



CVFS Product Release Notes

Release 5.0.11

Table of Contents

CVFS 5.0.11 version support matrix	5
CVFS 5.0.11 Release Documentation	6
Highlights of the CVFS 5.0.11 release	9
Highlights of the CVFS 5.0.10 release	9
Highlights of the CVFS 5.0.9 release	10
Highlights of the CVFS 5.0.8.3 release	11
Highlights of the CVFS 5.0.8.1 release	11
Highlights of the CVFS 5.0.8 release	11
Highlights of the CVFS 5.0.7 release	11
Highlights of the CVFS 5.0.6 release	12
Highlights of the CVFS 5.0.5 release	12
Highlights of the CVFS 5.0.3.3 release	12
Highlights of the CVFS 5.0.3.1 release	13
Highlights of the CVFS 5.0.2.2 release	13
Highlights of the CVFS 5.0.1.1 release	13
Highlights of the CVFS 4.7.14 release	13
Highlights of the CVFS 4.7.13 release	13
Highlights of the CVFS 4.7.12 release	14
Highlights of the CVFS 4.7.11 release	14
Highlights of the CVFS 4.7.10 release	14
Highlights of the CVFS 4.7.9 release	14
Highlights of the CVFS 4.7.8 release	14
Highlights of the CVFS 4.7.7 release	15
Highlights of the CVFS 4.7.6 release	15
Highlights of the CVFS 4.7.5 release	15
Highlights of the CVFS 4.7.4 release	16
Highlights of the CVFS 4.7.3 release	16
Highlights of the CVFS 4.7.2 release	17
Highlights of the CVFS 4.7 release	18
Highlights of the CVFS 4.6.3 release	18
Object Store Related Enhancements	18
User Management	19
CLI Commands for HSX	19
UI Related Enhancements	19
Highlights of the CVFS 4.6.2 release	20
Installer Enhancements for RWP support	20
Highlights of the CVFS 4.6.1 release	21

REST API enhancements for 1-minute RPO feature on HSX.....	21
REST API enhancements for S3 Object Storage	22
HyperscaleX optimizations	22
Process monitoring at the storage layer for ransomware protection	22
Swatch rule modifications.....	22
Highlights of the CVFS 4.6 release.....	23
Snapshot scheduler for Block Virtual Disks.....	23
CSV option for Block Virtual Disks renamed to SCSI-3 PR.....	24
Email and SNMP Alerts for Drive Enablement and RF3 Secondary Disk utilization	25
Hedvig WebUI supports Encryption@Rest	25
Hedvig WebUI supports S3 Bucket Operations	26
Hedvig WebUI adds Migration for Container Data Mover (CDM).....	26
Phase II support added for Copy Data Management (CDM/Database dev/test).....	29
Security Advisories	30
Description	30
Affected Products	30
Resolution.....	30
Improvements and bugs resolved (4.6).....	31
Highlights of the Hedvig CSI Driver 2.0 release.....	32

List of Figures

Figure 1: New CVS option for Block Virtual Disks, SCSI-3 PR.....16

Figure 2: Details for SCSI-3 PR option16

Figure 3: Example of RF3 Usage Storage Node Status Alert17

Figure 4: Enabling Encryption@Rest for a Virtual Disk in Hedvig WebUI.....17

Figure 5: Container Data Mover - Configuration Option18

Figure 6: Container Data Mover - Migration Locations and Add Migration Location.....19

Figure 7: Container Data Mover - Add Migration Location dialog.....19

Figure 8: Container Data Mover - Informative Tooltips20

Figure 9: Solution for CDM Phase II (HCL) - Showing a deployment for a minimum configuration21

CVFS 5.0.11 version support matrix

CVFS general release	Additional CVFS products
CVFS 5.0.11	CVFS CSI Driver version 2.0
CVFS 5.0.11	CVFS VAAI NAS Plugin version 4.2
CVFS 5.0.11	CVFS vCenter Plugin version 4.3.1

