Veeam Backup & Replication v12.3

What's new

Contents

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Introduction

Veeam Backup & Replication, the foundational component of Veeam Data Platform, delivers enterprise-grade data resilience for your entire hybrid estate, providing confidence in your protection, response, and recovery in the face of disaster and cyberthreats. The following is a list of the major new features and enhancements added in Veeam Backup & Replication v12.3. All capabilities here are transacted as the Veeam Data Platform, with certain features available only at the Advanced or Premium editions. Follow this link for a detailed edition comparison.

New Cloud Workload Support

Microsoft Entra ID

Microsoft Entra ID is a cloud-based Identity and Access Management (IAM) system that delivers access to your internal and external resources. But Entra ID is much more than just a directory of users and groups, and protecting this data and knowledge is paramount. Entra ID is at the core of nearly every organization and is essential to keep your business running, and Veeam can now give you peace of mind by protecting it. Key highlights of Veeam's Microsoft Entra ID support include:

Accelerate change detection — Quickly identify and revert changes created by human error, threat actors, automated attacks, and more when restoring Entra ID data. Bolster your forensic investigations with a point-in-time copy of your IAM data.

Simplify governance, risk and compliance — Reduce risk and stay compliant through fast, automated backup processes to reduce human error risks, ensuring consistent resiliency practices. Unlock cost-effective, long-term audit and sign-in log storage with unlimited retention to be able to easily go back in time during internal investigations of cybersecurity incidents.

Rapidly restore your business — Bring your business back online in seconds by pinpointing broken or missing app registrations and restoring them in seconds with comprehensive app registration recovery. Using object-level recovery empowers you to choose exactly what data you restore.

Role-based access for restores — Contrary to alternate solutions, which perform backup and restore operations under a single almighty account, Veeam relies on the native Entra ID permission system to ensure Entra ID administrators are unable to restore and/or overwrite data they do not have privileged access to.

New Data Center Workload Support

Microsoft Windows Server 2025

Upgrade to the latest Microsoft releases with confidence thanks to the official support for:

Microsoft Windows Server 2025 and Microsoft Windows 11 24H2 support — Included as a guest OS of protected machines, for installation of Veeam Backup & Replication components, and for agent-based backup with the Veeam Agent *for Microsoft Windows* 6.3 (included in V12.3).

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Microsoft Windows Server 2025 Hyper-V support — For host-based backup of virtual machines (VMs), allowing businesses to leverage the enhanced virtualization capabilities without compromising the ability to protect their production data.

Microsoft System Center Virtualization Machine Manager (SCVMM) 2025 support — For registering Hyper-V based virtualization infrastructure with Veeam Backup & Replication as a data source, streamlining VM management operations, and ensuring a robust backup strategy.

Microsoft SharePoint SE 24 H2 support — For application-aware processing with host-based and agent-based backup of machines running SharePoint SE 24 H2 and for application item-level recovery from such backups with Veeam Explorer *for Microsoft SharePoint*.

Nutanix AHV

Boost your Nutanix AHV protection with advanced Veeam capabilities you know and love, including application-aware processing for host-based backups, inline malware detection, and deeper data observability for Nutanix environments, all thanks to in-depth alerting and analytics with Veeam ONE.

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

Guest file system indexing — Adds visibility into Nutanix AHV VM guest OS files in the global catalog and enables accelerated self-service file-level recovery within the Veeam Enterprise Manager web UI.

Inline malware detection — Performs inline, low-impact entropy analysis of a data stream on the worker to immediately detect when previously unencrypted data becomes encrypted by ransomware, using a specially trained Machine Learning (ML) model, onion links, and certain ransom notes appearing on the protected machine.

Suspicious file system activity detection — Searches guest file system indexes for files with known malware file extensions, ransom notes, and similar flags of malware presence. Detects suspicious changes like bulk file deletes or renames, many new files with previously unknown extensions appearing, and other activities which can be a sign of a cyberattack or malicious insider activity. This functionality works independently from the inline malware detection and requires guest file system indexing enabled in the backup job settings.

Nutanix Guest Tools (NGT) — Added more granular control for environments leveraging NGT on protected VMs by allowing you to specify the application or crash consistent snapshot setting at the job level, as well as control database transaction log truncation or retention. However, NGT settings are ignored if application-aware processing is enabled.

