



storage pool commands

ONTAP 9.3 commands

NetApp
August 29, 2024

Table of Contents

- storage pool commands 1
 - storage pool add 1
 - storage pool create 3
 - storage pool delete 4
 - storage pool reassign 5
 - storage pool show-aggregate 6
 - storage pool show-available-capacity 7
 - storage pool show-disks 9
 - storage pool show 10

storage pool commands

storage pool add

Add disks to a storage pool

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `storage pool add` command increases the total capacity of an existing storage pool by adding the specified SSDs to the storage pool. The disks are split into four equal partitions and added to each of the allocation units of the storage pool. If any allocation units from the storage pool have already been allocated to an aggregate, the cache or usable capacity of that aggregate is increased depending on whether it is a Flash Pool or an All-Flash aggregate .

If capacity from a storage pool is already provisioned into a Flash Pool aggregate, the same storage pool cannot be used to provision an All-Flash aggregate and vice-versa.

For provisioning storage pool capacity into All-Flash aggregates, the Vserver option `raid.storagepool.data.enable` must be set to `true` . The storage pool data enabled mode of operation is not currently supported by OnCommand management software.

For example, if an SSD with a usable size of 745 GB is added to a storage pool that is part of four aggregates, each aggregate will grow its cache or usable capacity by 186.2 GB. If a different allocation is desired, create a new storage pool using the [storage pool create](#) command.

Parameters

-storage-pool <storage pool name> - Storage Pool Name

This parameter specifies the storage pool to which disks are to be added.

{ [-disk-count <integer>] - Number of Disks to Add in Storage Pool

This parameter specifies the number of disks that are to be added to the storage pool. The disks to be added come from the pool of spare disks.

[-nodes {<nodename>|local}] - Nodes From Which Spares Should be Selected

This parameter specifies a list of nodes from which SSD disks are selected for addition to the storage pool. If this parameter is not specified, disks to be added to the storage pool can be selected from both the nodes sharing the storage pool. Use this parameter to restrict the selection of spare disks to one particular node.

[-disk-list <disk path name>,...] - List of Spare Disks }

This parameter specifies a list of disks to be added to the storage pool. In an HA configuration, SSDs being added to a storage pool can be owned by either node in the HA pair.

{ [-quiet <true>] - Confirmations off (privilege: advanced)

When set to `true` , this parameter specifies the operation should be executed without pausing for confirmation.

[-simulate <true>] - Simulate Storage Pool Addition

When set to *true*, this parameter specifies the operation should be performed as a simulation. The command reports which aggregates would grow automatically as a result of adding the disks to the storage pool. The disks are not added to the storage pool.

Examples

In this example, the user requests a report detailing the changes that would occur if a new disk is added to the storage pool *SP1*. In this case, 186.2 GB of cache is added to the Flash Pool aggregates *nodeA_flashpool_1* and *nodeB_flashpool_1*. There are two unprovisioned allocation units in the storage pool and therefore the storage pool available capacity also grows by 372.5 GB.

```
cluster1::> storage pool add -storage-pool SP1 -disk-list 1.0.23 -simulate
```

This operation will result in capacity being allocated in the following way:

Container Name	Capacity To Be Added	Current Size	New Size
nodeA_flashpool_1	186.2GB	558.7GB	744.9GB
nodeB_flashpool_1	186.2GB	558.7GB	744.9GB
(Available Capacity)	372.5GB	1.09TB	1.45TB

The following example adds one disk to a storage pool named *SP1*. The spare disks are selected from either local node or its partner or both based on spare availability.

```
cluster-1::> storage pool add -storage-pool SP1 -disk-count 1
```

Info: The following disks will be added to storage pool "SP1":

Disk	Size	Type	Owner
1.0.12	744.9GB	SSD	cluster-1-01

New Allocation Unit Size: 744.8GB

Capacity will be allocated in the following way:

Container Name	Capacity To Be Added	Current Size	New Size
nodeA_flashpool_1	186.2GB	558.7GB	744.9GB
nodeB_flashpool_1	186.2GB	558.7GB	744.9GB
(Available Capacity)	372.5GB	1.09TB	1.45TB

Are you sure you want to continue with this operation?

{y|n}: y

[Job 48] Job succeeded: storage pool add job for "SP1" completed successfully

Related Links

- [storage pool create](#)

storage pool create

Create a new storage pool

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `storage pool create` command creates an SSD storage pool using a given list of spare SSDs.

