



Add (create) a local tier (aggregate)

ONTAP 9

NetApp
March 08, 2024

Table of Contents

- Add (create) a local tier (aggregate) 1
 - Add a local tier (create an aggregate) 1
 - Workflow to add a local tier (aggregate) 1
- Determine the number of disks or disk partitions required for a local tier (aggregate) 4
- Decide which method to use to create local tiers (aggregates) 5
- Add local tiers automatically (create aggregates with auto-provisioning) 7
- Add local tiers (create aggregates) manually 10

Add (create) a local tier (aggregate)

Add a local tier (create an aggregate)

To add a local tier (create an aggregate), you follow a specific workflow.

You determine the number of disks or disk partitions that you need for the local tier and decide which method to use to create the local tier. You can add local tiers automatically by letting ONTAP assign the configuration, or you can manually specify the configuration.

- [Workflow to add a local tier \(aggregate\)](#)
- [Determine the number of disks or disk partitions required for a local tier \(aggregate\)](#)
- [Decide which local tier \(aggregate\) creation method to use](#)
- [Add local tiers \(aggregates\) automatically](#)
- [Add local tiers \(aggregates\) manually](#)

Workflow to add a local tier (aggregate)

Creating local tiers (aggregates) provides storage to volumes on your system.


The workflow for creating local tiers (aggregates) is specific to the interface you use—System Manager or the CLI:

System Manager workflow

Use System Manager to add (create) a local tier

System Manager creates local tiers based on recommended best practices for configuring local tiers.


Beginning with ONTAP 9.11.1, you can decide to configure local tiers manually if you want a different configuration than the one recommended during the automatic process to add a local tier.



CLI workflow

Use the CLI to add (create) an aggregate

Beginning with ONTAP 9.2, ONTAP can provide recommended configurations when you create aggregates (auto-provisioning). If the recommended configurations, based on best practices, are appropriate in your environment, you can accept them to create the aggregates. Otherwise, you can create aggregates manually.



Determine the number of disks or disk partitions required for a local tier (aggregate)

You must have enough disks or disk partitions in your local tier (aggregate) to meet system and business requirements. You should also have the recommended number of hot spare disks or hot spare disk partitions to minimize the potential of data loss.

Root-data partitioning is enabled by default on certain configurations. Systems with root-data partitioning enabled use disk partitions to create local tiers. Systems that do not have root-data partitioning enabled use unpartitioned disks.

You must have enough disks or disk partitions to meet the minimum number required for your RAID policy and enough to meet your minimum capacity requirements.



In ONTAP, the usable space of the drive is less than the physical capacity of the drive. You can find the usable space of a specific drive and the minimum number of disks or disk partitions required for each RAID policy in the [Hardware Universe](#).

Determine usable space of a specific disk


The procedure you follow depends on the interface you use—System Manager or the CLI:

System Manager

Use System Manager to determine usable space of disks

Perform the following steps to view the usable size of a disk:

Steps

1. Go to **Storage > Tiers**
2. Click  next to the name of the local tier.
3. Select the **Disk Information** tab.

CLI

Use the CLI to determine usable space of disks

Perform the following step to view the usable size of a disk:

Step

1. Display spare disk information:

```
storage aggregate show-spare-disks
```

In addition to the number of disks or disk partitions necessary to create your RAID group and meet your capacity requirements, you should also have the minimum number of hot spare disks or hot spare disk partitions recommended for your aggregate:

- For all flash aggregates, you should have a minimum of one hot spare disk or disk partition.



The AFF C190 defaults to no spare drive. This exception is fully supported.

- For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions.
- For SSD storage pools, you should have a minimum of one hot spare disk for each HA pair.
- For Flash Pool aggregates, you should have a minimum of two spare disks for each HA pair. You can find more information on the supported RAID policies for Flash Pool aggregates in the [Hardware Universe](#).
- To support the use of the Maintenance Center and to avoid issues caused by multiple concurrent disk failures, you should have a minimum of four hot spares in multi-disk carriers.

Related information

[NetApp Hardware Universe](#)

[NetApp Technical Report 3838: Storage Subsystem Configuration Guide](#)

Decide which method to use to create local tiers (aggregates)

Although ONTAP provides best-practice recommendations for adding local tiers automatically (creating aggregates with auto-provisioning), you must determine whether the recommended configurations are supported in your environment. If they are not, you

must make decisions about RAID policy and disk configuration and then create the local tiers manually.

When a local tier is created automatically, ONTAP analyzes available spare disks in the cluster and generates a recommendation about how spare disks should be used to add local tiers according to best practices. ONTAP displays the recommended configurations. You can accept the recommendations or add the local tiers manually.