Note: Due to the significant expansion of Nutanix AHV protection capabilities, the above-mentioned new features are initially made available under experimental support terms. Enabling them does not impact full supportability for your existing Nutanix AHV backups; the experimental support SLA disclaimer applies only to the newly added capabilities.

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

MongoDB

MongoDB 8 support — Added support for protecting the new MongoDB version with MongoDB backup policies added in Veeam Backup & Replication v12.2.

PostgreSQL

PostgreSQL 17 support — Added support for application-aware processing with host-based and agent-based backup of machines running PostgreSQL 17 and for application item-level recovery from such backups with Veeam Explorer *for PostgreSQL*.

Oracle RMAN

Oracle 23ai support — Veeam Plug-in *for Oracle RMAN* now supports backing up and restoring databases running on Oracle 23ai with RMAN. In this release, version 23ai support is limited to virtual machines in cloud deployments (Oracle Cloud Infrastructure and Oracle Exadata Cloud) and On-Premises Engineered Systems (Oracle Exadata and Oracle Database Appliance).

SAP HANA

SLES 15 SP6 and RHEL 8.10 support — Veeam Plug-in *for SAP HANA* can now protect HANA databases running on SLES 15 SP6 (for x86-64 architecture only) and RHEL 8.10 Linux distributions.

Veeam Agent *for Linux*

Added support for agent-based backup with the Veeam Agent *for Linux* 6.3 (included in V12.3) for the following latest versions of Linux distributions:

x86_64 architecture — Added support for AlmaLinux 8.10, AlmaLinux 9.5, Rocky Linux 8.10, Rocky Linux 9.5, Red Hat Enterprise Linux (RHEL) 9.5, Oracle Linux 9.5, and Ubuntu 24.10.

IBM Power architecture — Added support for RHEL 8.8, RHEL 8.10, RHEL 9.4, SLES 15 SP5, and SLES 15 SP6.

Veeam Agent *for Mac*

macOS 15 (Sequoia) support — Added support for agent-based backup with the Veeam Agent *for Mac* 2.3 (included in V12.3). Ensure your Mac systems remain fully updated and secure while maintaining robust data protection by Veeam!

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Cyber Resiliency Enhancements

Indicators of Compromise (IoC) Detection

Stop cyberattacks right in their track with the built-in detection of early indicators of compromise (IoC) on protected machines. V12.3 leverages its file system indexing functionality to detect and report the sudden appearance of utilities from hacker's toolkit, which are commonly utilized by cybercriminals for lateral movement, data exfiltration, command and control, stored credential access, and more, with the list of tools constantly updated by Veeam.

Detecting the appearance of such tools significantly reduces the Mean Time to Detect (MTTD) threats, providing you with an opportunity to react before attackers can inflict significant damage. This lightweight and scalable detection of IoC on all protected machines is meant to draw your attention to potential issues. In cases when an attack is suspected, we recommend performing a more thorough scan of affected machines using the Recon Scanner available from Coveware by Veeam.

Veeam Threat Hunter

Many customers love the idea of using backups to identify potential dormant threats in their environment. Whether through periodic manual spot-checks, continuous scheduled scans (powered by SureBackup), or alert-driven scans, risks can be uncovered without adding overhead to production environments while allowing for a fast response from security teams.

However, one of the challenges with searching for threats is knowing what to look for. While YARA scans are fast, they can only search for a strictly defined list of signatures, which can be problematic when proactively looking for unknown threats. Although incredibly useful for a forensic investigation, it is not optimal as a defensive measure. On the other hand, purpose-built antivirus software avoids this issue by having millions of malware signatures in its database, but their speed and performance can be difficult to scale.

V12.3 brings the best of both worlds — the speed of YARA scans and the breadth of malware detection of a classic antivirus — with the new Veeam Threat Hunter. This advanced signature-based malware detection engine is integrated directly into Veeam Backup & Replication data processing engine for significantly faster scanning than with the Bring Your Own Antivirus approach, with the breadth of malware detection that YARA scans cannot touch.

Key benefits of Veeam Threat Hunter include:

- Built directly into Veeam Data Platform to offer highly optimized, accelerated signature-based backup content scans for malware while reducing costs and freeing up your critical IT resources from managing a third-party antivirus scanner on your mount hosts.
- Veeam Threat Hunter employs machine learning (ML) and heuristic analysis to identify advanced threats such as polymorphic malware, which are impossible to detect with YARA rules due to the dynamic nature of signatures of each malware instance.
- Updates to threat signatures and ML models used to detect polymorphic malware are delivered multiple times per day to quickly expand detection to newly developing threats.

