exported snapshot in a secure and protected way, while the live share continues operating uninterrupted.

Immutable backups - Snapshot lock enables restore points' immutability at the filesystem level.

Backup copy and tape offload - Backups stored in the Application Backup Repository can be further incrementally copied to a standard forever-forward incremental VBK/VIB chain to any generic Veeam Backup & Replication Repository and then offloaded to tape, enabling air-gapped long-term retention for application data.

Veeam Snap Scale Backup for Enterprise Applications

Large-scale database backup - Enterprise databases at the 100+ TB scale demand a fundamentally different approach to backup. This new backup policy is purpose-built for exactly that; delivering massive parallelization across both backup processing and target infrastructure to protect ultra-large databases within an optimized backup window.

At the start of each backup, applications are brought into a consistent state, and storage snapshots are created directly through the Veeam Universal Storage API (USAPI), with broad support for iSCSI and Fibre-Channel enterprise storage systems. From there, Veeam takes one of two paths: snapshots are either maintained in a retention chain through snapshot orchestration, or backups are created directly from the snapshots. During backup processing, snapshots are mounted to multiple Veeam Linux Proxies and read in high parallelization, leveraging Veeam's unstructured data backup engine under the hood. On the target side, Scale-Out Backup Repositories span multiple servers and storage systems to match the throughput capability of the backup source storage, ensuring optimized processing at peak performance of the underlaying hardware capabilities.

Backup efficiency follows the same source-side deduplication principle proven in image-level backup: in this case large database files are compared area by area against the last existing backup, and only changed data is transported to the target; keeping backup sizes lean without sacrificing recovery fidelity.

The first workload to leverage this policy is Epic EHR system protection, covered below.

Enterprise Application Plug-ins Enhancements

GFS retention for managed plug-ins - Long-term GFS (Grandfather-Father-Son) retention with Weekly, Monthly and Yearly flags is now available as part of the backup policy for Oracle RMAN, SAP HANA and Microsoft SQL Server, preserving restore points for the most critical workloads over extended periods. The configuration experience is identical to GFS retention elsewhere in the product. Note that a single backup file can carry multiple GFS flags, which should be considered during capacity planning, and flags are unassigned during backup runs whenever expired flags are detected. As an additional benefit, each GFS restore point is created as a standalone full backup whose .vab backup file is closed immediately upon completion, enabling immediate immutability.

Backup to Tape Support for Enterprise Application Plug-ins - Plug-in backup restore points can now be processed by Backup to Tape jobs, providing an additional layer of protection through long-term, air-gapped storage for organizations that require tape-based retention. Support is available for all remaining enterprise application plug-ins, including Oracle RMAN, Microsoft SQL Server, SAP HANA, SAP on Oracle, SAP MaxDB, and IBM Db2, whether deployed in managed or standalone mode. Tape backup operates on repository contents: full tape backups process the entire repository backup chain as it exists at the time the job runs, while incremental (forever) tape backups process only newly created plug-in backup files. Restore is performed in two stages—first to the backup repository, with automatic backup import, and then through the application itself or a Veeam Explorer. This functionality is available when a Veeam Universal License is installed on the backup server, regardless of the licensing mode used by the protected application workloads.

Plug-in backup properties - You can now review which databases were backed up directly from the backup console, with backup properties implemented for all plug-ins, including the option to view the list of backup files stored on the repository. Depending on the application type, additional details such as backup catalog content may also be available.

Individual Application and Database Enhancements

Epic EHR System Protection

Electronic Health Record (EHR) systems are among the most mission-critical workloads in healthcare environments, requiring reliable and automated data protection. Veeam now delivers dedicated protection for Epic environments running on the InterSystems IRIS Data Platform through the new Veeam Snap Scale Backup for Enterprise Applications backup method. By integrating application-aware orchestration with storage snapshots, this capability simplifies protection workflows while enabling fast, scalable backup and recovery for large healthcare deployments. Available with Advanced Edition and above of Veeam Universal License, it helps healthcare organizations strengthen resilience and ensure the availability of critical patient data. The following features are available:

Protection Group support - A dedicated Protection Group for InterSystems IRIS rolls out the necessary components onto the database servers and rescans the application topology automatically, discovering instances and their associated storage volumes with no manual mapping required.

Backup policy for the InterSystems IRIS databases - Purpose-built for the scale and performance demands of these environments, can operate in two modes: snapshot-only mode, which creates and retains storage snapshots, or backup mode, which captures data into a Veeam repository. In both cases, the policy leverages the Universal Storage API (USAPI) to communicate with the underlying storage system and create the storage snapshot. To guarantee application consistency, the IRIS instance is frozen only for the brief moment required to create the storage snapshot, then unfrozen immediately afterward. In backup mode, the snapshot clone is then mounted to a Linux-based proxy and the data is read by the Unstructured Data Backup engine. Backup efficiency follows the same source-side deduplication principle proven in image-level backup: in this case large database files are compared area by area against the last existing backup, and only changed data is transported to the target; keeping backup sizes lean without sacrificing recovery fidelity.

Short and simple restore wizard - Enables recovery to the original or to a different location, offering a quick restore from a storage snapshot by exporting its clone to the target, or a restore from a backup through physical data transfer with massive parallelization. Please refer to our technical documentation for additional pre-restore and post-restore recommendations to optimize the restore experience depending on your specific scenario.

