



# **Manage volume efficiency operations manually**

**ONTAP 9**

NetApp  
September 18, 2024

This PDF was generated from <https://docs.netapp.com/us-en/ontap/volumes/manage-volume-efficiency-operations-manual-concept.html> on September 18, 2024. Always check docs.netapp.com for the latest.

# Table of Contents

- Manage volume efficiency operations manually. . . . . 1
  - Manage volume efficiency operations manually overview . . . . . 1
  - Run an efficiency operation manually . . . . . 1
  - Checkpoints and efficiency operations. . . . . 2
  - Resume a halted efficiency operation . . . . . 2
  - Run an efficiency operation manually on existing data . . . . . 3

# Manage volume efficiency operations manually

## Manage volume efficiency operations manually overview

You can manage how the efficiency operations run on a volume by running efficiency operations manually.

You can also control how the efficiency operations run based on the following conditions:

- Use checkpoints or not
- Run efficiency operations on existing data or only new data
- Stop efficiency operations if required

You can use the `volume efficiency show` command with `schedule` as value for the `-fields` option to view the schedule assigned to the volumes.

## Run an efficiency operation manually

You can run efficiency operations on a volume manually. You might do this when scheduling efficiency operations is not appropriate.

### Before you begin

Depending on the efficiency operation you want to run manually, you must have enabled deduplication or both data compression and deduplication on a volume.

### About this task

This operation is performed using the `volume efficiency start` command. When temperature-sensitive storage efficiency is enabled on a volume, deduplication is run initially followed by data compression.

Deduplication is a background process that consumes system resources while it is running. If the data does not change often in a volume, it is best to run deduplication less frequently. Multiple concurrent deduplication operations running on a storage system lead to a higher consumption of system resources.

You can run a maximum of eight concurrent deduplication or data compression operations per node. If any more efficiency operations are scheduled, the operations are queued.

Beginning with ONTAP 9.13.1, if temperature-sensitive storage efficiency is enabled on a volume, you can run volume efficiency on existing data to take advantage of sequential packing to further improve storage efficiency.

## Run efficiency manually

### Steps

1. Start the efficiency operation on a volume: `volume efficiency start`

### Example

+ The following command allows you to manually start only deduplication or deduplication followed by logical compression and container compression on the volume `VolA`

+

```
volume efficiency start -vserver vs1 -volume VolA
```

## Repack existing data

To take advantage of sequential data packing introduced in ONTAP 9.13.1 on volumes with temperature-sensitive storage efficiency enabled, you can repack existing data. You must be in advanced privilege mode to use this command.

### Steps

1. Set the privilege level: `set -privilege advanced`
2. Repack existing data: `volume efficiency inactive-data-compression start -vserver vserver_name -volume volume_name -scan-mode extended_recompression`

### Example

```
volume efficiency inactive-data-compression start -vserver vs1 -volume  
voll -scan-mode extended_recompression
```

### Related information

- [Run efficiency operations manually on existing data](#)

## Checkpoints and efficiency operations

Checkpoints are used internally to log the execution process of an efficiency operation. When an efficiency operation is stopped for any reason (such as system halt, system disruption, reboot, or because the last efficiency operation failed or stopped) and checkpoint data exists, the efficiency operation can resume from the latest checkpoint file.

A checkpoint is created:

- in each stage or substage of the operation
- when you run the `sis stop` command
- when the duration expires

## Resume a halted efficiency operation

If an efficiency operation is halted due to a system halt, system disruption, or reboot, you can resume the efficiency operation from the same point it was halted. This helps to save time and resources by not needing to restart the operation from the beginning.

### About this task

If you enabled only deduplication on the volume, deduplication runs on the data. If you enabled both deduplication and data compression on a volume, then data compression runs first, followed by deduplication.

You can view the details of the checkpoint for a volume by using the `volume efficiency show` command.

By default, the efficiency operations resume from checkpoints. However, if a checkpoint corresponding to a previous efficiency operation (the phase when the `volume efficiency start -scan-old-data` command is run) is older than 24 hours, then the efficiency operation does not resume from the previous checkpoint automatically. In this case, the efficiency operation starts from the beginning. However, if you know that significant changes have not occurred in the volume since the last scan, you can force continuation from the previous checkpoint by using the `-use-checkpoint` option.

### Steps

1. Use the `volume efficiency start` command with the `-use-checkpoint` option to resume an efficiency operation.

The following command enables you to resume an efficiency operation on new data on volume VolA:

```
volume efficiency start -vserver vs1 -volume VolA -use-checkpoint true
```

The following command enables you to resume an efficiency operation on existing data on volume VolA:

```
volume efficiency start -vserver vs1 -volume VolA -scan-old-data true -use-checkpoint true
```

## Run an efficiency operation manually on existing data

You can run the efficiency operations manually on the data that exists in non-temperature sensitive storage efficiency volumes prior to enabling deduplication, data compression, or data compaction. You can run these operations with ONTAP versions earlier than ONTAP 9.8.

### About this task

This operation is performed using the `volume efficiency start` command with the `-scan-old-data` parameter. The `-compression` option does not work with `-scan-old-data` on temperature sensitive storage efficiency volumes. Inactive data compression runs automatically on pre-existing data for temperature sensitive storage efficiency volumes in ONTAP 9.8 and later.

If you enable only deduplication on a volume, then deduplication runs on the data. If you enable deduplication, data compression, and data compaction on a volume, then data compression runs first, followed by deduplication and data compaction.

When you run data compression on existing data, by default the data compression operation skips the data blocks that are shared by deduplication and the data blocks that are locked by Snapshot copies. If you choose to run data compression on shared blocks, then optimization is turned off and the fingerprint information is captured and used for sharing again. You can change the default behavior of data compression when compressing existing data.

You can run a maximum of eight deduplication, data compression, or data compaction operations concurrently per node. The remaining operations are queued.



Post process compression does not run on AFF platforms. An EMS message is generated to inform you that this operation was skipped.

### Steps

1. Use the `volume efficiency start -scan-old-data` command to run deduplication, data

compression, or data compaction manually on the existing data.

The following command enables you to run these operations manually on the existing data in volume VolA:

```
volume efficiency start -vserver vs1 -volume VolA -scan-old-data true [-  
compression | -dedupe | -compaction ] true
```

#### **Related information**

- [Run efficiency operations manually](#)

## Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

## Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.