When a storage pool is created, Data ONTAP splits the capacity provided by the SSDs into four equally-sized allocation units. In an HA configuration, two allocation units (containing 50% of the total capacity) are assigned to each node in the HA pair. This assignment can be modified using the [storage pool reassign](#) command.

After the storage pool is created, its allocation units can be provisioned into Flash Pool or All-Flash aggregates using the [storage aggregate add-disks](#) command and the `-storage-pool` parameter.

If capacity from a storage pool is already provisioned into a Flash Pool aggregate, the same storage pool cannot be used to provision an All-Flash aggregate and vice versa.

For provisioning storage pool capacity into All-Flash aggregates, the `vserver` option `raid.storagepool.data.enable` must be set to `true`. The storage pool data enabled mode of operation is not currently supported by OnCommand management software.

Parameters

-storage-pool <storage pool name> - Storage Pool Name

This parameter specifies the name of the storage pool that is to be created. The SSDs are partitioned and placed into the new storage pool.

{ [-nodes {<nodename>|local}] - Nodes Sharing the Storage Pool

This parameter specifies a list of nodes from which SSD disks are selected to create the storage pool. If two nodes are specified then they need to be in HA configuration. Spare disks are selected from either node or its partner or both. If this parameter is not specified, storage pool will be created by selecting disks from either the node or its partner or both from where command is run.

-disk-count <integer> - Number of Disks in Storage Pool

This parameter specifies the number of disks that are to be included in the storage pool. The disks in this newly created storage pool come from the pool of spare disks. The smallest disks in this pool are added to the storage pool first, unless you specify the `-disk-size` parameter.

[-disk-size {<integer>[KB|MB|GB|TB|PB]}] - Disk Size

This parameter specifies the size of the disks on which the storage pool is to be created. Disks with a usable size between 95% and 105% of the specified size are selected.

| -disk-list <disk path name>,... - Disk List for Storage Pool Creation }

This parameter specifies a list of SSDs to be included in the new storage pool. The SSDs must be spare

disks and can be owned by either node in an HA pair.

[`-simulate <true>`] - Simulate Storage Pool Creation

This option simulates the storage pool creation and prints the allocation unit size that would be used for the storage pool.

Examples

The following example creates a storage pool named SP1. The storage pool contains 3 SSD disks, the spare disks selected are from either local node, or its partner or both based on spare availability.

```
cluster1::> storage pool create -storage-pool SP1 -disk-count 3
```

The following example creates a storage pool named SP2. The storage pool contains 3 SSD disks, the spare disks selected are from either node0, or its partner node1 or both based on spare availability.

```
cluster1::> storage pool create -storage-pool SP2 -disk-count 3 -nodes  
node0,node1
```

The following example creates a storage pool named SP3 from four SSDs using disk list.

```
cluster1::> storage pool create -storage-pool SP3 -disk-list 1.0.13,  
1.0.15, 1.0.17, 1.0.19
```

Related Links

- [storage pool reassign](#)
- [storage aggregate add-disks](#)

storage pool delete

Delete an existing storage pool

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `storage pool delete` command deletes an existing SSD storage pool. At the end of the operation, the SSDs are converted back to spare disks.

Parameters

`-storage-pool <storage pool name>` - Storage Pool Name

This parameter specifies the storage pool that you want to delete. You can delete the storage pool only if all of the allocation units in the storage pool are available.

Examples

Verify that storage pool *SP3* is ready for deletion by confirming it has four available allocation units and then delete it.

```
cluster1::> storage pool show-available-capacity -storage-pool SP3
```

Node	Storage Pool	Storage Type	SyncMirror Pool	Allocation Unit size	Unit Count	Total Usable Size
node-a	SP3	SSD	Pool0	372.5GB	2	744.9GB
node-b	SP3	SSD	Pool0	372.5GB	2	744.9GB

```
2 entries were displayed.
```

```
cluster1::> storage pool delete -storage-pool SP3
```

Warning: Are you sure you want to delete storage pool "SP3"? {y|n}: y
[Job 313] Job succeeded: storage pool delete job for "SP3" completed successfully

storage pool reassign

Reassign capacity from one node to another node in storage pool

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `storage pool reassign` command changes the ownership of unprovisioned (available) storage pool allocation units from one HA partner to the other for an existing storage pool.

