



# Veeam Backup & Replication

## What's New in v8?

### Veeam Beats Legacy Backup

The results are clear – a recent ESG survey\* finds those who switch to Veeam increase their availability and find a reliable solution they can trust.

- 96% of Veeam recoveries are within RTO SLAs, compared to 78% for other solutions
- 83% of Veeam customers are more confident in their current backup compared to previous solutions
- 71% of Veeam customers report improved reliability of backups compared to other solutions
- 84% of Veeam Virtual Lab users save time by avoiding deployment problems
- 71% of customers say their risk awareness has improved using Veeam monitoring and reporting tools compared to other solutions

\*Source: ESG Customer Insights Research, February 2014

[Read how Veeam Beats Legacy Backup](#)

"Veeam has innovated numerous times over its eight-year history, and has earned its spot in the upper echelon of data protection vendors."

*Dave Simpson, 451 Research*

Veeam® Backup & Replication™ v8 provides new capabilities that deliver *Availability for the Modern Data Center™* through cutting-edge improvements in high-speed recovery and data loss avoidance.

### Backup from Storage Snapshots and Veeam Explorer for Storage Snapshots – now available for NetApp

Veeam's storage snapshot integration technologies have been extended to support NetApp storage. Achieve an unprecedented level of data protection through a combination of NetApp's low recovery point objective (RPO) capabilities and Veeam's low recovery time objective (RTO) enablement.

The integration of NetApp and Veeam allows you to:

- Create virtual machine (VM) backups from storage snapshots up to 20x faster than the competition
- Recover individual VMs, guest files and application items quickly and efficiently from NetApp Snapshot, SnapMirror and SnapVault
- Improve disaster recovery (DR) protection by creating instant, secondary backups from storage snapshots to SnapMirror and SnapVault
- Create frequent, no-impact storage snapshots with snapshot-only "backup" jobs managed directly in the Veeam Backup & Replication console, including support for application-aware processing

NetApp integration supports NFS, FC and iSCSI connectivity, and both 7-mode and cluster-mode deployment options. All Data ONTAP storage is supported, including NetApp FAS, FlexArray (V-Series), Edge VSA and IBM N Series.

### Veeam Explorer for Microsoft Active Directory

New Veeam Explorer™ for Microsoft Active Directory makes item-level recovery with Active Directory easier than ever by providing the following simple and powerful functionality:

- Instantly restore individual Active Directory (AD) object attributes, objects and entire containers
- Recover user and computer account passwords to avoid disruptions
- Perform 1-click compare between backed up and production Active Directory states on both object and object attribute levels
- Easily search for objects and attributes in the backed-up AD; includes plain text- and advanced LDAP filter-based search capabilities
- Export selected objects from backed-up AD state into LDIFDE format

### Veeam Explorer for Microsoft SQL Server

Veeam Explorer for Microsoft SQL Server enables you to achieve low RTOs and RPOs with Microsoft SQL Server database protection with the following all-new functionality:

- 1-click individual SQL database recovery for lowest RTO; no more searching for database files or spending time reattaching them to SQL server using native management tools

- Web-based restore portal for SQL database restores by delegated database administrators
- Agentless transaction log backup and replay for no-impact, low RPO protection of your SQL databases. Back up your databases as often as every five minutes!
- Point-in-time SQL database recovery down to individual transactions with the unique transaction logs viewer to ensure lowest possible RPOs
- Full support for SQL AlwaysOn Availability Groups for both backup and restore

### **Veeam Explorer for Microsoft Exchange**

Veeam Explorer for Microsoft Exchange brings additional functionality, making granular Exchange restores easier and more efficient than ever. New features include:

- Web-based restore portal for delegated help desk operators enables secure mailbox items recoveries without restore operators having access to the content of restored items
- Ability to recover hard-deleted items that you thought might have been lost forever
- Support for searching of and recovering to Online Archive mailboxes
- Restore of Public Folder permissions (in addition to Public Folder contents)
- Support for item-level recovery from secondary (passive) DAG node backups

### **EMC Data Domain Boost**

With EMC Data Domain Boost (DDBoost), you get source-side data deduplication, yielding faster, more efficient backups that enable lower RPOs and lower exposure to data loss. The benefits of combining the Veeam Backup & Replication with EMC Data Domain Boost include:

- Up to 50% faster backup performance with advanced source-side deduplication by DDBoost
- 10x faster in-place synthetic full backup file creation and transformation for dramatically shorter backup windows
- Fibre Channel (FC) connectivity enables LAN-free backup to Data Domain

DDOS 5.4 or later and DDBoost license from EMC required.

### **End-to-end encryption**

With Veeam's built-in AES 256-bit encryption you can secure your data end to end:

- At source (backup file encryption via backup proxy)
- In flight (network traffic encryption for replication and quick migration)
- At rest (via "at source" encryption, or using tape encryption)

Unique password loss protection enables IT managers to maintain the ability to recover backups even if the passwords set on backups by your IT staff are lost, all without compromising security.