IBM Db2

IBM Db2 on Windows - A new Enterprise Application plug-in enables native backups of Windows-based Db2 workloads, delivering the same DBA-driven, native backup experience already available on Linux and AIX using standard tools and external schedulers.

Microsoft SQL

Separate SQL account and SA authentication - With Veeam Application-aware image processing it is now possible to specify a separate account per machine rather than a single account per job. It leverages SQL Server authentication (SA) that is not tied to any Windows user accounts. Available for VMware and Hyper-V backup jobs via a new account dropdown and SA authentication checkbox in the processing settings.

Restore Amazon RDS for SQL Server backups to on-premises - Veeam Explorer for Microsoft SQL Server can now restore databases from backups created by Veeam Backup for AWS to on-premises target SQL servers, streamlining "reverse migration" and test-restore scenarios in hybrid environments.

MongoDB

New MongoDB platforms - Protection coverage at the Enterprise Application plug-in is expanded with official support for MongoDB Community Edition and Percona Server for MongoDB.

Oracle

Oracle Incremental Merge support with Veeam Application Backup Repository - One of the most popular DBA strategies which maintains an always-current image copy of the database by rolling forward only the latest incremental changes, minimizing both the backup window and storage usage, is now supported. DBAs can now perform Oracle Incremental Merge on our new Application Backup Repository (ABR), with ABR snapshots additionally enabling "time travel" to previous points in time. In this release, it is a manual process: DBAs specify the ABR as an NFS target in their RMAN script. Recovery options include switching to the image copy, or reverting to a snapshot (which requires a control file restore) and then switching back to production by moving the data files. Note that ABR snapshot creation and backup schedules must not overlap.

Oracle Database 26ai - Oracle 26ai is now fully supported by the Oracle RMAN plug-in and for application-aware processing of virtual and physical machines.

Extended Oracle HA, DR and Engineered Systems coverage - This release adds official support to the Oracle RMAN plug-in for additional Oracle high-availability and engineered platforms. Oracle Fail Safe based on Windows Failover Cluster and clusters based on Pacemaker, Corosync and the Red Hat HA add-on, as well as Oracle SEHA clusters, are now supported in standalone mode, while Oracle Database Appliance is supported in both standalone and managed modes.Oracle RMAN script export - Veeam Explorer for Oracle RMAN now enables export of the RMAN restore script to share with a DBA, making the restore procedure easier for DBA teams. Simply step through the familiar restore wizard in a password-less manner, the wizard performs no validations, and the script is generated at the end.

PostgreSQL

PostgreSQL on Windows - Windows-based PostgreSQL workloads can now be protected with application-aware processing on VMware and Hyper-V. In addition, Veeam Explorer for PostgreSQL now supports Windows-based restore targets. A new SSPI-based authentication mechanism, attempting Kerberos first and then failing over to NTLM, enables processing of Windows guests.

Patroni-managed PostgreSQL cluster support - Veeam now provides application-aware protection for PostgreSQL clusters managed by Patroni on Linux, delivering application-consistent backups with support for point-in-time recovery. Cluster topology is discovered and maintained automatically, with Veeam rescanning the environment during each backup to keep configuration information up to date while continuously processing WAL data from the primary node. Recovered databases are immediately accessible through Veeam Explorer for PostgreSQL, providing the same granular recovery capabilities available for standalone PostgreSQL deployments. Available with paid Veeam Universal License editions (Essentials, Advanced, and Premium), this capability simplifies protection and recovery of highly available PostgreSQL environments while maintaining consistent administrator experience.

Veeam Explorer for PostgreSQL individual database restore - In response to numerous customer requests, Veeam Explorer for PostgreSQL now supports selecting one or more individual databases and restoring them to the original or a different location, extending the scope of supported restore scenarios. This also maximizes the value of Patroni support, where granular recovery is often preferred over restoring an entire instance.. A new wizard step is now included in Veeam Explorer for PostgreSQL to authenticate to the target PostgreSQL instance. This functionality is supported for backups of Linux-based workloads only.

SAP HANA

Custom SAP HANA backup file prefixes - A new option in the policy settings lets you define custom prefixes for SAP HANA backup files, making backups easy to identify, sort and align with internal naming conventions, audits and multi-tenant environments.

Backup Infrastructure

Malware Detection

Entropy event investigation - Previously, investigating entropy-based malware detection events required switching to external tools, limiting visibility and slowing down response workflows. With this release, entropy detection events include built-in investigation context, allowing administrators to identify the specific files that triggered elevated entropy levels during block level scans. This capability improves visibility into suspicious activity, enabling faster and more accurate investigation directly within the backup environment.

Improved index analytics for failover cluster backup - File deletion tracking is now performed per volume instead of per backup job to significantly reduce false-positive malware events for clustered workloads.

Per-machine exclusions - With this release, administrators can define exclusions at the individual workload level. This capability allows precise tuning of malware detection, reducing false positives for specific systems while maintaining full protection across the rest of the environment.