Before you can accept ONTAP recommendations

If any of the following disk conditions are present, they must be addressed before accepting the recommendations from ONTAP:

- Missing disks
- Fluctuation in spare disk numbers
- Unassigned disks
- Non-zeroed spares
- Disks undergoing maintenance testing

The `storage aggregate auto-provision man` page contains more information about these requirements.

When you must use the manual method

In many cases, the recommended layout of the local tier will be optimal for your environment. However, if your cluster is running ONTAP 9.1 or earlier, or your environment includes the following configurations, you must create the local tier using the manual method.



Beginning with ONTAP 9.11.1, you can manually add local tiers with System Manager.

- Aggregates using third-party array LUNs
- Virtual disks with Cloud Volumes ONTAP or ONTAP Select
- MetroCluster system
- SyncMirror
- MSATA disks
- FlashPool tiers (aggregates)
- Multiple disk types or sizes are connected to the node

Select the method to create local tiers (aggregates)

Choose which method you want to use:

- [Add \(create\) local tiers \(aggregates\) automatically](#)
- [Add \(create\) local tiers \(aggregates\) manually](#)

Related information

[ONTAP 9 commands](#)

Add local tiers automatically (create aggregates with auto-provisioning)

If the best-practice recommendation that ONTAP provides for automatically adding a local tier (creating an aggregate with auto-provisioning) is appropriate in your environment, you can accept the recommendation and let ONTAP add the local tier.

Before you begin

Disks must be owned by a node before they can be used in a local tier (aggregate). If your cluster is not configured to use automatic disk ownership assignment, you must [assign ownership manually](#).

System Manager

Steps

1. In System Manager, click **Storage > Tiers**.
2. From the **Tiers** page, click [+ Add Local Tier](#) to create a new local tier:

The **Add Local Tier** page shows the recommended number of local tiers that can be created on the nodes and the usable storage available.

3. Click **Recommended details** to view the configuration recommended by System Manager.

System Manager displays the following information beginning with ONTAP 9.8:

- **Local tier name** (you can edit the local tier name beginning with ONTAP 9.10.1)
- **Node name**
- **Usable size**
- **Type of storage**

Beginning with ONTAP 9.10.1, additional information is displayed:

- **Disks:** showing the number, size, and type of the disks
- **Layout:** showing the RAID group layout, including which disks are parity or data and which slots are unused.
- **Spare disks:** showing the node name, the number and size of spare disks, and the type of storage.

4. Perform one of the following steps:

If you want to...	Then do this...
Accept the recommendations from System Manager.	Proceed to the step for configuring the Onboard Key Manager for encryption .
Manually configure the local tiers and not use the recommendations from System Manager.	<p>Proceed to Add a local tier (create aggregate) manually:</p> <ul style="list-style-type: none">• For ONTAP 9.10.1 and earlier, follow the steps to use the CLI.• Beginning with ONTAP 9.11.1, follow the steps to use System Manager.

5. (Optional): If the Onboard Key Manager has been installed, you can configure it for encryption. Check the **Configure Onboard Key Manager for encryption** check box.
 - a. Enter a passphrase.
 - b. Enter the passphrase again to confirm it.
 - c. Save the passphrase for future use in case the system needs to be recovered.
 - d. Back up the key database for future use.
6. Click **Save** to create the local tier and add it to your storage solution.

CLI

You run the `storage aggregate auto-provision` command to generate aggregate layout recommendations. You can then create aggregates after reviewing and approving ONTAP recommendations.

What you'll need

ONTAP 9.2 or later must be running on your cluster.

About this task

The default summary generated with the `storage aggregate auto-provision` command lists the recommended aggregates to be created, including names and usable size. You can view the list and determine whether you want to create the recommended aggregates when prompted.

You can also display a detailed summary by using the `-verbose` option, which displays the following reports:

- Per node summary of new aggregates to create, discovered spares, and remaining spare disks and partitions after aggregate creation
- New data aggregates to create with counts of disks and partitions to be used
- RAID group layout showing how spare disks and partitions will be used in new data aggregates to be created
- Details about spare disks and partitions remaining after aggregate creation

If you are familiar with the auto-provision method and your environment is correctly prepared, you can use the `-skip-confirmation` option to create the recommended aggregate without display and confirmation. The `storage aggregate auto-provision` command is not affected by the CLI session `-confirmations` setting.

The `storage aggregate auto-provision` [man page](#) contains more information about the aggregate layout recommendations.