CVFS 5.0.11 Release Documentation

Title	contents
<i>Hedvig Cluster Deployment Guide</i>	how to set up, install, and upgrade Hedvig Storage Cluster Nodes, Storage Proxies, and Deployment Servers. Also, how to set up for HCI, AWS, KVM, Hyper-V, Multipath I/O, and HTTPS support.
<i>Hedvig Command Reference</i>	how to use all commands in the Hedvig CLI to create and manage Virtual Disks, snapshots, and clones, as well as manage Hedvig Storage Proxies, Storage Pools, rereplication, SNMP, MIBs, and OST AIR.
<i>Hedvig CSI User Guide</i>	how to install, configure, and provision storage for the Hedvig CSI Driver, as well as how to create snapshots and clones and perform data migration.
<i>Hedvig Deduplication Design User Guide</i>	how Hedvig deduplication design, including with respect to GC (garbage collection), is implemented.
<i>Hedvig Encrypt360 User Guide</i>	how to use the Hedvig Encrypt360 feature, which provides software-based encryption that is AES-256, FIPS compliant encryption of data in flight and at rest.
<i>Hedvig MIB User Guide</i>	how to set up the Hedvig MIB and configure SNMP, including sending a test trap and all hedvigTrapsEntry traps and descriptions.
<i>Hedvig Operator User Guide</i>	how to use the Hedvig Operator, the official operator to deploy and manage Hedvig components in container orchestrators such as Kubernetes – including how to install, create a configuration for, and upgrade Hedvig components with.

Title	contents
<i>Hedvig Pensieve User Guide</i>	how to use Hedvig's Pensieve Service, which monitors the metrics for a system – providing a system snapshot to help troubleshoot Hedvig Cluster issues.
<i>CVFS Product Release Notes</i>	[this document] special notes on upgrading and installing the current release, as well as highlights of the current release and a few previous releases.
<i>Hedvig RBAC User Guide</i>	how to set up and manage Hedvig RBAC, including user, role, and tenant management.
<i>Hedvig: Resetting VMware ESXi NFS 4.1 Transfer Size</i>	how to use two hidden parameters, ForeMaxRqstSize and ForeMaxRespSize, in the ESXi advanced settings, to increase the maximum NFS 4.1 transfer size up to 1024 KB.
<i>Hedvig REST API User Guide</i>	how to use the Hedvig REST API to create and manage Virtual Disks, snapshots, and clones, as well as manage Hedvig Storage Clusters.
<i>Hedvig S3 Protocol-Compatible Object Storage User Guide</i>	how to use Hedvig's S3 Protocol-Compatible Object Storage, including how to set up configuration files, SSE, HTTPS, AWS CLI, and S3 API calls supported by Hedvig.
<i>Hedvig S3 Protocol-Compatible Object Storage User Interface Guide</i>	how to use Hedvig's S3 Protocol-Compatible Object Storage User Interface, which supports S3 Bucket operations. This UI enables basic S3 operations, such as bucket/object creation; bucket/object basic metadata management; listing objects; managing folders; uploading, downloading, and deleting objects; etc.

Title	contents
<i>Hedvig Troubleshooting Guide</i>	how to troubleshoot Hedvig Storage Proxy and Cluster Node service status, health status, and metrics, as well as changing log file settings, performing data collection, and recommended SSD specifications.
<i>Hedvig UI Installer User Guide</i>	how to install Hedvig products with a user interface-based installation process.
<i>Hedvig VAAI NAS Plugin User Guide</i>	how to use the Hedvig VAAI NAS Plugin, which enables the use of VMware vSphere Storage APIs – Array Integration (VAAI) in the Hedvig environment.
<i>Hedvig Virtualized Storage for Microsoft Hyper-V Clusters User Guide</i>	how to set up Hyper-V Failover Clusters with Hedvig, including how to configure Scale-out File Server (SOFS).
<i>CVFS vCenter Plugin User Guide</i>	how to use the CVFS vCenter Plugin, which integrates CVFS functionality into the VMware vSphere Client, a Web browser-based application that connects to the vCenter Server in order to manage vSphere infrastructure.
<i>Hedvig WebUI User Guide</i>	how to use the Hedvig WebUI to create and manage Virtual Disks, snapshots, and clones, as well as manage Hedvig Storage Clusters, Storage Proxies, Storage Pool Migration, rebalancing, and SNMP configurations.

Highlights of the CVFS 5.0.11 release

Following early access, several new CVFS features have been enabled to deliver improved resilience, performance, and operational efficiency:

Erase Coding Broadcast

Optimize CVFS erasure-coded writes through a broadcast-based data distribution method, reducing network overhead and improving write performance during high-throughput operations.

Attribute Scaling

Improves CVFS metadata performance by distributing attribute data across multiple Page replicas, minimizing hot-spots and maintaining consistent performance for workloads with large file and attribute counts.