Full Web UI coverage - With this release, the full malware detection control plane is available in the Web UI, including event management, granular exclusions, and end-to-end investigation workflows. This capability allows administrators to review detections, tune policies, and take action from a single interface, improving efficiency and simplifying security operations.

Performance improvements - With this release, we've decreased the amount of RAM consumed during guest index analysis up to 90% in some particular scenarios.

Security

Hybrid FIPS + Post-Quantum Cryptography (PQC) processing - With this release, Veeam Backup & Replication introduces a hybrid FIPS + PQC model that uses post-quantum algorithms (aligned with NIST FIPS 203, 204, and 205) for handshake and key exchange, while retaining FIPS-certified AES for data encryption. PQC support is enabled on existing deployments without requiring data migration or infrastructure changes. When strict FIPS compliance mode is enabled, PQC functionality is automatically disabled to ensure full adherence to FIPS requirements.

MFA enforcement for Users and Roles changes - Any change to Users and Roles settings now requires MFA confirmation (when enabled), adding a checkpoint to the most sensitive configuration area for your backup server.

Windows Event Log Integration for Security & Compliance - Security & Compliance Analyzer vfindings and status changes can now be published directly to the Windows Event Log. This enables integration with SIEM, Syslog, and other security monitoring platforms, allowing organizations to incorporate backup infrastructure compliance events into existing security operations and monitoring workflows.

SAML Backup server certificate - The SAML configuration page now uses the backup server certificate as the default Service Provider certificate, removing a manual selection step when configuring SAML single sign-on.

Exportable SAML certificate - Service Provider certificate can now be exported via the Download action in SAML configuration, simplifying IdP metadata exchange.

Encryption password change notifications - Changing backup encryption password now records a notification in the affected job's session log and generates an event, providing an audit trail for key rotation.

Configurable Web UI port for Windows Deployments - With this release the software allows administrators to configure a custom port for web services instead predefined of 443 port during backup server deployment or modify it later as needed.

Single-port transport - Previously, backup communication relied on a wide dynamic port range (2500–3300), increasing firewall complexity and expanding the potential attack surface in segmented environments. With this release, all backup communication uses a single port, eliminating the need for dynamic port ranges. This capability simplifies firewall configuration, reduces the attack surface, and enables easier deployment in secure and segmented network environments.

REST API Port - The REST API no longer requires a dedicated port and answers on port 443. Port 9419 remains open for backward compatibility so existing scripts aren't broken by this upgrade, and will be deprecated in a future release.

RBAC restore options refinement - Restore options for custom operator users now include grouping, descriptions, and search, making the right action easier to find.

.NET 10 adoption - The backup server components are migrated to .NET 10, delivering security patches, performance improvements, and continued long-term support alignment.

Veeam Intelligence

Veeam Intelligence - With V13.1, Veeam Intelligence evolves from an AI assistant that answers questions into an intelligent operational companion that can investigate real backup data and help perform actions on your behalf, while keeping administrators in full control. New AI-powered capabilities streamline troubleshooting, automate routine operations, improve support interactions, and provide deeper insight into backup environments.

Experienced Backup Admin Agent - provides expert-level troubleshooting directly within your environment. Administrators can investigate jobs using natural language or launch analysis directly from failed sessions in the Web UI. The agent automatically analyzes job and session logs and works through issues using the same information typically reviewed by Veeam Support engineers.

Support Case Integration - extends this capability to Veeam Support interactions. The agent can review existing support cases, add comments, and proactively offer to attach relevant diagnostic data, reducing the time required to collect and submit troubleshooting information.

Actions for Veeam Backup & Replication - enable Veeam Intelligence to execute approved operational tasks. Supported actions include job control, infrastructure maintenance operations, component updates, repository management, configuration backup execution, license management, malware scanning, and selected recovery operations. All actions remain subject to workflow-based approval controls to ensure administrators retain oversight of critical operations.

Capacity Planning - adds predictive storage analysis capabilities to the Experienced Backup Admin agent. Administrators can quickly understand current storage consumption trends, forecast future capacity requirements, and identify abnormal growth patterns without creating custom reports.

Additional enhancements - include multi-threaded conversations for parallel investigations, suggestion chips that accelerate common follow-up actions, an improved Data Resilience Daily Summary (Morning Coffee Report) with enhanced retrieval and feedback mechanisms, significantly more accurate and reliable responses, and expanded accessibility through French and German localization as well as light and dark themes.

These capabilities are available with paid Veeam Universal License editions (Essentials, Advanced, and Premium) and represent the first step toward a broader set of AI-driven operational and recovery workflows in Veeam Backup & Replication and the Veeam Data Platform.

Veeam Software Appliance and Veeam Infrastructure Appliance

Veeam continues to invest across the Veeam Software Appliance (VSA) and Veeam Infrastructure Appliance (VIA) platform, delivering targeted improvements in operational efficiency, security hardening, and deployment flexibility with each release.

General Improvements

Drive expansion - Administrators now can add additional drives to an existing Veeam Software Appliance deployment after installation to extend available storage capacity. This can be used for multiple purposes, including capacity extension within existing Scale-out Backup Repositories or creating new Repositories.

FC and iSCSI storage support - Veeam appliances now support attaching Fiber Channel (FC) and iSCSI storage devices, along with multipathing, enabling deployment in enterprise SAN environments.