Parameters

-storage-pool <storage pool name> - Storage Pool Name

This parameter specifies the storage pool within which available capacity is reassigned from one node to another.

-from-node {<nodename>|local} - Reassign Available Capacity from This Node

This parameter specifies the name of the node that currently owns the allocation units.

-to-node {<nodename>|local} - Reassign Available Capacity to This Node

This parameter specifies the name of the node that will now own the allocation units.

-allocation-units <integer> - Allocation Units

This parameter specifies the number of allocation units to be reassigned.

Examples

Move an available allocation unit from node-b to node-a in preparation for provisioning the allocation units on node-a.

```
cluster1::*> storage pool show-available-capacity -storage-pool SP2
```

Node	Storage Pool	Storage Type	SyncMirror Pool	Allocation Unit size	Unit Count	Total Usable Size
node-a	SP2	SSD	Pool0	744.9GB	2	1.45TB
node-b	SP2	SSD	Pool0	744.9GB	1	744.9GB

2 entries were displayed.

```
cluster1::*> storage pool reassign -storage-pool SP2 -from-node node-b -to  
-node node-a -allocation-units 1  
[Job 310] Job succeeded: storage pool reassign job for "SP2" completed  
successfully
```

```
cluster1::*> storage pool show-available-capacity -storage-pool SP2
```

Node	Storage Pool	Storage Type	SyncMirror Pool	Allocation Unit size	Unit Count	Total Usable Size
node-a	SP2	SSD	Pool0	744.9GB	3	2.18TB
node-b	SP2	SSD	Pool0	744.9GB	0	0B

2 entries were displayed.

storage pool show-aggregate

Display aggregates provisioned from storage pools

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `storage pool show-aggregate` command displays allocation information for SSD storage pools in the cluster. The command output depends upon the parameter or parameters specified with the command. If no parameters are specified, the command displays information about allocations of all storage pools in the cluster.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[`-storage-pool <storage pool name>`] - Name of Storage Pool

Selects the storage pools that match this parameter value.

[`-aggregate <aggregate name>`] - Aggregate

Selects the storage pools that match this parameter value.

[`-capacity {<integer>[KB|MB|GB|TB|PB]}`] - Capacity

Selects the storage pools that match this parameter value.

Capacity includes space provided by data and parity portions of each allocation unit. Only the data portions of each allocation unit contribute to the cache or usable capacity of Flash Pool or All-Flash aggregates respectively .

[`-allocated-unit-count <integer>`] - Number of AU's Assigned to This Aggregate

Selects the storage pools that match this parameter value.

[`-original-owner <text>`] - Original Owner Name

Selects the storage pools that match this parameter value.

[`-node {<nodename>|local}`] - Node

Selects the storage pools that match this parameter value.

Examples

Display information about the aggregate or aggregates using a storage pool called *SP2* :

```
cluster1::> storage pool show-aggregate -storage-pool SP2 -instance
Name of Storage Pool: SP2
                        Aggregate: node2_flashpool_1
                        Capacity: 744.9GB
Number of AU's Assigned to This Aggregate: 1
                        Original Owner Name: node2
                        Node: node2
```

storage pool show-available-capacity

Display available capacity of storage pools

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `storage pool show-available-capacity` command displays information about available capacity in SSD storage pools on each node in the cluster. The command output depends upon the parameter or parameters specified with the command. If no parameters are specified, the command displays information about available capacities in all shared pools in the cluster.

Storage pool available capacity is data storage space that has not yet been provisioned into Flash Pool or All-

Flash aggregates. Allocation units might be provisioned into aggregates using the [storage aggregate add-disks](#) command and the `-storage-pool` parameter.



All storage pool available capacity can be provisioned into aggregates. Available capacity within a storage pool is not used to protect against a disk failure. In the case of an SSD failure or predicted failure, Data ONTAP moves a suitable whole SSD spare disk from outside the storage pool into the storage pool and begins the recovery process (using either reconstruction or Rapid RAID Recovery, whichever is appropriate).

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-storage-pool <storage pool name>] - Name of Storage Pool

Selects the available capacities that match this parameter value.

[-node {<nodename>|local}] - Node

Selects the available capacities that match this parameter value.