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam Data Cloud Vault

Access to secure, affordable cloud storage has never been easier! Veeam Data Cloud Vault, our first-party cloud object storage offering, now features a more simplified onboarding experience, providing instant access to ultra-reliable yet competitively priced cloud storage. Key benefits of the updated Vault include:

Uncompromising security and reliability — Safeguard your offsite backups on cloud object storage that is always immutable and always encrypted, now with up to 12 nines of durability.

Unbeatable price and predictability — Choose between two new editions tailored specifically for primary and secondary backup use cases, both fully managed by Veeam with all-inclusive pricing lower than DIY solutions on leading hyperscalers.

Unbelievable ease — Provision and monitor your Vault directly from Veeam Data Platform for a straightforward and seamless cloud storage experience.

Unmatched flexibility — Vault can be used as a backup target in any product edition, including in the Community Edition, which does not normally allow backup to object storage.

Veeam Intelligence

Veeam Intelligence brings powerful, AI-driven insights directly into Veeam Data Platform. Based on the Veeam AI Assistant, Veeam Intelligence is the next iteration of Generative AI at Veeam and is always there to provide in-product assistance. Veeam Intelligence enables admins to quickly get answers about product capabilities, optimize backup performance, solve technical issues and proactively address potential risks — all without using any environment-specific information whatsoever.

Ready for even more AI power? Have AI help you make data-driven decisions, the new opt-in mode of advanced Veeam Intelligence can dramatically improve efficiency and reduce the burden on IT resources. Available in the Veeam ONE console, this new mode provides natural language-based reporting on your backup infrastructure and production environments, based on the monitoring data collected by Veeam ONE. For more information, please refer to the What's New in Veeam ONE v12.3 document.

Other Features and Enhancements

In addition to the above major new features, V12.3 includes many enhancements that are a response to customer feedback and ongoing R&D findings, the most significant of which are listed below:

Security

Malware Detection

More intelligent "Mark as Clean" — A backup server will now automatically raise the deleted files threshold for a particular machine once you mark it as clean to prevent further false positives from changes of comparable volume. This ensures that operations involving mass file deletions that are legitimate for the given server are recognized and allowed without further alerts.

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Default malware index retention — After validating a reduced retention policy with several customers through our Customer Support, we have enabled 14-day retention as the new default value for all installations. This should significantly reduce guest catalog size and backup server load during catalog data analysis.

Security & Compliance Analyzer

Backup encryption password strength check — This new Security & Compliance Analyzer test will validate the strength of passwords used for backup encryption. Long and complex encryption passwords are essential for backups because when copied to local storage, an attacker can probe a large number of passwords per second against encrypted data.

Security Events

Syslog filtering — You can now exclude some less important events from forwarding to the target syslog server. By specifying IDs and severity levels of unwanted events, users can take full control of their backup monitoring, ensuring only relevant information is captured. This capability was particularly important for our customers using cloud-based event management systems that charge them per event received.

Image-level backup

Backup Copy

Rotated drives support — Don't let true airgap get in the way of managing your backups copies! By popular demand, you can now create Backup Copy jobs from backups copies stored on repositories backed by rotated drives, as well as target such jobs to repositories backed by rotated drives.

Cloud Connect support as target — You can now target Backup Copy jobs that use backup copies as a source to cloud repositories, effectively allowing you to have a copy of your production backups stored both on-prem and with the Cloud Connect provider, without having to pull data from your production storage twice.

Recovery from image-level backups

Instant VM recovery

Instant recovery engine scalability — Multiple optimizations have been made to the instant recovery engine improving its scalability while dramatically reducing backup server load. Thanks to these improvements, V12.3 completes mass instant recovery of 200 VMs more than four times faster than the previous version (in 10 minutes vs. 45 minutes) while consuming up to three-times less CPU and seven-times less RAM on the backup server.

File-level recovery

Linux file-level recovery performance — Minimize downtime with up to five-times faster file-level recovery from Linux, Mac, AIX, and Solaris backups thanks to various file-level recovery engine optimizations, with restore performance of large files getting the biggest performance improvement.

OpenText Open Enterprise Server (OES) 24.1 support — The NSS file-level restore helper appliance has been updated to enable support for OES 24.1.

