

Veeam Backup & Replication

What's new in v10?

Veeam® Backup & Replication™ 10, part of Veeam Availability Suite™, delivers seriously powerful, modern data protection for NAS, ransomware prevention and advanced data portability. The following is a list of the major new features and enhancements added in v10.

Seriously Powerful NAS Backup

Modernize and simplify the protection of file shares with the powerful new capability of Network Attached Storage (NAS) data protection and recovery. Delivering a storage-agnostic architecture to back up your unstructured data at scale, Veeam helps you to achieve your NAS recovery objectives without additional hardware investments. There's no need to acquire and lock into expensive, specialized backup appliances, as Veeam enables you to use the backup storage you already have, and even offload backups to the cloud at no additional cost.

And when it comes to recovery, Veeam ensures your recovery time objectives (RTO) are met by providing flexible restore options tailored to typical disaster scenarios including: entire share recoveries to address complete hardware loss or NAS migration scenarios, point-in-time share state rollback for recoveries from ransomware attack, and file-level recovery with global search for day-to-day operational restores.

The unique benefits of Veeam's NAS backup engine include:

- **Scalable architecture** based on a proprietary distributed file system, designed and built from the ground up for protection of billions of files and PBs of unstructured data. The backup storage engine performance and capacity can be scaled seamlessly by leveraging additional compute and additional storage resources from scale-out backup repository extents based on commodity server hardware.
- **Extensive data source support** includes SMB file shares (any SMB protocol version), NFS file shares (NFS protocol versions 3.0 and 4.1), as well as files on Microsoft Windows or Linux based servers.
- **Efficient forever-incremental backup engine** that does not require periodic active full backups to be performed, as is the case with NDMP backups – making it possible to protect PB-size NAS environments with significantly reduced RPOs.
- **Innovative storage-agnostic changed file tracking** provides industry-leading incremental backup performance, allowing to achieve low RPOs on NAS devices without native changed file tracking capabilities.
- **Storage snapshots support** enables protection of locked files, backup of atomic file share states, and allows to implement zero impact NAS protection by retrieving file data from storage snapshots replicated to secondary storage arrays.

V10 – Better Backup! Faster. Stronger. Smarter.

With Veeam Backup & Replication™ v10, Veeam further extends its leadership in comprehensive data management by delivering powerful data protection across ALL your cloud, virtual, and physical workloads with a simple, flexible and reliable solution. Protect your data no matter the location or workload type. The latest released is packed with over 150 enhancements, including:

- Modernize and simplify the protection of data with seriously powerful NAS backup
- Recover faster than ever with our NEXT GENERATION Instant VM Recovery® engine
- Simplify off-site backup and get 100% bulletproof protection from ransomware and insider threat with immutable backups
- Put backups to work with greater platform extensibility with the NEW Veeam Data Integration API
- Broader platform and ecosystem support including new, advanced capabilities for Linux, Nutanix AHV, PostgreSQL, MySQL and more.

Veeam Backup & Replication works even better when it's combined with the powerful monitoring capabilities of Veeam ONE™ in Veeam Availability Suite v10 for the next generation of backup for what's next, including data protection capabilities that increase availability, portability, and extensibility.

Supported Environments

For a detailed list of supported environments, click [here](#) for the product Release Notes.

- **Simple offsite backup** with built-in ability to mirror primary NAS backup contents to other backup repositories with the same or different retention policy.
- **Object storage integration** for cost-effective, reliable archival of older file versions to Amazon S3, Azure Blob Storage, IBM Cloud Object Storage, as well as numerous S3-compatible service providers and on-premises storage solutions.

And it goes without saying that NAS backup jobs also include all the standard Veeam features including flexible scoping, integration with Microsoft VSS for consistent backups, data encryption, pre- and post-job scripting for automation, and multiple notification options.

Access to NAS backup requires the **Veeam Universal License**. When used alongside with a socket license, file version history archiving requires **Enterprise Plus** edition.

Ransomware Protection and Simplified Data Management

Protecting your backups against being encrypted by ransomware, or deleted by malware, hackers or malicious insiders is no longer a nice to have feature, but an essential part of any data protection strategy. While many Veeam customers have great success with leveraging tape to meet these requirements, other kept asking us for a simpler, disk-based solution. This is why we delivered 100% bulletproof protection for your backup copies in v10, with the level of backup isolation identical to the one provided by offline tapes!

The new **immutable backups** feature protects your recent backup copies from modification and deletions by leveraging advanced object storage functionality called S3 Object Lock. You can create immutable backup copies on Amazon S3, as well as any S3-compatible service provider or on-premises object storage solution which implements S3 Object Lock API. Uniquely, by leveraging object-level locking, we were able to preserve the forever-incremental object storage offload engine, thus avoiding periodic full backups requirement imposed by bucket-level retention lock. The immutable backup copies functionality is an optional extension of the new Capacity Tier data management policy.

