



Lightbits and Cinder Backup for OpenStack

Implementation Guide for OpenStack Cinder Backup with Lightbits Labs

Author: Rob Bloemendal

February 2026

Abstract

This white paper explores how Lightbits optimizes OpenStack environments by providing a high-performance, NVMe[®]-over-TCP storage layer that integrates natively with the standard Cinder driver—requiring zero custom hooks or manual configuration changes. For DevOps teams, this translates to a "storage-as-code" experience in which the entire volume lifecycle—including snapshots, clones, mounts, and automated resource reclamation—is managed through standard API calls without additional backend overhead. By abstracting the complexity of data management and offering a "zero-effort" restore path, Lightbits eliminates the traditional friction of storage operations, allowing engineers to focus on scaling infrastructure and accelerating CI/CD pipelines rather than managing storage silos.



Table of Contents

1. Introduction.....	3
1.1. Cinder Backup Configuration.....	3
1.1.1. Modifying cinder.conf.....	3
1.1.2. Service Initialization.....	3
2. Backup.....	4
3. Restore.....	6
4. Under the Hood: Lightbits Integration Logic.....	9
5. Conclusion.....	11
About Lightbits Labs™	12

1. Introduction

This white paper outlines the architectural implementation and operational procedures for backing up OpenStack Cinder volumes hosted on Lightbits NVMe-over-TCP storage to an S3-compatible object store (specifically utilizing Garage S3).

1.1. Cinder Backup Configuration

To facilitate the data transfer between the Lightbits storage backend and the S3 target, the cinder-backup service must be configured with the appropriate driver and endpoint metadata.

1.1.1. Modifying cinder.conf

Navigate to `/etc/cinder/cinder.conf` and ensure that the following parameters are defined within the `[DEFAULT]` section to initialize the S3 driver:

```
# Enable the S3 backup driver
backup_driver = cinder.backup.drivers.s3.S3BackupDriver
```

In this guideline, we are using the Garage S3 object store. The object store must also be configured. The configuration in the `[DEFAULT]` setting in this example is:

```
backup_s3_endpoint_url = http://192.168.1.216:3900
backup_s3_store_access_key = GK89bf439c5971daabb685ea75
backup_s3_store_secret_key =
32c1dab9605c18640781c1af91f7720d9684bb417a0db1aa7de9fa46167562b6
backup_s3_container = backup
backup_s3_region_name = garage
backup_s3_verify_ssl = False
driver_ssl_cert_verify = False
backup_s3_auth_version = 4
backup_s3_use_path_style = True
```

1.1.2. Service Initialization

Restart the Cinder services to apply the configuration changes. It is recommended to enable the backup service and restart the API to ensure full registration within the OpenStack environment.

```
Shell
sudo systemctl restart cinder-*
```



Now that the Cinder services have been restarted, the Cinder-backup must be started and enabled, and immediately after that, the Cinder API must be restarted.

```
Shell
sudo systemctl enable openstack-cinder-backup
sudo systemctl start openstack-cinder-backup

sudo systemctl restart openstack-cinder-api
```

Verify that the cinder-backup agent is successfully registered and reporting an up state:

```
Shell
openstack volume service list --service cinder-backup
```

Output:

```
Shell
+-----+-----+-----+-----+-----+-----+-----+-----+
| Binary      | Host   | Zone | Status | State | Updated At           | Cluster | Backend State |
+-----+-----+-----+-----+-----+-----+-----+-----+
| cinder-backup | devstack | nova | enabled | up   | 2026-02-20T12:44:00.000000 | None    | None           |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

2. Backup

First we need to select which volume we want to be backed up:

```
Shell
openstack volume list
```

Output:

```
Shell
+-----+-----+-----+-----+-----+-----+
| ID          | Name   | Status | Size | Attached to |
+-----+-----+-----+-----+-----+-----+
| 0cbd7795-aebe-4759-b9c9-a00a83eb9307 | vol1 | in-use | 10 | Attached to Demo on /dev/vdb |
+-----+-----+-----+-----+-----+-----+

```



To create the backup for the volume shown above:

```
Shell
openstack volume backup create --name lightbits-vol-backup vol1 --force
```

Output:

```
Shell
+-----+-----+
| Field   | Value                                     |
+-----+-----+
| id      | a530cf4e-6114-42fe-94c9-5c936c492f83 |
| name    | lightbits-vol-backup                   |
| volume_id | 0cbd7795-aebe-4759-b9c9-a00a83eb9307 |
+-----+-----+
```

The `--force` at the end of the command has to be used when the status of the volume is "In-use". With the `--force` option, the backup created is crash-consistent. To make a clean backup, the volume must be in status "Available". This guideline is not intended to describe the difference between crash-consistent and clean backups; rather, it explains how to make the backup happen.

To check that the backup is progressing, type:

```
Shell
openstack volume backup list
```

Output:

```
Shell
+-----+-----+-----+-----+-----+-----+-----+
| ID              | Name           | Description | Status  | Size | Incremental | Created At          |
+-----+-----+-----+-----+-----+-----+-----+
| 021fdecb-092e-4761-af81-f2974102f4b5 | lightbits-vol-backup | None       | creating | 10 | False | 2026-02-20T15:22:15.000000 |
+-----+-----+-----+-----+-----+-----+-----+
```

When the backup is complete, the status is changed to "available":

```
Shell
+-----+-----+-----+-----+-----+-----+-----+
| ID              | Name           | Description | Status  | Size | Incremental | Created At          |
+-----+-----+-----+-----+-----+-----+-----+
```

```
| 021fdecb-092e-4761-af81-f2974102f4b5 | lightbits-vol-backup | None | available | 10 | False | 2026-02-20T15:22:15.000000 |
```

3. Restore

In our example, the following directory structure is here on the mounted volume on `/dev/vdb`:

```
Shell
[almalinux@demo fio]$ ls -lah
total 2.5G
drwxr-xr-x. 3 almalinux almalinux 4.0K Feb 20 10:41 .
drwxr-xr-x. 3 root      root      17 Feb 20 10:23 ..
-rw-r--r--. 1 almalinux almalinux 2.5G Feb 20 10:42 fio.dsk
drwx-----. 2 root      root      16K Feb 20 10:22 lost+found
```

The `fio.dsk` file will be deleted:

```
Shell
[almalinux@demo fio]$ rm /mnt/fio/fio.dsk
[almalinux@demo fio]$ ls -lah
total 20K
drwxr-xr-x. 3 almalinux almalinux 4.0K Feb 20 15:38 .
drwxr-xr-x. 3 root      root      17 Feb 20 10:23 ..
drwx-----. 2 root      root      16K Feb 20 10:22 lost+found
```

To restore the volume, it must be detached from the instance. In this case, the instance is `Demo` and the volume is `vol1`.

```
Shell
openstack server remove volume Demo vol1
```

Verify that the volume has the status “available.”

```
Shell
openstack volume list
```



Output:

```
Shell
+-----+-----+-----+-----+
| ID                | Name | Status  | Size | Attached to |
+-----+-----+-----+-----+
| 0cbd7795-aebe-4759-b9c9-a00a83eb9307 | vol1 | available | 10 |             |
+-----+-----+-----+-----+
```

The restore can now begin. To find the id of the backup type:

```
Shell
openstack volume backup list
```

Output:

```
Shell
+-----+-----+-----+-----+-----+-----+-----+
| ID                | Name           | Description | Status  | Size | Incremental | Created At           |
+-----+-----+-----+-----+-----+-----+-----+
| 021fdecb-092e-4761-af81-f2974102f4b5 | lightbits-vol-backup | None       | available | 10 | False       | 2026-02-20T15:22:15.000000 |
+-----+-----+-----+-----+-----+-----+-----+
```

The restore command is as follows: (note that `--force` will be used, because the volume already exists):

```
Shell
openstack volume backup restore lightbits-vol-backup vol1 --force
```

Output:

```
Shell
+-----+-----+
| Field      | Value                |
+-----+-----+
| id         | 021fdecb-092e-4761-af81-f2974102f4b5 |
| volume_id  | 0cbd7795-aebe-4759-b9c9-a00a83eb9307 |
| volume_name | vol1                 |
+-----+-----+
```

Check the status of the backup:

```
Shell
openstack volume show vol 1 -c status
```

Output:

```
Shell
+-----+-----+
| Field | Value          |
+-----+-----+
| status | restoring-backup |
+-----+-----+
```

The volume must be in the “available” status before it can be re-attached to the instance.

