



Technical FAQ

FabricPool

ONTAP 9.2

John Lantz, NetApp
September 2017

Abstract

This technical FAQ is meant to assist NetApp SEs, PSEs, Support, and partners by providing answers to frequently asked questions surrounding the use and implementation of FabricPool. Topics include requirements, architecture, configuration, performance, and interoperability.

Version History

Version	Date	Document Version History
1.5	September	Removed free 10TB promotion information. Added questions to performance and interoperability sections.
1.4	July 2017	Updated process for obtaining free 10TB FabricPool license.
1.3	June 2017	GA release. Removed references to GA/RC differences. Added ONTAP Cloud supported volume types. Added recommended and max tiering ratios. FAS2620 added to supported controllers.
1.2	June 2017	Added instructions for determining used capacity tier size via GUI and CLI Clarified Volume Move use case Clarified why 8020 is not supported Removed FAS2620 from list of supported controllers until it becomes officially qualified as supported
1.1	May 2017	Clarified support for StorageGRID Webscale 10.3+ 90-day FabricPool license PVR Capacity reporting
1.0	May 2017	John Lantz Initial release

TABLE OF CONTENTS


Version History	2
1 Overview.....	6
What are the primary use cases for FabricPool in ONTAP 9.2?	6
Can I tier cold primary storage data using FabricPool?	6
Are AFF platforms the only platforms that support FabricPool?	7
How do I convince my customers to buy AFF for their secondary storage?	7
How is FabricPool different from AltaVault?	7
2 Requirements.....	8
Does FabricPool require a license?	8
How is FabricPool licensed?.....	8
How do I get a FabricPool license?	8
How do I install my FabricPool license?	8
What happens if the capacity limit is reached?	8
How is licensed capacity reported?	8
How do I add capacity to a license?	8
What is the fully qualified domain name (FQDN) for AWS S3?	8
How do I get the CA certificate required to add SGWS as an external capacity tier?	9
How do I install the CA certificate required by StorageGRID Webscale?	9
What NetApp controllers support FabricPool?.....	9
Why doesn't the FAS8020 support FabricPool?	9
What volume types are supported when using ONTAP Cloud?	10
How do I create a bucket in AWS S3?.....	10
How do I create a bucket in StorageGRID Webscale?	10
Why can't FabricPool attach to my AWS S3 bucket?	10
3 Architecture.....	11
3.1 Block Temperature.....	11
How is block temperature determined?	11
Can a user define block temperature?.....	11
How long does it take before a block turns cold?	11
Are blocks immediately tiered after they have been identified as cold?.....	11
Are storage efficiencies preserved when moving data to the capacity tier?	11
Is the complete 4MB object fetched when only one 4KB block is read?	11
In what scenario are objects removed/deleted from the capacity tier?	11
What happens if some, but not all, blocks in an object are deleted from the performance tier?	11
Does FabricPool prefetch sequential reads of blocks in the object store?	11
How does volume move work with FabricPool?	11
Is the capacity tier protected by ONTAP RAID?	12
3.2 Tiering.....	12
Why isn't my data tiering?	12

How much of an SSD aggregate must be used to trigger tiering?	12
Is the utilization parameter configurable?	12
Can tiering be scheduled?	12
Can tiering be associated with specific types of files?	12
What happens when tiering is disabled on a volume? Are capacity tier blocks moved to the performance tier?	12
Can data be moved from the performance tier to the capacity tier manually?	12
Can data be moved from the capacity tier to the performance tier manually?	12
Can an external capacity tier be disconnected from the aggregate to which it has been attached?	13
Can the bucket associated with an external capacity tier be changed after it has been attached to an aggregate?	13
3.3 Object Storage	13
What object storage providers does FabricPool support?	13
I don't see Amazon S3 as an option when trying to add an external capacity tier. Where is it?	13
How large are capacity tier objects?	13
Is object size configurable?	13
Can more than one type of object storage be connected to a cluster?	13
Can more than one type of object storage be connected to an aggregate?	13
Does FabricPool support information lifecycle (ILM) policies on object stores?	13
Can I remove a capacity tier from an aggregate after it has been attached?	13
4 Configuration	14
4.1 Sizing	14
How should I size the performance tier?	14
What happens if the performance tier doesn't have enough space to store objects being brought back from the capacity tier?	14
What is the maximum capacity of a FabricPool aggregate?	14
Why must FabricPool volumes set to use -space-guarantee none?	14
How can I determine the current size of the capacity tier?	14
4.2 Tiering Policies	16
What tiering policies does FabricPool support in ONTAP 9.2?	16
4.3 Security	16
Does FabricPool support NetApp Storage Encryption (NSE) and NetApp Volume Encryption (NVE)? ...	16
What about encryption in the capacity tier?	16
Who owns the keys when encryption is turned on in AWS?	16
How is object data protected over the wire?	16
Does FabricPool support aggregate granular encryption (AGE)?	16
5 Performance	17
What type of connectivity requirements does FabricPool need for optimal performance?	17
How does FabricPool affect SnapRestore performance??	17
What happens if the capacity tier goes offline? Does the aggregate go offline too?	17
What happens when the capacity tier has reached full capacity?	17