The new Capacity Tier **Copy policy** automatically duplicates backups to object storage as soon as they are created in the Performance Tier of the Scale-out Backup Repository (SOBR). This new policy fully automates the creation of backup copies for redundancy and off-site backup purposes, ensuring that you remain compliant with regulatory requirements and backup best practices – and are ready for ANY disaster. Highlights of the new SOBR functionality include:

- **Dramatically simplify off-site backup** and achieve compliance with 3-2-1 backup rule by immediately duplicating newly created backups to object storage with policy-based copy management.
- **Protect recent backup copies** in object storage against modification or deletion by malware or malicious actors by making them immutable for the required number of days, with the Compliance object locking mode ensuring that not even a root storage account is able to bypass the protection.
- Leverage new capabilities **without changing the way you use Veeam** today! The new copy policy supports forever-incremental source backup chains (without requiring a backup chain to be sealed) and works hand in hand with the existing Capacity Tier functionality of offloading older backup to object storage completely.

Access to SOBR requires **Veeam Universal License**, or the **Enterprise edition or higher** socket license.

Instant VM Recovery® – Reimagined

Veeam invented and pioneered Instant VM Recovery® in 2010, but now we're taking it to the whole new level! v10 enables additional data recovery and portability use cases by letting you instantly recover ANY Veeam backup created by ANY Veeam product to a VMware vSphere VM – finally making hybrid cloud DR a reality through meeting the RTO for any data or application regardless of its location. In addition, v10 sets the new instant recovery performance and scalability standards, making long downtime a distant memory. The new capabilities include:

- **Instant restore of ANYTHING to VMware:** Instantly recover from ANY image-level Veeam backup including physical servers or workstations, virtual machines or cloud instances – to a vSphere VM. No learning curve is required, as the recovery just works thanks to the built-in P2V/V2V conversion logic – enabling recoveries and migrations with new levels of speed and flexibility!
- **Next-generation Instant Recovery engine:** Recognizing that typical backup storage is inherently slow, instead of throwing SSDs at the problem, we focused on enhancing instant recovery performance for ANY storage. This effort resulted in the new Instant Recovery 2.0 engine, which uses workload-aware RAM caching to achieve unprecedented levels of recovery performance. Our v10 beta testers reported a few times faster VM boot times from typical backup storage!
- **Multi-VM Instant Recovery:** With the most advanced instant recovery engine on the market, you can now confidently initiate restores of multiple VMs at once – making this the perfect solution to bring those multi-server applications back online fast! To facilitate this use case, the new Multi-VM Instant Recovery wizard has been added, which makes it easy to perform mass instant restores in just a few clicks, reducing downtime associated with recovery and migration operations.
- **Instant restore of individual disks:** Reduce your instant recovery footprint by restoring only the required disks of large vSphere VMs, as opposed to restoring the entire VM (for example only OS disk, or only data disks). Instantly mount disks from backups to the selected VM for other use cases, for example comparing the disk content or performing mass file-level recoveries using 3rd party tools.

Veeam Data Integration API

We've enhanced data portability and platform extensibility to make it easier than ever for you to put your backups to work and reuse your data! v10 connects your entire data management ecosystem by enabling third-party applications and scripts to instantly access the content of any Veeam backup for data mining, security analysis, compliance checks, and other **data reuse** scenarios. By delivering the full automation needed to provide access to restore point content, v10 enables massive time savings compared with performing these tasks manually. The use cases for this new capability include:

- **Provide instant access** to the content of ANY Veeam backup or replica to specialized third-party mining and security analysis applications and scripts through the single PowerShell cmdlet that mounts the content of any restore point into the file system of the specified application server.
- **Perform data mining** of your production systems (including searching for Personally Identifiable Information, or PII, to aid in compliance processes and GDPR) without any impact on your production environment, by offloading overhead from such data scanning to backup storage hardware that usually remains idle during production hours.
- **Automated on-going security analysis** of your production environment by scanning recent backups for sleeping malware with additional antivirus applications, removing this burden from the production environment and reducing the number of antivirus licenses required.

Tighter Linux Ecosystem Integration

We've put a big focus on deeper integrating with Linux ecosystem technologies in v10 to provide our customers with more choices around their data management strategy. The new features include:

- **Advanced XFS integration:** By bringing fast clone technology to XFS, v10 delivers up to 20x faster synthetic full backup operations for shorter backup windows, decreased backup storage load and significantly reduced backup storage consumption. The space-less full backup technology enabled by this integration prevents duplication from occurring, resulting in raw disk space consumption of Grandfather-Father-Son (GFS) backup archives rivaling that of deduplicating storage appliances. But, unlike the scenario with

deduplicating storage, by integrating software-based deduplication and encryption with XFS block cloning capabilities, these storage savings exist even for compressed and encrypted backup files.

- **NFS backup repository:** v10 expands backup storage options with the native support for NFS backup repositories, simplifying the integration with next-gen hyper-converged storage appliances which commonly make their capacity available through the NFS protocol. v10 also provides a more reliable way to leverage low-end NAS devices without an iSCSI interface, safeguarding users against backup corruption issues caused by SMB protocol design peculiarities with non-continuously available (CA) SMB shares.
- **Linux backup proxy:** Achieve greater scale more efficiently with support for Linux servers as VMware HotAdd backup proxies, for example in "Linux-heavy" shops using VMware VSAN. Instead of forcing to use a pre-built proxy appliance, we help you to achieve compliance by leveraging your preferred Linux distributions – approved by your security team and having centralized monitoring and patching protocols already in place. Security aspects are critically important considering that backup proxies have direct and unrestricted access to the production VM data.
- **Network-less guest processing:** Seamlessly protect isolated environments with guest VM processing operations such as guest file system indexing, pre-freeze and post-thaw scripts, Oracle redo log backup as well as file-level restore even when direct network connection to the VM cannot be established. In such case, the connection into the guest OS will be established through the ESXi host instead, using VMware vSphere guest interaction API (formerly known as VIX API).
- **Expanded Linux compatibility.** Seamlessly connect to Linux machines with SSH server configured to use *keyboard-interactive* authentication type. This is supported for all Veeam capabilities including the registration of backup infrastructure servers, agent-based backups, application-aware processing for image-level backups and file-level recovery.