More importantly, because all Veeam technologies are tightly integrated, backup encryption does not impact the data reduction ratios of built-in compression and WAN Acceleration, as is the case when using third-party bandwidth reduction solutions.

### **Replication enhancements**

Improve offsite backup and recovery with additional replication enhancements from Veeam, allowing you to:

- Reduce the impact on your WAN link with Built-in WAN Acceleration for replication, resulting in up to 50x faster replication over slow links, and up to 20x reduction in bandwidth consumption on fast links
- Reduce the impact on primary storage with replication from backup files, and replication from NetApp storage snapshots
- Lower failover complexity with new Failover Plans, providing built-in orchestration for easy 1-click entire-site failover
- Lower RTOs to ensure zero data loss with Planned Failover functionality designed to facilitate data center migrations using replication jobs
- Reduce the time difference between replicated states when replicating a group of VMs by processing retention policy at the end of the replication job for all VMs, as opposed to doing this after each replicated VM
- Eliminate bandwidth loss on network disconnects with true resume of data transfer upon network disconnects; available for both standard and WAN-accelerated replication jobs

## Veeam Cloud Connect

Avoid the risk of catastrophic data loss in case of a site disaster through hosted backup repositories with Veeam Cloud Connect. It's the easy, efficient way to get backups off site to a service provider through:

- **Hosted offsite backups:** Get your backups off site to a cloud repository hosted by the service provider of your choice through a single port over a secure connection protected by SSL authentication and network traffic encryption.
- **Complete visibility and control:** Access hosted backup repository content and perform recoveries directly from a cloud repository—no need to copy backup files locally first!
- **Modern backup architecture:** Leverage Built-in WAN Acceleration, forever incremental backups, Backup Copy jobs, GFS retention and MORE with cloud repositories, just as if they were your local backup repositories.

Veeam Cloud Connect also gives service providers an easy way to host offsite backups for Veeam customers by providing:

- **Cloud gateways:** Virtualize all traffic between the customer and service provider across SSL—no need for VPN; deploy multiple cloud gateways to enable high availability (HA) and automatic load balancing.
- **Multi-tenant support:** Segment a single backup repository into multiple cloud repositories, and manage resource allocations and expiration dates for each tenant.
- **Service provider licensing:** Monthly subscription-based, per-VM licensing with no upfront investment or commitment.
- **Service provider lookup:** Web portal where Veeam customers and resellers can find YOU!

## Backup I/O Control PATENT PENDING

Backup I/O Control ensures your workload's availability by reducing the impact of backup and replication jobs on production VMs running on the same storage where backed-up VMs reside by monitoring the production datastore read latency and controlling job I/O to keep the latency within user-defined, acceptable limits. This ensures availability of production workloads during the backup window without unnecessarily sacrificing backup job performance when the datastore load is low (as is often the case with legacy solutions offering static per-datastore task limits).

Backup I/O Control works in conjunction with Automatic Intelligent Load Balancing to ensure no new I/O-intensive tasks (such as virtual disk backup, restore or snapshot removal) are assigned to the same datastore if the datastore's latency is above a user-defined first latency threshold.

Additionally, Backup I/O Control actively throttles Data Mover I/O for existing tasks if the production datastore's latency goes above a user-defined second latency threshold (e.g., due to a production VM starting to execute an I/O intensive operation).

To ensure that backups can complete under any circumstances, Backup I/O Control will always allow at least one active task per datastore, and will never throttle any given task to less than 10% of full I/O capacity.

## Self-restore portal for application owners PATENT PENDING

Veeam Backup Enterprise Manager v8 includes a new, dedicated self-service, file-level restore portal for application owners. The portal will automatically detect what VM the user logs on from, find the given VM backup, and check if the user was a member of a Local Administrators group on the given VM at the time when the backup was executed. If this is the case, the user is immediately provided with the ability to browse, search and restore files and folders from the latest restore point of this VM or select another restore point if necessary.

The self-restore portal simplifies restores for application owners, as they no longer have to guess the VM name of the server from which they need to restore files (and thus increase their restore times). Since delegation is fully automated and leverages the existing Microsoft Windows permissions already in place, the portal also reduces TCO by eliminating the need for portal administrators to constantly tweak the restore delegation settings as new VMs are added to the environment, or new users requiring privileges to be granted for performing file-level restores.

## 200+ additional enhancements

As a result of ongoing research and development efforts and in response to customer feedback, Veeam Backup & Replication v8 includes an extensive range of additional features and enhancements, the most significant of which are listed below.

### General

- **Forever forward incremental backup mode:** This new, default backup mode for primary backup jobs creates forward incremental backups for all the following runs after the initial full backup, which always remains the oldest restore point in a full backup chain. When the retention policy needs to remove the oldest restore point, the job merges the oldest incremental backup into the full backup file, discarding any data replaced in the full backup file by this process. This is identical to how Backup Copy jobs store the backup files.