Does FabricPool require the use of AWS Direct Connect?	17
6 Interoperability.....	18
Are any ONTAP features unsupported by FabricPool?	18
How is FabricPool different from Flash Pool?.....	18
Does FabricPool replace SnapMirror or SnapVault® data protection technologies?	18
How does FabricPool impact concurrent SnapMirror operations?	18
Does FabricPool work with SVMDR?	18
How does volume move (DataMotion™ for Volumes) work with FabricPool?	18
Additional Resources	20
Contact Us	20

1 Overview

FabricPool is a NetApp® ONTAP® feature that allows any workload (SAN or NAS) on an ONTAP system to retain hot data on an all-flash tier and tier inactive cold data onto low-cost object storage, either on the premises or in the cloud. Studies suggest that inactive data can compose up to 80% of an overall storage footprint, and tiering this data provides significant savings on the total cost of ownership. Unlike hierarchical storage management that offers automated tiering capabilities, FabricPool is transparent to applications and allows enterprises to take advantage of cloud economics without sacrificing performance or having to rearchitect solutions to leverage storage efficiency.



What are the primary use cases for FabricPool in ONTAP 9.2?

Answer: For primary storage:

- Snapshot® data

For secondary storage:

- Snapshot data
- Secondary data (Backup and DR; SnapVault and SnapMirror)

Can I tier cold primary storage data using FabricPool?

Answer: Yes, provided that data is unique to Snapshot copies and is not shared with the active file system. We anticipate that tiering additional types of cold data from primary storage will be available in future releases.

Are AFF platforms the only platforms that support FabricPool?

Answer: No. AFF and FAS with all-SSD aggregates are supported FabricPool in ONTAP 9.2.

Note: Flash Pool™ and HDD aggregates are not supported.

How do I convince my customers to buy AFF for their secondary storage?

Answer: By taking advantage of FabricPool, your customers do not need to purchase 1:1 primary to secondary storage. Using FabricPool, only 10% to 20% of your customers' secondary data needs to be accessed using the SSD tier. The rest can be stored as low-cost object storage.

How is FabricPool different from AltaVault?

Answer: AltaVault™ is a complete backup solution. Although FabricPool can significantly reduce storage footprints, it is not a complete backup solution because WAFL® metadata stays on the performance tier. FabricPool's capacity tier does not include WAFL metadata. Likewise, although AltaVault is supports ONTAP-integration using SnapMirror, it is not an ONTAP-integrated solution for tiering hot and cold data like FabricPool.

2 Requirements

Does FabricPool require a license?

Answer: Yes, but only when using Amazon S3 as the capacity tier in ONTAP. No license is required when using StorageGRID® Webscale as the capacity tier or Amazon S3 as the capacity tier in ONTAP Cloud.

How is FabricPool licensed?

Answer: FabricPool uses a capacity-based license. Capacity is equal to amount of data (used capacity) that is stored on Amazon S3.

Note: No license is required when tiering to StorageGRID Webscale.

How do I get a FabricPool license?

Answer: When ordering a new cluster, sales can include a free 10TB FabricPool license for use with AWS S3. For existing systems, a FabricPool license needs to be purchased.

How do I install my FabricPool license?

Answer: See the [ONTAP FabricPool Licensing Overview](#). The NetApp License File (NLF) is installed using System Manager (Configurations/Licenses/Add).