Other Enhancements

In addition to the above-mentioned major areas of improvements, v10 includes over one hundred other enhancements which are a response to customer feedback and on-going R&D learnings, the most significant of which are listed below.

General

Engine

- **Enhanced backup format:** Metadata bank format was optimized and now consumes 20% less space, resulting in faster incremental backup file chain mounting, smaller backup files and reduced RAM consumption by data movers.
- **On-demand metadata loading:** When opening an incremental backup chain for the next incremental backup, or to perform a restore, metadata from each backup file will now be loaded on-demand, as opposed to immediately for all backup files. This approach accelerates restore process initialization and further reduces RAM consumption by data movers considerably.
- **Secure certificate:** Backup server certificate used in communication between backup infrastructure components and agent management functionality has been changed from SHA-1 to SHA-256 to align with industry security best practices.
- **Reduced port range:** To avoid overlap with the RDP port, v10 reduces the default port range for data movers to TCP 2500 to 3300. If your deployment uses extremely high number of concurrent tasks, please expand the port range manually following the v10 upgrade.

- **Continuous public API expansion:** We keep expanding our PowerShell and RESTful API to cover all newly added functionality, but also to address gaps in the existing cmdlets and API calls based on your requests and use cases! While there are too many new additions and updates to list in this document – close to 200 new PowerShell cmdlets alone – please refer to the official documentation for the corresponding APIs.

Application-aware processing

- **SQL Server 2019 support:** Added full support of SQL Server 2019, including application-aware processing, transaction log backups, and all restore types with Veeam Explorer for SQL Server.
- **Oracle log retention:** Redo log retention period can now be set to values higher than 60 hours to prevent logs being purged before the backup takes place, in extended site downtime scenarios.
- **Enhanced log backup metadata format:** New transaction log backup metadata was reduced in size and includes support for incremental updates – providing for significantly increased performance and scalability of log backup jobs. For example, transaction log backups from 100 SQL Servers with 20 databases each can now be done with 5 min RPO, comparing to 20 min RPO before the change.

Backup

- **GFS retention:** Primary backup jobs with periodic full backups enabled can now use GFS retention policy for those backups. This functionality is designed to work hand in hand with the SOBR Capacity Tier using the copy+move policy, thus dramatically simplifying off-site backup while optimizing your long-term retention costs.
- **Daily retention:** You can now choose between restore-point based retention (for more predictable repository disk space usage) and the new retention based on the number of days (for when you need to guarantee your fixed retention policy is met even with manual ad hoc job runs). In the latter case, to protect against scenarios when backup job has not been running for a long time, Veeam will always keep at least 3 latest restore points by default. You can use *RetentionDaysBackupsCountMin* (DWORD) registry value under the *HKLM\SOFTWARE\Veeam\Veeam Backup and Replication* key on the backup server to configure this parameter.

Backup Copy

- **Immediate copy mode:** This new Backup Copy job mode copies every restore point created by selected primary backup jobs as soon as it is created in the primary backup repository. This mode is designed to help meet strict off-site backup RPOs, and even comes with the built-in RPO monitor functionality that will notify you when the restore point was not copied within the specified period of time after its creation.
- **Transaction log backup processing:** In addition to image-level backups, Backup Copy job in the immediate copy mode can now be configured to copy transaction log backups, thus enabling point-in-time database recovery from off-site backup copies.
- **Enhanced per-VM backup chains support:** Backup Copy jobs in both immediate and periodic copy mode will no longer wait for the entire primary backup job to finish the execution before starting the copy process when per-VM backup chains are used, thus improving off-site RPO.

Backup Infrastructure

Scale-out Backup Repository

- **Sealed mode:** SOBR extents can now be marked as Sealed, which prevents new backups from being added to it, however still allows restores and retention processing. This is useful when you want to safely decommission storage hardware, without putting it through much stress associated with the backup evacuation. And with object storage extents, this provides an economical way to vacate a cloud provider without extra egress charges associated with downloading all backups back to on-premises – by simply allowing those backups to age out.

- **Backup placement enhancements:** SQL Server and Oracle transaction log backups are now treated as incremental backups and placed on extents dedicated to incremental backups. This helps to ensure meeting the required RPOs for your databases even when SOBR includes extents backed by slower backup storage (such as deduplicating storage appliances).
- **Improved disk space reservation logic:** Extent disk space reservations for new backups will now use previous full and incremental backup sizes for the given job, instead of static percentages of source machine size. This should improve SOBR extent scheduling in situations when SOBR is nearing its capacity.
- **Tiering job status:** To make object storage copy and move tasks easier to monitor and troubleshoot, SOBR tiering jobs now have the additional "Waiting" status. This status means there are tasks to process, but no backup repository slots are currently available (as object storage offload is the lowest priority task).

