




# Veeam Backup & Replication 12.2

What's New



---

## Introduction

Veeam Backup & Replication®, the workhorse of Veeam Data Platform, delivers enterprise-grade resiliency capabilities that ensure confidence in your protection, response, and recovery in the face of both classic disaster and modern cyberattacks across the hybrid cloud. The following is a list of the major new features and enhancements added in Veeam Backup & Replication v12.2. All capabilities here are transacted as the Veeam Data Platform with certain features available only at the Advanced or Premium editions.

---

## New Datacenter Workloads Support

### Proxmox VE

Native support for host-based backups of Proxmox VE virtual machines (VMs) gives users even greater freedom of choice to best fit changing business needs for both virtualization and cloud. Key highlights of Veeam's Proxmox VE integration include:

**Immutable backups** — keep on-premises and cloud backups of Proxmox VE VMs safe and protected from encryption by ransomware, accidental deletions, or malicious destruction during cyberattacks with storage-based immutability. Or leverage native support for tape and for rotated drives to achieve true airgap at the lowest possible cost.

**Uncompromised backup performance** — powered by advanced changed block tracking (CBT) integration and HotAdd backup mode, Proxmox VE users can now leverage fast and efficient virtual machine (VM) backups, minimizing disruptions from long-running backups and ensuring smooth day-to-day operations. Uniquely, Veeam is the first to bring support for CBT backups of even powered-off VMs to market.

**Veeam BitLocker** — by automatically excluding unused disk blocks, such as those belonging to deleted guest OS files, this intelligent feature significantly enhances backup performance while conserving valuable storage space.

**Storage flexibility** — support for all Veeam Backup & Replication backup repository types, including hardened repositories and object storage, empowers businesses to easily create backup repositories from any storage that suits their RTO/RPO requirements and budget.

**Freedom of full VM recovery** — VMs from leading virtualization platforms (VMware vSphere, Microsoft Hyper-V, Nutanix AHV, Red Hat Virtualization, Oracle Linux Virtualization Manager) or public cloud (AWS, Azure and Google Cloud) can be promptly restored to Proxmox VE — and the other way around too, reducing downtime and maximizing productivity. You can also restore your physical servers backups to Proxmox VMs, simplifying disaster recovery and P2V migrations.

**Advanced granular recovery capabilities** — quickly recover all changed and deleted files, due to a ransomware attack for example, ensuring business continuity and peace of mind in the face of unexpected data loss or unplanned disruption.

## Nutanix AHV

The new major update of our Nutanix AHV integration represents a fundamental evolution of our support for enterprise AHV environments through improved centralized management, expanded functional capabilities and backup architecture simplification.

**Prism Central integration** — support your backup operations across Prism Central-managed clusters from a single backup appliance yielding a unified, easy to use backup architecture.

**Dynamic backup jobs scope** — leverage Prism Central VM categories (aka “tags”) to facilitate automatic inclusion of VMs into backup jobs based on Prism Central VM category affiliation.

**Backup from secondary VM copy** — for Nutanix replication-enabled environments, backup from the VM replica may be enabled via a single checkbox. This allows for backup operations to be offloaded from production clusters, reducing the impact on production environments. If the replica copy is out of RPO compliance or otherwise unavailable, backup processing fails back to the production VM instance.

**Multiple network adapter support** — both AHV appliance and workers may now leverage multiple network interfaces for improved data processing efficiency.

**Advanced job settings** — you can now customize job compression level and block size, as well as configure BitLocker options to further reduce backup size while improving performance.

**Gmail and Microsoft 365 support email notifications** — in addition to basic SMTP servers, V6 now supports Google Gmail and Microsoft 365 with their OAuth 2.0 protocol-based secure authorization and access-token-based authentication.

**Resources view** — the new VM inventory view of the web interface summarizes protection posture across managed clusters and enables adding VMs to existing protection policies with a single click. The view is customizable per cluster, protection status and protection type.

