

Manage the use of local tiers (aggregates) ONTAP 9

NetApp September 18, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap/disks-aggregates/manage-use-local-tiers-overview-task.html on September 18, 2024. Always check docs.netapp.com for the latest.

Table of Contents

Manage the use of local tiers (aggregates)	. 1
Manage the use of local tiers (aggregates)	. 1
Rename a local tier (aggregate).	. 1
Set media cost of a local tier (aggregate)	. 2
Manually Fast zero drives	. 2
Manually assign disk ownership.	. 3
Determine drive and RAID group information for a local tier (aggregate)	. 6
Assign local tiers (aggregates) to storage VMs (SVMs).	. 7
Determine which volumes reside on a local tier (aggregate)	. 8
Determine and control a volume's space usage in a local tier (aggregate)	. 8
Determine space usage in a local tier (aggregate)	10
Relocate ownership of a local tier (aggregate) within an HA pair.	. 11
Delete a local tier (aggregate)	14
Commands for aggregate relocation	14
Commands for managing aggregates	14

Manage the use of local tiers (aggregates)

Manage the use of local tiers (aggregates)

After you have created local tiers (aggregates), you can manage how they are used.

You can perform the following tasks:

- Rename a local tier (aggregate)
- Set the media cost of a local tier (aggregate)
- Determine drive and RAID group information for a local tier (aggregate)
- Assign local tiers (aggregates) to storage VMs (SVMs)
- Determine which volumes reside on a local tier (aggregate)
- Determine and control a volume's space usages in a local tier (aggregate)
- Determine space usage in a local tier (aggregate)
- Relocate local tier (aggregate) ownership within an HA pair
- Delete a local tier (aggregate)

Rename a local tier (aggregate)

You can rename a local tier (aggregate). The method you follow depends on the interface you use—System Manager or the CLI:

System Manager

Use System Manager to rename a local tier (aggregate)

Beginning with ONTAP 9.10.1, you can modify the name of a local tier (aggregate).

Steps

- 1. In System Manager, click Storage > Tiers.
- 2. Click next to the name of the local tier.
- 3. Select Rename.
- 4. Specify a new name for the local tier.

CLI

Use the CLI to rename a local tier (aggregate)

Step

1. Using the CLI, rename the local tier (aggregate):

```
storage aggregate rename -aggregate aggr-name -newname aggr-new-name
```

The following example renames an aggregate named "aggr5" as "sales-aggr":

> storage aggregate rename -aggregate aggr5 -newname sales-aggr

Set media cost of a local tier (aggregate)

Beginning with ONTAP 9.11.1, you can use System Manager to set the media cost of a local tier (aggregate).

Steps

- In System Manager, click Storage > Tiers, then click Set Media Cost in the desired local tier (aggregate) tiles.
- 2. Select active and inactive tiers to enable comparison.
- 3. Enter a currency type and amount.

When you enter or change the media cost, the change is made in all media types.

Manually Fast zero drives

On systems freshly installed with ONTAP 9.4 or later and systems reinitialized with ONTAP 9.4 or later, *fast zeroing* is used to zero drives.

With *fast zeroing*, drives are zeroed in seconds. This is done automatically before provisioning and greatly reduces the time it takes to initialize the system, create aggregates, or expand aggregates when spare drives are added.

Fast zeroing is supported on both SSDs and HDDs.



Fast zeroing is not supported on systems upgraded from ONTAP 9.3 or earlier. ONTAP 9.4 or later must be freshly installed or the system must be reinitialized. In ONTAP 9.3 and earlier, drives are also automatically zeroed by ONTAP, however, the process takes longer.

If you need to manually zero a drive, you can use one of the following methods. In ONTAP 9.4 and later, manually zeroing a drive also takes only seconds.

CLI command

Use a CLI command to fast-zero drives

About this task

Admin privileges are required to use this command.

Steps

1. Enter the CLI command:

storage disk zerospares

Boot menu options

Select options from the boot menu to fast-zero drives

About this task

- The fast zeroing enhancement does not support systems upgraded from a release earlier than ONTAP 9.4.
- If any node on the cluster contains a local tier (aggregate) with fast-zeroed drives, then you cannot revert the cluster to ONTAP 9.2 or earlier.

