



Technical Report

Migrating an SVM into a MetroCluster solution using SnapMirror asynchronous replication

Tony Ansley and Stephen Galla, NetApp
April 2023 | TR-4966

Abstract

This technical report provides a step-by-step methodology for migrating a storage virtual machine (SVM) from a standalone ONTAP® storage cluster into a MetroCluster solution using SnapMirror®.

TABLE OF CONTENTS

Introduction	3
Configuration overview	3
Prerequisites	4
Limitations	4
Terminology	4
Before you begin	5
Overview of steps involved in migrating SVM to MCC	5
Migration planning worksheet	5
Migration process	5
Before you begin	6
Create peer relationship	6
Create a replication schedule	7
Create the SnapMirror relationship	7
Migrate the data using SnapMirror	7
Completing the migration	8
Start mcsvm1	9
Post migration verification	10
Support	10
Appendix A: SVM volume worksheets	10
Where to find additional information	12
Version history	12

LIST OF TABLES

Table 1) Naming conventions for migration process	5
---	---

LIST OF FIGURES

Figure 1) SVM migration between a stand-alone cluster and a MetroCluster solution	3
---	---

Introduction

NetApp® MetroCluster is a free feature of ONTAP software that provides real-time protection of your storage infrastructure. When deployed, MetroCluster automatically protects existing and new data and SVMs by synchronously replicating data volumes and SVMs between two physically disparate sites.

This implementation guide is useful if you are already using an existing NetApp storage cluster for your data solution, but your data protection requirements have expanded to providing site-level, real-time data protection for all your workloads – new and existing. This document explains how you can get your existing data and SVMs into this new MetroCluster infrastructure.

This implementation guide provides step-by-step instructions to use SnapMirror SVM disaster recovery (SVM DR) technology to migrate individual SVMs currently running on a standalone ONTAP cluster into a new production cluster within your MetroCluster deployment.

Note: Beginning with ONTAP 9.10.1, SVM DR has extended its support to include replication of MetroCluster SVMs to independent, external ONTAP clusters for disaster recovery scenarios. Although this use case is fully supported, it falls outside the scope of this document.

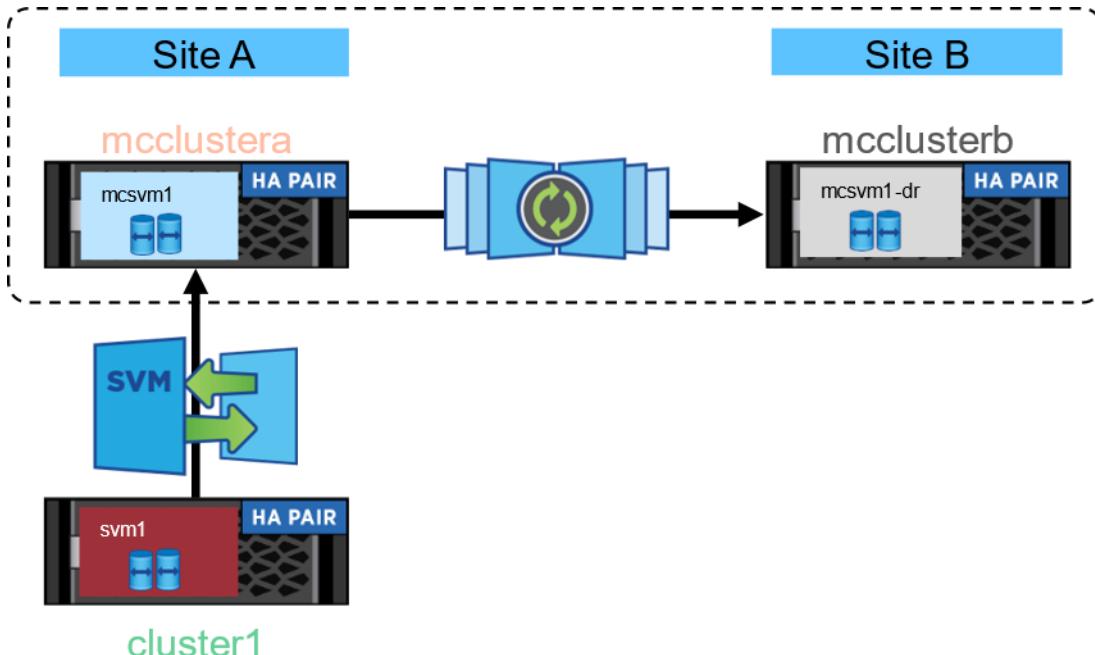
Note: This procedure is restricted by the supported scalability currently provided by SnapMirror for SVM-DR implementation. To understand the current SVM-DR scalability limitations, refer to [TR-4015: SnapMirror configuration and best practices guide for ONTAP 9](#).