Because forward incremental backup creation involves mostly sequential write I/O, the forever forward incremental backup mode reduces the time a VM runs off of a snapshot by up to 3x when compared to the reverse incremental backup mode. This prevents a VM snapshot from growing large and making its commit much faster, thus reducing both backup window and load on production storage.

- **Quick rollback:** When restoring an entire VM, the new quick rollback option can be used to restore only the data blocks known to be changed since the last backup was taken. This option is not selected by default, and should only be enabled in case of application-level issues. This option should not be used in scenarios with storage corruption or power loss, because CBT data cannot be trusted in such cases, which is why the entire VM restore is highly recommended. However, quick rollback can significantly improve RTOs in disaster scenarios that were not caused by hardware problems.
- **Certificate-based authentication:** Support for certificate-based authentication for Linux servers used as backup repositories, and for the guest OS for application-aware processing and guest file system indexing has been added. Supported SSH keys formats include: OpenSSH RSA, OpenSSH DSA, OpenSSL PEM, Open SSL PKCS#8, SSH.com and Putty Private Key (PPK).
- **Native SQL authentication:** Native Microsoft SQL Server authentication is now supported for the configuration database. When native authentication is used, Veeam Backup & Replication user interface users are no longer required to have configuration database owner privileges.
- **Pre-job script:** In addition to the existing post-job script, backup and replication jobs can now be configured to run pre-job scripts, for example to stop certain VMs before backup, using a PowerCLI script to ensure cold backup.
- **Update notification:** Veeam Backup & Replication can now optionally notify users about new updates that are available by checking the current version against the Veeam website. This includes updates for Veeam Backup & Replication, as well as critical Microsoft Hyper-V updates and patches. When the backup server does not have an Internet connection, an updated **vbrupdate.vdf** definition file can be downloaded and copied to **C:\ProgramData\Veeam\** manually, replacing the existing file.

### Engine

- **Preferred networks selection:** Enables the selection of preferred networks for multi-homed backup proxy and backup repository servers. This helps to ensure the backup and replication traffic is transferred over dedicated backup networks, without impacting the production network.
- **Multi-streaming of restores:** Veeam Data Movers now leverage multiple TCP/IP streams when performing restores, in addition to backups. This helps to improve performance by more fully leveraging available network bandwidth of high latency links.
- **Concurrent jobs:** The limit of 100 concurrent jobs has been removed to allow running even more jobs at once.
- **Parallel processing for host discovery:** Infrastructure rescan now scans multiple hosts and servers at once, thus dramatically improving operations such as registering large virtual infrastructures.
- **Infrastructure rescan:** The default infrastructure scanning period was changed from 24 hours to four hours to be able to detect any virtual environment or storage changes faster, thus preventing failed jobs.

## Automatic Intelligent Load Balancing

- **Job start time priority:** The backup infrastructure resource scheduler will now prioritize all jobs according to their start times, and attempt to finish any particular job as soon as possible once it has been started. This means that if you start multiple jobs at once, the scheduler will assign newly appearing processing resources to the job that started earlier, as opposed to assigning them between all running jobs. This approach ensures that the backed-up state of all VMs within the same job remains as close together as possible.
- **Retry cycle priority:** The retry cycle inherits the priority of the job it belongs to, as opposed to getting the same priority as other running jobs.
- **Periodic job priority:** The resource scheduler will now give highest priority to jobs scheduled to run periodically over daily and monthly jobs to ensure required RPOs are being met.
- **Snapshot removal priority:** Pending snapshot removal tasks get highest priority among all other tasks in the event that Backup I/O Control starts restricting task assignment to a certain datastore.
- **Backup window enforcement:** Activities isolated to backup repositories, such as full backup transformation, should no longer be affected by the backup window setting.
- **Periodic jobs improvements:** Scheduled, periodic jobs will now have more predictable start time behavior upon reaching their defined time periods, which have been limited to factors of 24 hours (i.e., 1, 2, 3, 4, 6, 8, 12, and 24).

## Application-aware processing

- **Custom pre-freeze and post-thaw scripts:** You can now specify custom pre-freeze and post-thaw scripts for both Windows and Linux in the advanced application-aware processing settings to prepare applications requiring custom processing for backup, which completely removes the hurdle of deploying and updating said VMware Tools scripts in vSphere and brings this otherwise missing functionality to Hyper-V.
- **Copy-only VSS processing:** Application-aware processing can now be configured to perform “copy-only” backups, which allows application-consistent images to be created without touching transaction logs. This can be required if transaction logs are processed with a third-party application.
- **Microsoft Exchange VSS improvements:** If standard VSS processing of a Microsoft Exchange Server times out, the job will retry processing using persistent in-guest VSS snapshot, which should prevent VSS processing timeouts commonly observed with Microsoft Exchange 2010. Persistent snapshots are removed from the production VMs once backups are finished. All VM, file-level and item-level restore options have been enhanced to detect and properly handle restores from Exchange backups created with the help of persistent VSS snapshots.
- **Oracle VSS improvements:** Oracle databases that do not meet the requirements for successful VSS processing are automatically excluded from application-aware processing to avoid job failures.
- **Guest logging improvements:** Improved guest logging significantly reduces disk space consumption in the guest OS. This should prevent situations with guests running out of disk space on the system volume due to too many logs being created.

