



Provisioning storage through aggregates

System Manager Classic

NetApp
January 13, 2022

Table of Contents

- Provisioning storage through aggregates 1
 - Provisioning storage by creating an aggregate based on storage recommendations 1
 - Provisioning storage by creating an aggregate manually 2
 - Provisioning storage by creating a Flash Pool aggregate manually 3
 - Provisioning storage by creating a SnapLock aggregate manually 4
 - Provisioning storage by creating a FabricPool-enabled aggregate manually 5

Provisioning storage through aggregates

You can create an aggregate based on storage recommendations or manually depending on your requirement. You can create Flash Pool aggregates, SnapLock aggregates, and a FabricPool-enabled aggregates to provide storage for one or more volumes by using System Manager.

Before you begin

You must have enough spare disks to create an aggregate.

About this task

You cannot perform the following actions by using System Manager:

- Combine disks of different sizes even if there are enough spare disks of different sizes.
You can initially create an aggregate with disks of the same size and then add disks of a different size later.
- Combine disks with different checksum types.
You can initially create an aggregate with a single checksum type and add storage of a different checksum type later.

Related information

[Aggregates window](#)

[Storage Tiers window](#)

Provisioning storage by creating an aggregate based on storage recommendations

You can use System Manager to create an aggregate based on storage recommendations. System Manager analyzes the configuration of your storage system and provides storage recommendations such as the number of aggregates that will be created, the available nodes, and the available spare disks.

About this task

- You cannot create an aggregate based on storage recommendations in Cloud Volumes ONTAP, ONTAP Select, and MetroCluster configurations.
- Errors, if any, are displayed on the screen.

You can fix these errors and then create an aggregate based on the storage recommendations, or you can create an aggregate manually.

Steps

1. Create an aggregate by using one of the following methods:
 - Click **Applications & Tiers > Storage Tiers > Add Aggregate**.
 - Click **Storage > Aggregate & Disks > Aggregates > Create**.

2. Review the storage recommendations, and then click **Submit**.

The Information dialog box displays the status of the aggregates.

3. Click **Run in Background** to navigate to the **Aggregates** window.
4. Click **Refresh** to view the aggregates that are created.

Provisioning storage by creating an aggregate manually

You can manually create an aggregate that consists of only HDDs or only SSDs by using System Manager.

Before you begin

All of the disks must be of the same size.

About this task

- If you are creating an aggregate on a four-node cluster in ONTAP Select, the mirrored aggregate option is selected by default.
- Starting with ONTAP 9.0, you can create aggregates with disk size equal to or larger than 10 TB.
- If the disk type of the aggregate disks is FSAS or MSATA, and the disk size is equal to or larger than 10 TB, then RAID-TEC is the only available RAID type.

Steps

1. Create an aggregate by using one of the following methods:
 - Click **Applications & Tiers > Storage Tiers > Add Aggregate**.
 - Click **Storage > Aggregate & Disks > Aggregates > Create**.
2. Enable the **Manually Create Aggregate** option to create an aggregate.
3. To create an aggregate:
 - a. Specify the name of the aggregate, the disk type, and the number of disks or partitions to include in the aggregate.

The minimum hot spare rule is applied to the disk group that has the largest disk size.

- b. Modify the RAID configuration of the aggregate:
 - i. Click **Change**.
 - ii. In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.

Shared disks support two RAID types: RAID DP and RAID-TEC.

- iii. Click **Save**.
- c. If you want to mirror the aggregate, select the **Mirror this aggregate** check box.

For MetroCluster configurations, creating unmirrored aggregates is restricted. Therefore, the mirroring option is enabled by default for MetroCluster configurations.

4. Click **Create**.

Results

The aggregate is created with the specified configuration, and is added to the list of aggregates in the Aggregates window.

Provisioning storage by creating a Flash Pool aggregate manually

You can use System Manager to create a Flash Pool aggregate manually, or to convert an existing HDD aggregate to a Flash Pool aggregate by adding SSDs. When you create a new HDD aggregate, you can provision an SSD cache to it and create a Flash Pool aggregate.

Before you begin

- You must be aware of the platform-specific best practices and workload-specific best practices for the Flash Pool aggregate SSD tier size and configuration.
- All of the HDDs must be in the zeroed state.
- If you want to add SSDs to the aggregate, all of the existing SSDs and dedicated SSDs must be of the same size.

About this task

- You cannot use partitioned SSDs while creating a Flash Pool aggregate.
- You cannot mirror the aggregates if the cache source is storage pools.
- Starting with ONTAP 9.0, you can create aggregates with disk size equal to or larger than 10 TB.
- If the disk type of the aggregate disks is FSAS or MSATA, and the disk size is equal to or larger than 10 TB, then RAID-TEC is the only option available for RAID type.

Steps

1. Create an aggregate by using one of the following methods:
 - Click **Applications & Tiers > Storage Tiers > Add Aggregate**.
 - Click **Storage > Aggregate & Disks > Aggregates > Create**.
2. Enable the **Manually Create Aggregate** option to create an aggregate.
3. In the **Create Aggregate** window, specify the name of the aggregate, the disk type, and the number of disks or partitions to include for the HDDs in the aggregate.
4. If you want to mirror the aggregate, select the **Mirror this aggregate** check box.