Latency-Aware Reads

Now enabled by default, CVFS employs a data-driven latency tracker to dynamically select the fastest replica for each read, ensuring more consistent performance and reducing tail latency for backups and restores.

Space Reservation for RBL/SPM

Dynamically reserves capacity on destination pools with a configurable buffer to ensure smooth Rebalance and Storage Pool Migration. When space is constrained, CVFS intelligently skips containers to prevent disk-full issues and maintain stability.

Additional Enhancements

- Various memory optimizations and service-level tuning improve stability under heavy workloads.
- Resolved interface detection issues on Rocky8 caused by unreliable IP resolution that prevented DThrottle from initializing properly. This resulted in missing throttling on commit log operations and potential clog buildup under load. The update ensures robust IP detection and restores throttling behavior, stabilizing commit log throughput.
- Other CVFS upgrade and node rebuild bug fixes for improved reliability.

Highlights of the CVFS 5.0.10 release

- Addressed potential data bloat during inter-node rebalancing when substantial read I/O is present on the rebalancing container. This fix improves the stability and efficiency of the Read Repair mechanism under high-load conditions.
- Resolved an issue where the root user could forcibly unmount CvFS directories via alias commands. This fix enhances filesystem integrity and prevents unauthorized unmount operations.
- Addressed critical security vulnerabilities that allowed remote code execution. This

update strengthens system defenses and mitigates potential exploitation vectors.

Highlights of the CVFS 5.0.9 release

- Automatically enable Read Repair after upgrade completion and performs repair writes during I/O reads when fragments fail due to specific error classes.
- Resolved memory bloat in `ganesha.nfsd` by replacing inode-level metric tracking with fixed-size buckets.
- Supports concurrent execution of SPM and Rebalance tasks once a configurable SPM completion threshold is reached.
- Allows initiation of new Rebalance jobs even when an SPM task is active or pending.
- Automatically excludes containers from Rebalance if they have already been rebuilt by SPM.
- Fixes issues where Rebalance tasks could get stuck due to disk failures by ensuring affected containers remain eligible for reconstruction.
- Addressed bugs in the early access space reservation feature to improve prioritized space allocation for containers.
- Enhanced internal tooling for Rebalance and SPM to support diagnostics and performance tuning.
- Supports seamless end-to-end node replacement while preserving hostnames and cluster metadata.
- Enhancements and bug fixes in the rebuild and refresh node processes for HSX NFS.
- Supports end-to-end node replacement while preserving hostname and cluster attributes
- Improved RPM installation performance when upgrading from version 4.7 to 5.x.
- Pages and hblock services now wait for `local-fs.target` to ensure all drives are mounted before startup.
- Delays service startup until the network is fully available after a reboot.
- Ganesha now binds to `127.0.0.1` to optimize for local-only mounts.
- Prevents the root user from forcibly unmounting `cvfs` directories via aliases.
- Ensures hblock `systemdb` backups maintain accurate generation numbers.
- `whichCommit` output now reflects consistent release versions across all CVFS modules.
- Non-disruptive upgrades prioritize updating nodes that are currently down, minimizing operational impact.

Highlights of the CVFS 5.0.8.3 release

- Fixed the problem with Hblock service not coming up after abrupt shutdown.

Highlights of the CVFS 5.0.8.1 release

- Restoration of correct RWP context during drive failure handling.

Highlights of the CVFS 5.0.8 release

- Enhanced write performance using broadcast (Early Adopter).
- Clog is now synced using O_DIRECT.
- Reduced hblock logging.
- Health Check Service enhancements to monitor soft disabled and unmounted pools, check if SPMs are triggered for unmounted disks and monitor RF3 growth.
- Enhancements to visualize progress of housekeeping tasks such as SPM and RBL.
- Enhancement to detect batch disk failure.
- Enhancement to include CVFS version number in rpms.
- Enhancement to maintain CVFS upgrade history.
- Enhancement to ensure n/w interfaces are up during boot before starting CVFS services.
- Enhancement to perform read repair fragments in case of a missing fragment.
- Implemented a fix for add node operation to existing HSX cluster. Adds host IP instead of SID when hostname resolution fails.
- Implemented a fix for nested disk failures that cause SPMs to get stuck.
- Implemented a fix for updatecontainerreplicas that should validate storage sid before updating container replicas.