## Guest file system indexing

- **Parallel indexing:** Guest file system indexing is now performed in parallel with virtual disk data backup, as opposed to holding backups until indexing finishes.
- **Linux indexing:** Support for Linux guest file system indexing has been added. This feature requires that the Linux server has mlocate (which is standard for the vast majority of Linux distributions).
- **NTFS indexing improvements:** NTFS indexing logic has been redesigned for up to 10x better indexing performance on file servers with large amount of files.
- **ReFS indexing improvements:** ReFS indexing logic has been optimized for over 2x better indexing performance.
- **Indexing from VSS snapshots:** The indexing will now be performed from a VSS snapshot to make the catalog more precisely reflect the backed-up state in cases where the guest OS supports Microsoft VSS.
- **Indexing priority:** The indexing process now runs with a lower priority to reduce impact from its CPU consumption on existing applications running on the backed-up VM.

## Instant File-Level Recovery

- **Overwrite/keep options:** When restoring to the original location, users performing restores can now choose whether the original file should be overwritten or renamed and preserved. This functionality is available in both Windows UI and Veeam Backup Enterprise Manager web UI.
- **Enhanced LVM support:** Logical volume names are now fetched from the LVM configuration and displayed to the user, as opposed to showing /dev/mapper device names (dm0...dmN) to help users find the location of the files they need to restore more easily.
- **Solaris UFS support:** Support for Solaris UFS to multi-OS file level-recovery wizard has been added.
- **Extended ZFS version support:** ZFS pools are now supported up to version 5000. Note that ZFS pools from [proprietary Oracle Solaris distributions](#) are still not supported due to Oracle licensing limitations.
- **Improved performance:** Windows file-level recovery performance was increased for scenarios in which multiple files are being restored over the network or locally to a backup server. Performance improvements should be especially noticeable when the backup server is protected with antivirus software with real-time scanning enabled.
- **Inactive IFLR session shutdown:** After 30 minutes of inactivity, an Instant File-Level Recovery (IFLR) session will be considered abandoned, and will be automatically closed after displaying the corresponding warning. This will remove the locks from backup files and prevent potential failures of daily backups.
- **Volume quantity limit removed:** The maximum of 20 mounted volumes limit has been removed from the native file-level recovery engine. This engine is leveraged for both file and item recovery, so the change should help with item-level recovery scenarios from application servers with large amount of volumes hosting application databases.
- **4K native disk support:** File-level recovery is now supported from native 4 KB formatted disk with. This requires switching the IFLR engine to an alternative mounter. To enable one, create the **EnableHvVDK** (DWORD) = 0 registry value under **HKLM\SOFTWARE\Veeam\Veeam Backup and Replication** key.

## VMware vSphere

- **Snapshot Hunter:** Detects and automatically consolidates hidden VM snapshots to prevent production VMs from stopping due to datastores overflowing with snapshot files. To detect hidden snapshots before exiting, and in addition for checking "Consolidation Needed" VM status, the jobs will physically scan datastores for snapshot files belonging to snapshots not registered in vSphere. If hidden snapshots are found, the job will attempt to automatically consolidate them by following the procedure described in VMware support KB article KB1002310. If the job fails to remove hidden snapshots due to snapshot files being locked, the job will log an informational event and exit, however a background system process will attempt to perform consolidation three more times every four hours in the hope that file locks will be removed. If these attempts fail too, Snapshot Hunter will stop trying, and a warning email will be sent to the global notification recipients urging them to take manual action.
- **Direct SAN restore:** This enables thick virtual disk restore directly through the SAN fabric for faster performance and no impact on the production network. Using thick disks is common among users who choose to do thin-provisioning on a storage level. This functionality is explicitly disabled for thin disks, because it impacts vCenter Server and does not provide acceptable performance due to vSphere storage architecture. We recommend using hot add for fast restore of VMs with thin disks.
- **Hot add improvements:** Backup proxies can now process multiple virtual disks of the same VM in parallel using hot add. This includes hot add backup, restore and replication (on both source and target proxy). Additionally, performance of said operations can see an increase of up to a few times depending on the environment.
- **Tag-based job population:** Jobs can now be populated using VM tags as the dynamic selection scope. This provides incredible flexibility around job management. Consider using Veeam Business View (part of Veeam Availability Suite™ v8) to automate the assignment of tags to newly created VMs in your environment based on pre-defined rules.
- **Tag backup and restore:** VM tags are now backed up and restored when restoring VMs back to the original location.
- **DRS cluster-based job population:** Jobs can now be populated using DRS clusters as the dynamic VM selection scope. This job population method gives greater predictability of job size compared to other methods.