**RestAPI enhancements** — include Prism Central management scope enhancements and multiple networks management.

## MongoDB

Native support for backup and recovery of MongoDB on Linux, the world's most popular NoSQL database, enables you to perform backups at the per-replica set level without the need to create and manage pre-freeze and post-thaw scripts. We've designed this feature to provide a seamless, simple, and intuitive Veeam user experience while maintaining a connection with the native tool's experience (OpsManager). This means that both DBAs and backup admins can effortlessly work with our solution without extensive knowledge of either Veeam Backup & Replication or MongoDB. Key highlights of Veeam's MongoDB integration include:

**Simple discovery** — just add your MongoDB replica sets to the new, dedicated type of protection group and Veeam will automatically identify all replica set members, discover the application topology, and install the necessary backup components on each replica set node.

**Backup flexibility** — the backup policy can be populated with either individual replica sets or with entire protection groups. It also offers a unique capability to select the preferred node for data retrieval, either manually or automatically. This capability ensures that the protected data is retrieved only once, minimizing impact from backup activities on production environment.

**Storage flexibility** — Support for all Veeam Backup & Replication backup repository types including hardened repositories and object storage, empowers businesses to easily create backup repositories from any storage that suits their RTO/RPO requirements and budget.

**Proven backup engine** — under the hood, MongoDB protection uses our existing volume-level backup engine from Veeam Agent for Linux. This makes many of the features like immutable backups, synthetic full backups, GFS retention policy, encryption, backup copy jobs and more available in the first release of this new capability.

**Veeam Explorer for MongoDB** — the new addition to the Veeam Explorer's family provides the flexibility to restore individual collections and databases for day-to-day operational restores, or entire MongoDB instances for disaster recovery, both to their original or alternative locations.

## IBM Db2

**Linux on Power support** — seeking a reliable solution to safeguard Db2 workloads running on Linux on Power? Our updated Veeam Plug-in for IBM Db2 now offers its complete functionality on Power platforms as well. This plug-in operates in standalone mode and supports the same deployment types as our existing plug-in for AIX and Linux x86\_64 platforms.

**SLES 15 SP5 support** — added full support for the plug-in to operate on SLES 15 SP5 distribution. For a complete list of supported Operating System distributions, please refer to the compatibility matrix available in our Help Center.

## SAP HANA

**RHEL 9.4 support** — added full support for the plug-in to operate on RHEL 9.4 for SAP HANA and RHEL 9.4 for SAP Solutions distributions on both Linux x86 64 and Linux on Power platforms.

## Veeam Agent for Linux

**Latest Linux distributions versions support** — the new agent version adds full support for OpenSuse 15.6, SLES 15 SP6 (for x86\_64), Debian 12.4 and Fedora 40.

## VMware Cloud Director

**Full VMware Cloud Director 10.6 support** — in addition to the basic compatibility of version 12.1.2 for backup and restore only, this new release delivers full support of this Cloud Director version including the updated vSphere Web Client plug-in and Veeam CDP support.

## VMware vSphere

**Full vSphere 8.0 Update 3 support** — as opposed to compatibility-level support of vSphere 8.0 U3 provided by version 12.1.2, version 12.2 delivers full support of this vSphere version without the need to apply workarounds.

---

# New Cloud Workloads Support

## Amazon AWS

**Amazon Redshift Cluster protection** — manage your Redshift Cluster protection alongside other AWS workloads you're protecting with Veeam. Benefit from flexible policy-based scheduling and experience seamless Redshift Cluster recovery in case of a disaster.

**Amazon FSx protection** — manage your FSx file system protection alongside other AWS workloads you're protecting with Veeam. Benefit from flexible policy-based scheduling and experience seamless file system recovery to either the original or a new location.

## Microsoft Azure

### Data Lake

Protect production data residing in Microsoft Azure Data Lake (Gen2) with the powerful backup and recovery functionality built on a storage-agnostic architecture. Veeam helps you to achieve your data recovery objectives without additional hardware investments. The unique benefits of Veeam's backup engine include:

**Scalable, storage-agnostic architecture** — based on a proprietary distributed file system specifically built for the protection of billions of objects, of PBs in size, to a storage target of your choice. You get complete freedom to direct your backup to object storage or a scale-out backup repository built on standard server hardware with internal or directly attached storage.

**Efficient forever-incremental backup** — innovative "forever-incremental" engine that eliminates the need for periodic active full backups. This efficiency enables the protection of petabyte-sized data buckets with significantly reduced recovery point objectives (RPOs).