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Continuous Data Protection (CDP)

Extended short-term retention — The maximum short-term retention period has been expanded from one day (24 hours) to seven days (168 hours) to provide low-RPO coverage with the I/O journal to cover not only weekends but also extended public holidays. Note that the effective retention may be limited by a maximum I/O journal size of 2TB per disk, but we're planning to remove this limitation in the next updates.

CDP Policy cloning — Now added is the ability to clone CDP policies. This functionality is particularly useful for creating new CDP policies from "template" policies based on different SLAs, target locations, and CDP proxy settings, speeding up new policies provisioning and ensuring their consistency.

Unstructured Data Backup

NetApp SnapDiff integration — Perform NetApp filers backup faster and more efficiently while reducing impact on the storage system during incremental backup runs through the integration with NetApp SnapDiff. This changed file tracking API allows backup jobs to query a list of changed or removed files without the need to scan a file share's content.

NetApp FSx support — NetApp FSx is now officially supported as a data source by file share backup jobs, providing a robust solution for protecting your data stored on Amazon FSx for NetApp ONTAP, ensuring that your critical data located in public cloud is reliably protected and easily recoverable.

Agents

Agent Management

Computers with preinstalled agents — Windows backup agents distributed via this protection group type can now utilize the new *matchbackupserver* parameter for the *setVBRsettings* command. This eliminates the need to reapply configurations to agents that are already configured, provided the backup server specified in the configuration file remains unchanged.

Veeam Agents *for AIX* and *Solaris*

File-level recovery direct to target — Streamline file recoveries by transferring restored files from backups directly to the target host. This feature minimizes recovery time, enhances efficiency, and improves security by reducing potential vulnerabilities in helper appliances.

Publish backup — Instantly access backup content by mounting disks from backup directly to the target system and letting system administrators interact with it directly. This enables several use cases, such as performing bulk file recoveries using native tools, scanning backup content with third-party tools, and more with minimal involvement from backup administrators. Any accidental changes made to the content of published disks are discarded once the backup is unpublished.

Veeam Agent *for Mac*

Hotspot and Low Data Mode detection — The agent was made aware of the Low Data Mode macOS setting and is now capable of detecting hotspots, allowing it to avoid sending traffic over metered networks and perform backups over more suitable network connections only.

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam Agent *for Microsoft Windows*

Automated Bare Metal Recovery — Automate the recovery process with support for answer files, streamlining the restoration of systems to the original or dissimilar hardware. By completely eliminating the need for manual input during recovery or limiting users to only basic settings like restore point selection, this capability eliminates the requirement for end-user training or detailed guidance during the recovery process while avoiding input errors. This makes bare metal recoveries effortless for backup admins and helps to get your end users back to being productive faster.

Bare Metal Recovery from manual copies — Due to popular demand, we are adding the ability to perform bare metal recovery from a locally attached storage device containing backups manually copied from a backup repository. This provides added flexibility and convenience by allowing you to restore systems without establishing a direct network connection to the backup repository, which may not always be feasible. **Note:** To enable this functionality, create recovery media once agent is upgraded.

Backup network detection — Optimize backup data flow and reduce backup window by creating the *AgentDirectConnectionPriority* (DWORD, 1) registry value under the *HKLM\SOFTWARE\Veeam\Veeam Endpoint Backup* key on the protected machine to instruct the backup agent to prioritize the backup repository IP address within the same subnet as the agent, optimizing network efficiency and improving backup performance through reduced latency and increased data transfer speeds. **Note:** This option is not enabled by default because it changes the current behavior, and we try to avoid such disruption in minor releases.

Hidden repository size for managed agents — We hid the repository size and remaining capacity from the backup job wizard and from the Control Panel on the managed backup agents. This should eliminate questions and concerns from end users about backup repository space which many Veeam backup administrators told us they struggle with.

Expanded GFS configuration — GFS settings now offer additional options, allowing you to specify the second, third, and fourth week for the monthly GFS setting. This helps backup administrators align long-term retention with business needs and organizational policies.

GFS restore point notification — The job action log will now specifically highlight the creation of a GFS restore point, providing better visibility into GFS restore point selection logic and facilitating troubleshooting of long-term retention policy behavior.

Enterprise Applications

General

Preferred network selection — Configuration tools for all application plug-ins now support preferred network selection and prioritization, allowing DBAs to route all backup traffic from an application server to a backup repository over a dedicated network. The preferred network settings are applied to application plug-ins only; they do not interfere with any other workloads like agent-based backups. This also helps to achieve lower RPO through improving transaction log backup frequency, as plug-ins can now identify which network interface to use much quicker.