Configuration overview

This migration process uses SnapMirror asynchronous to migrate one or more SVMs – including configuration and FlexVol volumes – from a standalone ONTAP storage cluster to a cluster participating in a MetroCluster infrastructure, with the goal of having the migrated SVMs act as a source SVM for MetroCluster replication to a remote MetroCluster DR site (Figure 1).

Note: This process can be used to migrate an SVM into a MetroCluster IP or MetroCluster FC solution.

Figure 1) SVM migration between a stand-alone cluster and a MetroCluster solution.



Prerequisites

The following list of prerequisites must be met before starting the migration process. Any aspect of the standalone source ONTAP cluster and the source and destination MetroCluster infrastructure that does not fit these prerequisites invalidates the supportability of this migration process:

- All nodes in the external cluster (migration source) and MetroCluster clusters (migration destination) are running ONTAP 9.5 or later.
- The migrating SVM is hosting NAS services only.
- The MetroCluster infrastructure is already properly configured for operation as described in [ONTAP 9 MetroCluster IP Installation and Configuration](#).
- MetroCluster utilizing either IP or FC site-to-site infrastructure.
- The source and destination cluster ONTAP versions is supported for SnapMirror interoperability as defined in [Compatible ONTAP versions for SnapMirror relationships \(netapp.com\)](#).
- All required intercluster LIF IP addresses and administrator passwords for each cluster are available.
- The MetroCluster inter-switch links are sized to support the data I/O volume of the migrated SVMs.
- The MetroCluster replication switch infrastructure is dedicated to the MetroCluster configuration.

Limitations

The following limitations restrict the supportability of this implementation guide for moving an SVM from a standalone cluster to a MetroCluster cluster.

- This process only supports SVMs hosting NFS, SMB, or both.
Although it might be possible to migrate SAN LUNs and namespaces, SVM DR does not currently support migration of SAN configuration information. Any LUNs or namespaces hosted by the migrating SVM must be reconfigured when migrated to the MetroCluster infrastructure.
- This process assumes that only a single SVM is migrated at any time.
Although it might be possible to migrate multiple SVMs simultaneously, this restriction makes sure that each SVM migrates as fast as possible, and without any possibility of outside influences – such as other migrating SVMs – that might impact the SVM migration.
- This process is not supported for migrating SVMs between two MetroCluster infrastructures.
ONTAP currently restricts migration between two clusters that are part of a MetroCluster infrastructure.
- The MetroCluster ecosystem (NetApp OnCommand®, Active IQ® Unified Manager, Health Monitor, ConfigAdvisor, OnCommand Performance Monitor, and so on) are not supported.
- There is limited support for AutoSupport.

Terminology

The migration process uses the following naming and color conventions to reference various source and destination clusters, SVMs, and volumes. This aids in properly identifying each cluster in the various CLI commands. Actual names of clusters, SVMs, volumes, and aggregates can be of any value appropriate for the customer's production storage environment (use the provided worksheets to document cluster, volume, and SVM information before executing migrations).

Table 1) Naming conventions for migration process.