```
Shell
+-----+-----+
| Field | Value          |
+-----+-----+
| status | available      |
+-----+-----+
```

The volume can now be attached to the server:

```
Shell
openstack server add volume Demo vol1
```

Output:

```
Shell
+-----+-----+
| Field          | Value          |
+-----+-----+
| ID             | 0cbd7795-aebe-4759-b9c9-a00a83eb9307 |
+-----+-----+
```

```
| Server ID           | 426a074d-91e5-4cdf-97ee-847d39c95756 |
| Volume ID          | 0cbd7795-aebe-4759-b9c9-a00a83eb9307 |
| Device              | /dev/vdb                               |
| Tag                 | None                                    |
| Delete On Termination | False                                  |
+-----+-----+
```

On the server, the volume needs to be mounted again. Check the directory:

```
Shell
ls -lah /mnt/fio
```

Output

```
Shell
[almalinux@demo fio]$ ls -lah
total 2.5G
drwxr-xr-x. 3 almalinux almalinux 4.0K Feb 20 10:41 .
drwxr-xr-x. 3 root      root      17 Feb 20 10:23 ..
-rw-r--r--. 1 almalinux almalinux 2.5G Feb 20 10:42 fio.dsk
drwx-----. 2 root      root      16K Feb 20 10:22 lost+found
```

The file has been restored.

4. Under the Hood: Lightbits Integration Logic

When the OpenStack volume backup create command is issued, the Cinder driver coordinates with the Lightbits cluster to perform the following automated steps.

1. **Snapshot Creation:** A temporary snapshot is created on the Lightbits cluster to serve as a point-in-time source.

```

demo@lb-01:~
X demo@lb-01:~ (ssh)
Every 2.0s: lbcli list snapshots --project-name acme                               lb-01: Fri Feb 20 17:02:14 2026

Name                               UUID                               Source volume UUID                               State
for_clone_748bc257-2318-480c-9b5d-e2fcaefd2368  6601c462-f5ac-4107-aeef-aaaffcd59035  2857de2a-9dca-48de-af54-ccb6f9992878  Creating

X demo@lb-01:~ (ssh)
Every 2.0s: lbcli list volumes --project-name acme                               lb-01: Fri Feb 20 17:02:15 2026

Name                               UUID                               State      Protection State  NSID   Size  Repli
cas  Compression  ACL                               UUID                               State      Protection State  Rebuild Progress  NSID   Size  Repli
volume-0cbd7795-aebe-4759-b9c9-a00a83eb9307  2857de2a-9dca-48de-af54-ccb6f9992878  Available  FullyProtected    13     10 GiB  2
true                                         values:"nqn.2014-08.org.nvmeexpress:uuid:2873ff4b-5947-49a9-b891-24cc79f15f12"  None

```

2. **Transient Cloning:** A hidden, transient volume is cloned from this snapshot. The snapshot is immediately deleted to minimize metadata overhead.
3. **Data Streaming:** The cinder-backup service streams the blocks from this transient volume directly to the S3 object store.

```

demo@lb-01:~
X demo@lb-01:~ (ssh)
Every 2.0s: lbcli list snapshots --project-name acme                               lb-01: Fri Feb 20 17:02:20 2026

Name      UUID      Source volume UUID  State

X demo@lb-01:~ (ssh)
Every 2.0s: lbcli list volumes --project-name acme                               lb-01: Fri Feb 20 17:02:21 2026

Name                               UUID                               State      Protection State  NSID   Size  Repli
cas  Compression  ACL                               UUID                               State      Protection State  Rebuild Progress  NSID   Size  Repli
volume-0cbd7795-aebe-4759-b9c9-a00a83eb9307  2857de2a-9dca-48de-af54-ccb6f9992878  Available  FullyProtected    13     10 GiB  2
true                                         values:"nqn.2014-08.org.nvmeexpress:uuid:2873ff4b-5947-49a9-b891-24cc79f15f12"  None
volume-748bc257-2318-480c-9b5d-e2fcaefd2368  c72ca91b-ac00-4277-a3f9-d9f467024915  Available  FullyProtected    15     10 GiB  2
true                                         values:"nqn.2014-08.org.nvmeexpress:uuid:2873ff4b-5947-49a9-b891-24cc79f15f12"  None

```



4. **Cleanup:** Upon completion, the transient Lightbits volume is purged.