Restore performance improvements — Minimize downtime with up to two-times faster database restores from image-level backups thanks to the implementation of asynchronous read logic specifically optimized for reading data from FUSE mounts. This enhancement should significantly improve database restore and publish performance with Veeam Explorers for Oracle, MongoDB, and PostgreSQL.

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Microsoft SQL

Microsoft SQL Server Management Studio 20 support — The application plug-in for Microsoft SQL Server is now fully compatible with the latest version of SQL Server Management Studio, allowing you to run Configuration, Backup, and Restore wizards directly from it.

Microsoft OLE DB Driver for SQL Server 19 — This new version of the OLE DB driver is now supported for application-aware image processing and Veeam Explorer *for Microsoft SQL Server* when configured with default settings in the guest OS. For the advanced options enabling connections to SQL Servers operating in the Strict mode, please refer to the technical documentation.

Oracle

Multiple Oracle homes support — Windows-based Oracle database servers with multiple Oracle homes are now supported for agent-based and host-based backups; as well as application-level backups with RMAN plug-in in both standalone and managed modes. These enhancements, combined with the same existing support for Linux-based deployments delivered in version 12.1, enable universal support for Oracle database servers with multiple homes.

Corrupted block level recovery — Following extensive QA testing, restoring corrupted blocks from RMAN plug-in backups is now officially supported. No additional configuration is required; simply run the recover block command directly in the RMAN console and the corrupted block in a data file on the production Oracle server will be replaced with a healthy one from backup.

Database authentication in Veeam Explorer — A new authentication method in Veeam Explorer *for Oracle RMAN* allows you to completely disable OS-based authentication on the production Oracle server. This method additionally improves security by enabling "per-database" level authentication, so there's no longer a need to add users to the ORA_DBA or OSDBA groups. To accommodate the new authentication type, Veeam Explorer's restore wizard will now check the OS authentication status on the target Oracle server, and if it is disabled, prompt the user to enter the SYS user password for the database.

Backup appliances

Google Cloud

Imageless appliance deployment — Deploy backup appliances directly from the backup console using standard Marketplace images to reduce complexity, improve security, and enhance the efficiency of managing your infrastructure. This method ensures a consistent environment, significantly reducing the likelihood of configuration errors and improving the reliability of your deployments.

Expanded workload support — With a new release, we expanded Cloud SQL protection support to SQL Enterprise edition and Cloud SQL on PostgreSQL v16.

Enhanced security — Core backup appliance components have been updated to the latest major versions. This includes Microsoft .NET v8, PostgreSQL v16, and Ubuntu OS 22.04.

oVirt KVM

This section applies to Red Hat Virtualization (RHV) and Oracle Linux Virtualization Manager (OLVM) platforms.

Cross-platform VM restores support — You can now effortlessly restore your physical, virtual, and cloud machines from any image-level Veeam backup directly to RHV or OLVM hypervisors. This expanded support closes the only remaining gap in our Data Portability compatibility matrix, making it all green!

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Backup Infrastructure

Scale-out Backup Repository

Enhanced migration experience — Extents from different object storage systems and/or extents with different immutability settings can now coexist within the same Performance and Capacity tier, unlocking many migration scenarios. For example, this allows you to easily move your backup from one cloud object storage provider to another, or from non-immutable bucket to a bucket with immutability enabled. To add an extent of a different type, you need to first put existing extents into the sealed mode.

Archive Tier encryption — You can now configure backup encryption directly on the Archive Tier, which can be helpful when a scale-out repository is configured to move data from Performance Tier directly to Archive Tier, bypassing Capacity Tier. Whereas previously, this scenario required enabling encryption in the backup job itself, which may not be desired on storage systems with deduplication capabilities for Performance Tier, now you have the flexibility to encrypt backups only as they are offloaded to Archive Tier.

Improved support log collection — Gathering support logs capturing scale-out repository activities such as evacuation, offloading, and archiving has been simplified. Whereas previously, this required collecting logs from all involved infrastructure servers and jobs individually, now you have an option to collect all relevant logs for the selected scale-out repository at once, reducing the number of turnarounds with your support engineer.

Storage integrations

Object Storage

Automatic bucket provisioning — Many S3-compatible object storage solutions use dedicated metadata databases for each bucket, so creating multiple buckets helps to spread the load. While a scale-out backup repository supports adding multiple buckets to Performance and Capacity tiers, there has been no solution for customers using regular object storage repositories, which routinely causes performance issues as buckets grow large, due to the metadata database becoming too large to promptly serve S3 API requests.