Object Storage Repository

- **Backup import:** Backups residing in an object storage repository can now be imported with the click of a button, similarly to how you import local backup files today. This capability dramatically simplifies disaster recovery following the complete loss of on-premises backup infrastructure. You can start restoring your environment from backup copies in object storage – including directly to the public cloud IaaS – in no time by simply downloading and installing the free Veeam Backup & Replication *Community Edition* on your laptop!
- **Amazon S3 One Zone-IA support:** Customers who want a lower-cost option for infrequently accessed backups (such as GFS restore points) and do not require the availability and resilience of S3 Standard or S3 Standard-IA storage classes now have an option to leverage S3 One Zone-Infrequent Access storage class.
- **Microsoft Azure Data Box support:** Added native support and dedicated UI for registering Microsoft Azure Data Box storage solutions, enabling seamless seeding of SOBR Capacity Tier backups to Microsoft Azure Blob Storage.
- **S3 operations performance improvements:** Amazon S3 and S3-compatible storage should see much faster rescan and retention processing thanks to these operations becoming multi-threaded.
- **Veeam Backup for Microsoft Azure support:** Register Microsoft Azure blob storage containers with backups created by Veeam Backup for Microsoft Azure as external repositories, enabling you to perform all types of restores, and copy your Azure VM backups to on-prem backup repositories for disaster recovery purposes and for compliance with the 3-2-1 rule.

Tape

- **Backup to Tape performance:** We've made multiple changes under-the-hood to significantly accelerate Backup to Tape synthetic full backup offload performance when using enterprise-grade backup storage with Windows-based backup repositories.
- **Backup to Tape job scheduling:** List of tasks for an already running tape offload job will no longer be re-built when the new restore points appears for the already processed task while the job is still in progress. This logic often prevented the jobs from finishing as expected, which resulted in dependent jobs or scripts not being triggered.
- **File to Tape performance:** File to Tape backup performance from SMB/NFS file shares have been accelerated through direct share access to file shares registered in the Inventory. Existing File to Tape jobs will be automatically modified to leverage the direct share access, with the corresponding file shares automatically registered in the Inventory – please review your job settings after upgrading to v10 to make sure our upgrade scripts got it all right.
- **NDMP to Tape performance:** NDMP to Tape backup performance has been accelerated up to 3 times depending on NAS and tape hardware.

- **Multi-drive inventoring:** Tape inventory process will now use all available tape drives, instead of being limited to a single drive.
- **Tape management performance enhancements:** Various performance optimizations of mass tape operations, including Mark as Free and mass removal.
- **Email notification enhancements:** By popular demand, tape backup job email report now includes tape library slot information.

Veeam Cloud Connect

- **Log backup copy support:** Tenants can now use Backup Copy jobs in immediate copy mode to copy transaction log backups to a cloud repository.
- **Delete tenant backups:** When a service provider deletes the tenant, the backup console will now offer to automatically delete all tenant's backups respectively.
- **Application-item recovery from agent backups:** Tenants can now perform application item recovery from agent-based backups located in cloud repositories.
- **Improved scalability:** Added new request-based cloud repository disk space allocation logic for backups direct to cloud repositories which significantly increases backup performance to fast cloud repositories for tenants with high upload bandwidth. The new behavior can be enabled using `CloudConnectQuotaAllocationMode (DWORD) = 1` registry value under the `HKLM\SOFTWARE\Veeam\Veeam Backup and Replication` key on the Cloud Connect server.

Access to Veeam Cloud Connect *for Service Providers* requires a Cloud Connect Rental license.

For access to Veeam Cloud Connect *for the Enterprise*, please contact your Veeam sales representative.

WAN Acceleration

- **High-bandwidth mode:** New high-performance data processing mode for built-in WAN accelerators significantly reduces bandwidth usage while providing transfer times comparable to direct transfer mode for WAN links under 1Gbps. High-bandwidth mode will be used if both source and target WAN accelerators have the corresponding option selected, otherwise low-bandwidth mode will be used.
- **Low-bandwidth mode enhancements:** The existing data processing mode has been optimized and now provides up to 2x faster data processing throughput than with the previous version. As such, this mode is now recommended for WAN links up to 100 Mbps, or whenever minimal bandwidth consumption is required – as this mode provides noticeably better data reduction than high-bandwidth mode.

Access to built-in WAN accelerators requires the **Veeam Universal License**, or the **Enterprise Plus edition** socket license (the **Enterprise edition or higher** for Veeam Cloud Connect tenants).

Platforms

Microsoft Windows

- **Windows Server version 1909 support:** Added support as guest OS for application-aware VM processing, as server OS for protection with agent-based jobs, and for the installation of backup infrastructure components.
- **Windows 10 version 1909 support:** Added November 2019 Update support as guest OS for application-aware VM processing, as server OS for protection with agent-based jobs, and for the installation of Veeam Backup & Replication and its backup infrastructure components.

Microsoft Hyper-V

- **Automatic reference point clean-up:** Jobs will now automatically clean up obsolete reference points to prevent the issue with slow VM start up due to VMCX file size when a large amount of reference points have been accumulated. This logic does not apply to VMs with VHDS disks, because removing reference points breaks VHDS changed block tracking.
- **Disk exclusion:** Removed limitation that only first 16 virtual disks on each SCSI controller could be excluded from processing in backup and replication jobs.
- **SureBackup on Windows 10 Hyper-V:** SureBackup Virtual Lab now supports Microsoft Windows 10 Hyper-V hosts for SureBackup jobs with Hyper-V VM backups. This was added based on the feedback from MSPs who build Veeam-powered all-in-one DR appliances based on Windows 10.