Custom log path configuration - A custom file system path now can be set for Veeam log files using Host Management Console to redirect logs to a dedicated mount point or storage volume on the Logs and Services page.

Extended unattended deployment control - Two additional flags are now available: DataCollection to allow automatic joining to Veeam Service Provider Console and Veeam ONE and High Availability to skip high availability-related steps during automated deployment.

Upgrade to Veeam JeOS 9.6 - All Veeam Appliances are now upgraded to Veeam Enterprise Linux JeOS 9.6, extending hardware compatibility and incorporating the latest kernel improvements and platform updates.

Updates Mirror via Enterprise Manager - Enterprise Manager can now act as a centralized update mirror, caching product updates from the official Veeam repository and distributing them to managed Veeam Infrastructure Appliance (VIA) and Veeam Software Appliance (VSA) deployments. This capability is particularly valuable in environments with restricted, proxied, or tightly controlled internet connectivity, where direct access to external update services is limited. By centralizing update distribution and reducing external connectivity requirements, organizations can simplify update management while maintaining greater control over how updates are delivered across their backup infrastructure.

Observability

Prometheus-compatible metrics - Veeam appliances now expose a broad set of system metrics through Node Exporter, enabling seamless integration with Prometheus-compatible monitoring platforms and observability stacks. This allows organizations to collect, visualize, and analyze appliance health and performance data using their existing monitoring tools and processes. This capability is available when a Veeam Universal License is installed on the backup server.

Syslog support for appliance events - Veeam appliances now can forward system logs and operational events to an external syslog server for centralized log aggregation, SIEM integration and compliance-driven audit trails without additional agents.

Resource consumption dashboard - Host Management Console now includes a live resource monitoring view that shows CPU, memory and disk utilization. This helps administrators quickly identify resource bottlenecks and correlate workload spikes with infrastructure load.

Email notifications for storage load - Veeam Backup & Replication now sends an email notification whenever appliance storage utilization exceeds configured thresholds, providing early warning before capacity becomes critical.

User management and MFA

Streamlined MFA configuration - Multi-factor authentication (MFA) initialization is now available in the Web UI in addition to text user interface (TUI).

Unified user management - User creation includes optional assignment of Veeam Backup & Replication roles during onboarding. MFA configuration is also unified, with settings defined in Host Management automatically applied to Veeam Backup & Replication access. This capability simplifies user management, reduces administrative overhead, and ensures consistent access control and security policies across the environment.

Optional MFA at deployment - Multi-factor authentication (MFA) was enforced during deployment, making it difficult to automate installations or deploy in environments without access to authenticator apps. With this release, Veeam Software Appliance (VSA) and Veeam Infrastructure Appliance (VIA) can be installed without enforcing MFA for the veeamadmin and veeamso accounts. MFA can be enabled later from the Host Management Console.

Longer MFA secret length support - MFA now supports 160-bit secrets to improve compatibility with modern authenticator applications and enterprise MFA solutions, helping organizations align with current security standards.

Localization

Host Management localization - The Host Management Console adds support for German, Japanese, and French, allowing each user to select their preferred display language independently.

Keyboard layout selection - Terminal User Interface initialization wizard and menu now include keyboard layout selection.

Security and compliance

Improved certificate management - It's now possible to import workload certificates into the system certificate store and apply custom TLS certificates to the Host Management Console Web UI without manual certificate operations at the OS level.

Password expiration configuration - Administrators can define password expiration policies directly in Host Management to enforce password lifecycle requirements, improving security posture and ensuring compliance with internal and regulatory standards.

Password reset restriction - Administrators can disable password reset through the Host Management Console, limiting password reset operations to the Host Manager Terminal User Interface only.

Control over FIPS-compliance - Host Management Console introduces a strict, global FIPS compliance mode that enforces consistent cryptographic standards across both system-level and application-level components to simplify compliance management for regulated environments, ensure consistent security configuration while reducing administrative overhead and configuration errors.

Reduced permissions on log folders - Log directories and files now use stricter permissions, reducing exposure of sensitive log data.

Veeam Software Appliance

Extended HotAdd support - VSA now supports HotAdd transport mode for VMware vSphere backups in all-in-one deployments. This capability enables higher backup performance and removes previous limitations, allowing organizations to deploy Veeam Software Appliance in VMware Cloud environments while achieving efficient data transfer.

Hyper-V VM image - A ready-to-use Hyper-V virtual machine image is available for Veeam Software Appliance.

PostgreSQL security hardening - The internal PostgreSQL database on Veeam Software Appliance is hardened following industry-standard database security best practices, with a focus on audit logging and log-file security.

Automatic database disk quota expansion - Veeam Software Appliance now automatically expands the database quota when its free space becomes low and alerts administrators if the underlying storage volume is running out of space.

High Availability

Support for cross-subnet HA clusters - Designed for real-world enterprise architectures, v13.1 enhances High Availability with support for cross-subnet HA clusters, enabling HA nodes to be deployed across different networks or sites - meeting common enterprise segmentation requirements without added deployment complexity.

Flexible Deployment - High Availability now supports flexible deployment approaches: a standard cluster for straightforward environments, or a cross-subnet cluster for larger organizations that want to place nodes in a geographically separate data center or a separate network segment - while still presenting a single HA cluster identity.