**Changed object tracking** — a unique approach to monitoring modifications delivers industry-leading incremental backup performance. You can achieve low RPOs without the need for native changed object tracking in your object storage.

And naturally, Data Lake backup jobs also come with all standard Veeam features, such as encrypted backups, backup copies, pre- and post-job scripting for automation, and multiple notification options.

### Cosmos DB

**Cosmos DB protection** — Manage your Microsoft Azure Cosmos DB protection alongside your other Azure workloads with Veeam. Enhance your PostgreSQL database protection by adding an extra backup layer on top of native Azure capabilities: safeguard your data against various types of outages and data loss by utilizing hot, cool, and archive repositories, with optional immutability through Azure Storage. Benefit from flexible policy-based scheduling and experience seamless database recovery to either the original or a new location.

## Other Features and Enhancements

In addition to the above major new features, V12.2 includes a lot of enhancements that are a response to customer feedback and ongoing R&D findings, the most significant of which are listed below:

### Direct to Archive

By popular request, scale-out backup repositories (SOBR) now support offloading aging backups directly from Performance Tier to Archive Tier, without requiring offloaded backups to land on Capacity Tier first, which can be relatively expensive in comparison to archive-class object storage. Whereas previously, Direct to Archive was only supported when Performance Tier was backed by Amazon S3 or Microsoft Azure Blob Storage-based extents, with 12.2 this configuration is supported for all on-prem backup repository types.

This new SOBR capability provides additional flexibility when configuring a scale-out backup repository for long-term backup archival in scenarios when the 3-2-1 rule compliance is achieved with a different process, such as by copying backups to another data center or exporting backups to tape. In such scenarios, you can save significant cloud object storage costs by archiving backups directly to Archive Tier and skipping Capacity Tier altogether.

However, Capacity Tier remains the recommended way of implementing the immediate copying of newly created backups to object storage for disaster recovery purposes. If you are currently leveraging Capacity Tier as your offsite backup copy strategy, do not remove it without first replacing your offsite backups copy process with a different one.

### CDP I/O Filter Cross Compatibility

V12.2 brings multiple I/O filter protocol version support, allowing I/O filters to continue communicating with upgraded backup server of later versions. This allows CDP policies to continue functioning in special compatibility mode following the product update, without requiring immediately placing all ESXi hosts into maintenance mode to perform I/O filters upgrade.

While placing ESXi hosts into maintenance mode for I/O filter upgrade is unavoidable and cannot be controlled by Veeam as the filter upgrade process is managed entirely by vCenter, now you can postpone this upgrade until the next infrastructure maintenance window — for example, when a security patch or a periodic update needs to be deployed to your ESXi hosts.

As it comes to Veeam Backup & Replication 12.2 specifically, it also supports I/O filters of versions 12.0 and 12.1. This allows organizations using backup server versions 12 and 12.1 to upgrade to version 12.2 with no disruption and choose a convenient time for updating CDP I/O filters later.

### Veeam App for Splunk

With this new Splunk extension, customers can now monitor the health and security status of their Veeam backup infrastructure using Splunk capabilities. The app enables Splunk customers to analyze Veeam events, monitor Veeam backup environments and provides access to alerts, dashboards and reports while integrating seamlessly with Splunk user roles and location management.

The app processes events sent by Veeam Backup & Replication™ and provides Splunk users with built-in dashboards and reports to monitor job statuses and security events, built-in alerts with severity level management, role-based permissions for locations, app configuration backup and more. And with multiple Veeam Backup & Replication servers and multiple data source locations support, the App is ready for the largest Veeam deployments out there.

Veeam App for Splunk supports Splunk Enterprise 9.1.0 or Splunk Cloud Platform 9.1.2308 and later versions and is available via [Splunkbase](#). The app requires Veeam Data Platform Advanced or Premium edition license.

---

## Security

### Malware Detection

**Improved Onion links detection** — based on user feedback, the onion links detection logic has been improved to reduce false positives in some corner cases such as browser cache and OS swap files. The inline malware detection engine now intelligently determines the type of file where onion links are found and excludes non-text based files like cache or swap files from consideration, as they are unlikely to represent a usable ransomware note.