Names	Description
cluster1	The cluster that hosts the SVM at the beginning of the migration process. This cluster is a standalone cluster outside of the MetroCluster solution that the SVM is being migrated to.
mccluster1a	A cluster that is located at one of the two sites that make up the MetroCluster infrastructure. This cluster will be the post-migration primary host of the migrating SVM.
mccluster1b	A cluster that is located at one of the two sites that make up the MetroCluster infrastructure. This cluster will be the MetroCluster DR destination of the SVM once the migration is complete.
svm1	The source SVM that is being migrated from cluster1 to mccluster1a.
mcsvm1	The destination SVM that svm1 is being migrated to. This SVM is hosted by mccluster1a
mcsvm1-mc	The destination SVM that MetroCluster is replicating from source mcsvm1 is being migrated to. This SVM is hosted by mccluster1b
mcsvm1_root	The root volume of mcsvm1 on mccluster1a

Before you begin

Overview of steps involved in migrating SVM to MetroCluster

1. Verify that all prerequisites are met.
2. Understand and accept the current limitations of this process.
3. Create the migration target SVM and volumes on the MetroCluster cluster that will be the post-migration production cluster.
4. Create the peering relationship between the external cluster (SVM migration source) and the MetroCluster cluster (SVM migration target destination)
5. Create the SVM peering relationship between the source SVM (external cluster) and the destination SVM (MetroCluster cluster)
6. Create and initialize the SnapMirror relationship between the source SVM and the destination SVM
7. Failover the SVM from the source cluster to the MetroCluster SVM
8. Post-migration cleanup

Migration planning worksheet

Appendix A: SVM volume worksheets provides worksheets that can be used to document each SVM that will be migrated.

Migration process

The following steps are provided to migrate an SVM from a standalone ONTAP cluster source to a destination cluster that is a member of a MetroCluster infrastructure.

Before you begin

It is recommended that you test the process using a non-production SVM hosting at least one volume storing some miscellaneous random data to make sure that the process is understood and successfully migrates an SVM within your specific deployment environment.

Note: All command input parameters and resulting outputs are for reference only. Substitute actual names and IP addresses as appropriate for the specific solution deployment requirements.

Create peer relationship

1. Create cluster peer relationship between `cluster1` and `mccluster1`.

```
mccluster1::> cluster peer create -peer-addrs cluster1 -username admin
Remote Password:
mccluster1::> cluster peer show -instance
Peer Cluster Name: cluster1
    Remote Intercluster Addresses: mccluster1
    Availability of the Remote Cluster: Available
        Remote Cluster Name: cluster1
        Active IP Addresses: 10.1.1.246, 10.1.1.243
        Cluster Serial Number: 1-80-123456
    Address Family of Relationship: ipv4
    Authentication Status Administrative: no-authentication
    Authentication Status Operational: absent
    Last Update Time: 02/05 21:05:41
    IPspace for the Relationship: Default
Encryption for Inter-Cluster Communication: none
```

2. Create `mcsvm1` on `mccluster1`. This is the new production SVM post-migration.

```
mccluster1::> vserver create -vserver mcsvm1
[Job 96] Job succeeded: Vserver creation completed.
```

3. Take the root volume of `mcsvm1` offline.

```
mccluster1::> volume offline -vserver mcsvm1 -volume mcsvm1_root
mccluster1::> vserver show -vserver mcsvm1 -fields rootvolume, rootvolume-security-style,
aggregate
vserver rootvolume aggregate rootvolume-security-style
-----
mcsvm1  mcsvm1_root
        aggr1      unix
```

4. Delete the root volume of `mcsvm1`.

Note: This required advanced level privileges.

```
mccluster1::> set advanced
Warning: These advanced commands are potentially dangerous; use them only when directed to do so
by NetApp personnel.
Do you want to continue? {y|n}: y
mccluster1::> volume delete -vserver mcsvm1 -volume mcsvm1_root -force true
[Job 98] Job succeeded: Successful
```