Synchronization Experience Enhanced - To make High Availability more transparent and resilient, we've improved the day-to-day maintenance experience across HA nodes by reducing manual coordination and keeping key configuration elements aligned automatically.

Users synchronization - User synchronization is now automatic, keeping the secondary aligned with the primary and avoiding the inconsistencies that can occur with manual user management.

Private fix synchronization - Private fixes (hotfixes delivered outside the standard update channel) are now automatically synchronized between HA cluster nodes, eliminating manual application to the secondary node.

REST API coverage - makes HA management easier to automate and operate at scale, with support for key HA lifecycle actions.

Data lag observability - See exactly how much data lag exists before a switchover or failover, so administrators can make an informed decision before committing to the operation.

Version observability - See when a secondary node is in the process of being upgraded, so administrators have full visibility into replica health before relying on it for failover. Overall, more granular statistics about HA are now collected.

Modify cluster settings - Existing High Availability clusters can now be edited, allowing changes to cluster endpoint parameters: the virtual IP address and cluster DNS name for a regular cluster, or both external IP addresses and the cluster DNS name for a cross-subnet cluster.

Veeam Infrastructure Appliance

Single-disk deployment option - Veeam Infrastructure Appliance can be deployed on a single-disk server without a dedicated backup storage disk. This capability supports use cases where Veeam Infrastructure Appliance operates as a backup proxy with remote storage or when Application Backup Repository should be used, simplifying deployment and allowing more flexible use of available infrastructure.

Unified boot options - Veeam Infrastructure Appliance now uses a unified boot menu for all roles, including Infrastructure Server, iSCSI & NVMe/TCP, and Hardened Repository. The appliance role is selected later during the initialization wizard.

Expanded repository path support - Backup repositories can be created at or under /var/lib/veeamdata/backup on Veeam Infrastructure Appliance deployments for greater flexibility in storage design, supporting single-disk configurations and simplifying deployments where additional storage devices are not available.

Application Backup Repository deployment - Veeam Infrastructure Appliance in Hardened Repository role can now act as an Application Backup Repository (ABR). Storage Pool management for ABR is available directly from Host Management Console. For more details, see the Application backup repository section of this document.

Direct SAN mode support - Veeam Infrastructure Appliance now supports Direct SAN access over Fibre Channel and iSCSI, letting VIA read virtual machine data directly from SAN storage for faster, production LAN-free backups.

Backup from Storage Snapshots over NVMe/FC - BoSS jobs can now retrieve snapshot data over NVMe/FC, giving VIA a lower-latency path to the storage array and reducing backup windows for NVMe-based arrays.

Backup Repositories

Linux based Hardened Repository

In addition the strict Veeam Hardened Repository Appliance, we are opening the same concept for less strict Linux Repository environments that want to benefit from Immutability while keeping the flexibility of deleting backup at the system administrator level with Immutability in governance-mode.

Governance-mode immutability - Standard Linux repositories now support governance-mode immutability using the same underlying immutability engine as we use with the Hardened Repository appliance to enable flexible, policy-driven data protection, allowing organizations to benefit from immutability while simplifying deployment and allowing additional roles which are not allowed on Hardened Repository.

Authorization for immutability changes - Reducing or disabling the immutability period on Hardened Repositories and governance-mode Linux repositories now requires approval from a second authorized user whenever four-eyes approval is enabled.

Configuration backup immutability - Veeam Backup & Replication configuration backups can be stored immutably on Hardened Repositories and governance-mode Linux repositories.

Object Storage

Object storage repository read-only mode - V13.1 adds an "Enable read-only access" option when connecting to an immutable object storage repository, letting a second backup server (such as a DR-site instance) attach to a repository already owned by your production server. In read-only mode the second server never writes to or modifies the data and performs restore operations only, so you can run restores, recoverability checks and Data Integration API based forensics/ransomware scans from a separate location without taking ownership away from production or disrupting running jobs. This turns regular restore testing, previously impractical because of the coordination overhead of transferring ownership, into a native, supported capability rather than a workaround.

Veeam Data Cloud Vault

Veeam Data Cloud Vault guided AWS onboarding - Azure-based Veeam Data Cloud Vault has always offered a dedicated, guided onboarding experience directly in the Veeam Backup & Replication console. For customers with specific use cases for AWS based processing, V13.1 brings the same experience to AWS-based Veeam Data Cloud Vault: instead of adding it as a generic Amazon S3 object storage repository and pasting shared keys, you simply select your Vault in the unified wizard and the repository is provisioned automatically, with immutability enabled by default.

Setup also moves away from long-lived shared keys to federated, role-based access, the same secure approach used for Azure Vault, so there are no static credentials to manage, rotate, or risk leaking. For customers already running Azure Vault, adding AWS is now an identical workflow, one wizard, both clouds covered.

Veeam Data Cloud Vault Archive - Veeam extends the fully managed Vault experience to archive-class cloud storage. With Azure Blob Archive tier, we are delivering the lowest-cost option for long-term retention without requiring you to manage your own cloud storage infrastructure. Vault Archive can serve as a primary backup target for unstructured data (NAS) backups, or as the Archive Tier of a scale-out backup repository, automatically moving or copying GFS restore points from Vault to Vault Archive.

