



Red Hat Enterprise Linux 9.0 Beta

Deduplicating and compressing logical volumes on RHEL

Using VDO to increase LVM storage capacity

Red Hat Enterprise Linux 9.0 Beta Deduplicating and compressing logical volumes on RHEL

Using VDO to increase LVM storage capacity

Legal Notice

Copyright © 2021 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document explains how to use the Virtual Data Optimizer (VDO) feature in LVM to manage deduplicated and compressed logical volumes on RHEL.

Table of Contents

RHEL BETA RELEASE	3
MAKING OPEN SOURCE MORE INCLUSIVE	4
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	5
CHAPTER 1. INTRODUCTION TO VDO ON LVM	6
CHAPTER 2. LVM-VDO REQUIREMENTS	7
2.1. VDO MEMORY REQUIREMENTS	7
2.2. VDO STORAGE SPACE REQUIREMENTS	7
2.3. EXAMPLES OF VDO REQUIREMENTS BY PHYSICAL SIZE	8
2.4. PLACEMENT OF LVM-VDO IN THE STORAGE STACK	9
CHAPTER 3. CREATING A DEDUPLICATED AND COMPRESSED LOGICAL VOLUME	10
3.1. LVM-VDO DEPLOYMENT SCENARIOS	10
3.2. THE PHYSICAL AND LOGICAL SIZE OF AN LVM-VDO VOLUME	12
3.3. SLAB SIZE IN VDO	12
3.4. INSTALLING VDO	13
3.5. CREATING AN LVM-VDO VOLUME	13
3.6. MOUNTING AN LVM-VDO VOLUME	14
3.7. CHANGING THE COMPRESSION AND DEDUPLICATION SETTINGS ON AN LVM-VDO VOLUME	15
CHAPTER 4. TRIM OPTIONS ON AN LVM-VDO VOLUME	17
4.1. ENABLING DISCARD MOUNT OPTION ON VDO	17
4.2. SETTING UP PERIODIC TRIM OPERATION	17

RHEL BETA RELEASE

Red Hat provides Red Hat Enterprise Linux Beta access to all subscribed Red Hat accounts. The purpose of Beta access is to:

- Provide an opportunity to customers to test major features and capabilities prior to the general availability release and provide feedback or report issues.
- Provide Beta product documentation as a preview. Beta product documentation is under development and is subject to substantial change.

Note that Red Hat does not support the usage of RHEL Beta releases in production use cases. For more information, see [What does Beta mean in Red Hat Enterprise Linux and can I upgrade a RHEL Beta installation to a General Availability \(GA\) release?](#).

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Please let us know how we could make it better. To do so:

- For simple comments on specific passages:
 1. Make sure you are viewing the documentation in the *Multi-page HTML* format. In addition, ensure you see the **Feedback** button in the upper right corner of the document.
 2. Use your mouse cursor to highlight the part of text that you want to comment on.
 3. Click the **Add Feedback** pop-up that appears below the highlighted text.
 4. Follow the displayed instructions.
- For submitting more complex feedback, create a Bugzilla ticket:
 1. Go to the [Bugzilla](#) website.
 2. As the Component, use **Documentation**.
 3. Fill in the **Description** field with your suggestion for improvement. Include a link to the relevant part(s) of documentation.
 4. Click **Submit Bug**.

CHAPTER 1. INTRODUCTION TO VDO ON LVM

The Virtual Data Optimizer (VDO) feature provides inline block-level deduplication, compression, and thin provisioning for storage. You can manage VDO as a type of LVM logical volumes (LVs), similar to LVM thinly provisioned volumes.

VDO volumes on LVM (LVM-VDO) are composed of the following LVs:

VDO pool LV

This is the backing physical device that stores, deduplicates, and compresses data for the VDO LV. The VDO pool LV sets the physical size of the VDO volume, which is the amount of data that VDO can store on the disk.