What happens if the capacity limit is reached?

Answer: Tiering stops, including SnapMirror® copies to volumes using the [backup tiering policy](#). Additional data cannot be tiered until the license capacity is increased.

Note: Although tiering will stop, data remains accessible from the capacity tier. Additional cold data will remain on SSDs until the licensed capacity is increased.

How is licensed capacity reported?

Answer:

- **Using ONTAP:**
`system license show-status`
- **Using OnCommand System Manager:**
Configurations / Licenses / FabricPool License / Details

How do I add capacity to a license?

Answer: Additional capacity can be purchased in 1TB increments.

What is the fully qualified domain name (FQDN) for AWS S3?

Answer: s3.amazonaws.com.

How do I get the CA certificate required to add SGWS as an external capacity tier?

Answer:

Using SGWS

1. Select Configuration > Grid Options.
2. From the Grid Options menu, click Overview.
3. Under API Server Certificates, expand CA Certificate.
4. Copy the CA certificate, including -----BEGIN CERTIFICATE----- and -----END CERTIFICATE-----.

Using OpenSSL

1. `openssl s_client -connect <FQDN> -showcerts.`
2. Copy the CA certificate, including -----BEGIN CERTIFICATE----- and -----END CERTIFICATE-----.

How do I install the CA certificate required by StorageGRID Webscale?

Answer: `security certificate install -type server-ca`

Note: When adding SGWS as an external capacity tier, the server name (FQDN) must match the FQDN used by the CA certificate. An admin-provided certificate common name (CN) might not work.

Storage Tiers

Add External Capacity Tier

External capacity tiers such as StorageGRID Webscale object stores can be used to store infrequently used data. [Learn more](#)

External capacity Provider: S3 CloudStorage

Server Name (FQDN):

Port:

Access Key ID:

Secret Key:

Container Name:

SSL: ☒

Encryption: Enabled

What NetApp controllers support FabricPool?

Answer:

- A200, A300, A700, A700s
- AFF8040, AFF8060, AFF8080
- FAS2620, FAS2650, FAS8040, FAS8060, FAS8080, FAS8200, FAS9000

Note: FabricPool is only supported on SSD aggregates in FAS controllers.

Why doesn't the FAS8020 support FabricPool?

The FAS8020 doesn't meet the performance requirements associated with the encryption and decryption of tiered data.

What volume types are supported when using ONTAP Cloud?

Answer:

- Performance tier: gp2 and st1 volumes
- Capacity tier: AWS S3

How do I create a bucket in AWS S3?

Answer: FabricPool cannot create a bucket in AWS S3. To create a bucket in AWS S3, use the Amazon Management Console as documented here:

<http://docs.aws.amazon.com/AmazonS3/latest/gsg/CreatingABucket.html>

How do I create a bucket in StorageGRID Webscale?

Answer: FabricPool cannot create a bucket in SGWS. An S3 browser (such as CloudBerry Explorer) is required to create a bucket in SGWS.

Why can't FabricPool attach to my AWS S3 bucket?

Answer: ONTAP might be unable to attach to AWS S3 buckets in some regions. This issue is documented in public report 1081807.

3 Architecture

3.1 Block Temperature

How is block temperature determined?

Answer: When a block is first written to SSD, it is assigned a temperature value. Over time, its temperature cools provided it is not read. If read, it immediately becomes hot.

Can a user define block temperature?

Answer: No. Block temperature cannot be defined by users.

How long does it take before a block turns cold?

Answer: Blocks are cooled by cooling scans. During this process, blocks that have not been used have their block temperature moved/cooled to the next lower value. Assuming no activity, a block becomes cold within 48 hours.

Note: ONTAP Cloud must be active for the cooling scan to work.

Are blocks immediately tiered after they have been identified as cold?

Answer: No, cold blocks are not immediately tiered. Two conditions must first be met:

- The aggregate must be at least 50% full.
- FabricPool's tiering scan needs to identify a block as cold and mark it for tiering. When enough 4KB blocks from the same volume have been collected, they are concatenated into a 4MB object and moved to the capacity tier.

Are storage efficiencies preserved when moving data to the capacity tier?