5. Stop `mcsvm1`. This prevents any access to `mcsvm1` while performing the next steps.

```
mccluster1::> vserver stop -vserver mcsvm1
[Job 99] Job succeeded: DONE
mccluster1::> vserver show -vserver mcsvm1
Vserver      Type      Subtype      Admin      Operational      Root
           state      state      Volume      Aggregate
-----
mcsvm1.  data      default      stopped      stopped      root_vs1      aggr0
mccluster1::> set admin
```

6. Create the SVM peer relationship process between `svm1` and `mcsvm1`.

```
cluster1::> vserver peer create -vserver svm1 -peer-vserver mcsvm1 -applications snapmirror -
peer-cluster mccluster1
```

```
Info: [Job 505] 'vserver peer create' job queued. Vserver creation completed.
cluster1::> vserver peer show -vserver svm1
      Peer          Peer          Peering          Remote
Vserver  Vserver  State  Peer Cluster  Applications  Vserver
-----
svm1      mcsvm1  peered  mccluster1  snapmirror  mcsvm1
```

7. Complete the SVM peer relationship process between **svm1** and **mcsvm1** by accepting the peer relationship.

```
mccluster1::> vserver peer accept -vserver mcsvm1 -peer-vserver svm1 -application snapmirror
Info: [Job 505] 'vserver peer create' job queued. Vserver creation completed.
mccluster1::> vserver peer show -vserver mcsvm1
      Peer          Peer          Peering          Remote
Vserver  Vserver  State  Peer Cluster  Applications  Vserver
-----
mcsvm1      svm1  peered  cluster1  snapmirror  svm1
```

Create a replication schedule

The following step can be used to create a schedule for the replication process that will migrate the SVM and data to the destination MetroCluster cluster. This might not be necessary if you have an appropriate SnapMirror schedule configured that meets your migration timeline.

1. Create a schedule on **mccluster1** to control the migration process. It is recommended to use a minimum schedule of no less than 15 minutes for the migration process.

```
mccluster1::> job schedule cron create -name 15min_schedule -minute "00,15,30,45"
(job schedule cron create)

mccluster1::> job schedule cron show -name 15min_schedule
Cluster Vserver  Name          Description
-----
mccluster1      mcsvm1  15min_schedule
                  @:10,:25,:40,:55
```

Note: The schedule might need to be longer depending on the amount of data that might change on **svm1** before the final cutover.

Create the SnapMirror relationship

The following steps create the SnapMirror relationship that migrates the SVM configuration and data to the new MetroCluster cluster.

1. Create an asynchronous SnapMirror relationship between **svm1** and **mcsvm1** using the schedule defined in Create a replication schedule.

```
mccluster1::> snapmirror create -source-path svm1: -destination-path mcsvm1: -identity-preserve
true -schedule 15min_schedule
```

Note: This relationship must use the **-identity-preserve true** parameter to migrate all NAS settings to the destination **mcsvm1** SVM.

Migrate the data using SnapMirror

1. Verify that the SnapMirror relationship shows the Mirror State as Broken-off.

```
mccluster1::> snapmirror show
Source          Destination  Mirror  Relationship  Total          Last
Path           Type    Path      State  Status    Progress  Healthy  Updated
-----
svm1:          XDP    mcsvm1:  Broken-off
                           Idle      -       true      -
```

2. Perform a SnapMirror resync operation to start the data migration.

```
mccluster1::> snapmirror resync mcsvm1:
```

Note: The resync operation changes the SVM type for `mcsvm1` from default to dp-destination and the state from running to stopped.

3. Verify and monitor the progress of the SnapMirror data migration. During the transfer, the Mirror State continues to show Broken-off, the Relationship Status is Transferring, and the Total Progress field displays the amount of data replicated.