[-allocation-unit-size {<integer>[KB|MB|GB|TB|PB]}] - Allocation Unit Size

Selects the available capacities that match this parameter value.

Allocation units are the units of storage capacity that are available to be provisioned into aggregates.

[-storage-type <SSD>] - Type of Storage Pool

Selects the available capacities that match this parameter value. Only the *SSD* type is supported for this version of Data ONTAP.

[-allocation-unit-count <integer>] - Number of Allocation Units Available

Selects the available capacities that match this parameter value.

Allocation units are the units of storage capacity that are available to be provisioned into aggregates. Each allocation unit is one minimum unit of allocation (MUA) and its capacity is given as `allocation-unit-size`.

[-syncmirror-pool <text>] - Syncmirror Pool

Selects the available capacities that match this parameter value.

The SyncMirror pool of an allocation unit must match the SyncMirror pool of the disks of the aggregate when adding allocation units into an aggregate.

Mirroring of aggregates that are provisioned from SSD storage pools is not supported.

[-available-size {<integer>[KB|MB|GB|TB|PB]}] - Total Usable Available Size

Selects the available capacities that match this parameter value.

The `available-size` is the sum of the capacities of the allocation units that are assigned but not yet provisioned. The amount of `available-size` that is contributed to the cache or usable capacity of an aggregate depends upon the RAID type used when provisioning the allocation units.

Examples

In this example, two nodes of an HA pair share available capacity from two storage pools, *SP1* and *SP2* . There are a total of 5 allocation units that have not yet been provisioned.

```
cluster1::> storage pool show-available-capacity
```

Node	Storage Pool	Storage Type	SyncMirror Pool	Allocation Unit size	Unit Count	Total Usable Size
node-a	SP1	SSD	Pool0	558.7GB	1	558.7GB
node-b	SP1	SSD	Pool0	558.7GB	1	558.7GB
node-a	SP2	SSD	Pool0	744.9GB	2	1.45TB
node-b	SP2	SSD	Pool0	744.9GB	1	744.9GB

Related Links

- [storage aggregate add-disks](#)

storage pool show-disks

Display disks in storage pools

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `storage pool show-disks` command displays information about disks in storage pools in the cluster. The command output depends on the parameter or parameters specified with the command. If no parameters are specified, the command displays information about all disks in all storage pools in the cluster.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-storage-pool <storage pool name>] - Name of Storage Pool

Selects the storage pools that match this parameter value.

[`-disk <disk path name>`] - Name of the disk

Selects the storage pools with the disks that match this parameter value.

[`-disk-type {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SSD | VMDISK | SSD-NVM}`] - Disk Type

Selects the storage pools with the disks that match this parameter value. Only the *SSD* type is supported for this version of Data ONTAP.

[`-usable-size {<integer>[KB|MB|GB|TB|PB]}`] - Disk Usable Size

Selects the storage pools with the disks that match this parameter value.

In this command, `usable-size` refers to the sum of the capacities of all of the partitions on the disk.

[`-total-size {<integer>[KB|MB|GB|TB|PB]}`] - Total Size

Selects the storage pools with the disks that match this parameter value.

[`-node-list <nodename>,...`] - List of Nodes

Selects the storage pools with the disks that are visible to all of the specified nodes.

Examples

Show information about SSDs in a storage pool called *SP2*.

```
cluster1::> storage pool show-disks -storage-pool SP2
```

Storage Pool Name: SP2

Storage

Disk	Type	Usable Size	Total Size
-----	-----	-----	-----
1.0.16	SSD	745.0GB	745.2GB
1.0.18	SSD	745.0GB	745.2GB
1.0.20	SSD	745.0GB	745.2GB
1.0.22	SSD	745.0GB	745.2GB

storage pool show

Display details of storage pools

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `storage pool show` command displays information about SSD storage pools in the cluster. By default, the command displays information about all storage pools in the cluster. You can specify parameters to limit the output to a specific set of storage pools.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `-fields ?` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-storage-pool <storage pool name>] - Storage Pool Name

Selects the storage pools that match this parameter value.

[-storage-pool-uuid <UUID>] - UUID of Storage Pool

Selects the storage pools that match this parameter value.

[-nodes {<nodename>|local}] - Nodes Sharing the Storage Pool

Selects the storage pools that match this parameter value.

In an HA pair, either node name may be specified.

[-disk-count <integer>] - Number of Disks in Storage Pool

Selects the storage pools that match this parameter value.