**Malware detection index improvements** — the CPU consumption during the guest file system index analysis has been lowered significantly to reduce the backup server load. Further, you can now reduce the index retention period if your guest catalog storage is running out of space with the `IndexRetentionDays` (DWORD) registry value under the `HKLM\SOFTWARE\Veeam\Veeam Backup and Replication` key on the backup server with your preferred retention duration in days.

### Role-Based Access (RBAC)

**Incident API Operator role** — this new user role is designed exclusively for interaction with Veeam Incident API REST endpoint. It's ideal for automated systems or users who only need to create or manage incidents without having broader access to the backup server, which enhances the overall security posture by adhering to the principle of least privilege.

**Security Administrator role** — this new role is tailored towards security teams, enabling businesses to delegate certain sensitive tasks, such as managing saved credentials, backup encryption passwords and four-eyes authorization requests, without providing the ability to manage other backup server settings, backups and restores. By empowering designated personnel to perform these functions, you can ensure internal compliance with security best practices such as zero trust.

### Security & Compliance Analyzer

**Additional security checks** — in our ongoing effort to bolster security, we have expanded the analyzers checks to include LSASS and NetBIOS configuration on network interfaces of the backup server. These enhancements ensure your system adheres to best practices for protecting credentials in memory, securing against unauthorized access, and managing legacy network protocols.

**Additional product configuration checks** — added the recommendation against using hardened repositories as backup proxy servers due to expanded attack surface from the addition of VMware VDDK and outgoing network connections established to protected ESXi hosts.

**Auto-apply script** — the script referenced in KB4525 has been upgraded to simplify the enforcement of security and compliance measures.

## Security Events

**Backup server name** — we added the originating backup server's fully qualified domain name (FQDN) to both Windows Event Log and Syslog event parameters to facilitate event tracking in infrastructures with multiple backup servers.

---

# Image-level backup

## Backup Copy

**Direct to object from hardened repository** — Backup Copy jobs from hardened repositories will now also transfer data directly to object storage, as opposed to looping through a gateway server, thus eliminating potential bottleneck.

**Enhancements for agent-based backups** — backup copy task builder and resource locking manager have been redesigned to enable faster and more robust processing of agent-based backups, thus dramatically increasing scalability of such Backup Copy jobs.

## SureBackup

**Continuous schedule** — added a new scheduling option enabling continuous operation for SureBackup jobs within specified time windows. In conjunction with VM randomizer and parallelization limiter, this new capability enables new interesting use cases, such as continuously testing random VMs for recoverability or scanning their disks for threats or unwanted content with YARA rules during hours where the backup infrastructure is idling.

**Enhanced NSX-T support** — we've improved network mapping logic in a number of corner cases where NSX-T networks were involved.

---

# Unstructured Data Backup

**Reduced RAM consumption** — file share and object storage backup jobs processing data sources with several billions of objects should now consume significantly less RAM on cache repositories and backup proxies.



---

# Agents

## Agent Management

**Nosnap agents support** — pre-installed Linux backup agents deployed with the x86 Veeam-nosnap package are now fully supported by all agent management functionality except application processing. This allows for controlling backup settings on Linux systems that cannot use the Veeam snapshot module for whatever reason and therefore have to rely on LVM snapshots.

**Enhanced audit trail retention** — agent-based backup job session history will now be preserved for backup agents which have been removed from the backup server configuration. These historical sessions will adhere to the global session retention settings specified in the backup server properties.

## Veeam Agents for AIX and Solaris

**Flexible bare metal recovery (BMR)** — you can now restore your backup to a system with a different combination of disks from the original system or restore only selected volumes.

**IPv6 support** — the new version of Agents now supports operation in IPv6-enabled and IPv6-only networks.

**List backup content** — you can now easily list all files contained in the selected restore point using CLI, without having to mount a backup first.

**New compression algorithm** — agents will now use a modern and highly efficient compression algorithm that was first adopted in Veeam Backup & Replication v12 for High and Extreme compression levels to improve data reduction ratios by up to 20%, lower CPU usage by up to 3x and increase restore performance by up to 2x over the previously used algorithm.

**Backup performance improvements** — parallel compression and other under-the-hood tweaks should make your AIX and Solaris backups run significantly faster.

**Read-write backup mount [AIX only]** — the new ability to publish AIX backup content directly from the backup repository enables several interesting use cases, such as performing database restores using native tools directly from the mounted backup.