- **DRS cluster support.** DRS clusters can now be selected as the target for restore or replication. The actual VM placement will be determined by the restore or replication job automatically by querying DRS recommendations at VM creation time.
- **Append to VM attribute:** Instead of overwriting the existing contents of a selected VM attribute (such as VM Notes) with the last backup status string, last VM backup information can now be optionally appended to the current attribute's value. (Shout out to Michael White for this one – we appreciate your continuous support!)
- **Support for vSphere Flash Read Cache:** Restore jobs and replication jobs will now correctly configure VMs when restoring or replicating to target hosts without SSD storage installed.
- **Building VMs list optimization:** The performance of the Building VMs list job operation was enhanced by optimizing virtual infrastructure data queries.
- **Snapshot release improvements:** VM snapshot interaction logic in corner cases (such as job hangs and crashes) has been reworked to reduce the chance of snapshot files being locked by data movers that may still be running, thus preventing potential snapshot removal issues.
- **Backup server connectivity:** The backup server no longer requires connection to source and target ESX(i) host on port 902, as the operation requiring this connection is now performed by the source and target backup proxies respectively. This requirement was causing issues in environments with high network isolation, such as in service provider environments.

## Microsoft Hyper-V

- **Shared VHDX support:** Backup jobs can now back up shared VHDX virtual disks in a crash-consistent state. The full VM restore wizard has also been enhanced to account for shared VHDX restore specifics.
- **Consistency groups:** Jobs can now process more than one VM from the volume snapshot at a time instead of creating a separate volume snapshot for each processed VM. For reliability reasons caused by Hyper-V backup architecture, the maximum amount of VMs is limited to four per snapshot by default in case of software VSS, and eight per snapshot in case of hardware VSS, but can be increased on fast primary storage with the **MaxVmCountOnHvSoftSnapshot** (DWORD) and **MaxVmCountOnHvHardSnapshot** (DWORD) registry values under **HKLM\SOFTWARE\Veeam\Veeam Backup and Replication** key.
- **Volume-based job population:** Jobs can now be populated using volumes as dynamic VM selection scope.
- **Checkpoint-based replicas:** Replica restore points are now stored as VM checkpoints for faster recovery to any replica restore point, better reliability and the ability to power on any required restore point manually if required.
- **VLAN ID support:** SureBackup® network mapping rules for production networks and multi-OS file-level recovery helper appliance configurations now support pointing to specific VLANs in a multi-VLAN environment.
- **VSS failover control:** Failover from hardware-based VSS snapshots to software-based VSS snapshots has been made optional in order to prevent undesired usage of software-based VSS snapshots when a hardware-based snapshot fails.
- **Hyper-V resource pool support:** Hyper-V resource pool settings are now preserved when restoring VM backups to their original location, and are automatically removed when restoring a VM to a different location.
- **Smart Paging File support:** The Smart Paging File location can now be customized in replication jobs to control the file's placement on the target host, thus preventing potential issues upon failover.
- **CBT driver certification:** The proprietary changed block tracking (CBT) driver has passed Microsoft Windows Hardware Certification testing for Windows Server 2012 R2.
- **CBT driver improvements:** Optimizations to further reduce the CBT driver's CPU consumption have been added.
- **Bitlocker support:** Support for CBT on CSV volumes encrypted with Bitlocker has been added.
- **Differencing VHDX improvements:** Backup and replication performance has been improved significantly for differencing virtual disks, especially for VMs running on Windows Server 2012 R2.

## Built-in WAN Acceleration

- **Performance improvements:** Data processing performance has been improved up to 3x compared to v7 with the latest patch. The enhanced data processing engine includes additional optimizations specifically aimed to better leverage faster CPUs and SSD cache.
- **Cache optimizations:** The WAN accelerator cache will now consume the entire allocated space immediately upon creation to prevent fragmentation and avoid unexpected lack of disk space issues at a later time as the cache grows.
- **Cache management:** The WAN accelerator cache population can now be triggered without having to run the job by selecting one or more backup repositories to crawl for OS data. Cache reset was also made available directly in the user interface for troubleshooting purposes.

## Rotated drives

Rotated drives functionality can now be enabled with a checkbox in the advanced backup repository settings, as opposed to being controlled by registry settings (which are still supported for upgrade compatibility reasons). Additionally, when a backup repository backed by rotated drives is running on the Windows server, the following additional functionality is supported:

- **Intelligent cleanup:** When a rotated drive that was used during previous cycles is attached, the backup job will only delete backup files if they are no longer under retention policy. This prevents unwanted deletion of backup files still under retention due to accidentally attaching an incorrect rotated drive.
- **Drive letter tracking:** Backup jobs will now track rotated drives, and will follow them without failing even when the drive letter changes. Drive letters may change when adding new volumes or adding additional storage hardware, such as CD-ROM.