[-allocation-unit-size {<integer>[KB|MB|GB|TB|PB]}] - Allocation Unit Size

Selects the storage pools that match this parameter value.

Allocation units represent the unit of storage allocated to aggregates from this storage pool.

[-allocation-unit-data-size-raid4 {<integer>[KB|MB|GB|TB|PB]}] - Allocation Unit Data Size for RAID4

This parameter shows the amount of additional data capacity provided if an allocation unit from this storage pool was added to an aggregate with `-raidtype` as *raid4*.

[-allocation-unit-data-size-raid-dp {<integer>[KB|MB|GB|TB|PB]}] - Allocation Unit Data Size for RAID-DP

This parameter shows the amount of additional data capacity provided if an allocation unit from this storage pool was added to an aggregate with `-raidtype` as *raid_dp*.

[-allocation-unit-data-size-raid-tec {<integer>[KB|MB|GB|TB|PB]}] - Allocation Unit Data Size for RAID-TEC

This parameter shows the amount of additional data capacity provided if an allocation unit from this storage pool was added to an aggregate with `-raidtype` as *raid_tec*.

[-storage-type <SSD>] - Storage Type

Selects the storage pools that match this parameter value.

Only the *SSD* type is supported for this version of Data ONTAP.

[-pool-usable-size {<integer>[KB|MB|GB|TB|PB]}] - Storage Pool Usable Size

Selects the storage pools that match this parameter value.

The `pool-usable-size` is the sum of the capacities of the allocation units that are assigned to nodes but not yet provisioned. The amount of `pool-usable-size` that is contributed to the cache or usable capacity of an aggregate depends upon the RAID type used when provisioning the allocation units.

`[-pool-total-size {<integer>[KB|MB|GB|TB|PB] }]` - Storage Pool Total Size

Selects the storage pools that match this parameter value.

The `pool-total-size` is the sum of the capacities of allocation units belonging to this storage pool.

`[-is-healthy {true|false}]` - Is Pool Healthy?

Selects the storage pools that match this parameter value.

For storage pools with `is-healthy`false`` , the ``unhealthy-reason`` parameter provides more information.

`is-healthy` must be `true` to provision allocation units from a storage pool into an aggregate.

`[-pool-state <State of the Storage Pool>]` - State of the Storage Pool

Selects the storage pools that match this parameter value. Possible states are:

- normal - the storage pool is operating normally.
- degraded - the storage pool has one or more failed disks.
- creating - the storage pool is being created.
- deleting - the storage pool is being deleted.
- reassigning - allocation units are being reassigned from one node to another.
- growing - allocation units in the storage pool are expanding due to the addition of new capacity into the storage pool.

`[-unhealthy-reason <text>]` - Reason for Storage Pool Being Unhealthy

Selects the storage pools that match this parameter value.

The message provided gives additional details about why the storage pool is unhealthy.

`[-current-operation-job-id <integer>]` - Job ID of the Currently Running Operation

Selects the storage pools that match this parameter value.

Long-running operations associated with storage pools will be managed via jobs. For example, if you provision allocation units from a storage pool into an aggregate and the disks associated with the storage pool need to be zeroed, the operation will be completed via a job.

Examples

Display the storage pools in the cluster.

```
cluster1::> storage pool show
```

Storage Pool	Type	#Disks	Nodes	Total Size
LargeSP	SSD	10	noda-a,node-b	7.27TB
SmallSP	SSD	2	noda-a,node-b	1.45TB

2 entries were displayed.

The following example displays the details of a storage pool named SmallSP. Only one of its four allocation unit has been provisioned, so 75% of its size is available (usable).

```
cluster1::> storage pool show -storage-pool SmallSP
```

```
Storage Pool Name: SmallSP
                UUID of Storage Pool: 60f2f1b9-e60f-11e3-a5e7-
00a0981899a2
        Nodes Sharing the Storage Pool: node-a, node-b
    Number of Disks in Storage Pool: 2
        Allocation Unit Size: 372.5GB
            Storage Type: SSD
    Storage Pool Usable Size: 1.09TB
    Storage Pool Total Size: 1.45TB
        Is Pool Healthy?: true
    State of the Storage Pool: normal
    Reason for storage pool being unhealthy: -
    Job ID of the Currently Running Operation: -
```

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.