Steps

- 1. From the boot menu, select one of the following options:
 - (4) Clean configuration and initialize all disks
 - (9a) Unpartition all disks and remove their ownership information
 - \circ (9b) Clean configuration and initialize node with whole disks

Manually assign disk ownership

Disks must be owned by a node before they can be used in a local tier (aggregate).

About this task

- If you are manually assigning ownership in an HA pair that is not being initialized and does not have only DS460C shelves, use option 1.
- If you are initializing an HA pair that has only DS460C shelves, use option 2 to manually assign ownership for the root drives.

Option 1: Most HA pairs

For an HA pair that is not being initialized and does not have only DS460C shelves, use this procedure to manually assigning ownership.

About this task

- The disks you are assigning ownership for must be in a shelf that is physically cabled to the node you are assigning ownership to.
- If you are using disks in a local tier (aggregate):
 - Disks must be owned by a node before they can be used in a local tier (aggregate).
 - You cannot reassign ownership of a disk that is in use in a local tier (aggregate).

Steps

1. Use the CLI to display all unowned disks:

storage disk show -container-type unassigned

2. Assign each disk:

storage disk assign -disk disk name -owner owner name

You can use the wildcard character to assign more than one disk at once. If you are reassigning a spare disk that is already owned by a different node, you must use the "-force" option.

For an HA pair that you are initializing and that only has DS460C shelves, use this procedure to manually assign ownership for the root drives.

About this task

• When you initialize an HA pair that has only DS460C shelves, you must manually assign the root drives to conform to the half-drawer policy.

After HA pair initialization (boot up), automatic assignment of disk ownership is automatically enabled and uses the half-drawer policy to assign ownership to the remaining drives (other than the root drives) and any drives added in the future, such as replacing failed disks, responding to a "low spares" message, or adding capacity.

Learn about the half-drawer policy in the topic About automatic assignment of disk ownership.

• RAID needs a minimum of 10 drives for each HA pair (5 for each node) for any greater than 8TB NL-SAS drives in a DS460C shelf.

Steps

- 1. If your DS460C shelves are not fully populated, complete the following substeps; otherwise, go to the next step.
 - a. First, install drives in the front row (drive bays 0, 3, 6, and 9) of each drawer.

Installing drives in the front row of each drawer allows for proper air flow and prevents overheating.

b. For the remaining drives, evenly distribute them across each drawer.

Fill drawer rows from front to back. If you don't have enough drives to fill rows, then install them in pairs so that drives occupy the left and right side of a drawer evenly.

The following illustration shows the drive bay numbering and locations in a DS460C drawer.



2. Log into the clustershell using the node-management LIF or cluster-management LIF.

3. Manually assign the root drives in each drawer to conform to the half-drawer policy using the following substeps:

The half-drawer policy has you assign the left half of a drawer's drives (bays 0 to 5) to node A, and the right half of a drawer's drives (bays 6 to 11) to node B.

- a. Display all unowned disks: storage disk show -container-type unassigned`
- b. Assign the root disks: storage disk assign -disk disk_name -owner owner_name

You can use the wildcard character to assign more than one disk at a time.

Determine drive and RAID group information for a local tier (aggregate)

Some local tier (aggregate) administration tasks require that you know what types of drives compose the local tier, their size, checksum, and status, whether they are shared with other local tiers, and the size and composition of the RAID groups.

Step

1. Show the drives for the aggregate, by RAID group:

storage aggregate show-status aggr_name

The drives are displayed for each RAID group in the aggregate.