To address this issue, newly created standalone object storage repositories can automatically provision a new bucket for every given number of workloads protected. You can also enable this setting for existing repositories, however new backup placement rules will apply to newly protected workloads only. Managing buckets automatically requires additional CreateBucket and DeleteBucket permissions on object storage, whose presence will be validated by the Backup Repository wizard if the corresponding option is selected.

While we expect most object storage to benefit from this new capability, some advanced storage solutions perform similar load-balancing natively under the hood. Please check with your storage vendor whether they recommend enabling this option.

Background checkpoint removal enhancements — The reporting around background checkpoint removal activities added in version 12.2 resulted in many upgraded customers receiving errors caused by object storage infrastructure issues they were not aware of. V12.3 approaches this problem from a few angles. First, it significantly increases timeouts on operations that were common root causes for object storage connection failures simply due to object storage being too slow to respond, such as certificate retrieval. Second, the failure reports now offer more insights into object storage issues encountered and possible underlying causes. Finally, in light of checkpoint removal being a retriable problem as opposed to a critical backup issue, V12.3 now lowers the severity of this event from Error to Warning and adds an ability for users to make it an information event by creating *ObjectStorageCheckpointRemovalSeverity* (DWORD, 0 — info, 1 — warning, 2 — error) registry value under the *HKLM\SOFTWARE\Veeam\Veeam Backup and Replication* key on the backup server.

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Smart Object Storage API (SOSAPI) workload expansion — Nutanix AHV, Red Hat Virtualization and Oracle Linux Virtualization Manager backup jobs will now interact with SOSAPI-enabled object storage using extended SOSAPI capabilities instead of regular S3 API.

11:11 Cloud Object Storage integration — Cloud object storage from 11:11 is now integrated directly into the backup console UI to streamline the management experience for its users.

Secondary Storage

Dell Data Domain integration — V12.3 adds support for DD OS versions up to 8.1 and updates the included DD Boost SDK to version 7.13.

HPE StoreOnce Catalyst Copy for Proxmox — Catalyst-powered backup copy jobs can now copy Proxmox VM backups between Catalyst Stores. This functionality is available under experimental support terms.

Tape

Enhanced NAS Backup to Tape performance — The performance of Backup to Tape jobs has been significantly boosted for NAS backups through optimizing read patterns and caching, improving the processing speed up to a few times depending on source backup repository type, with deduplicating storage getting the biggest benefit. Note that these changes affected minimum system requirements for backup repositories, and you may not experience their full benefits if your backup infrastructure does not meet updated requirements.

Granular retries for NAS Backup to Tape jobs — Instead of retrying the entire backup job, backup jobs will now retry only failed tasks, thereby reducing load on the backup repository and improving tape utilization.

General NAS Backup to Tape improvements — Now addressed are multiple corner cases reported by customers which could lead to job failures and unexpected behaviors, making 12.3 the recommended upgrade for all NAS Backup to Tape users.

Resizable tape properties dialog — Due to popular demand, the tape media properties dialog has been made resizable to help reduce the need for scrolling.

Backup Console

Performance improvements — Several optimizations were made based on support cases to improve backup console performance and overall stability while reducing its compute resources consumption.

Configuration database logs collection — A dedicated option to include all relevant local PostgreSQL instance logs has been added to the Export Logs wizard for more efficient troubleshooting of database-related issues.

Enhanced navigation to technical documentation — The Help > Online Help option of the main menu now leads users directly to the Veeam Backup & Replication help center section, providing quicker access to relevant product documentation. This is also a good opportunity to remind users that pressing F1 anywhere in the backup console provides context-sensitive help about the current dialog.

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Enterprise Manager

Restore To option — Previously available only in the Backup Browser of the backup console, the "Restore To" option has made its way to the Veeam Backup Enterprise Manager web UI. This feature enables users to specify any Windows-based server in the environment to restore to, enhancing flexibility and operational efficiency.

Restore point selector improvements — The restore point selection control now includes a tooltip with a job name to provide users with additional information to help them make the correct selection and reduce errors.

Secure LDAP connections support — Connections from Enterprise Manager to domain controllers now support secure LDAP for added security of traffic encryption.