Currently, each VDO pool LV can hold only one VDO LV. As a result, VDO deduplicates and compresses each VDO LV separately. In other words, VDO cannot deduplicate or compress a piece of data that is shared between several VDO LV.

VDO LV

This is the virtual, provisioned device on top of the VDO pool LV. The VDO LV sets the provisioned, logical size of the VDO volume, which is the amount of data that applications can write to the volume before deduplication and compression occurs.

Table 1.1. A comparison of components in VDO on LVM and LVM thin provisioning

	Physical device	Provisioned device
VDO on LVM	VDO pool LV	VDO LV
LVM thin provisioning	Thin pool	Thin volume

Since VDO is thinly provisioned, the file system and applications only see the logical space in use and are not aware of the actual physical space available. Use scripting to monitor the actual available space and generate an alert if use exceeds a threshold: for example, when the VDO pool LV is 80% full.

CHAPTER 2. LVM-VDO REQUIREMENTS

VDO on LVM has certain requirements on its placement and your system resources.

2.1. VDO MEMORY REQUIREMENTS

Each VDO volume has two distinct memory requirements:

The VDO module

VDO requires a fixed 38 MB of RAM and several variable amounts:

- 1.15 MB of RAM for each 1 MB of configured block map cache size. The block map cache requires a minimum of 150MB RAM.
- 1.6 MB of RAM for each 1 TB of logical space.
- 268 MB of RAM for each 1 TB of physical storage managed by the volume.

The UDS index

The Universal Deduplication Service (UDS) requires a minimum of 250 MB of RAM, which is also the default amount that deduplication uses. You can configure the value when formatting a VDO volume, because the value also affects the amount of storage that the index needs.

The memory required for the UDS index is determined by the index type and the required size of the deduplication window:

Index type	Deduplication window	Note
Dense	1 TB per 1 GB of RAM	A 1 GB dense index is generally sufficient for up to 4 TB of physical storage.
Sparse	10 TB per 1 GB of RAM	A 1 GB sparse index is generally sufficient for up to 40 TB of physical storage.

The UDS Sparse Indexing feature is the recommended mode for VDO. It relies on the temporal locality of data and attempts to retain only the most relevant index entries in memory. With the sparse index, UDS can maintain a deduplication window that is ten times larger than with dense, while using the same amount of memory.

Although the sparse index provides the greatest coverage, the dense index provides more deduplication advice. For most workloads, given the same amount of memory, the difference in deduplication rates between dense and sparse indexes is negligible.

Additional resources

- [Examples of VDO requirements by physical size](#)

2.2. VDO STORAGE SPACE REQUIREMENTS

You can configure a VDO volume to use up to 256 TB of physical storage. Only a certain part of the physical storage is usable to store data. This section provides the calculations to determine the usable size of a VDO-managed volume.

VDO requires storage for two types of VDO metadata and for the UDS index:

- The first type of VDO metadata uses approximately 1 MB for each 4 GB of *physical storage* plus an additional 1 MB per slab.
- The second type of VDO metadata consumes approximately 1.25 MB for each 1 GB of *logical storage*, rounded up to the nearest slab.
- The amount of storage required for the UDS index depends on the type of index and the amount of RAM allocated to the index. For each 1 GB of RAM, a dense UDS index uses 17 GB of storage, and a sparse UDS index will use 170 GB of storage.

Additional resources

- [Examples of VDO requirements by physical size](#)
- [Slab size in VDO](#)

2.3. EXAMPLES OF VDO REQUIREMENTS BY PHYSICAL SIZE

The following tables provide approximate system requirements of VDO based on the physical size of the underlying volume. Each table lists requirements appropriate to the intended deployment, such as primary storage or backup storage.

The exact numbers depend on your configuration of the VDO volume.

Primary storage deployment

In the primary storage case, the UDS index is between 0.01% to 25% the size of the physical size.