## Tape

- **Remote tape server:** This enables leveraging tape libraries and standalone drives connected to any Windows server in the environment, as opposed to the requirement of attaching them directly to the backup server. Adding multiple tape devices connected to different servers is supported.
- **Extended backup modes support:** Backup to Tape jobs now support backup jobs using backup modes without periodic full backups as the source. For such jobs, on the full backup day, Backup to Tape jobs will synthesize the full backup file on the fly, using data from source backup files on disk. The synthesized full backup file will be compatible with the following incremental backup files created by the primary backup job.
- **Backup Copy job support:** Backup to Tape can now leverage Backup Copy jobs as the source.
- **Performance enhancements:** The performance of File to Tape backup has been increased up to 50x when backing up workloads consisting of a large number of small files. Performance of backing up large files should see minor improvement as well.
- **Media Vault:** Enables tracking the physical tape medium location, such as storage shelves or offsite locations. Tape mediums are automatically assigned to the specified vault as they go offline (removed from tape library). Different media vaults can be specified for each media pull.
- **Software overwrite protection:** Enables the ability to protect certain tape mediums from being overwritten according to the data retention policy. This functionality can be useful in certain situations, such as legal hold.
- **Media pool retention:** Changing the retention policy on existing media pools is now supported.
- **Job scripts:** File to Tape and Backup to Tape jobs now allow configuring pre- and post-job scripts to perform auxiliary actions with the tape library.
- **Granular email notification:** File to Tape and Backup to Tape jobs now allow setting up per-job email notification recipients and settings.
- **Catalog auto-refresh:** The tape catalog is now subscribed to change notification from all jobs, and performs automatic granular refresh as required (e.g., when a new tape backup is created).
- **Tape library rename:** Tape libraries and tape drives can now be renamed in the user interface, in cases when the native name supplied by the driver is unreadable.



- **Last chain backup:** When creating a new Backup to Tape job that uses an existing backup job as the source, and that has multiple full backup chains already created, users can choose whether all backup files, or only the latest full backup chain files, should be copied to tape.
- **Job timeouts:** Tape backup and restore jobs will now time out and fail after 72 hours of waiting for user interaction (e.g., inserting tape), as opposed to waiting for user action forever and not reporting any issues.
- **Amazon Storage Gateway support:** Added support for AWS Storage Gateway in virtual tape library (VTL) configuration for seamless tape-like backup experience to Amazon S3 and Amazon Glacier. Requires AWS Storage Gateway version 23-Oct-2014 or later and using Sony TSL-A500C Autoloader medium changer emulation.

## Primary storage integration

- **HP 3PAR StoreServ iSCSI support:** Veeam Explorer for Storage Snapshots and Backup from Storage Snapshots are now supported for HP 3PAR StoreServ with iSCSI connectivity.
- **Processing failover:** Jobs using Backup from Storage Snapshots can now be set to automatically fail over to standard vStorage API-based backup, should primary storage integration fail. The failover is disabled by default to prevent unexpected impact on primary storage during production storage and needs to be enabled manually if desired.

## Backup storage integration

In addition to EMC Data Domain integration described earlier, v8 integrates with the following deduplicating storage:

- **ExaGrid:** Joint development between Veeam and ExaGrid R&D enables the Veeam Data Movers to run directly on the ExaGrid appliance, resulting in improved backup performance due to enabling local access to backup files. Note that integration must be configured in the ExaGrid management UI before registering ExaGrid appliance with Veeam. Additionally, backup jobs will now automatically change the default settings to one recommended by ExaGrid when pointed at the ExaGrid-backed repository.
- **HP StoreOnce:** Backup jobs will now automatically change the default settings to those recommended by HP when pointed at a StoreOnce-backed repository. The integration does NOT include HP Catalyst support. (We are working on it though!)

## Ad hoc backups

- **Quick Backup:** Enables fast, incremental backups of one or more selected VMs. Quick Backup looks up the job to which the selected VM belongs, and runs incremental backup for the selected VM only. If multiple jobs exist for the same VM, the one with most recent restore point will be used to ensure smallest incremental backup and fastest backup. Restore points produced by Quick Backup do not affect overall job retention; they are "glued" to the previous restore point produced by a standard job run, and are removed during regular job retention processing.
- **VeeamZIP™ auto retention:** VeeamZIP now includes optional retention on produced backup files. When enabled, created backup files will be automatically deleted after the specified amount of time.
- **vSphere Web Client integration:** Both Quick Backup and VeeamZIP are now available directly from the vSphere Web Client, placing excellent alternatives to classic VM snapshots at your fingertips.