Highlights of the CVFS 5.0.7 release

- CVFS S3 bug fixes in the bucket/object deletion code paths.
- Process stop/restart/os-shutdown with systemd KillSignal SIGTERM are caught by hblock and pages services and appropriately initiates graceful shutdown of core components absorbing incoming IO and commit log cleanup.
- Read path improvements to achieve loading and reading data of containers distributed across multiple storage pools.
- Active Read Repair feature to read and repair data from all erasure coded fragments.
- Miscellaneous enhancements to developer cli tools and HSX collect.

- Latency aware read improvements.
- Support for cvbackupadmin user commands.
- CVFS log file archiving improvements.
- hv_deploy command fixes for returning cluster information.

Highlights of the CVFS 5.0.6 release

- Fixed a non-disruptive upgrade bug where hblock might hang on gossip.
- Fixed a log disk full issue caused by daily archival of old pages and hblock log files.
- Feature: Space reservation on target pools for RBL & SPM. This feature is disabled by default and can be enabled by CLI.
- Send out an alert to customers for upgrading smartpqi driver on HP.

Highlights of the CVFS 5.0.5 release

- Add extra security headers to the CvFS REST API response.
- Implemented stack canary and read-only relocation in binaries for increased security.
- Fix for one way thrift issue causing improper updates to container replicas in metadata subsystem causing container data to spread out onto non-replica nodes **incorrectly**.
- Passwords in CvFS are hashed with a salt for enhanced security.
- Fix security vulnerability due insecure password handling in scripts/libraries.
- Optimized CvFS pages and hblock service memory usage under heavy system load.
- Handle error cases in lock/unlock operations from zookeeper more efficiently
- Fix security vulnerabilities in SNMP API.
- Additional rebalance tools options for dev usability.
- Minor bug fixes and optimizations in rebalance and storage pool migration workflows.
- CvFS S3 - Copy object API functionality improvements.
- Support for Rebuild node workflow to rebuild hblock data in the event of node refresh/replacement.
- CvFS S3 - APIs to query bucket status and error code differentiation for - marked for delete, exists, and quota full cases.

Highlights of the CVFS 5.0.3.3 release

- Fixed NDU upgrade failure.

Highlights of the CVFS 5.0.3.1 release

- Improved drive replacement handling when slot cannot be identified.
- Implemented opt-in feature for scaling out the attributes on metadata nodes for better performance.
- Implemented multi-staged delete handling to process delete requests in parallel.

Highlights of the CVFS 5.0.2.2 release

- Implemented a hook to reduce memory footprint in memory pressure situations.
- Implemented optimized method to identify and correct duplicate indices on disk.
- Deleter is now set to 2 nodes by default.
- Free up space on source disk during rebalances in batches.
- Implemented switch to disable fetching rebalance history.
- Implemented multi staged delete in block service for parallel delete processing.

Highlights of the CVFS 5.0.1.1 release

- Support for Rocky Linux 8 and RHEL 7.9 Operating systems.
- Python version upgraded to 3.11.
- Implemented updates to underlying libraries boost (1.72), openssl (1.11), and gcc (8.3).

Highlights of the CVFS 4.7.14 release

- Dynamic throttling service is auto started.
- Implemented a switch to perform duplicate index validation on service start.

Highlights of the CVFS 4.7.13 release

- In case of drive failure, storage pool migration will be retried for 24 hours in case of errors.
- Implemented optimization to fingerprint and send unique files during rebalancing restarts.
- Fixed replication redundancy block bloat in case of write failures when the cluster is running in degraded mode.
- Improved disk rebuild in case of disk failure.
- Fixed the issue with duplicate index that cause rebalancing to stall due to restart of rebalancing.
- Fixed the issue with data bloat during read repairs of read modified write conditions.

- Implemented memory optimizations in the Metadata system.

Highlights of the CVFS 4.7.12 release

- Fixed read failures when more than one fragment is inaccessible.
- Fixed Read failures when the target node for rebalancing data is down.
- Fixed backup timeouts on unaligned writes.