**ZFS encryption for BMR [Solaris only]** — when restoring a backup of an encrypted ZFS volume, we will automatically enable ZFS encryption on the restored volume as well.

## Veeam Agent for Mac

**Recovery token support** — a simplified way to provide end users with access to a particular backup has made its way to our Mac agents as well. Backup administrators are now able to generate time-limited access keys, or recovery tokens, that can be shared with users, enabling them to connect to a Veeam repository to perform a restore without the need for backup server credentials.

**Library selection option** — an additional predefined selection option is available for inclusion or exclusion from backup under the User Profiles Data section. Library typically contains fonts and other items used by apps that are available to all users of your Mac.

## Veeam Agent for Microsoft Windows

**File restore target selection** — users can now leverage the new Restore To option during file-level recoveries to specify a different Windows machine as the restore target. This dramatically simplifies restore operations when the original machine is no longer available and can facilitate complex migration scenarios by providing the flexibility to easily extract all data from backup to the desired destination.

**ReFS support in cloud machines** — file-level recovery is now supported from cloud-native agent backups of AWS EC2 instances and Microsoft Azure virtual machines with ReFS disks.

**UI branding** — in response to popular demand, we have introduced the ability to customize the Control Panel logo image and the tray icon. This feature allows you to align the user interface with your corporate aesthetics, reinforcing your brand's identity. To perform the customization, create the Logolcon (DWORD, 1) registry value under the HKLM\SOFTWARE\Veeam\Veeam Endpoint Backup key on each endpoint and place your custom logo.png and logo.ico files into the %ProgramFiles%\Veeam Endpoint Backup\Resources folder.

---

# Application plug-ins

## General

**Plug-in throttling** — you can now reduce CPU usage by application plug-ins by instructing the agent to start with a lower process priority. This grants the production application preferential access to the compute resources, thus reducing impact on production environment from backup activities during periods of high system load. Because the impact on backup performance may be noticeable, we recommend utilizing throttling only in highly loaded environments.

## Oracle

**Database authentication** — the Oracle RMAN plug-in now supports database authentication in both standalone and managed modes. This capability is particularly beneficial for Oracle environments where OS authentication is disabled. In addition to enhanced authentication capabilities, it also improves Oracle RAC processing because the previous requirement of adding the grid user to the oradba group is no longer necessary.

**Application policy enhancements** — you can now select the desired authentication type for application policies, allowing you to leverage the authentication method that best fits your environment by connecting to the Oracle instance using the specified database user, further enhancing security and control.

**Improved configuration experience** — the new, flexible configuration wizard allows you to easily specify database user credentials at per-database level and either using the same set of common credentials for multiple databases, or specifying unique credentials for each database.

**RAC processing visibility** — in response to user feedback, we expanded the application policy action log to display the specific node being utilized as a source during Oracle RAC processing, providing a better visibility into the backup process.

---

# Backup appliances

## General

**Imageless deployment** — you can now deploy cloud backup appliances directly from the Veeam Backup & Replication server to reduce complexity 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.

## AWS

**Immutability architecture improvements** — we have made some optimizations allowing users to save up to 30% of costs on immutability extensions, compared to the previous version. To further reduce costs, we will now use a default generation period up to 25 days for immutability instead of 10 days.

## Microsoft Azure

**Backup copy for virtual network configuration** — enhance your virtual network configuration protection by adding an extra layer of protection with backup copies to hot or cool repositories, with support for immutability to prevent data loss or corruption. Easily import backup copies to any Veeam Backup for Microsoft Azure appliance for hassle-free restores.

**Worker placement flexibility** — simplify your resource management by selecting the tenant, subscription, and resource group for workers deployment. For example, you can deploy workers in a dedicated subscription for added isolation, or within the same subscription as your protected resources for improved efficiency. This can help to streamline resource and cost management for backup and restore operations while ensuring adherence to internal security standards and regulations.

## oVirt KVM

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

**Restore to original location** — when performing file-level recoveries, users now have an option to restore files directly to their original location, thus avoiding many extra clicks to specify the destination. When using this option, you can also choose whether to Keep or Replace the existing files, should they still be present at the original location.

**Improved backup performance** — multiple under-the-hood enhancements for workers have been made to improve backup performance and reduce compute resources consumption.