Nutanix AHV

The following new Veeam Backup & Replication v10 functionality requires using the Nutanix AHV **v2** backup proxy appliance:

- **Centralized management:** You can now centrally deploy and manage AHV backup proxy appliances directly from the backup console. We have even reduced the appliance image size significantly to make its deployment and updates very fast!
- **Nutanix AHV integration:** Added ability to register Nutanix AHV clusters in the backup consoles, which enables direct restore of ANY Veeam backup into an AHV VM, file-level recovery from backups and native AHV snapshots, VeeamZIP of AHV VMs, Veeam ONE reporting and centralized license management.
- **AHV backup proxy v2:** New appliance functionality includes snapshot-only backup jobs, file-level and disk-level restore from AHV snapshots, support for deduplicating storage appliances, email notifications and more. For the complete list of new features, please refer to the corresponding What's New document.

VMware vSphere

- **VMware NSX-T support:** VMs that use VMware NSX-T networking are now correctly replicated and restored with full and instant VM recovery. Previously, such VMs were restored and powered on with their vNICs disconnected from the network, and thus not immediately usable.
- **Automatic CBT reset option:** You can now choose to automatically reset VMware Changed Block Tracking (CBT) on VMs upon Active Full backup. This feature was suggested by our users as a safeguard against possible undiscovered CBT data corruptions. Enabling this option increases the backup window due to CBT reset requiring the creation and removal of an additional VM snapshot.
- **vCloud Director 10 support:** Added full support for vCloud Director 10 including backup and restore of vCloud VMs and vApps, and self-service portal for vCloud Director.

VMware Cloud on AWS

- **VMC VDDK usage:** Backup job will now detect if the processed VM resides in VMware Cloud on AWS and, if so, initialize the VMware VDDK with the special flag that lets VMware track the data protection vendor and software version in use, and leverage this information for VMware Cloud updates planning.

Restore

File Level Recovery

- **Remote mount:** By default, backups will now be mounted exclusively to the mount server associated with the given backup repository, with the backup console receiving file system contents from the mount server. This new architecture radically improves FLR mount performance in ROBO environments. The ability to mount backups directly to the backup console is still available for the appropriate use cases through the corresponding new option in the Backup Browser.
- **File search:** The redesigned Backup Browser now includes file search functionality to help quickly locate required files on the selected volume from the mounted restore point. You can search any file system supported for the file-level restore.
- **File restore audit:** All file-level restore operations are now logged into the comprehensive audit trace that includes details on every restored file, user performing the restore and so on. Audit logs are stored as plain text CSV files on the backup server, with path to file included in the restore session. If desired, audit files storage location can be redirected to a specialized WORM storage – otherwise, older audit logs will be automatically compressed and archived to preserve disk space.
- **Restore under user context:** File-level recovery to Linux can now be performed using standard (non-elevated) user accounts as long as the selected account has rights to write to the given files/directories. In addition, you can now enable the prompt for single-use target machine credentials using *LinuxFLRSingleUseCreds* (DWORD) = 1 registry value under the *HKLM\SOFTWARE\Veeam\Veeam Backup and Replication* key on the backup server.
- **NSS file restore:** NSS file-level restore helper appliance was updated to use OES 2018 SP1 to support new file system features.

Veeam Explorer™ for Microsoft SQL Server

- **Publish database to cluster:** Instant database publish functionality now supports Microsoft SQL Server failover cluster as the target.

Veeam Explorer™ for Oracle

- **Publish database:** Added the ability to instantly publish a point-in-time state of any backed-up database to the selected Oracle server for dev/test purposes by running the database directly from the backup file. Changed database state can be exported or discarded with no impact to the actual backup.
- **Database export:** Added the ability to export a point-in-time state of any backed up database to a native RMAN backup file, or as regular database files, to simplify the process of providing the database backup to Oracle developers, BaaS clients or Oracle Support.
- **Out of place restore enhancements:** Support for restore of database files to alternative paths, and for restore of RMAN backups to another server.
- **No sudo requirement:** Restore from RMAN backups no longer require elevating rights with sudo on the target Oracle server.
- **RESETLOGS option:** Added the new restore option to allow Oracle Database Administrators to easily restore the database in a state that allows applying additional logs manually following the restore from RMAN backups.

- **Restore blocker for Windows:** In addition to similar existing capability for Linux-based Oracle servers, we added support for blocking Veeam Explorer from being able to restore to the particular Oracle server running on Microsoft Windows. To block restores, an empty *disablerestore* file should be created on the system volume in the `%ProgramData%\Veeam\` folder.
- **Oracle Data Guard experimental support for RMAN backups:** See Veeam [KB2976](#) for more information.

Direct Restore to Cloud

- **Security improvements:** We made public IP allocation and restored instance power on operations optional to improve security.
- **Amazon EC2 enhancements:** Improved restore performance of machines large number of disks through accelerating EC2 import step.
- **Azure IaaS enhancements:** Added support for Managed Disks and 8TB disks, as well as the ability to select disk type when performing the restore.
- **PowerShell interface enhancements:** We enhanced our PowerShell coverage for the Direct Restore functionality based on feedback of our customers and managed service providers who fully automate disaster recovery to public cloud with Veeam.