Answer: Storage efficiencies such as compression and deduplication are preserved when moving data to the capacity tier. Compaction and aggregate inline deduplication are only supported on the performance tier. Third-party object tier deduplication has not been qualified by NetApp.

Is the complete 4MB object fetched when only one 4KB block is read?

Answer: No. Only the requested 4KB blocks in the object are fetched. The exception to this is if ONTAP performs sequential readahead to increase performance.

In what scenario are objects removed/deleted from the capacity tier?

Answer: If all blocks of an object are moved back or otherwise no longer referenced, then the object is deleted from the object store.

What happens if some, but not all, blocks in an object are deleted from the performance tier?

Answer: There is no change to the object. The object is deleted only when there is no reference to any of the blocks.

Does FabricPool prefetch sequential reads of blocks in the object store?

Answer: Yes, data blocks are prefetched if ONTAP determines it is a sequential read.

How does volume move work with FabricPool?

Answer: If the volume has data in the capacity tier, all the data is pulled back to the buffer cache on the performance tier before being written to disk.

Is the capacity tier protected by ONTAP RAID?

Answer: No. Protection against disk failures in the capacity tier is provided by StorageGRID Webscale or AWS S3.

3.2 Tiering

Why isn't my data tiering?

Answer: Two conditions must first be met:

- The aggregate must be at least 50% full.
- FabricPool's tiering scan needs to identify a block as cold and mark it for tiering. When enough 4KB blocks from the same volume have been collected, they are concatenated into a 4MB object and moved to the capacity tier.

Note: ONTAP Cloud must be active for the cooling scan to work.

How much of an SSD aggregate must be used to trigger tiering?

Answer: Tiering to the capacity tier only happens after 50% of the performance tier has been utilized.

Is the utilization parameter configurable?

Answer: No.

Can tiering be scheduled?

Answer: No.

Can tiering be associated with specific types of files?

Answer: No. [Tiering policies](#) are set at the volume level.

What happens when tiering is disabled on a volume? Are capacity tier blocks moved to the performance tier?

Answer: No. Blocks on the capacity tier remain there until they become hot or a volume move is performed.

Can data be moved from the performance tier to the capacity tier manually?

Answer: Yes. Performing a volume move with the tiering policy set to backup moves all the data in the volume to the capacity tier. A volume move (using CLI or the Volume Move wizard) is the only time when a backup tiering policy can be set on a non-dp volume.

```
vol move start -vserver <name> -volume <name> -destination-aggregate <name>
-tiering-policy backup
```

Note: After the volume move has completed, the tiering policy on the volume automatically reverts from backup to snapshot-only.

Can data be moved from the capacity tier to the performance tier manually?

Answer: Yes. Performing a volume move with the tiering policy set to snapshot-only moves all the data in the volume to the performance tier.

Note: You need to set the tiering policy to none to prevent cold data moved this way from returning to the capacity tier.

Can an external capacity tier be disconnected from the aggregate to which it has been attached?

Answer: No.

Can the bucket associated with an external capacity tier be changed after it has been attached to an aggregate?

Answer: No.

3.3 Object Storage

What object storage providers does FabricPool support?

Answer: FabricPool supports AWS S3 and StorageGRID Webscale 10.3+.

I don't see Amazon S3 as an option when trying to add an external capacity tier. Where is it?

Answer: A FabricPool license must first be installed before Amazon S3 can be selected as an external capacity tier.

How large are capacity tier objects?

Answer: 4MB.

Is object size configurable?

Answer: No. The object size is fixed based on performance recommendations from Amazon and StorageGRID Webscale.

Can more than one type of object storage be connected to a cluster?

Answer: Yes. The type of object storage used is attached at the aggregate level. One aggregate could use StorageGRID Webscale, and another aggregate could use Amazon S3.

Can more than one type of object storage be connected to an aggregate?

Answer: No. One aggregate could use StorageGRID Webscale, and another aggregate could use Amazon S3. One aggregate cannot use both (and cannot attach to two different buckets on the same object storage provider).

Does FabricPool support information lifecycle (ILM) policies on object stores?

Answer: No. ILM policies are not supported by FabricPool. ILM typically includes various movement and deletion policies, which would be disruptive to FabricPool capacity tier data.

Note: FabricPool has no knowledge of ILM policies or configurations set on the object store. Misconfiguration of ILM policies can result in data loss.