**BitLocker enhancement** — BitLocker performance has been improved by skipping processing of small disk image chunks altogether to avoid backup slowdown on highly fragmented disks.

# Storage integrations

## Object Storage

V12.2 is a recommended upgrade to all object storage users due to multiple under-the-hood improvements and optimizations. The new version should put significantly less stress on object storage through reducing the frequency of certain API calls up to 10 times, which therefore also reduces costs with cloud object storage providers that do charge for API calls. We've also reduced gateway server memory consumption and enhanced our code's reliability in several corner cases observed in Customer Support.

**Background checkpoint removal** — checkpoint removal operation has been decoupled from the backup job and will now run as a system session after backup jobs and offload processes finish and there are no conflicting tasks (checkpoint removal has the lowest priority for the task scheduler). In addition, this process is triggered daily at 3 a.m. to perform the same activity. This change will remove additional load from object storage during the time when it is already extremely busy accepting incoming backups, and also prevent backup jobs from appearing to "hang" at the end while processing checkpoints removal.

**Intelligent extent selection** — unstructured data backups pointed to scale-out backup repositories (SOBR) will now consider both running task count and the available free space during extent selection. This more intelligent allocation will improve load distribution across SOBR extents compared to selecting them solely based on available free space.

## Primary Storage

**Immutable snapshots integration for HPE Storage Arrays** — HPE 3PAR, Primera, Alletra, and Alletra MP storage users now have one additional immutability option in their toolbox. V12.2 brings an option to utilize HPE's native Virtual Lock functionality to make the created storage snapshots immutable, offering an extra layer of security against data tampering and ransomware attacks by creating read-only storage snapshots.

**USAPI Plug-in for IBM FlashSystem** — the built-in storage snapshot integration for IBM FlashSystem has been re-implemented into the dedicated USAPI plug-in, enabling IBM to continue managing it on their own and provide more timely updates that are not dependent on Veeam Backup & Replication release vehicles. The full set of features that our users have come to rely on has been preserved in this transformation, and so is the continuity of job and snapshot history.

## Tape

**Hardware compression** — to ensure optimal use of tape media by all backup job types, hardware compression is now enabled by default for all new tape jobs. This change should enhance data transfer efficiency and maximize storage capacity without offload performance degradation.

**Source restore point monitoring** — based on a popular demand, the "no new restore points" warning will now only be logged when the source job actually fails to create a new restore point. This change should eliminate unnecessary warnings when the source backup job simply did not run yet or had been intentionally disabled.



**File Backup to Tape enhancements** — in addition to significantly improving performance of file backups to tape jobs, we've also refined the job retry logic and conflict handling in several corner case scenarios reported by our customers.

## Backup Console

**Elevated AI capabilities** — the Veeam AI Assistant has been migrated to the new GPT-4o model. This upgrade should improve response times and the quality of responses, especially as it comes to REST code creation and analysis. With GPT-4o, navigating Veeam's features, troubleshooting issues with your custom Veeam PowerShell scripts, and exploring best practices becomes more enjoyable and more intuitive than ever, with all your questions answered with significantly improved precision.

**Redesigned dashboards** — Managed Servers, Protection Groups and Object Storage dashboards have been redesigned to add the necessary grouping as to accommodate the ever-increasing number of available options due to the continuous expansion of supported platforms and storage devices.

**Improved UI performance** — we have reduced configuration databases load and improved load time and responsiveness of the following UI views and dialogs: the Home > Backup view, the Home > Jobs view, the Backup Properties dialog, the Tape Media Details dialog and the Entire Machine Recovery wizard.

**Configuration Backup and Restore improvements** — enhanced our code's reliability in several corner cases observed in Customer Support particularly around large-scale deployments of File to Tape backup jobs. These enhancements should facilitate seamless large-scale configuration database migrations from Microsoft SQL Server to PostgreSQL.

## Enterprise Manager

**One-click restore enhancements** — Enterprise Manager web UI as well as its self-service backup and restore portals now support file-level recovery for backups that do not include a system disk.

**Upgrade performance** — Enterprise Manager database upgrade performance has been improved by an order of magnitude.

This document includes features and enhancements first introduced in version 12.2 only. If you are looking for the same information on the previous V12 releases, refer to the following documents:

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

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