```
demo@lb-01:~  
X demo@lb-01:~ (ssh)  
Every 2.0s: lbcli list snapshots --project-name acme lb-01: Fri Feb 20 17:10:59 2026  
Name      UUID      Source volume UUID  State  
  
X demo@lb-01:~ (ssh)  
Every 2.0s: lbcli list volumes --project-name acme lb-01: Fri Feb 20 17:10:58 2026  
Name      UUID      State      Protection State  NSID      Size      Repli  
cas      Compression  ACL      Rebuild Progress  
volume-0cbd7795-aebe-4759-b9c9-a00a83eb9307  2857de2a-9dca-48de-af54-ccb6f9992878  Available  FullyProtected  13      10 GiB  2  
true      values:"nqn.2014-08.org.nvmeexpress:uuid:2873ff4b-5947-49a9-b891-24cc79f15f12"  None
```

5. **Restore Logic:** Unlike the backup process, the restoration is performed via a direct stream from the S3 object store back to the existing Lightbits volume, requiring no intermediate snapshots or clones within the storage cluster.

5. Conclusion

The integration of OpenStack Cinder with Lightbits and S3-compatible object storage provides a robust, scalable framework for high-performance data protection. By leveraging Lightbits' native NVMe-over-TCP capabilities, the backup workflow minimizes impact on production workloads by using transient snapshots and clones. This ensures that even "in-use" volumes can be backed up with crash consistency, while the transition to S3 storage, such as Garage, provides a cost-effective, durable repository for long-term data retention outside the primary storage cluster.

Ultimately, this architecture simplifies the lifecycle management of block storage by automating data movement between high-performance flash tiers and archival object tiers. The ability to perform direct-to-volume restorations without requiring intermediate storage staging reduces Recovery Time Objectives (RTO) and operational complexity. For enterprise environments running mission-critical applications on OpenStack, this synergy between Lightbits and S3 creates a resilient disaster recovery strategy that balances high-speed local performance with cloud-native storage flexibility.



About Lightbits Labs™

Lightbits Labs (Lightbits) is leading the digital data center transformation by making high-performance elastic block storage available to any cloud. Creators of the NVMe® over TCP (NVMe/TCP) protocol, Lightbits software-defined storage is easy to deploy at scale and delivers performance equivalent to local flash to accelerate cloud-native applications in bare metal, virtual, or containerized environments. Backed by leading enterprise investors including Cisco Investments, Dell Technologies Capital, Intel Capital, JP Morgan Chase, Lenovo, and Micron, Lightbits is on a mission to make high-performance elastic block storage simple, scalable and cost-efficient for any cloud.

www.lightbitslabs.com

US Offices
1830 The Alameda,
San Jose, CA 95126, USA

info@lightbitslabs.com

Israel Office
17 Atir Yeda Street,
Kfar Saba 4464313, Israel

The information in this document and any document referenced herein is provided for informational purposes only, is provided as is and with all faults and cannot be understood as substituting for customized service and information that might be developed by Lightbits Labs Ltd for a particular user based upon that user's particular environment. Reliance upon this document and any document referenced herein is at the user's own risk.

The software is provided "As is", without warranty of any kind, express or implied, including but not limited to the warranties of merchantability, fitness for a particular purpose and non-infringement. In no event shall the contributors or copyright holders be liable for any claim, damages or other liability, whether in an action of contract, tort or otherwise, arising from, out of or in connection with the software or the use or other dealings with the software.

Unauthorized copying or distributing of included software files, via any medium is strictly prohibited.

COPYRIGHT© 2026 LIGHTBITS LABS LTD. - ALL RIGHTS RESERVED

LBWP22/2026/2