Highlights of the CVFS 4.7.11 release

- Optimize caching of attributes in case of HyperScale deployment
- Implement CLI to configure streamer threads to allow tuning of intake of the rebalanced data
- Implement CLI to configure designated nodes to participate in delete operation
- Fix rebalancing stall in case of PVC disk with null mount location
- Fix storage pool migration stall in case of null virtual disk information

Highlights of the CVFS 4.7.10 release

- Fix hedvigfsc service crashes due to scanning of port 4321 by external port scanners
- Fix null check for PVC disks in RBL to prevent RBL stall
- Do not allow ingest and other housekeeping operations when metadata drives exceed 90% capacity usage in order to prevent metadata corruptions in certain conditions
- Fix issues with disk rebalancing operation due to hung socket connections

Highlights of the CVFS 4.7.9 release

- Implement CLI tool command to enable/disable rebalancing throttle parameter
- Implement CLI tool command to resume a particular rebalance ID forcefully
- Implement CLI tool command to rebalance RF3 containers by count
- Implement read operation optimization for EC 8:4 deployments

Highlights of the CVFS 4.7.8 release

- Fix for NFS server, where when NFS Server is loaded and the system is under stress, in a higher latency situation, certain files may report the file size to be smaller than the actual size.

- Implemented the TTL logic for hblock commit logs using the priority queue for memtable flushes.
- Fixed a corner condition where healing of data from RF3 to EC may result in data consistency issue with very small amount of healed data.
- Added the ability to preserve site-specific tunable for memory across upgrades.

Highlights of the CVFS 4.7.7 release

- The problem related to Hedvigfsc service crash is fixed.
- Implemented the CLI tool to rebuild the metadata disk.

Highlights of the CVFS 4.7.6 release

- Unicode characters (except for certain special characters) in file names are now supported.
- The ability to override default memory caps for CVFS processes has been added.
- Boundary conditions for pending deletes are now handled.
- Validation of usernames in CVFS processes has been made more secure.
- Stronger passwords for CVFS users are now required.
- The existing refresh node now supports *full node rebuild* when a complete machine is lost.

Highlights of the CVFS 4.7.5 release

- Improvements have been made to aggressively enforce-flush and async-flush to prevent the commit-log directory from becoming full and recover if full.
- Storage pool migration target selection has been improved.
- RF5 for pages is now supported by the Ansible file.
- Previously, the disk replacement process would keep the hblock service down in case replacement itself failed (and needed intervention). The hblock process is now kept up.
- Changes have been made to libthrift for the hedvigfsc process to handle security scanners.
- A problem with the templates used for the CVFS user interface has been resolved.

Highlights of the CVFS 4.7.4 release

- Enhancements were made to the internal housekeeping tools for more fine-grained configuration control.
- Memory caps are now automatically set for CVFS processes as part of the upgrade.

Highlights of the CVFS 4.7.3 release

- The CLI Tool command was added to handle the Deleter Task.
- Improvements were made in read operations done during rebalancing.

Highlights of the CVFS 4.7.2 release

- The NVMe replacement for HSX has been automated.

To run this playbook, do the following, as user admin.

```
env HV_PUBKEY=1 /opt/hedvig/bin/hv_deploy --replace_hsx_nvme
--cluster_name $cluster
```

This script must be run after NVMe replacement is done and the drives formatted.

- The HSX debug playbook collection has been automated using hv_deploy.

To run the HSX debug playbook, type:

```
env HV_PUBKEY=1 /opt/hedvig/bin/hv_deploy --hsx_debug_playbook
--cluster_name $cluster
```

This runs a series of commands on each individual node in the cluster. These commands are listed in:

```
/u/tools/bin/mk_hsx_debug_playbook.pl
```

When run, this pipes (into the current directory) the output files corresponding with each node, for example:

```
hostname_hsx_debug_report.smc2300-8d.nj.hedvig.commvault.com.txt
```

and displays the names of all the files that it has generated.

- Support has been added for Copy Data Management Deployment.

Copy Data Management (CDM) deployment is supported with the following command:

```
export HV_PUBKEY=1; /opt/hedvig/bin/hv_deploy
--cdm_deploy ansible-filename
```

For detailed syntax, refer to the sample file for CDM deployments.

- Drive replacement fixes have been made for RWP environments with corrected return codes.
- The `fstab` options have been synced for replace drives with HSX standard.
- Robust handling for the HSX refresh node process has been added to save state on other disks.
- Scripts for non-HSX environments have been added to clean up the Commvault registry after vdisk deletion (`rm_hs_vdisk.pl`).
- Security vulnerability for port 8778 / 8777 has been fixed as part of the start-up process.

Highlights of the CVFS 4.7 release

- Web Server-Generated CORS Headers have been disabled for HSX deployments.
- Timeout values have been optimized for synchronization of finger-printed files.
- Overrides have been added to the config files to lower the memory usage of the Metadata process.
- The `start-server.sh` script has been updated to leverage default settings for capping the Direct Byte Buffer limits.