For MetroCluster configurations, creating unmirrored aggregates is restricted. Therefore, the mirroring option is enabled by default for MetroCluster configurations.

5. Click **Use Flash Pool Cache with this aggregate**.
6. Specify the cache source:

If you want to select the cache source as...	Then...
Storage pools	<ol style="list-style-type: none"> a. Select Storage pools as the Cache Source. b. Select the storage pool from which the cache can be obtained, and then specify the cache size. c. Modify the RAID type, if required.
Dedicated SSDs	<ol style="list-style-type: none"> a. Select Dedicated SSDs as the Cache Source. b. Select the SSD size and the number of SSDs to include in the aggregate. c. Modify the RAID configuration, if required: <ol style="list-style-type: none"> i. Click Change. ii. In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size. iii. Click Save.

7. Click **Create**.

Results

The Flash Pool aggregate is created with the specified configuration, and is added to the list of aggregates in the Aggregates window.

Related information

[How storage pool works](#)

[NetApp Technical Report 4070: Flash Pool Design and Implementation](#)

Provisioning storage by creating a SnapLock aggregate manually

You can use System Manager to create a SnapLock Compliance aggregate or a SnapLock Enterprise aggregate manually. You can create SnapLock volumes on these aggregates, which provide “write once, read many” (WORM) capabilities.

Before you begin

The SnapLock license must have been added.

About this task

- In MetroCluster configurations, you can create only SnapLock Enterprise aggregates.
- For array LUNs, only SnapLock Enterprise aggregates are supported.
- Starting with ONTAP 9.0, you can create aggregates with disk size equal to or larger than 10 TB.
- If the disk type of the aggregate disks is FSAS or MSATA, and the disk size is equal to or larger than 10 TB, then RAID-TEC is the only option available for RAID type.

- Starting with ONTAP 9.1, you can create a SnapLock aggregate on an AFF platform.

Steps

1. Create a SnapLock aggregate by using one of the following methods:
 - Click **Applications & Tiers > Storage Tiers > Add Aggregate**.
 - Click **Storage > Aggregate & Disks > Aggregates > Create**.
2. Enable the **Manually Create Aggregate** option to create an aggregate.
3. To create a SnapLock aggregate:
 - a. Specify the name of the aggregate, the disk type, and the number of disks or partitions to include in the aggregate.

You cannot change the name of a SnapLock Compliance aggregate after you create the aggregate.

The minimum hot spare rule is applied to the disk group that has the largest disk size.

- b. Modify the RAID configuration of the aggregate:
 - i. Click **Change**.
 - ii. In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.

Shared disks support two RAID types: RAID-DP and RAID-TEC.
 - iii. Click **Save**.
- c. Specify the SnapLock type.
- d. If you have not initialized the system ComplianceClock, select the **Initialize ComplianceClock** check box.

This option is not displayed if the ComplianceClock is already initialized on the node.



You must ensure that the current system time is correct. The ComplianceClock is set based on the system clock. Once the ComplianceClock is set, you cannot modify or stop the ComplianceClock.

- e. If you want to mirror the aggregate, select the **Mirror this aggregate** check box.

For MetroCluster configurations, creating unmirrored aggregates is restricted. Therefore, the mirroring option is enabled by default for MetroCluster configurations.

By default, the mirroring option is disabled for SnapLock Compliance aggregates.

4. Click **Create**.

Provisioning storage by creating a FabricPool-enabled aggregate manually

You can use System Manager to create a FabricPool-enabled aggregate manually or to convert an existing SSD aggregate to a FabricPool-enabled aggregate by attaching a cloud tier to the SSD aggregate.

Before you begin

- You must have created a cloud tier and attached it to the cluster in which the SSD aggregate resides.
- An on-premises cloud tier must have been created.
- A dedicated network connection must exist between the cloud tier and the aggregate.

About this task

The following object stores can be used as cloud tiers:

- StorageGRID
- Alibaba Cloud (Starting with System Manager 9.6)
- Amazon Web Services (AWS) Simple Storage Service (S3)
- Amazon Web Services (AWS) Commercial Cloud Service (C2S)
- Microsoft Azure Blob storage
- IBM Cloud
- Google Cloud



- Azure Stack, which is an on-premises Azure services, is not supported.
- If you want to use any object store as a cloud tier, other than StorageGRID, you must have the FabricPool capacity license.

Steps

1. Create a FabricPool-enabled aggregate by using one of the following methods:
 - Click **Applications & Tiers > Storage Tiers > Add Aggregate**.
 - Click **Storage > Aggregate & Disks > Aggregates > Create**.
2. Enable the **Manually Create Aggregate** option to create an aggregate.
3. Create a FabricPool-enabled aggregate:
 - a. Specify the name of the aggregate, the disk type, and the number of disks or partitions to include in the aggregate.



Only all flash (all SSD) aggregates support FabricPool-enabled aggregates.

The minimum hot spare rule is applied to the disk group that has the largest disk size.

- b. Modify the RAID configuration of the aggregate:
 - i. Click **Change**.
 - ii. In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.

Shared disks support two RAID types: RAID-DP and RAID-TEC.
 - iii. Click **Save**.
4. Select the **FabricPool** checkbox, and then select a cloud tier from the list.
 5. Click **Create**.

Copyright Information

Copyright © 2022 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.