Unlike the standard Archive Tier, Vault Archive requires no archiver appliance. Backup data is consolidated natively in the cloud, so there is no extra compute to deploy, manage, or pay for. As with any archive-class storage, restores require a data retrieval step and are therefore not instant, making Vault Archive ideal for data retained for years and accessed very rarely; where lowest cost matters more than immediate availability.

Veeam Data Cloud Vault built-in cost guardrails - Veeam Data Cloud Vault is built around all-inclusive, predictable pricing, and we now introduce edition-aware guardrails that keep usage aligned with that model, so you are never surprised by cloud charges that fall outside of it. For Vault instances on the Foundation edition, the product now automatically prevents configurations that would generate unplanned egress or retrieval costs, keeping all scale-out repository tiers within the same cloud region, blocking offload or archiving to non-Vault storage that would incur egress, and disabling the higher-cost priority (Azure) retrieval options for Vault Archive. To honor the cloud providers' minimum storage commitment periods, immutability is always on and enforced for at least 30 days on Vault, and 180 days on Vault Archive. Vault Advanced is unaffected, as its unlimited-egress commercial model carries no such restrictions.

Scale-Out Backup Repository

Archive Tier copy policy - This new Copy policy for the scale-out backup repository Archive Tier was added in addition to the existing Move policy. Instead of waiting for GFS restore points to age out of the archive window before they move to the Archive Tier, the Copy Policy archives the GFS points upon creation while keeping the local copy on the Performance Tier until it naturally expires. Copy and Move policies can also be combined to apply different retention per tier. For example, shorter on the Performance Tier and longer on the Archive Tier which enables to keep a local copy for fast restores during the archive window and an archive copy available for long-term retention, compliance, or ransomware resilience. The Copy policy applies to Direct Archival configurations only, a scale-out backup repository with a Performance Tier and an Archive Tier. Configurations that include a Capacity Tier (Capacity Tier → Archive Tier) do not support Copy mode.

Actual backup size reporting - Shows how much space your backups truly occupy on object storage, not just their logical size before optimizations such as block reuse. Both values are displayed side by side wherever you review backups, "Backup size" for the logical size and "Actual size" for the space consumed after optimizations. Wherever a backup resides, on a standalone repository or any tier of a scale-out backup repository, you can see the actual space its chain segment occupies on that specific tier, or the total consumed across all tiers at once. This takes the guesswork out of capacity planning across tiered storage.

Archive Tier minimum immutability period - Archive repositories can now be configured with a minimum immutability period. A fixed number of days during which data cannot be deleted or modified can be set as an alternative to keeping backups immutable for their entire retention period. . A minimum immutability window gives you protection where it matters most: the critical period right after a backup is created is protected from alteration or attacks, while leaving the flexibility to manage or delete archive data once that window has passed.

Storage Integrations

Primary Storage Integrations

Fleet management solutions support - Adds native support for storage fleet based single management systems that manage multiple storage systems holistically. In a single wizard the fleet management system can be added once and all underlying storage system hierarchies with their volumes and snapshots are discovered automatically for Veeam's storage-based processing. Functionality available via Universal Storage API (USAPI) allows vendors providing fleet capabilities to implement it in their plug-ins.

Independent VM processing in Backup from Storage Snapshot jobs - VMs are now automatically grouped by datastore, volume, or replication group, with each group processed in parallel instead of sequentially, allowing large jobs with 1,000+ VMs to cut preparation stage by 2x and more and reduce individual VM snapshot hold time by 10x and more.

Secondary storage integration

Dell Data Domain Managed File Replication (MFR) support - In addition to the existing HPE Catalyst Copy support we are delivering the same Veeam Backup Copy Job concept to DDBoost based deduplication storage, enabling efficient server-side data movement without intermediate staging through backup gateway server for processing.

HPE Cloud Bank direct backup for more workloads - Direct to HPE Cloud Bank backup and backup copy jobs are now supported for agent-based, enterprise plug-ins and unstructured data, not just VM backup.

Multiple backup repositories in a single HPE StoreOnce Catalyst Store - Configuring multiple backup repositories within a single Catalyst Store (using subfolders to maximize deduplication efficiency) previously required PowerShell as documented in KB2987. This version adds full UI support: repositories sharing a single Catalyst Store can now be created, edited, and rescanned directly from the backup console without scripting.

Configuration backup immutability on HPE Catalyst (StoreOnce) and Dell DDBoost (Data Domain) - Backup server configuration backups (.bco) can now be stored with immutability enabled on Dell Data Domain DDBoost and HPE StoreOnce Catalyst, protecting the backup server configuration from ransomware or accidental deletion.

Updated firmware and SDK support - This version adds support for HPE StoreOnce OS 5.2.0 with Catalyst Client 5.2.0 and Dell Data Domain OS 8.7 with DDBoost SDK 8.7.

Accelerated Restore from Dell Data Domain performance improvement - The default number of DDBoost connections used during full VM restore has been reduced to 1, improving restore throughput by up to 30% by letting Data Domain optimize read ahead caching.