Secure Restore

- **Kaspersky Security 10 support:** Added out-of-the-box support for Kaspersky antivirus.

Storage Integrations

Primary Storage Integrations

The following functionality was designed for Tier 1 workloads, which are impacted by classic backups from VM snapshots, and leverages primary or secondary storage array snapshots as "backups". When using this functionality, please **remember to take real backups** of your VMs periodically to another storage, for example during the scheduled service downtime, to achieve media break required by the 3-2-1 rule.

- **Transaction log backups from storage snapshots:** You can now perform Microsoft SQL Server and Oracle transaction log backups with storage snapshot-only backup jobs. In this case, image-level backup of a VM (that includes database files) will be stored in the primary and/or secondary storage snapshot, while transaction log backups will be performed to a Veeam repository normally. Veeam Explorers have been updated to support point-in-time database restore transparently for such "split backup" scenarios.
- **VM snapshot requirement removal:** Storage snapshot-only backup job will no longer leverage the VM snapshot when performing application-aware processing of VMs with virtual disks located on the dedicated datastores, as it is often the case with Tier 1 workloads.

Access to backup from storage snapshots functionality requires **Veeam Universal License**, or the **Enterprise Plus edition** socket license.

Other primary storage integration enhancements include:

- **Volume exclusions:** Added the ability to include or exclude volumes from/to VMFS rescan process based on volume name mask/wildcard.
- **Mount server selection:** When registering a storage array, you can now specify the mount server to be used by Veeam Explorer™ for Storage Snapshots (for example, during file-level recovery from storage snapshots). Previously, backup console itself was always used as the mount server.

- **VM configuration files backup:** When backing up from storage snapshots, VMX, NVRAM and VMDK descriptor files will now be retrieved through the backup proxy, as opposed to the backup server.
- **Enhanced NFS storage support:** All primary storage integrations now support deployments with NFS datastores registered via DNS names, but without the associated PTR records in the DNS server, which causes reverse lookup failure. In this case, matching will now be performed through DNS A record.
- **HPE Primera support:** Added HPE Primera primary storage array support for all storage snapshot integration capabilities.
- **NetApp ONTAP 9.7 support:** Added support for ONTAP 9.7 for all storage snapshot integration capabilities. v10 is a mandatory update due to the breaking API change in ONTAP 9.7.
- **NetApp SVM support:** NetApp cDot cluster can now be registered with Veeam by adding the required vServer (SVM) only, as opposed to the requirement to add the entire cluster.

Secondary Storage Integrations

- **New integration architecture:** New deduplication storage support architecture moves storage-specific interaction logic from the source agent or proxy to the target agent (backup repository component), which enables all current and future Agents and Plug-ins for Enterprise Applications to utilize all deduplicating storage appliances supported by Veeam Backup & Replication.
- **HPE StoreOnce Catalyst Copy support:** HPE StoreOnce users can now leverage Catalyst Copy-based Backup Copy jobs for backups produced by VM and Server backup jobs managed by the backup server. Such jobs can only use StoreOnce-based backup repositories as both the source scope and the destination, and operate in the immediate copy mode – copying all backups to the target repository as soon as they appear in the source repository.
- **Quantum DXi block cloning integration:** Added experimental support for Quantum DXi native block cloning functionality that is based on v10 advanced XFS integration. This functionality is pending Quantum's internal validation, and must be enabled on the storage side by Quantum support.
- **DELL EMC DD OS 7.0 support:** Added support for DELL EMC Data Domain storage using DD OS 7.0.

Access to secondary storage integrations requires the **Veeam Universal License**, or the **Enterprise edition or higher** socket license.

Agents and Enterprise Application Plug-ins

Agent Management

- **WSFC on S2D support:** Backup of Microsoft Windows Server Failover Clusters based on Storage Spaces Direct is now supported by "managed by server" agent-based jobs.
- **Veeam Cloud Connect support:** You can now configure cloud repositories hosted by the Veeam Cloud Connect service provider as targets for all agent-based backup job types.
- **Backup cache support:** Enable the usage and configure the placement of Veeam Agent for Microsoft Windows backup cache centrally for "managed by agent" image-level backup jobs.
- **Job management improvements:** Added abilities to control (start/stop), perform Active Full backups, and view job statistics interactively (include transaction log backup statistics) for "managed by agent" jobs.
- **Granular scripts:** Agent-based backup jobs for Linux now support granular (per-machine) settings for pre/post job scripts.

Veeam Agents

- **Veeam Agent for Microsoft Windows v4:** New functionality includes new OS versions support, parallel disk processing for servers, backup job and cache upload resume following sleep, hibernate or short network disconnect events, backup mapping and more. For the complete list of new features, please refer to the corresponding What's New document.
- **Veeam Agent for Linux v4:** New functionality includes new OS versions support, application-aware processing for PostgreSQL and MySQL, radically increased file-level backup performance (up to 20x), and more. For the complete list of new features, please refer to the corresponding What's New document.

Access to agent-based backups requires the **Veeam Universal License**.