Table 2.1. Storage and memory requirements for primary storage

Physical size	RAM usage: UDS	RAM usage: VDO	Disk usage	Index type
10GB-1TB	250MB	472MB	2.5GB	Dense
2-10TB	1GB	3GB	10GB	Dense
	250MB		22GB	Sparse
11-50TB	2GB	14GB	170GB	Sparse
51-100TB	3GB	27GB	255GB	Sparse
101-256TB	12GB	69GB	1020GB	Sparse

Backup storage deployment

In the backup storage case, the UDS index covers the size of the backup set but is not bigger than the physical size. If you expect the backup set or the physical size to grow in the future, factor this into the index size.

Table 2.2. Storage and memory requirements for backup storage

Physical size	RAM usage: UDS	RAM usage: VDO	Disk usage	Index type
10GB-1TB	250MB	472MB	2.5 GB	Dense
2-10TB	2GB	3GB	170GB	Sparse
11-50TB	10GB	14GB	850GB	Sparse
51-100TB	20GB	27GB	1700GB	Sparse
101-256TB	26GB	69GB	3400GB	Sparse

2.4. PLACEMENT OF LVM-VDO IN THE STORAGE STACK

You must place certain storage layers under a VDO logical volume and others above it.

You can place thick-provisioned layers on top of VDO, but you cannot rely on the guarantees of thick provisioning in that case. Because the VDO layer is thin-provisioned, the effects of thin provisioning apply to all layers above it. If you do not monitor the VDO volume, you might run out of physical space on thick-provisioned volumes above VDO.

The supported placement of the following layers is under VDO. Do not place them above VDO:

- DM Multipath
- DM Crypt
- Software RAID (LVM or MD RAID)

The following configurations are **not supported**:

- VDO on top of a loopback device
- Encrypted volumes on top of VDO
- Partitions on a VDO volume
- RAID, such as LVM RAID, MD RAID, or any other type, on top of a VDO volume
- Deploying Ceph Storage on LVM-VDO

Additional resources

- [Stacking LVM volumes knowledgebase article](#)

CHAPTER 3. CREATING A DEDUPLICATED AND COMPRESSED LOGICAL VOLUME


You can create an LVM logical volume that uses the VDO feature to deduplicate and compress data.

3.1. LVM-VDO DEPLOYMENT SCENARIOS

You can deploy VDO on LVM (LVM-VDO) in a variety of ways to provide deduplicated storage for:

- block access
- file access
- local storage
- remote storage

Because LVM-VDO exposes its deduplicated storage as a regular logical volume (LV), you can use it with standard file systems, iSCSI and FC target drivers, or as unified storage.




NOTE

Deploying Ceph Storage on LVM-VDO is currently not supported.

KVM


You can deploy LVM-VDO on a KVM server configured with Direct Attached Storage.



RHEL_462492_1117

File systems


You can create file systems on top of a VDO LV and expose them to NFS or CIFS users with the NFS server or Samba.



RHEL_466924_0218

iSCSI target


You can export the entirety of the VDO LV as an iSCSI target to remote iSCSI initiators.



RHEL_466924_0218

Encryption

Device Mapper (DM) mechanisms such as DM Crypt are compatible with VDO. Encrypting a VDO LV volumes helps ensure data security, and any file systems above the VDO LV are still deduplicated.



RHEL_466924_0218



IMPORTANT

Applying the encryption layer above the VDO LV results in little if any data deduplication. Encryption makes duplicate blocks different before VDO can deduplicate them.

Always place the encryption layer below the VDO LV.

3.2. THE PHYSICAL AND LOGICAL SIZE OF AN LVM-VDO VOLUME

This section describes the physical size, available physical size, and logical size that VDO can utilize.

Physical size

This is the same size as the physical extents allocated to the VDO pool LV. VDO uses this storage for:

- User data, which might be deduplicated and compressed
- VDO metadata, such as the UDS index

Available physical size

This is the portion of the physical size that VDO is able to use for user data.

It is equivalent to the physical size minus the size of the metadata, rounded down to a multiple of the slab size.

Logical Size

This is the provisioned size that the VDO LV presents to applications. It is usually larger than the available physical size. VDO currently supports any logical size up to 254 times the size of the physical volume with an absolute maximum logical size of 4 PB.