Highlights of the CVFS 4.6.3 release

- *Object Store Related Enhancements.*
- *User Management*
- *CLI Commands for HSX*
- *UI Related Enhancements*

Object Store Related Enhancements.

- Object lock functionality now has the ability to take `Year` as a parameter.
- There is a new REST API for `ResizeBucketQuota`.

User Management

- You are now able to add an individual LDAP/AD user with @ in the username. Also, you are not restricted to having to specify an email to begin the process.
- There is a fix for LDAP users to show the Edit Profile page after logging in.
- There is a fix for the display when multiple users are added to a tenant.

CLI Commands for HSX

A new CLI command has been added to get container sizes.

Here is an example:

```
Welcome to CVDS.
```

```
Type 'showcommands' or '?' for help. Type 'exit' or 'quit' to quit.
prodmsds> getallcontainersize
```

Container Name	Container Size
__ecBackingVDisk_NFSFILE_HedvigStorageEC\$1009	101.4 GB
__ecBackingVDisk_NFSFILE_HedvigStorageEC\$1031	108.5 GB
__ecBackingVDisk_NFSFILE_HedvigStorageEC\$1040	102.0 GB
__ecBackingVDisk_NFSFILE_HedvigStorageEC\$1050	101.9 GB

UI Related Enhancements

- The CVLT Theme is now the default color scheme.
- The default for Cluster Watch is now List View.

Highlights of the CVFS 4.6.2 release

Installer Enhancements for RWP support

- The `cron` utility previously used to check *process out of context* has been removed. (This check existed only in 4.6.1, not in any releases prior to 4.6.1 or after 4.6.1.)
- The script `fix_post.sh` can be used to correct existing environments that have 4.5.3 or 4.6.1 installed.
- The following files – in `/usr/local/hedvig/hblock` – can have context set in the case of 4.5.3 or 4.6.1:

```
/usr/local/hedvig/hblock/start-server.sh
```

```
/usr/local/hedvig/hblock/start-server.sh.tmpl
```

```
/usr/local/hedvig/hblock/start-server.sh.prod.tmpl
```

```
/usr/local/hedvig/hblock/start-server.sh.prod.cvlt.tmpl
```

Highlights of the CVFS 4.6.1 release

- *REST API enhancements for 1-minute RPO feature on HSX*
- *REST API enhancements for S3 Object Storage*
- *HyperscaleX optimizations*
- *Process monitoring at the storage layer for ransomware protection*
- *Swatch rule modifications*

REST API enhancements for 1-minute RPO feature on HSX

In HSX (and Hedvig as a backup target) deployments, the **hedvigfsc** service exposes a REST API interface, which includes capabilities such as the distributed lock service.

This REST API interface has been enhanced with the following APIs for the 1-minute RPO feature:

1. **GetFolderList(path)** - Returns all the directory names in the given path.
2. **GetFileList(path)** - Returns all the files (including directories) in the given path.
3. **GetFileDetailsList(path)** - Returns all the files (including directories) and the corresponding file attributes in the given path.
4. **GetFileInfo(path)** - Returns the file/directory attributes corresponding to the given path.

For 1, 2 and 3, a paginated version of the REST API has also been provided

REST API enhancements for S3 Object Storage

SendSecurityToken REST API backwards compatibility fixes have been included in this release.

As a part of these fixes, the response will include the active security token in the HTTP response. In addition, the security token will be sent to the given user's email address.

HyperscaleX optimizations

In clusters running close to full capacity, the healing process (for RF3 data) has been optimized.

Process monitoring at the storage layer for ransomware protection

In scenarios where processes at the storage layer lose context, resulting in I/O errors, cron-based monitoring has been implemented for process contexts when ransomware protection is enabled.

This monitoring guards against currently open issues with ransomware protection at the MM layer, which will be fixed as a part of the April release.

Swatch rule modifications

Swatch rules have been modified to ignore spurious, non-critical xfs error messages.