Veeam Plug-ins for Enterprise Applications

General

- **SOBR placement enhancements:** Oracle RMAN or SAP HANA Backint channels will now be spread across all available SOBR extents (based on task slots availability) to maximize throughput when SOBR Performance Mode is used.
- **Garbage collector:** Plug-in backup files can now be force-deleted by backup server after the specified time. This provides automated backup repository clean up from orphaned backups, which can happen when some databases are dropped because Oracle RMAN or SAP HANA Backint backup catalogs no longer maintains the inventory of the associated backup files in this case.
- **Restore tasks display:** Running plug-in restore tasks are now displayed in the backup console as active tasks and can be canceled by the user as necessary.

Backup Copy

- **Plug-in backup copy support:** You can now create Backup Copy jobs for your plug-in backups, which will operate in the immediate copy mode – copying backups to the target repository as soon as they appear in the source repository. Backup copies use an independent retention policy, thus protecting primary backups from unintentional deletion, and enabling the segregation of duties between the Database Administrator and the Backup Administrator.
- **Dedicated RPO monitor:** You can choose to be notified when the newly created plug-in backup file was not copied to the target repository within the specified period of time after its creation.
- **Read-only mount:** Mount backup copies with the plug-ins as read-only to restore from them without the risk of accidental deletes.
- **Disaster recovery mount:** Should the primary backup repository be lost, the repository with plug-in backup copies can be easily promoted to the primary repository to continue backup using the existing backup file chains. A new backup copy job to another backup repository can then be created for additional copies.

Veeam Plug-in for Oracle RMAN

- **Oracle Solaris support:** Added Oracle RMAN plug-ins for Oracle Solaris versions 10 and 11 running on Intel x86_64 and SPARC server hardware.
- **Out of place restore enhancements:** Restore to other servers can now be processed without the Veeam XML configuration change, by instead defining the original server in the channel definition with SEND "*srcSrv=<originalserver>*" command.
- **Autobackup set up:** The plug-in configuration wizard now automatically enables Autobackup setting in the RMAN configuration to ensure fast and reliable restores.

- **Backup file naming:** Application backup files extension is now hardcoded to .vab to avoid issues with backup file management functionality in Veeam Backup & Replication (for example, SOBR).
- **Improved Oracle RAC support:** Added support for concurrent plug-in backups from multiple RAC cluster nodes.
- **Resource usage optimizations:** Reduced Oracle server CPU load from plug-in backup activities and added the ability to disable network traffic compression (--compression parameter) in the plug-in configuration to further reduce CPU load if needed.
- **Oracle Data Guard experimental support:** See Veeam [KB2976](#) for more information on experimental support.

Veeam Plug-in for SAP HANA

- **Improved SAP HANA scale-out cluster support:** Added support for concurrent plug-in backups from multiple scale-out cluster nodes.
- **Performance and resource usage optimizations:** Multiple under the hood improvements for up to 40% faster single-channel database backups and improved scalability in cluster environments.
- **SAP HANA 1.0 experimental support:** See Veeam [KB2997](#) for more information on deploying the plug-in for SAP HANA 1.0.

Access to Enterprise Application plug-ins requires the **Veeam Universal License**, or the **Enterprise Plus edition** socket license.

User Interface

Backup Console

- **UI cache service:** Backup service will now cache configuration database query results for reusing and invalidate them in case of related database table changes. The new architecture dramatically reduces configuration database load and improves UI responsiveness in large environments, especially when multiple backup administrators are using their backup consoles concurrently.
- **Scalable views:** Stored procedures behind all view have been refactored for performance. In addition, certain "heavy" views now load their content asynchronously and limit the number of items displayed to 1000 by default. You can then choose to see all items or apply a search filter to reduce the number of items returned.
- **Job filters:** Define flexible job filters, and quickly switch between them to make the jobs grid only display the jobs you need to see at the given moment.
- **New backup properties view:** We've redesigned this commonly used dialog to make it more functional and scalable.
- **Console startup time:** Backup console should now open more than twice as fast. Anything counts for the TCO and RTO!
- **Fewer annoying questions:** We hate those too, so our wizards will no longer ask whether you want to exit without applying changes even when you did not actually make any edits.
- **Daily emails time setting:** The time when daily status emails are sent is now configurable in the UI.
- **Require restore reason:** Based on the popular demand, we added the ability to disable hiding the Restore Reason step of restore wizards and require at least one symbol be provided as the reason to prevent

clever operators from simply clicking through this step. You can enable this functionality by creating *UIRequireRestoreReason* (DWORD = 1) registry value on the backup server. But if the step was previously selected by the user already, then you will have to do edit *user.config* file manually in the Veeam folder under *%LOCALAPPDATA%* path on each affected backup console.

- **Session history:** Default session history retention has been reduced to 3 months to reduce configuration database size, which is important for customers trying to stick with SQL Express due to its 10GB database size limit.
- **High DPI support:** All wizards, file-level recovery Backup Browser and dashboards were migrated to Windows Presentation Foundation (WPF) UI framework for enhanced support of high-resolution monitors.

Enterprise Manager

- **SAML 2.0 support:** Allows the usage of 3rd party identity providers to authenticate users into the Enterprise Manager, adding support for non-Active Directory users, multi-factor authentication etc.
- **Individual disk restore:** Enterprise Manager and all self-service portals now provide individual disk restore functionality for when the entire VM restore is not desired.
- **vCenter scoping:** Self-service backup and restore portal for VMware vSphere now supports selection of vCenter servers in the quota settings to restrict users to certain parts of the virtual environment only.
- **Backup server audit report:** Enterprise Manager can now export backup server audit reports. This report lists all user activities in the backup console during the specified time frame.
- **Large file systems support:** Web UI should no longer time out when browsing very large file systems during file-level restore. New browse dialog will fetch results asynchronously, and display contents of large directories in portions.
- **RESTful API security:** HTTP access to the REST API is now disabled for enhanced security. Please update your integrations to use HTTPS instead.