ISO

Streamlined configuration database management — By popular request, we've included the Database Configuration Utility Tool directly on the product ISO to simplify managing database settings, providing greater convenience, especially when troubleshooting database configuration issues.

ISO-based cumulative patches — Starting from V12.3, future cumulative patches will be delivered as ISO files to significantly increase patching speed and remove the additional disk space requirement on the backup server for unpacking a patch prior to its installation. **Note:** This may affect existing Veeam Backup & Replication patching processes.

Service Providers

VMware Cloud Director

Self-Service disaster recovery — In addition to existing backup and restore capabilities of our Cloud Director Self-Service portal, this version adds self-service disaster recovery for both VM snapshot-based replication jobs and CDP policies, enabling Cloud Director tenants to initiate failovers independently from service providers. Additionally, the improved portal enables self-service configuration of VM snapshot-based replication jobs, giving tenants greater control and flexibility in managing their disaster recovery processes.

Malware detection — Backup content scan functionality is now unlocked for Cloud Director VM backups as well, with both YARA rule-based and signature-based antivirus scans options available, including Veeam Threat Hunter. This enables service providers to deliver additional services to clients who want to ensure that their VMs are free from active malware and dormant threats.

Veeam Cloud Connect

CDP replica recoverability verification — V12.3 enables service providers and tenants to initiate test failover plans to validate their CDP replicas and ensure failover will be successful. Notably, test failovers do not affect existing CDP replication, which continues to run seamlessly during the replica testing.

Expanded restore options — In addition to performing Instant Recovery of a tenant backup to a vSphere VM, this new version allows service providers to restore those backups directly to AWS EC2 or Microsoft Azure IaaS, providing an opportunity to deliver new types of services to their clients and putting committed consumption contracts with hyperscalers to a good use.

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

API Enhancements

PowerShell

VMware vCenter migration tool — Cmdlets for mapping moRef IDs following vCenter redeployment to enable existing jobs to continue their incremental backup chains based on new moRef IDs is now a part of Veeam Backup & Replication PowerShell module.

Microsoft Entra ID backup and restore — Automate Entra ID tenant and audit log backup management with dedicated Powershell cmdlets enabling you to perform backup, browse restore points for items, compare them to production, validate, and execute restores.

Disk mapping for Instant Recovery to Hyper-V — Automatically configure disk mapping while performing VMware vSphere VMs to Microsoft Hyper-V migration using instant VM recovery capabilities.

Instant Recovery to Hyper-V cluster — Get-VBRHvServerNetworkInfo cmdlet now supports targeting Hyper-V cluster networks when performing Instant Recovery to Hyper-V.

NAS backup chain mapping — You can now point your NAS backup jobs to existing backups using the newly extended NAS backup job cmdlets. This can be helpful to automate reconfiguration of multiple NAS backup jobs in case of migration to a new backup repository.

Linux file-level recovery helper host — Automate helper host selection when performing Linux file-level recovery from your backups, snapshot-based replicas and CDP replicas.

REST API

Microsoft Entra ID backup & restore — Automate Entra ID tenant and audit log backup management with REST API endpoints enabling you to perform backup, browse restore points for items, compare them to production, validate, and execute restores.

Data Integration API — Effortlessly access data in backups from any server in the environment without the need to install any Veeam components by using new REST API endpoints to initiate iSCSI and FUSE mount of selected restore point to any specified host.

License management — Automate assigning, updating, installing, and reporting on your license consumption with new REST API endpoints for license management.

Quick backup — Easily trigger Quick Backup via REST API, for example as an automated response to receiving a malware alert from a third-party cybersecurity system.

Authorization events — Track and audit all authorization-related events, such as 4-eye authorization tasks or roles assignment through REST API.

Syslog events filtering — Automate events exclusion management to ensure only relevant events are forwarded to your syslog provider, for example, to automatically stop most events from being forwarded during backup infrastructure maintenance windows.

Backup repositories availability — Use the new dedicated repositories view to see whether your backup repositories are online and discover possible connectivity issues.

Enhanced backup statistics — Access backup job processing statistics through REST to conveniently monitor your backup health, performance and disk space occupied in the target repository.

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Missing Something?

This document includes features and enhancements first introduced **in V12.3 only**. If you are looking for information on the previous V12 releases, please refer to the following documents:

What's New in 12.2 (released August 28, 2024)

What's New in 12.1 (released December 5, 2023)

What's New in 12.0 (released February 14, 2023)

© 2024 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.