When you set up a VDO logical volume (LV), you specify the amount of logical storage that the VDO LV presents. When hosting active VMs or containers, Red Hat recommends provisioning storage at a 10:1 logical to physical ratio, that is, if you are utilizing 1 TB of physical storage, you would present it as 10 TB of logical storage.

If you do not specify the `--virtualsize` option, VDO provisions the volume to a **1:1** ratio. For example, if you put a VDO LV on top of a 20 GB VDO pool LV, VDO reserves 2.5 GB for the UDS index, if the default index size is used. The remaining 17.5 GB is provided for the VDO metadata and user data. As a result, the available storage to consume is not more than 17.5 GB, and can be less due to metadata that makes up the actual VDO volume.

Additional resources

- [Examples of VDO requirements by physical size](#)

3.3. SLAB SIZE IN VDO

The physical storage of the VDO volume is divided into a number of slabs. Each slab is a contiguous region of the physical space. All of the slabs for a given volume have the same size, which can be any power of 2 multiple of 128 MB up to 32 GB.

The default slab size is 2 GB in order to facilitate evaluating VDO on smaller test systems. A single VDO volume can have up to 8192 slabs. Therefore, in the default configuration with 2 GB slabs, the maximum allowed physical storage is 16 TB. When using 32 GB slabs, the maximum allowed physical storage is 256 TB. VDO always reserves at least one entire slab for metadata, and therefore, the reserved slab cannot be used for storing user data.

Slab size has no effect on the performance of the VDO volume.

Table 3.1. Recommended VDO slab sizes by physical volume size

Physical volume size	Recommended slab size
10–99 GB	1 GB
100 GB – 1 TB	2 GB
2–256 TB	32 GB

You can control the slab size by providing the `--config 'allocation/vdo_slab_size_mb=size-in-megabytes'` option to the `lvcreate` command.

3.4. INSTALLING VDO

This procedure installs software necessary to create, mount, and manage VDO volumes.

Procedure

- Install the `vdo` and `kmod-kvdo` packages:

```
# yum install vdo kmod-kvdo
```

3.5. CREATING AN LVM-VDO VOLUME

This procedure creates an VDO logical volume (LV) on a VDO pool LV.

Prerequisites

- Install the VDO software. For more information, see [Installing VDO](#).
- An LVM volume group with free storage capacity exists on your system.

Procedure

1. Pick a name for your VDO LV, such as `vdo1`. You must use a different name and device for each VDO LV on the system.
In all the following steps, replace `vdo-name` with the name.
2. Create the VDO LV:

```
# lvcreate --type vdo \
  --name vdo-name \
  --size physical-size \
  --virtualsize logical-size \
  vg-name
```

- Replace `vg-name` with the name of an existing LVM volume group where you want to place the VDO LV.
- Replace `logical-size` with the amount of logical storage that the VDO LV will present.

- If the physical size is larger than 16TiB, add the following option to increase the slab size on the volume to 32GiB:

```
--config 'allocation/vdo_slab_size_mb=32768'
```

If you use the default slab size of 2GiB on a physical size larger than 16TiB, the **lvcreate** command fails with the following error:

```
ERROR - vdoformat: formatVDO failed on '/dev/device': VDO Status: Exceeds maximum number of slabs supported
```

Example 3.1. Creating a VDO LV for container storage

For example, to create a VDO LV for container storage on a 1TB VDO pool LV, you can use:

```
# lvcreate --type vdo \
  --name vdo1 \
  --size 1T \
  --virtualsize 10T \
  vg-name
```



IMPORTANT

If a failure occurs when creating the VDO volume, remove the volume to clean up.

3. Create a file system on the VDO LV:

- For the XFS file system:

```
# mkfs.xfs -K /dev/vg-name/vdo-name
```

- For the ext4 file system:

```
# mkfs.ext4 -E nodiscard /dev/vg-name/vdo-name
```

Additional resources

- [lvmvdo\(7\)](#) man page

3.6. MOUNTING AN LVM-VDO VOLUME

This procedure mounts a file system on an LVM-VDO volume, either manually or persistently.