Licensing

General

- **No double licensing:** Whenever the same machine is protected by more than one Veeam product and the protection is managed by the same backup server, only a single license will be consumed (previously, each product would consume a license separately). For example, using an agent-based backup job *and* Veeam Plug-in for Oracle RMAN to protect the same server now requires only 1 Veeam Universal License (VUL). Note that this does not apply to NAS backup jobs, which are licensed based on protected front-end capacity, as opposed to per machine/workload.
- **In-product license merge:** You can now install Socket and Instance licenses with matching Company Name alongside one another without having to merge them in the Veeam Customer Portal first – and manage them separately going forward. As with the previous product version, product functionality level continues to be determined by the Socket license edition, while Support and License expiration dates of the merged license continue to be inherited from the license that expires sooner.
- **New license file format:** v10 uses the new and improved license file format that no longer includes personal information, which is replaced with the unique license ID. In addition, the new format is no longer restricted to the specific major version (but is compatible with v10 and later versions only).
- **License auto update:** When performing an in-place upgrade to v10, the setup wizard will offer to download your v10 license automatically. This requires uploading your currently installed license to Veeam servers. If your backup server has no Internet connection, or if you prefer not to have your current license uploaded – you can download your v10 license from the Veeam Customer Portal.

- **Perpetual licenses:** Perpetual licenses no longer have the License Expiration field. Customers can therefore install and use indefinitely any product version shipped before the Support Expiration date. However, as per the existing Veeam support policy, customers without an active support agreement will not be able to install hotfixes, patches or updates shipped after the Support Expiration date of the installed license.
- **Subscription licenses:** Subscription licenses no longer have the Support Expiration field, because support expires along with the license itself. We've retained the 1-month grace period following the License Expiration date.
- **Rental licenses:** Rental licenses now leverage the same Point system that is used by Veeam Cloud & Service Providers (VCSP) for usage consumption reporting. The product will update Point weights periodically from Veeam servers to take into account any recent changes if the license auto-update functionality is enabled.

Instance Licenses

- **Veeam Universal License (VUL):** v10 uses portable, universal license which comes in the single fully-featured edition and no longer uses a weighting table, with every workload consuming one instance. This dramatically simplifies license management and portability.
- **Veeam Instance License (VIL) conversion:** Existing VIL will be converted to VUL, at no charge, for the remainder of the contract. The conversion is done at 1:1 ratio regardless of VIL edition, to ensure that each customer can continue protecting the same number of workloads.
- **NAS backup capability:** The new v10 file share backup functionality requires VUL, with 1 instance consumed for each 250GB of protected front-end capacity. For example: up to 250GB = 1 VUL, 251GB to 500GB = 2 VUL, etc.

Socket Licenses

- **Mandatory socket license usage for VM protection:** *In presence of a Socket license*, protecting vSphere or Hyper-V VMs with agents or enterprise application plug-ins requires and consumes a socket license. This is the existing Veeam licensing policy that is now enforced at the product level. To avoid backup interruption after v10 upgrade, please ensure your socket license is sufficient to cover all hypervisor hosts where protected VMs reside, no matter of which Veeam product is used for protecting those VMs.
- **Expanded socket license coverage:** VMs on hypervisor hosts licensed with a Socket license can now be protected with any Veeam product without requiring additional Instance license, so long as the protection is managed by the same backup server. For example, protecting a Windows failover cluster vSphere VM with Veeam Agent for Windows no longer requires an Instance license for the agent, as the VM protection is already covered by the socket license for the corresponding host. Again, no double charging.
- **Gifted instances are now VUL instances:** Up to 6 gifted instances that are provided according to your licensed Socket count have been converted to VUL instances. In particular, this means that Standard and Enterprise edition Socket license users will be able to protect slightly more workloads requiring an Instance license than with the previous product version.
- **Gifted instance count can no longer be exceeded:** Unlike with subscription Instance license, gifted Instances usage is now limited to the granted amount.
- **Gifted instances are now persistent:** Merging a Socket license with an Instance license will no longer remove gifted instances, with their number added to the total number of usable Instances. Such license usage can be exceeded normally, based on the total number of Instances.

Community Edition

- **Expanded free license capacity:** The 10 Instances provided by the built-in free license are now VUL Instances. This allows users to protect up to 10 physical servers (or up to 30 workstations) with the Community Edition – a significant increase from 3 servers (or 10 workstations) with the previous version. The number of protected VMs remains unchanged at up to 10 VMs.
- **NAS backup capability:** Yes, Community Edition can protect NAS too, as it shares the same 10 VUL instance pool provided by the built-in free license.
- **Mandatory check for product updates:** You can no longer disable the option to check for product updates periodically. In addition, Community Edition users will be notified of the new version availability sooner and are expected to become the early adopters. However, you're not required to upgrade to the new versions immediately, and can continue to do so as you see fit.



Learn more
veeam.com



Download free trial
vee.am/backup