You can see the RAID type of the drive (data, parity, dparity) in the Position column. If the Position column displays shared, then the drive is shared: if it is an HDD, it is a partitioned disk; if it is an SSD, it is part of a storage pool.

```
cluster1::> storage aggregate show-status nodeA fp 1
Owner Node: cluster1-a
Aggregate: nodeA fp 1 (online, mixed raid type, hybrid) (block checksums)
 Plex: /nodeA_fp_1/plex0 (online, normal, active, pool0)
  RAID Group /nodeA fp 1/plex0/rg0 (normal, block checksums, raid dp)
                                      Usable Physical
    Position Disk
                                       Size
                                               Size Status
                   Pool Type
                                RPM
    _____
                              ----- ------ ------ ------
                     ____ ____
    shared 2.0.1
                         SAS
                               10000 472.9GB 547.1GB (normal)
                      0
    shared 2.0.3
                      0 SAS
                              10000 472.9GB 547.1GB (normal)
                              10000 472.9GB 547.1GB (normal)
    shared 2.0.5
                      0 SAS
    shared 2.0.7
                      0 SAS
                              10000 472.9GB 547.1GB (normal)
                      0 SAS
    shared 2.0.9
                              10000 472.9GB 547.1GB (normal)
    shared 2.0.11
                         SAS
                              10000 472.9GB 547.1GB (normal)
                      0
  RAID Group /nodeA flashpool 1/plex0/rg1
  (normal, block checksums, raid4) (Storage Pool: SmallSP)
                                      Usable Physical
    Position Disk
                    Pool Type
                                       Size
                                               Size Status
                               RPM
    shared 2.0.13
                   0 SSD
                               - 186.2GB 745.2GB (normal)
                               - 186.2GB 745.2GB (normal)
    shared 2.0.12 0 SSD
8 entries were displayed.
```

Assign local tiers (aggregates) to storage VMs (SVMs)

If you assign one or more local tiers (aggregates) to a storage virtual machine (storage VM or SVM, formerly known as Vserver), then you can use only those local tiers to contain volumes for that storage VM (SVM).

What you'll need

The storage VM and the local tiers you want to assign to that storage VM must already exist.

About this task

Assigning local tiers to your storage VMs helps you keep your storage VMs isolated from each other; this is especially important in a multi-tenancy environment.

Steps

1. Check the list of local tiers (aggregates) already assigned to the SVM:

```
vserver show -fields aggr-list
```

The aggregates currently assigned to the SVM are displayed. If there are no aggregates assigned, "-" is displayed.

2. Add or remove assigned aggregates, depending on your requirements:

If you want to	Use this command
Assign additional aggregates	vserver add-aggregates
Unassign aggregates	vserver remove-aggregates

The listed aggregates are assigned to or removed from the SVM. If the SVM already has volumes that use an aggregate that is not assigned to the SVM, a warning message is displayed, but the command is completed successfully. Any aggregates that were already assigned to the SVM and that were not named in the command are unaffected.

Example

In the following example, the aggregates aggr1 and aggr2 are assigned to SVM svm1:

vserver add-aggregates -vserver svm1 -aggregates aggr1,aggr2

Determine which volumes reside on a local tier (aggregate)

You might need to determine which volumes reside on a local tier (aggregate) before performing operations on the local tier, such as relocating it or taking it offline.

Steps

1. To display the volumes that reside on an aggregate, enter

volume show -aggregate aggregate name

All volumes that reside on the specified aggregate are displayed.

Determine and control a volume's space usage in a local tier (aggregate)

You can determine which FlexVol volumes are using the most space in a local tier (aggregate) and specifically which features within the volume.

The volume show-footprint command provides information about a volume's footprint, or its space usage within the containing aggregate.

The volume show-footprint command shows details about the space usage of each volume in an aggregate, including offline volumes. This command bridges the gap between the output of the volume show-space and aggregate show-space commands. All percentages are calculated as a percent of aggregate size.

The following example shows the volume show-footprint command output for a volume called testvol:

Vserver · theus
Volume : testvol
Feature Used Used%
Volume Data Footprint120.6MB4%
Volume Guarantee 1.88GB 71%
Flexible Volume Metadata 11.38MB 0%
Delayed Frees 1.36MB 0%
Total Footprint2.01GB76%

The following table explains some of the key rows of the output of the volume show-footprint command and what you can do to try to decrease space usage by that feature:

Row/feature name	Description/contents of row	Some ways to decrease
Volume Data Footprint	The total amount of space used in the containing aggregate by a volume's data in the active file system and the space used by the volume's Snapshot copies. This row does not include reserved space.	 Deleting data from the volume. Deleting Snapshot copies from the volume.
Volume Guarantee	The amount of space reserved by the volume in the aggregate for future writes. The amount of space reserved depends on the guarantee type of the volume.	Changing the type of guarantee for the volume to none.
Flexible Volume Metadata	The total amount of space used in the aggregate by the volume's metadata files.	No direct method to control.
Delayed Frees	Blocks that ONTAP used for performance and cannot be immediately freed. For SnapMirror destinations, this row has a value of 0 and is not displayed.	No direct method to control.
File Operation Metadata	The total amount of space reserved for file operation metadata.	No direct method to control.
Total Footprint	The total amount of space that the volume uses in the aggregate. It is the sum of all of the rows.	Any of the methods used to decrease space used by a volume.