When the baseline data migration is complete, the Total Progress field shows Idle and the Relationship Status field shows Snapmirrored.

```
mccluster1::> snapmirror show
Source          Destination Mirror  Relationship  Total          Last
Path           Type     Path      State   Status    Progress  Healthy Updated
-----  -----
svm1:          XDP     mcsvm1:  Broken-off
                           Transferring   -       true      -
                                         Total Progress
                                         Last
mccluster1::> snapmirror show
Source          Destination Mirror  Relationship  Total          Last
Path           Type     Path      State   Status    Progress  Healthy Updated
-----  -----
svm1:          XDP     mcsvm1:  Snapmirrored
                           Idle      -       true      -
```

Completing the migration

1. When the baseline data replication has been completed, quiesce the SnapMirror relationship.

```
mccluster1::> snapmirror quiesce mcsvm1:
mccluster1::> snapmirror show
Source          Destination Mirror  Relationship  Total          Last
Path           Type     Path      State   Status    Progress  Healthy Updated
-----  -----
svm1:          XDP     mcsvm1:  Snapmirrored
                           Quiesced   -       true      -
```

2. Stop the `svm1` SVM. This prevents access to the SVM by application clients.

```
cluster1::> vserver stop svm1
[Job 791] Job succeeded: DONE
```

3. Verify that `svm1` has stopped

```
cluster1::> vserver show
Vserver      Type     Subtype  Admin      Operational Root
           State   State    Volume    Aggregate
-----  -----
svm1:        data    default  stopped  stopped  root_svm1  aggr1
```

4. Restart the SnapMirror relationship from the quiesced state.

```
mccluster1::> snapmirror resume mcsvm1:
```

5. Verify that `mcsvm1` has restarted SnapMirror replication. Status should show Idle.

```
mccluster1::> snapmirror show
Source          Destination Mirror  Relationship  Total          Last
Path           Type     Path      State   Status    Progress  Healthy Updated
-----  -----
svm1:          XDP     mcsvm1:  Snapmirrored
                           Idle      -       true      -
```

6. Execute a final update to the SnapMirror relationship.

```
mcclustera::> snapmirror update mcsvm1:
```

7. Monitor the SnapMirror progress. Once the update has completed, the Status will change from Transferring to Idle.

```
mcclustera::> snapmirror show
Source          Destination Mirror  Relationship  Total          Last
Path           Type    Path      State   Status    Progress  Healthy Updated
-----
svml:          XDP    mcsvm1:   Snapmirrored
                           Transferring   -       true    -
                           Idle          -       true    -
```



```
mcclustera::> snapmirror show
Source          Destination Mirror  Relationship  Total          Last
Path           Type    Path      State   Status    Progress  Healthy Updated
-----
svml:          XDP    mcsvm1:   Snapmirrored
                           Idle          -       true    -
```

8. Break the SnapMirror relationship between **svml** and **mcsvm1**. This results in **mcsvm1** volumes switching from read-only to read-write.

```
mcclustera::> snapmirror break mcsvm1:
```

```
mcclustera::> snapmirror show
Source          Destination Mirror  Relationship  Total          Last
Path           Type    Path      State   Status    Progress  Healthy Updated
-----
svml:          XDP    mcsvm1:   Broken-off
                           Idle          -       true    -
```

Start mcsvm1

1. Start **mcsvm1** on **mcclustera**. This restarts all NFS and SMB servers using the same configuration as the original **svml** SVM.

```
mcclustera::> vserver start -vserver mcsvm1
[Job 519] Job succeeded: DONE
```

2. Verify that the new **mcsvm1** is running.

```
mcclustera::*> vserver show
Vserver      Type     Subtype   Admin      Operational Root
             State     State     Volume    Aggregate
-----
mcsvm1       data     default   running   running   root_mcsvm1
                                         aggr1
```

3. Verify that **mcsvm1-mc** is present on **mcclusterb**.

```
mcclusterb::*> vserver show
Vserver      Type     Subtype   Admin      Operational Root
             State     State     Volume    Aggregate
-----
mcsvm1-mc   data     default   running   running   root_mcsvm1
                                         aggr1
```