## Backup Copy jobs

- **Backup Copy suspend:** A Backup Copy job is now suspended when source backup files are locked by another job or process, such as a restore or full backup file transformation, instead of just failing.
- **Backup mapping:** In addition to full backup, you can now map Backup Copy jobs to any existing backup chain, as long as it does not contain reversed incremental (VRB) files. This dramatically simplifies Backup Copy seeding, and makes it possible to re-point Backup Copy jobs after moving backup files to another storage target.
- **Backup Copy as a source:** Backup Copy can now use backups produced by another Backup Copy job as a source. While the Backup Copy job itself cannot be the source due to its continuously running nature with no stop event, you can point the "child" Backup Copy to a backup repository used by the "parent" Backup Copy job to store backups in, to achieve the desired result.
- **Backup repository throttling:** Backup Copy jobs now respect the I/O throttling rule on the source backup repository. For example, if more than one Backup Copy job reads from repository with I/O throttling enabled, all Backup Copy jobs will split the allowed bandwidth equally.

- **Forward lookup:** Backup Copy jobs can now be configured to wait for the new restore point to be created by a primary backup job, instead of using the latest non-copied restore point already available. To enable the new behavior, create the **BackupCopyLookForward** (DWORD) = 1 registry value under **HKLM\SOFTWARE\Veeam\Veeam Backup and Replication** key.

## Configuration backup

- **Cross-version configuration restore:** Configuration restore now supports restore of configuration backups created by any product version supported for upgrade to the given product version. In case of v8, this means support of configuration backups created by versions 6.5, 7.0 and 8.0.
- **Encryption:** Configuration backups can now be encrypted. When encryption is enabled, configuration backup will include all credentials, which in turn accelerates the configuration recovery process.
- **Tape catalog backup:** Configuration backup now catalogs all tapes written by File to Tape and Backup to Tape jobs. This significantly reduces backup server recovery times, as re-cataloging all tapes can take significant time to complete.
- **Session history backup:** Configuration backup now includes all job session history, including restore operator activity.
- **New restore wizard:** Configuration restore has been streamlined with the new restore wizard, allowing users to choose between two distinct restore scenarios: regular configuration restore and backup server migration.
- **Flexible restore options:** Configuration restore now supports restoring only certain parts of product configurations (i.e., only backup infrastructure, only jobs configurations and restore points catalog or only session history).
- **Flexible scheduling options:** Configuration backup now provides more flexible scheduling options, similar to what is available for regular jobs.
- **Additional notification options.** Configuration backup jobs now provide SNMP and email notifications to global notification recipients.

## SureBackup

- **Enhanced SQL Server script:** The pre-defined Microsoft SQL Server script has been replaced with a more advanced script that checks each database's availability, instead of just SQL Server service availability.
- **BIOS UUID removal:** SureBackup jobs can now automatically remove BIOS UUID from VMX file of tested VMs to prevent problems with VMware SRM and similar tools that track VMs using this parameter. To enable UUID removal, create the **RemoveBiosUuid** (DWORD) = 1 registry value under **HKLM\SOFTWARE\Veeam\Veeam Backup and Replication\SureBackup** key.

## User interface

- **Guest account testing:** Guest accounts used for application-aware processing and guest file system indexing can now be automatically tested on the corresponding job wizard step. Testing is performed building the list of all VMs added to the job, and attempting to connect to each guest using the provided guest credentials settings (including any overrides in the advanced guest credentials settings dialog).
- **Job-level email notifications:** Email notification settings can now be configured on the job level. Job-level notification can include a custom set of recipients and be triggered by different status than global notifications.
- **Save as Default:** The advanced job settings pages now provide the ability to save the current settings as defaults to simplify the configuration of future jobs. Note that clicking Save as Default saves the settings of the currently displayed dialog page only.
- **VM snapshot removal progress:** The VM snapshot removal progress percentage is now included in the job statistics to provide users with a better idea of how much longer the process will take.
- **Password peeker:** Password peeker control is now included everywhere where passwords need to be supplied to save users from typing the same password twice for input verification purposes.
- **Minute-based periodic job start time:** Start times within an hour can now be set for minute-based periodic jobs in addition to hourly jobs.