Prerequisites

- An LVM-VDO volume exists on your system. For more information, see [Creating an LVM-VDO volume](#).

Procedure

- To mount the file system on the LVM-VDO volume manually, use:

```
# mount /dev/vg-name/vdo-name mount-point
```

- To configure the file system to mount automatically at boot, add a line to the `/etc/fstab` file:
 - For the XFS file system:

```
/dev/vg-name/vdo-name mount-point xfs defaults 0 0
```

- For the ext4 file system:

```
/dev/vg-name/vdo-name mount-point ext4 defaults 0 0
```


If the LVM-VDO volume is located on a block device that requires network, such as iSCSI, add the `_netdev` mount option. For iSCSI and other block devices requiring network, see the `systemd.mount(5)` man page for information on the `_netdev` mount option.

Additional resources

- `systemd.mount(5)` man page

3.7. CHANGING THE COMPRESSION AND DEDUPLICATION SETTINGS ON AN LVM-VDO VOLUME

This procedure enables or disables the compression and deduplication of a VDO pool logical volume (LV).



NOTE

Compression and deduplication are enabled by default.

Prerequisites

- An LVM-VDO volume exists on your system.

Procedure

1. To find out if the compression and deduplication is enabled or disabled on your logical volumes:

```
# lvs -o+vdo_compression,vdo_deduplication
```

2. Find out status of the compression and status of the deduplication index of your running active VDOPoolLV:

```
# lvs -o+vdo_compression_state,vdo_index_state
```

The `vdo_index_state` can show as **error**, **close**, **opening**, **closing**, **online**, and **offline**.

3. To enable or disable the compression for VDOPoolLV:

```
# lvchange --compression y|n vg-name/vdopoolname
```

4. To enable or disable the deduplication for VDOPoolLV:

```
█ # lvchange --deduplication y|n vg-name/vdopoolname
```

Additional resources

- **lvmvdo(7)** man page

CHAPTER 4. TRIM OPTIONS ON AN LVM-VDO VOLUME

You can mount your file system with the **discard** option, which informs the VDO volume of the unused space. Another option is to use the **fstrim** application, which is an on-demand discarding, or **mount -o discard** command for immediate discarding.

When using the **fstrim** application, the admin needs to schedule and monitor an additional process, while using **mount -o discard** command allows for immediate recovery of space when possible.

Note that it is currently recommended to use **fstrim** application to discard unused blocks rather than the **discard** mount option because the performance impact of this option can be quite severe. For this reason, **nodiscard** is the default.

4.1. ENABLING DISCARD MOUNT OPTION ON VDO

This procedure enables the **discard** option on your VDO volume.

Prerequisites

- An LVM-VDO volume exists on your system.

Procedure

- Enable the **discard** on your volume:

```
# mount -o discard /dev/vg-name/vdo-name mount-point
```

Additional resources

- **xfs(5)**, **mount(8)**, and **lvmvdo(7)** man pages

4.2. SETTING UP PERIODIC TRIM OPERATION

This procedure enables a scheduled TRIM operation on your system.

Prerequisites

- An LVM-VDO volume exists on your system.

Procedure

- Enable and start the timer:

```
# systemctl enable --now fstrim.timer
```

Verification

- Verify that the timer is enabled:

```
# systemctl list-timers fstrim.timer
```

Example 4.1. Possible output of the verification procedure

```
# systemctl list-timers fstrim.timer
NEXT          LEFT    LAST PASSED UNIT    ACTIVATES
Mon 2021-05-10 00:00:00 EDT 5 days left n/a  n/a    fstrim.timer fstrim.service
```



NOTE

You will not see any reference to a VDO volume, because the **fstrim.timer** runs across all mounted filesystems.

Additional resources

- **fstrim(8)** man page