NetApp Technical Report 3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment

Determine space usage in a local tier (aggregate)

You can view how much space is used by all volumes in one or more local tiers (aggregates) so that you can take actions to free more space.

WAFL reserves a percentage of the total disk space for aggregate level metadata and performance. The space used for maintaining the volumes in the aggregate comes out of the WAFL reserve and cannot be changed.

In aggregates smaller than 30 TB, WAFL reserves 10% of the total disk space for aggregate level metadata and performance.

Beginning with ONTAP 9.12.1, in aggregates that are 30 TB or larger, the amount of reserved disk space for aggregate level metadata and performance is reduced, resulting in 5% more usable space in aggregates. The availability of this space savings varies based on your platform and version of ONTAP.

Disk space reserved by ONTAP in aggregates 30 TB or greater	Applies to platforms	In ONTAP versions
5%	All AFF and FAS platforms	ONTAP 9.14.1 and later
5%	AFF platforms and FAS500f platforms	ONTAP 9.12.1 and later
10%	All platforms	ONTAP 9.11.1 and later

You can view space usage by all volumes in one or more aggregates with the aggregate show-space command. This helps you see which volumes are consuming the most space in their containing aggregates so that you can take actions to free more space.

The used space in an aggregate is directly affected by the space used in the FlexVol volumes it contains. Measures that you take to increase space in a volume also affect space in the aggregate.



Beginning with ONTAP 9.15.1, two new metadata counters are available. Together with changes to several existing counters, you can get a clearer view of the amount of user data allocated. See Determine space usage in a volume or aggregate for more information.

The following rows are included in the aggregate show-space command output:

Volume Footprints

The total of all volume footprints within the aggregate. It includes all of the space that is used or reserved by all data and metadata of all volumes in the containing aggregate.

Aggregate Metadata

The total file system metadata required by the aggregate, such as allocation bitmaps and inode files.

Snapshot Reserve

The amount of space reserved for aggregate Snapshot copies, based on volume size. It is considered

used space and is not available to volume or aggregate data or metadata.

Snapshot Reserve Unusable

The amount of space originally allocated for aggregate Snapshot reserve that is unavailable for aggregate Snapshot copies because it is being used by volumes associated with the aggregate. Can occur only for aggregates with a non-zero aggregate Snapshot reserve.

Total Used

The sum of all space used or reserved in the aggregate by volumes, metadata, or Snapshot copies.

Total Physical Used

The amount of space being used for data now (rather than being reserved for future use). Includes space used by aggregate Snapshot copies.

The following example shows the aggregate show-space command output for an aggregate whose Snapshot reserve is 5%. If the Snapshot reserve was 0, the row would not be displayed.

cluster1::> storage aggregate show-space				
Aggregate : wqa_gx106_aggr1				
Feature	Used	Used%		
Volume Footprints	101.0MB	 0%		
Aggregate Metadata	300KB	0%		
Snapshot Reserve	5.98GB	5%		
Total Used	6.07GB	5%		
Total Physical Used	34.82KB	0%		

Related Information

- Knowledge Base article: Space Usage
- Free up 5% of your storage capacity by upgrading to ONTAP 9.12.1

Relocate ownership of a local tier (aggregate) within an HA pair

You can change the ownership of local tiers (aggregates) among the nodes in an HA pair without interrupting service from the local tiers.

Both nodes in an HA pair are physically connected to each other's disks or array LUNs. Each disk or array LUN is owned by one of the nodes.

Ownership of all disks or array LUNs within a local tier (aggregate) changes temporarily from one node to the other when a takeover occurs. However, local tiers relocation operations can also permanently change the ownership (for example, if done for load balancing). The ownership changes without any data-copy processes

or physical movement of the disks or array LUNs.

About this task

• Because volume count limits are validated programmatically during local tier relocation operations, it is not necessary to check for this manually.

If the volume count exceeds the supported limit, the local tier relocation operation fails with a relevant error message.