Can I remove a capacity tier from an aggregate after it has been attached?

Answer: No. Once attached to an aggregate a capacity tier cannot be removed.

4 Configuration

4.1 Sizing

How should I size the performance tier?

Answer: Refer to the SPM tool for FabricPool sizing recommendations. A general best practice for sizing is a 1:4, performance tier:capacity tier ratio.

What happens if the performance tier doesn't have enough space to store objects being brought back from the capacity tier?

Answer: Capacity tier blocks are not moved and made hot in the performance tier if space is not available. Instead, they are accessed directly from the capacity tier.

Note: A capacity warning is triggered when the capacity tier reaches 85% (default) of the maximum capacity set by the capacity-based license.

What is the maximum capacity of a FabricPool aggregate?

Answer: Maximum capacity for the performance tier is set by the platform on which FabricPool is enabled.

FabricPool has a 1:20 performance tier:capacity tier ratio. Cold data will stop tiering to the capacity tier if it reaches 20 times the size of the performance tier.

For example, 16PB can be tiered from an 800TB aggregate: in total, a 16.8PB FabricPool.

Maximum capacity of the capacity tier is set by the object store. Possibly it is unlimited.

Why must FabricPool volumes set to use `-space-guarantee none`?

Answer: Using `-space-guarantee none` assures thin provisioning of the volume. The amount of space consumed by volumes with this guarantee type grows as data is added instead of being determined by the initial volume size, which is important for FabricPool because the volume needs to support capacity tier data that becomes hot and is brought back to the performance tier.

How can I determine the current size of the capacity tier?

- **OnCommand System Manager**

1. Click **Configurations**.
2. Click **Licenses**.
3. Click **Details**.
4. Click **FabricPool License**.
5. Current capacity is listed in the Current Capacity column.

- **ONTAP**

When looking at an aggregate:

```
storage aggregate object-store show-space
```

When looking at a volume:

```
volume show-footprint
```

4.2 Tiering Policies

What tiering policies does FabricPool support in ONTAP 9.2?

Answer: Three tiering policies are available at the volume level:

- **Snapshot only** (default)
Cold Snapshot blocks in the volume that are not shared with the active file system are moved to the capacity tier.
- **Backup**
All data blocks (does not include WAFL metadata) in the volume are moved to the capacity tier until read and made hot. Hot blocks are moved to the performance tier.

The backup policy can only be set on data protection (DP) volumes (destination targets).

- **None**
Cold blocks in the volume are not moved to the capacity tier.

Setting the tiering policy to none prevents new tiering. Volume data that had previously been moved to the capacity tier remains in the capacity tier until it is accessed and moved back to the performance tier.

Note: Cloned volumes inherit the policy of their parent volume by default. After volume creation, the tiering policy can be changed using `volume modify`.

4.3 Security

Does FabricPool support NetApp Storage Encryption (NSE) and NetApp Volume Encryption (NVE)?

Answer: Yes, on the performance tier using ONTAP.

Note: When enabled on FabricPool volumes, NVE on the performance tier uses the encryption key generated by the capacity tier's object store.

What about encryption in the capacity tier?

Answer: FabricPool's capacity tier uses S3, so NSE and NVE are not applicable. Instead, all objects are encrypted server-side using AES256-SHA by default (x-amz-server-side-encryption is sent in the header).

Who owns the keys when encryption is turned on in AWS?

Answer: Amazon owns the keys for the encryption that is enabled on the capacity tier.

How is object data protected over the wire?

Answer: Objects moving between performance and capacity tiers are encrypted using TLS 1.2.

Does FabricPool support aggregate granular encryption (AGE)?

Answer: No. AGE is exclusive to ONTAP Cloud.

5 Performance

What type of connectivity requirements does FabricPool need for optimal performance?

Answer: FabricPool read latency is a function of connectivity to the external capacity tier. LIFs using 10Gbps ports and round trip latencies under 30ms provide adequate performance. NetApp recommends validating the latency and throughput of your specific network environment to determine the impact it has on FabricPool performance.