4. Verify that the volume is in the online state and is of type RW

```
mcclustera::*> volume show -vserver mcsvm1
Vserver      Volume   Aggregate  State     Type  Size   Available Used%
-----
mcclustera   mcsvm1   aggr1     online   RW    2GB    1.9GB   5%
```

5. Verify that MetroCluster is healthy and replicating to **mcclusterb**.

```
mcclustera::> metrocluster vserver show
Cluster: mcclustera
          Partner
          Vserver
  -----
  mcsvm1      mcsvm1-mc
  vs1         vs1-mc
Cluster: mcclusterb
          Partner
          Vserver
  -----
  vs2         vs2-mc
3 entries were displayed.
```

Repeat the entire Migration process for any additional SVMs that need to be migrated from **cluster1** to **mcclustera**.

Post migration verification

Perform the following steps to delete the SVM from **cluster1** for each SVM that has been migrated to the MetroCluster solution:

1. Verify that the data hosted by **mcsvm1** matches the data on original source cluster **svm1**.
2. Perform any steps on the client systems to redirect access from volumes on **svm1** and redirect access to SMB shares or NFS exports on **mcsvm1**.
3. Delete the SnapMirror relationship between **svm1** and **mcsvm1**.

```
mcclustera::> snapmirror delete -source-path svm1: -destination-path mcsvm1:
```

4. Delete the SVM peer relationship between **svm1** and **mcsvm1**.

```
cluster1::> vserver peer delete -vserver svm1 -peer-vserver mcsvm1
```

5. Delete the volumes hosted on **cluster1://svm1** when no longer needed.

```
cluster1::> volume delete -vserver svm1 -volume vol1_old
```

Note: Repeat step 6 for each volume hosted by **svm1**, including the **svm1_root** volume.

6. Delete the original, external source SVM **svm1**.

```
cluster1::> vserver delete -vserver svm1
```

7. Delete the cluster peer relationship between **cluster1** and **mcclustera**.

```
mcclustera::> cluster peer delete -cluster cluster1
```

Support

Hardware or software support for host or storage systems during this migration process should use the standard interfaces as they are defined in your support contract with NetApp.

NetApp CSS handles host environment hardware and software issues. Escalations should be directed to the appropriate CSS specialist or NetApp support communities on communities.netapp.com or the NetApp Discord community

Appendix A: SVM volume worksheets

Migration target MetroCluster information

Cluster name	
Cluster intercluster LIF IP address	
Cluster management LIF IP address	
Administrator account (keep passwords hidden)	

Migration source information (non-MetroCluster cluster)

Cluster name	
Cluster intercluster LIF IP address	
Cluster management LIF IP address	
Administrator account (keep passwords hidden)	

SVM

SVM name	
Cluster intercluster LIF IP address	
Cluster management LIF IP address	
Administrator account (keep passwords hidden)	

Volume list

Source	Destination	Notes

SVM migration information

	Source	Destination
SVM name		
Cluster intercluster LIF IP address(es)		
Cluster management LIF IP address		

	Source	Destination
Administrator account (keep passwords hidden)		

Volume list

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- MetroCluster Documentation (netapp.com)
<https://docs.netapp.com/us-en/ontap-metrocluster/>
- Cluster and SVM peering overview with the CLI (netapp.com)
<https://docs.netapp.com/us-en/ontap/peering/index.html>
- About SnapMirror SVM replication (netapp.com)
<https://docs.netapp.com/us-en/ontap/data-protection/snapmirror-svm-replication-concept.html>
- Compatible ONTAP versions for SnapMirror relationships (netapp.com)
<https://docs.netapp.com/us-en/ontap/data-protection/compatible-ontap-versions-snapmirror-concept.html>

Version history

Version	Date	Document version history
Version 1.0	April 2023	Initial release of document

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 © 2023 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—with 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.