• You should not initiate local tier relocation when system-level operations are in progress on either the source or the destination node; likewise, you should not start these operations during the local tier relocation.

These operations can include the following:

- Takeover
- · Giveback
- Shutdown
- Another local tier relocation operation
- Disk ownership changes
- · Local tier or volume configuration operations
- · Storage controller replacement
- ONTAP upgrade
- ONTAP revert
- If you have a MetroCluster configuration, you should not initiate local tier relocation while disaster recovery operations (*switchover, healing*, or *switchback*) are in progress.
- If you have a MetroCluster configuration and initiate local tier relocation on a switched-over local tier, the operation might fail because it exceeds the DR partner's volume limit count.
- You should not initiate local tier relocation on aggregates that are corrupt or undergoing maintenance.
- Before initiating the local tier relocation, you should save any core dumps on the source and destination nodes.

Steps

1. View the aggregates on the node to confirm which aggregates to move and ensure they are online and in good condition:

storage aggregate show -node *source-node*

The following command shows six aggregates on the four nodes in the cluster. All aggregates are online. Node1 and Node3 form an HA pair and Node2 and Node4 form an HA pair.

cluster::> Aggregate	storage Size	aggregate Available	show Used%	State	#Vols	Nodes	RAID Status
aggr_0	239.0GB	11.13GB	95%	online	1	node1	raid_dp, normal
aggr_1	239.0GB	11.13GB	95%	online	1	node1	raid_dp, normal
aggr_2	239.0GB	11.13GB	95%	online	1	node2	raid_dp, normal
aggr_3	239.0GB	11.13GB	95%	online	1	node2	raid_dp, normal
aggr_4	239.0GB	238.9GB	0%	online	5	node3	raid_dp,
aggr_5	239.0GB	239.0GB	08	online	4	node4	raid_dp, normal
6 entries were displayed.							

2. Issue the command to start the aggregate relocation:

storage aggregate relocation start -aggregate-list aggregate-1, aggregate-2...
-node source-node -destination destination-node

The following command moves the aggregates aggr_1 and aggr_2 from Node1 to Node3. Node3 is Node1's HA partner. The aggregates can be moved only within the HA pair.

cluster::> storage aggregate relocation start -aggregate-list aggr_1, aggr_2 -node node1 -destination node3 Run the storage aggregate relocation show command to check relocation status. node1::storage aggregate>

3. Monitor the progress of the aggregate relocation with the storage aggregate relocation show command:

storage aggregate relocation show -node source-node

The following command shows the progress of the aggregates that are being moved to Node3:

When the relocation is complete, the output of this command shows each aggregate with a relocation status of "Done".

Delete a local tier (aggregate)

You can delete a local tier (aggregate) if there are no volumes on the local tier.

The storage aggregate delete command deletes a storage aggregate. The command fails if there are volumes present on the aggregate. If the aggregate has an object store attached to it, then in addition to deleting the aggregate, the command deletes the objects in the object store as well. No changes are made to the object store configuration as part of this command.

The following example deletes an aggregate named "aggr1":

```
> storage aggregate delete -aggregate aggr1
```

Commands for aggregate relocation

There are specific ONTAP commands for relocating aggregate ownership within an HA pair.

If you want to	Use this command
Start the aggregate relocation process	storage aggregate relocation start
Monitor the aggregate relocation process	storage aggregate relocation show

Related information

• ONTAP command reference

Commands for managing aggregates

You use the storage aggregate command to manage your aggregates.

If you want to	Use this command
Display the size of the cache for all Flash Pool aggregates	storage aggregate show -fields hybrid- cache-size-total -hybrid-cache-size -total >0
Display disk information and status for an aggregate	storage aggregate show-status
Display spare disks by node	storage aggregate show-spare-disks
Display the root aggregates in the cluster	storage aggregate show -has-mroot true
Display basic information and status for aggregates	storage aggregate show
Display the type of storage used in an aggregate	storage aggregate show -fields storage- type
Bring an aggregate online	storage aggregate online
Delete an aggregate	storage aggregate delete
Put an aggregate into the restricted state	storage aggregate restrict
Rename an aggregate	storage aggregate rename
Take an aggregate offline	storage aggregate offline
Change the RAID type for an aggregate	storage aggregate modify -raidtype

Related information

ONTAP command reference

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.