Steps

1. Run the `storage aggregate auto-provision` command with the desired display options.
 - no options: Display standard summary
 - `-verbose` option: Display detailed summary
 - `-skip-confirmation` option: Create recommended aggregates without display or confirmation
2. Perform one of the following steps:

If you want to...	Then do this...
-------------------	-----------------

Accept the recommendations from ONTAP.

Review the display of recommended aggregates, and then respond to the prompt to create the recommended aggregates.

```
myA400-44556677::> storage aggregate auto-
provision
Node                               New Data Aggregate
Usable Size
-----
-----
myA400-364                         myA400_364_SSD_1
3.29TB
myA400-363                         myA400_363_SSD_1
1.46TB
-----
-----
Total:                             2      new data aggregates
4.75TB

Do you want to create recommended
aggregates? {y|n}: y

Info: Aggregate auto provision has
started. Use the "storage aggregate
      show-auto-provision-progress"
command to track the progress.

myA400-44556677::>
```

Manually configure the local tiers and **not** use the recommendations from ONTAP.

Proceed to [Add a local tier \(create aggregate\) manually](#).

Related information

[ONTAP 9 Commands](#)

Add local tiers (create aggregates) manually

If you do not want to add a local tier (create a aggregate) using the best-practice recommendations from ONTAP, you can perform the process manually.

Before you begin

Disks must be owned by a node before they can be used in a local tier (aggregate). If your cluster is not configured to use automatic disk ownership assignment, you must [assign ownership manually](#).

System Manager

Beginning with ONTAP 9.11.1, if you do not want to use the configuration recommended by System Manager to create a local tier, you can specify the configuration you want.

Steps

1. In System Manager, click **Storage > Tiers**.
2. From the **Tiers** page, click [+ Add Local Tier](#) to create a new local tier:

The **Add Local Tier** page shows the recommended number of local tiers that can be created on the nodes and the usable storage available.

3. When System Manager displays the storage recommendation for the local tier, click **Switch to Manual Local Tier Creation** in the **Spare Disks** section.

The **Add Local Tier** page displays fields that you use to configure the local tier.

4. In the first section of the **Add Local Tier** page, complete the following:
 - a. Enter the name of the local tier.
 - b. (Optional): Check the **Mirror this local tier** check box if you want to mirror the local tier.
 - c. Select a disk type.
 - d. Select the number of disks.
5. In the **RAID Configuration** section, complete the following:
 - a. Select the RAID type.
 - b. Select the RAID group size.
 - c. Click RAID allocation to view how the disks are allocated in the group.
6. (Optional): If the Onboard Key Manager has been installed, you can configure it for encryption in the **Encryption** section of the page. Check the **Configure Onboard Key Manager for encryption** check box.
 - a. Enter a passphrase.
 - b. Enter the passphrase again to confirm it.
 - c. Save the passphrase for future use in case the system needs to be recovered.
 - d. Back up the key database for future use.
7. Click **Save** to create the local tier and add it to your storage solution.

CLI

Before you create aggregates manually, you should review disk configuration options and simulate creation.

Then you can issue the `storage aggregate create` command and verify the results.

What you'll need

You must have determined the number of disks and the number of hot spare disks you need in the aggregate.

About this task

If root-data-data partitioning is enabled and you have 24 solid-state drives (SSDs) or fewer in your

configuration, it is recommended that your data partitions be assigned to different nodes.

The procedure for creating aggregates on systems with root-data partitioning and root-data-data partitioning enabled is the same as the procedure for creating aggregates on systems using unpartitioned disks. If root-data partitioning is enabled on your system, you should use the number of disk partitions for the `-diskcount` option. For root-data-data partitioning, the `-diskcount` option specifies the count of disks to use.



When creating multiple aggregates for use with FlexGroups, aggregates should be as close in size as possible.

The `storage aggregate create` man page contains more information about aggregate creation options and requirements.

Steps

1. View the list of spare disk partitions to verify that you have enough to create your aggregate:

```
storage aggregate show-spare-disks -original-owner node_name
```

Data partitions are displayed under `Local Data Usable`. A root partition cannot be used as a spare.

2. Simulate the creation of the aggregate:

```
storage aggregate create -aggregate aggregate_name -node node_name  
-raidtype raid_dp -diskcount number_of_disks_or_partitions -simulate true
```

3. If any warnings are displayed from the simulated command, adjust the command and repeat the simulation.

4. Create the aggregate:

```
storage aggregate create -aggregate aggr_name -node node_name -raidtype  
raid_dp -diskcount number_of_disks_or_partitions
```

5. Display the aggregate to verify that it was created:

```
storage aggregate show-status aggregate_name
```

Related information

[ONTAP 9 commands](#)

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.