Operation	Primary	Secondary (Backup)
PUT	100MBps	300MBps
GET	32KB/IO: 10K IOPS = 320MBps 8KB/IO: 20K IOPS = 160MBps	32KB/IO: 10K IOPS = 320MBps 8KB/IO: 20K IOPS = 160MBps

Note: External capacity tiers do not provide performance similar to that found on the performance tier (typically GB per second). Recovery time objectives (RTOs) should be adjusted accordingly.

How does FabricPool affect SnapRestore performance??

Answer: SnapRestore® takes about the same amount of time to restore as it does without enabling FabricPool. Slower performance occurs if blocks existing only on the capacity tier need to be accessed.

Blocks stored on the capacity tier are moved back to the performance tier for access. The process of doing so takes longer than accessing data that is already located on the performance tier.

What happens if the capacity tier goes offline? Does the aggregate go offline too?

Answer: No. The aggregate (or volume) remains online, but the user/application receives an error message because cold blocks that exist exclusively in the capacity tier are unavailable until it comes back online.

Additionally, cold performance tier data (including SnapMirror transfers to volumes set to use the backup tiering policy) fails to be written to the capacity tier if the capacity tier is offline.

What happens when the capacity tier has reached full capacity?

Answer: Performance tier data is not written to the capacity tier until the additional capacity becomes available.

Does FabricPool require the use of AWS Direct Connect?

Answer: No. Although AWS Direct Connect will provide better performance and lower data transfer rates, it is not required by FabricPool.

6 Interoperability

Are any ONTAP features unsupported by FabricPool?

Answer: Yes. FabricPool does not currently support:

- 7-Mode transition tools
- Flash Pool
- FlexGroup
- MetroCluster™
- NDMP
- ONTAP Select
- SMTape
- SnapLock®
- SVMDR
- SyncMirror®

How is FabricPool different from Flash Pool?

Answer: Flash Pool uses SSDs to provide a high-performance cache for the active dataset of the data volumes provisioned on the Flash Pool aggregate, offloading I/O operations from HDDs to the SSDs.

Environments using FabricPool are already using SSD as part of their aggregate capacity (the performance tier), so there is no need for a Flash Pool aggregate. FabricPool additionally use object storage to store inactive data in a lower-cost capacity tier.

Does FabricPool replace SnapMirror or SnapVault® data protection technologies?

Answer: No. FabricPool does not replace disaster recovery or backup requirements. Data in the capacity tier should provide the same level of availability as ONTAP.

How does FabricPool impact concurrent SnapMirror operations?

Answer: Because concurrent SnapMirror and SnapVault replication operations share bandwidth to the external capacity tier, performance degradation will occur when connectivity resources have been saturated.

To maintain optimal performance and utilization of connectivity resources, NetApp recommends initiating no more than eight concurrent SnapMirror or SnapVault operations, and no more than 20 concurrent SnapRestore operations, with external capacity tiers.

Does FabricPool work with SVMDR?

Answer: No. FabricPool volumes cannot be created or moved to an SVM using SVMDR. Likewise, SVMDR cannot be set on an SVM containing a FabricPool volume.

How does volume move (DataMotion™ for Volumes) work with FabricPool?

Answer: By default, destination volumes inherit the policy of their parent volume. A different tiering policy can be set when initiating a volume move using:

```
vol move start -vserver <name> -volume <name> -destination-aggregate <name>
-tiering-policy <policy>
```

After the volume has been moved, the tiering policy can be changed using `volume modify`.

Note: Because all FabricPool volumes must use volume guarantee = none, attempts to move volumes using volume guarantee = volume to a FabricPool fail. Volume guarantee must be set to none prior to moving the volume into a FabricPool aggregate. Additionally, moving a volume to an aggregate that uses a different capacity tier (for example, SGWS rather than AWS) changes the location of that volume's capacity tier.

Additional Resources

- FabricPool Field Portal Collection
<https://fieldportal.netapp.com/collections/533113>
- Cluster Management Using OnCommand System Manager
https://library.netapp.com/ecm/ecm_download_file/ECMLP2674480
- StorageGRID Webscale 10.4 Administrator Guide
https://library.netapp.com/ecm/ecm_download_file/ECMLP2574683
- ng-fabricpool-info@netapp.com

Contact Us

Do you have a question that you think we should include in this FAQ?

Ask us at docfeedback@netapp.com.

Include FABRICPOOL TECHNICAL FAQ in the subject line.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2017 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.