- **Monthly job schedule improvement:** Monthly jobs can now be configured to run on a particular day of the month (e.g., 1st or 31st), in addition to the first or last day of the week of a month.
- **Job abort reasons:** The exact reason for job stopping is now displayed in the job's session when it is stopped by a user (username is included), another job (name of the job is included) or when it runs over the forced backup window.
- **Job size calculation improvements:** The total VM size now takes into account excluded disk sizes. (We are sorry it took us a while to fix this one!)
- **TCP/IP thread control:** The number of concurrent TCP/IP connections can now be set in the UI on the Network Traffic Rules dialog, as opposed to being a registry value. Increasing the value helps to saturate high latency links, but setting it too high may increase backup proxy CPU usage significantly as well as affect reliability due to intermediary network equipment misbehaving under heavy load.
- **Session type display:** Session window caption and session list views now show the job session type (i.e., full, active full, synthetic full, incremental or reversed incremental).
- **Disk exclusion improvements:** IDE and SATA disks exclusions can now be configured.
- **Multi-select support:** Multi-select is now supported in the advanced application-aware processing and Guest OS Credentials dialogs.
- **Guest OS details:** Guest OS details are now displayed in the Virtual Machines view as an additional column.
- **Capacity details:** Backup repository capacity is now displayed in the Backup Repositories view.
- **Session report sort order:** By popular demand, the sort order for the multiple jobs session report was changed to display the latest session at the top.
- **Hide unneeded tabs:** The ability to hide unneeded tabs in the management tree has been added. For example, you can now hide the Storage Infrastructure and Tape Infrastructure tabs completely if you don't have supported primary storage and don't use tape.
- **New repository type:** The backup repository wizard now allows adding backup repositories backed by deduplicating storage through the corresponding designated option. This option lets you can choose between deduplication appliances featuring appliance-specific integration with Veeam Backup & Replication, as described earlier.
- **Backup Copy jobs with an empty source:** You can now create Backup Copy jobs with an empty source to be used as the secondary destination job later when configuring primary backup jobs. This helps to streamlines job configuration in large environments.
- **Hide splash screen:** The ability to hide the splash screen has been added based on requests from terminal users experiencing color display issues as well as users wishing to hide name and email address information of the product's licensee. To hide the splash screen, create the **ShowSplashScreen** (DWORD) = 0 registry value under **HKLM\SOFTWARE\Veeam\Veeam Backup and Replication** key.

## Veeam Backup Enterprise Manager web UI

- **1-Click Linux file restore:** Enables 1-Click file restores from Linux guests. Just like with Windows guests, this functionality requires guest file system indexing to be enabled in the backup job settings.
- **Failover initiation:** Enables VM failover from the web UI using Failover Plans. Delegated users can only initiate failover if they have full VM restore permission for all VMs included in the Failover Plan they choose.
- **Self-dismissing full VM restore warning:** Initiating full VM restore will first show the warning explaining that the original VM will be deleted. The warning will dismiss itself and the full VM restore will proceed if the user provides no input within 20 seconds, so as to not impact any existing web UI integration logic.
- **Backup to Tape jobs display:** Backup to Tape jobs can now be monitored with the Veeam Backup Enterprise Manager web UI.
- **Job disable:** Disabling jobs is now supported with the Veeam Backup Enterprise Manager web UI.
- **New design:** The Veeam Backup Enterprise Manager web UI has been refreshed to have a more modern look. (We hope you like it!)

## PowerShell API

- **Multi-OS IFLR:** By popular demand, multi-OS IFLR can now be managed through the newly added PowerShell cmdlet.
- **Improved cmdlet architecture:** All-new v8 cmdlets, as well as some existing cmdlets, have been implemented using the new architecture to allow access to certain internal objects without "dirty hacks." This should help to ensure compatibility of newly developed scripts referring to internal objects with future Veeam Backup & Replication versions.

## RESTful API

- **Multiple enhancements:** We've added many improvements and enhancements based on feedback from real-world implementations. (Special thanks go to iland web portal developers.)

## Setup

- **Streamlined setup:** The setup wizard has been streamlined to enable the installation of Veeam Backup & Replication in default configuration with fewer clicks. This includes support for running system services under LOCAL SYSTEM (new v8 installations only). Setup will now also automatically detect and use the volume with most free disk space for the default backup repository and guest file system catalog location, as opposed to defaulting to the C: drive.
- **Automatic upgrade:** The new VBR\_AUTO\_UPGRADE="YES" unattended setup option has been added. When enabled, the setup wizard initiates all Veeam components to upgrade automatically and make truly unattended upgrade a possibility for environments with dozens of backup servers.
- **SQL Server 2012 Express:** Veeam Backup & Replication v8 ships with Microsoft SQL Server 2012 Express Edition as the default local database option. Remote SQL Server databases and other SQL Server versions are still supported as before.

## Tools and utilities

- **Configuration database management wizard:** The new wizard has been added to allow users to easily update database connection settings (i.e., database name, authentication type, account or password).
- **Backup Validator enhancements:** To simplify scripting, Veeam Backup Validator no longer requires verified backup files to be imported into the backup server and can be pointed at standalone backup files (e.g., VBM, VLB, VBK+VBIs chain, VBK+VRBs chain). Encrypted backups are also supported, but require backup passwords to be present in the Password Manager of the corresponding backup server. Additionally, Veeam Backup Validator can now create nicely formatted HTML reports at the specified location that contain backup verification results and details.



Learn more   
[www.veeam.com](http://www.veeam.com)

Download free trial   
[www.veeam.com/backup](http://www.veeam.com/backup)