Automatic transport TCP port assignment for deduplication repositories - ExaGrid, Quantum DXi, Infinidat Infiniguard, and Fujitsu Eternus CS800 deduplication repositories now support automatic transport TCP port assignment, removing the need for manual port configuration when adding these repositories to the backup infrastructure.

Immutability for NFS/SMB Repositories - Starting from this release, you can enable immutability for NFS and SMB backup repositories backed by WORM-capable storage solutions, such as NetApp SnapLock. To signal the underlying storage to lock the files, Veeam Backup & Replication sets the read-only file attribute and updates the last access date on backup files, which the WORM solution uses to commit the files to an immutable state. The immutability settings for these repositories are configured in the same way as for other supported repository types. Default share and volume based global WORM settings are not supported and should be disabled in order to let Veeam fully manage the immutability settings.

Tape

NetApp SMTape (SnapMirror to Tape) Support - Veeam Backup & Replication now supports NetApp SMTape, enabling block-level backups of NAS volumes directly to tape using storage-side snapshots. Designed for large-scale NetApp environments that rely on onboard snapshot retention for operational recovery, SMTape provides an efficient way to maintain an independent disaster recovery copy on tape for long-term retention and compliance purposes. Available in Advanced Edition and above and licensed using the same 500 GB per-instance model as Unstructured Data Backup, this capability delivers high-performance, cost-efficient archival for large NAS deployments. To ensure recoverability, restore operations remain available regardless of license status.

IBM LTO-10 Read Past Read Permanent Error handling - Veeam Backup & Replication now supports handling IBM LTO-10 Read Past Read Permanent Errors, allowing tape operations to continue past damaged sectors and improving data recovery from partially corrupted media.

Enterprise Manager

Search history extended to 10 years - The maximum search history window in Enterprise Manager has been increased to 10 years to support long-term compliance and discovery workflows.

Search now includes folders - Enterprise Manager search now returns repository folders alongside backup objects.

SureBackup jobs listed in grids - SureBackup content scan jobs now appear in Enterprise Manager job grids, giving unified visibility across all verification job types.

Default Linux mount server - Enterprise Manager now honors the default Linux mount server set in backup server for Linux file-level restores.

Veeam Cloud Connect

Restore encrypted tenant backups on the service provider side - Service providers can now restore encrypted tenant backups directly on the Veeam Cloud Connect server, for tenants that opt-in and set the scope the provider can restore. This enables managed recovery services for tenants and provides an additional disaster recovery option.

Additional restore options on the service provider side - In addition to already existing capabilities, several new restore capabilities are now available for service providers directly on the Veeam Cloud Connect server: entire VM restore to Nutanix AHV, Proxmox VE, oVirt KVM, Scale Computing HyperCore, and HPE Morpheus VM Essentials, plus Instant VM Recovery to Microsoft Hyper-V.

Malware scan of tenant backups - Service providers can run on-demand Threat Hunter malware scans and YARA rule scanning against tenant backups hosted on the Veeam Cloud Connect server, pinpointing the last clean restore point and enabling pre-recovery malware verification without requiring tenant compute resources.

Universal CDP for Linux - Enables service providers to offer Disaster Recovery as a Service (DRaaS) to tenants protecting Linux VMs and physical servers, replicating to Cloud Connect targets backed by VMware Cloud Director or VMware vSphere hosts.

RBAC support for Cloud Connect backups - Custom roles on the tenant Veeam Backup & Replication server now extend their scope to Cloud Connect-hosted backups, so a scoped-role user can restore from backups stored on the service provider as well, giving tenants consistent role-based access across all their backups.

HPE StoreOnce - Can now be used as Veeam Cloud Connect repository, enabling an additional repository type for an off-site offering for tenants.

Web UI

Instant Recovery from any platform - Instant Recovery can now restore any image-level backup from any hypervisor through Web UI, allowing workloads to be recovered directly to VMware vSphere, Hyper-V, or Nutanix AHV targets, streamlining recovery through a single, unified interface.

UI personalization & localization - Web UI introduces a refreshed default theme with per-user customization and expanded language support (Japanese, German, French), with preferences applied consistently across sessions, login, and related components.

Veeam ONE Alarm Overview integration - The dashboard now displays Veeam ONE alarms in a dedicated panel, providing immediate visibility into backup health and alerts without requiring a separate console.

Email Settings - Mail server and email notification settings can now be configured directly in the Web UI, eliminating the need to switch to the Remote Console for setup or updates.

Backup Copy Job Detail View - The updated Backup Copy job details view in the Web UI consolidates status, restore points, and diagnostic insights into a single interface, enabling faster troubleshooting with visibility from high-level status to individual workers and objects.

Application item-level recovery - For Microsoft Active Directory and Microsoft SQL Server are now available directly in the Web UI.

Advanced RBAC - Full role-based access control is now available in the Web UI, including custom role assignment, scope-based visibility, and centralized role management, ensuring users can access and manage only the resources within their assigned scope.

General Improvements

Veeam Updater on Windows - Windows-based Veeam Backup & Replication deployments, enabling administrators to automatically download and silently install product updates without manually obtaining ISO images or running upgrade wizards. By streamlining the update process and reducing operational overhead, this capability helps organizations keep their backup infrastructure current with less administrative effort and improved consistency across deployments.