Highlights of the CVFS 4.6 release

- *Snapshot scheduler for Block Virtual Disks*
- *CSV option for Block Virtual Disks renamed to SCSI-3 PR*
- *Email and SNMP Alerts for Drive Enablement and RF3 Secondary Disk utilization*
- *Hedvig WebUI supports Encryption@Rest*
- *Hedvig WebUI supports S3 Bucket Operations*
- *Hedvig WebUI adds Migration for Container Data Mover (CDM)*
- *Phase II support added for Copy Data Management (CDM/Database dev/test)*
- *Security Advisories*
- Systemic enhancements moving to NFS Ganesha release 3.3
- Enhancements to support DR at clone for Container Data Mover (CDM) functionality
- Performance Optimizations for K8S and OpenShift Filer workloads
- Compression no longer implicitly set when enabling Deduplication feature
- *Improvements and bugs resolved (4.6)*

Snapshot scheduler for Block Virtual Disks

CVFS now supports the scheduling of snapshots for Block Virtual Disks.

The snapshots taken by the scheduler are **crash-consistent** snapshots.

If a Virtual Disk is scheduled for snapshots, **manual create, revert, or delete snapshot operations are not supported.**

To enable this feature, you must create an *SLA (service level agreement) Domain* and then associate a Virtual Disk with that SLA Domain.

CSV option for Block Virtual Disks renamed to SCSI-3 PR

The CSV (Cluster Shared Volumes) option for Block Virtual Disks has been renamed to **SCSI-3 PR**. This option should be enabled if the application that uses this disk relies on SCSI-3 Persistent Reservations to manage the ownership of the disk.

Add New Virtual Disk

Batch

Name

hyprCSV1

Size

10

TB

Disk Type

BLOCK

Migration

Encryption

Enable Deduplication

Clustered File System

Description

Compressed

Client-side Caching

SCSI-3 PR

Run

Replication Policy

RackAware

Replication Factor

3

Block Size

4096

Residence

HDD

Flash

Figure 1: New CVS option for Block Virtual Disks, SCSI-3 PR

Details

Containers

Name: testscsi3pr

Block Size: 4k

Cache Enabled: No

Compressed: No

Controllers: Unassigned

Created By: HedvigAdmin

Creation Time: 31-12-2021 21:33:04:976

Deduplication: No

Description:

Disk Type: block

Encryption: No

Encryption@Rest: No

Exported Block Size: 4k

File System: Non-clustered

Frozen: No

Migration: No

Protection Policy: Replication

Replication Factor: 3

Replication Policy: RackAware

Residence: HDD

SCSI Serial Number: 242

SCSI-3 PR: Yes

Size: 10 GB

Tenant: Hedvig

Version Counter: 1

Workload Type: NONE

Figure 2: Details for SCSI-3 PR option

Email and SNMP Alerts for Drive Enablement and RF3 Secondary Disk utilization

Email and SNMP alerts for drive enablement and RF3 secondary disk utilization are now automatically provided.



Figure 3: Example of RF3 Usage Storage Node Status Alert

Hedvig WebUI supports Encryption@Rest

You can now enable Encryption@Rest when you create a Virtual Disk in the Hedvig WebUI.

For more details on creating Virtual Disks, see the *Hedvig WebUI User Guide*.




Figure 4: Enabling Encryption@Rest for a Virtual Disk in Hedvig WebUI

Hedvig WebUI supports S3 Bucket Operations

The Hedvig WebUI now supports S3 Bucket Operations for Hedvig's S3 Protocol-Compatible Object Storage. The following S3 operations are now enabled:

- Bucket and object creation
- Bucket and object basic metadata management
- listing objects
- managing folders
- uploading, downloading, and deleting objects

Hedvig WebUI adds Migration for Container Data Mover (CDM)

You can now enable the Migration option for the Container Data Mover when you create a Virtual Disk in the Hedvig WebUI.

Note: For more details on creating Virtual Disks, see the *Hedvig WebUI User Guide*.

When this option is enabled, extensive information regarding the selected Migration Location is displayed in a tooltip.

Here are the steps to use the Container Data Mover feature.

1. Each Virtual Disk that will be migrated must be associated with a Migration Location. To create a Migration Location, select the **Configuration** menu, and then select **Migration Locations** and **Add Migration Location**.




Figure 5: Container Data Mover - Configuration Option




Figure 6: Container Data Mover - Migration Locations and Add Migration Location

- 2. In the **Add Migration Location** dialog, chose the Migration Location **Name** and **Replication Policy** to be used in the destination cluster. The cluster composition is chosen automatically from the seed names provided during cluster setup.




Figure 7: Container Data Mover - Add Migration Location dialog

- When you are creating the related Virtual Disk, select the **Migration** option, and a dropdown list of locations is displayed. Extensive information regarding the selected **Migration Location** is displayed in a tooltip.