Automatic Component Upgrade Distribution - In addition to the existing manual process, a Veeam Backup & Replication upgrade can now automatically update remote backup infrastructure components. By automating the distribution and deployment of component updates across the backup infrastructure, this capability reduces administrative effort, accelerates upgrade rollouts, and helps maintain version consistency throughout the environment. It is particularly beneficial for large-scale deployments with numerous managed servers, proxies, and repositories, where manual upgrade operations can be time-consuming and operationally complex.

Move job-to-job during job operation - "Move backup" operations no longer disable the entire primary source job when run on individual child objects for VMware and Hyper-V workloads, improving flexibility when moving large backups.

Container and tag-based exclusions - Global exclusions configuration now supports container-level and tag-based exclusion rules, enabling policy-driven exclusion management in dynamic virtualized environments.

Multi-string description in Credentials Manager - Credential entries in Credentials Manager now support multi-line descriptions, improving documentation in large credential inventories.

Remote Console performance improvements - Windows-based remote console connection and rendering performance have been improved for large-scale environments.

VeeamZIP user experience enhancements - VeeamZIP now remembers the most recently selected backup repository and pre-selects it as the default target for subsequent VeeamZIP operations, eliminating repetitive repository selection for ad-hoc, on-demand backups.

Empty Job Creation Support - Backup jobs can now be created from preconfigured templates in the Web UI without selecting a workload upfront, allowing administrators to define job settings first and add protected objects later.

Job & Backup ownership edit - Jobs and backups now include a Permissions window for setting object ownership and defining which roles can access them, giving administrators role-based control to delegate ownership and manage access with precision.

REST API

V13.1 continues progress toward full REST API coverage with additions across RBAC, malware detection, High Availability, deduplication repositories, and backup operations.

RBAC - REST API now covers complete Users and Roles management including creating and modifying custom roles and their ACL-based scope definitions.

Malware Detection - Malware detection settings, event management, scan triggering, and investigation actions are now fully accessible via REST API.

High Availability - Cluster management, failover, and switchover are now available via REST API for automation of HA workflows.

Deduplication repository management - Deduplication appliance repositories (Dell Data Domain, HPE StoreOnce, ExaGrid, Quantum DXi, Infinidat Infiniguard and Fujitsu CS800) can now be added, configured, and managed via REST API.

Instant Recovery to Azure - Helper appliance templates that assist with facilitating instant recoveries to Azure can now be fully managed via REST API.

Quick Backup for Hyper-V and Agents - Ad-hoc quick backup sessions can be triggered via REST API for Hyper-V and Agent managed machines, enabling post-malware-event backup automation.

Any-to-vSphere and Any-to-Hyper-V Instant Recovery - Scripted Instant Recovery from any image-level backup source to VMware vSphere or Hyper-V can now be initiated and managed via REST API.

Proxmox VE and Nutanix AHV - New methods to get Proxmox and Nutanix backups, restore points and sessions along with querying infrastructure and starting / stopping jobs are now supported.

PowerShell

Epic EHR backup & restore - This release adds PowerShell support for Epic EHR protection on InterSystems IRIS, enabling scripting and automation of deploying the required components, automatically discovering IRIS instances/volumes, and configuring the dedicated IRIS backup policy. With PowerShell now available, you can also automate restores (to the original or custom location) from either storage snapshots or repository backups.

GFS for Application Backups - Now you can script and manage GFS long-term retention for SQL Server, Oracle RMAN, and SAP HANA backups with new dedicated GFS settings in PowerShell.

Veeam Data Cloud Vault Archive - Configure new Veeam Data Cloud Vault Archive repository via PowerShell and use it in your workflows to backup and restore from.

Application-aware processing for Universal CDP - Automate configuring application-aware processing settings for Universal CDP policies in PowerShell to allow restoring application items from replica's application-consistent restore points.

Malware detection analysis - Start malware detection analysis with a dedicated PowerShell cmdlet specifying event id and target folder path to save analysis results to.

Infrastructure Metrics - Manage external monitoring capabilities with a new set of PowerShell cmdlets that allow you to configure both Node Exporter and Syslog server settings.

SMTape integration - Adds end-to-end management of NetApp NDMP with SMTape for file-to-tape workflows, including enabling NDMP on the NetApp side, registering and administering NDMP servers in Veeam Backup & Replication tape infrastructure, and selecting NDMP volumes for backup. It will also allow retrieving NetApp NDMP volume restore points from tape and initiating volume restores.

Manage agent backups ownership - Use PowerShell to configure ownership of managed by Agent backups when a user's laptop is replaced, so the new device can be granted access to backups created by the old device. This includes listing which backups are associated with a given agent and attaching/detaching same-platform backups.

Capacity tier support in evacuation cmdlet - Extends the existing evacuation cmdlet workflow so it can be started against capacity tier extents as well, enabling support to script repeatable evacuations when large moves fail due to timeouts or transient infrastructure issues. The approach is to unify performance and capacity extents under a common interface, so the same cmdlet can target either tier without requiring a separate parameter set.

Custom Role Assignment - Administrators can now assign custom roles to users for the remote console and WebUI directly through PowerShell, closing a gap from previous releases.