Figure 8: Container Data Mover - Informative Tooltips

Phase II support added for Copy Data Management (CDM/Database dev/test)




Figure 9: Solution for CDM Phase II (HCL) - Showing a deployment for a minimum configuration

Phase II support for Copy Data Management (CDM/Database dev/test) has been added.

Technical assumptions include:

- CDM workload: iSCSI [RMAN data in native format], NFS [Transaction Logs in CV Chunk Format].
- Encryption@Rest support added.
- Backup: NFS, EC (4,2) + Compression [scope limited to just the Transaction logs as part of CDM solution, and cluster does not account for any other type of backup workloads].
- iSCSI Data will have support for RF2 and RF3 at a per Virtual Disk level.
- Instant Mount support irrespective of Size/Version depth will be provided via Shallow/Thin clones.
- Fast restore from a snapshot will be provided via Full clone implementation. Dedicated test setup with real world datasets to test out efficacy of solution.
- H/W Configuration supported HSX43504. [168 TB [12*14] RAW capacity, 512 GB RAM, 2*6.4TB NVMe]
- Storage tier will not enable Deduplication on any workloads. In use cases, E2E encryption is a requirement. Compression will not be enabled.
- Encryption support with any KMIP compliant vendor. Out-of-the-box integration with KeySecure-KV190, Fortanix.
- Minimum configuration is 4 nodes [NFS will run on 2 nodes with MA and 1 iSCSI HA pair].

Technical assumptions (continued):

- CDM release for 11.26 will be released with CVFS 4.6.
- Dedicated test setup with real world datasets to test out efficacy of solution.
- Solution barring the install process will be common for both CVLT as backup target and HSX deployments.

Security Advisories

CV_2021_12_1: Vulnerability in Apache Log4j Logging Libraries Impacting Commvault Products

Advisory ID: CV_2021_12_1

External Reporting IDs: CVE-2021-44228, CVE-2021-45046

Issued On: December 11, 2021

Updated On: December 11, 2021

Severity: Critical

Version: 2.0

Description

A critical vulnerability has been found on Apache Log4j logging libraries. For more information about this vulnerability, refer to the following report:

[CVE-2021-44228](#): Apache Log4j2 JNDI features do not protect against attacker-controlled LDAP and other JNDI related endpoints.

[CVE-2021-45046](#): Apache Log4j 2.15.0 was incomplete in certain non-default configurations.

Affected Products

This vulnerability may affect the following products:

- Commvault Distributed Storage
- Commvault Hyperscale X

Resolution

We have upgraded to the Apache Log4j release 2.17.1.

Improvements and bugs resolved (4.6)

Category	Description
Bug	Graceful handling during freeing of file attribute when the NFS server receives an unlink() for a file that was open and being read.

Highlights of the Hedvig CSI Driver 2.0 release

This is a summary of new features in the Hedvig CSI Driver 2.0 release. For more detail, see the *Hedvig CSI User Guide*.

- **Secrets Support** – The Hedvig CSI Driver authenticates with a Hedvig Cluster on startup using *Kubernetes Secrets*.
- **Volume Expansion** – The Hedvig CSI Driver enables users to expand dynamically provisioned persistent volumes for **hedvig-block** and **hedvig-nfs** backend types.
- **Raw Block Volume** – The Hedvig CSI Driver enables users to dynamically provision persistent volumes that will appear as Block devices inside containers. This feature applies only to the **hedvig-block** backend type.
- **Scheduled Snapshots** – The *Snapshot Scheduler* enables users to schedule periodic snapshots of dynamically provisioned persistent volumes through schedule policies. This feature applies only to the **hedvig-block** backend type.
- **Provisioning of Agnostic/Rack Aware volume with single data center** – If a Hedvig Storage Cluster is configured to support both Rack Aware and Data Center Aware replication policies, you can create an Agnostic or Rack Aware Virtual Disk with replicas confined to a *single* data center.
- **Container Data Mover** – For Kubernetes environments, the *Container Data Mover* feature enables the streaming of persistent volume data from a source cluster to a target cluster.

Commvault Systems, Inc., believes the information in this publication is accurate as of its publication date. The information is subject to change without notice. The information in this publication is provided as is. Commvault Systems, Inc., makes no representations or warranties of any kind with respect to the information in this publication and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any Commvault Systems, Inc., software described in this publication requires an applicable software license. All trademarks are the property of their respective owners. Revision date: 030423.

Software-defined AES-256, FIPS compliant encryption of data in flight and at rest.