



Solutions for Hyper-V with NetApp

NetApp Solutions

NetApp
September 18, 2024

Table of Contents

Solutions for Hyper-V with NetApp 1

 SnapMirror Active Sync with Microsoft Stretch Clusters 1

Solutions for Hyper-V with NetApp

SnapMirror Active Sync with Microsoft Stretch Clusters

This paper documents SnapMirror active sync technology synchronous bidirectional replication between Microsoft stretches failover clusters, allowing multisite application data for example MSSQL and Oracle to be actively accessible and in sync across both sites.

Introduction

Beginning with ONTAP 9.15.1, SnapMirror active sync supports symmetric active/active deployments, enabling read and write I/O operations from both copies of a protected LUN with bidirectional synchronous replication. A Windows Stretch Cluster is an extension of the Windows Failover Cluster feature that spans multiple geographic locations to provide high availability and disaster recovery. With SnapMirror active sync symmetric active/active and clustered applications like Windows failover clustering, we can achieve continuous availability for Microsoft Hyper-V business-critical applications to achieve zero RTO and RPO during unexpected incidents. This solution provides the following benefits:

- **Zero Data Loss:** Ensures data is replicated synchronously, achieving zero Recovery Point Objective (RPO).
- **High Availability and Load Balancing:** Both sites can actively handle requests, providing load balancing and high availability.
- **Business continuity:** Implement a symmetric active/active configuration to ensure that both data centers are actively serving applications and can seamlessly take over in case of a failure.
- **Improve performance:** Use symmetric active/active configuration to distribute the load across multiple storage systems, improving response times and overall system performance.

This paper documents SnapMirror active sync technology synchronous bidirectional replication between Microsoft stretches failover clusters, allowing multisite application data for example MSSQL and Oracle to be actively accessible and in sync across both sites. If a failure occurs, applications are immediately redirected to the remaining active site, with no loss of data and no loss of access, providing high availability, disaster recovery and geographic redundancy.

Use Cases

In the event of a disruption such as a cyber-attack, power outage, or natural disaster, a globally connected business environment demands rapid recovery of business-critical application data with zero data loss. These demands are heightened in areas such as finance and those adhering to regulatory mandates such as the General Data Protection Regulation (GDPR). Deploy a symmetric active/active configuration to replicate data between geographically dispersed locations, providing local access to data and ensuring continuity in case of regional outages.

SnapMirror active sync provides the following use cases:

Application deployment for zero recovery time object (RTO)

In a SnapMirror active sync deployment, you have a primary and mirror cluster. A LUN in the primary cluster (\L1P) has a mirror (L1S) on the secondary; Read and Writes are served by the site local to the Hosts based on Hot Proximity settings.

Application deployment for zero RTO or TAF

Transparent Application Failover (TAF) is based on host MPIO software-based path failover to achieve non-disruptive access to the storage. Both LUN copies—for example, primary (L1P) and mirror copy (L1S) --have the same identity (serial number) and are reported as read-writable to the host.

Clustered applications

Clustered applications including VMware vSphere Metro Storage Cluster(vMSC), Oracle RAC, and Windows Failover Clustering with SQL require simultaneous access so the VMs can fail over to the other site without any performance overhead. SnapMirror active sync symmetric active/active serves IO locally with bidirectional replication to meet the requirements of clustered applications.

Disaster scenario

Synchronously replicate multiple volumes for an application between sites at geographically dispersed locations. You can automatically failover to the secondary copy in case of disruption at the primary, thus enabling business continuity for tier-one applications.

Windows failover


SnapMirror active sync provides flexibility with easy-to-use application-level granularity and automatic failover to achieve high data availability and fast data replication for your business-critical applications such as Oracle, Microsoft SQL Server, and so on, in both virtual and physical environments.

Solution Architecture

The Microsoft failover stretch cluster has two Hyper-V nodes on each site. These two nodes share NetApp storage and use SnapMirror active sync symmetric active-active to replicate the volumes between the two sites. A consistency group ensures all volumes of a dataset are quiesced and then snapped at precisely the same point in time. This provides a data-consistent restore point across volumes supporting the dataset. The ONTAP Mediator receives health information about peered ONTAP clusters and nodes, orchestrating between the two and determining if each node/cluster is healthy and running.

Solution Components:

- Two NetApp storage systems ONTAP 9.15.1: first and second failure domain
- A Rethat 8.7 VM for ONTAP mediator
- Three Hyper-V failover clusters on Windows 2022:
 - site1, site 2 for the applications
 - site 3 for mediator
- VM on Hyper-V: Microsoft Domain Controller, MSSQL Always On Failover cluster instance, ONTAP Mediator



Install a Microsoft Stretch Failover Cluster

You can use Windows Admin Center, PowerShell, or the Server Manager console to install the Failover Clustering feature and its associated PowerShell cmdlets. For details on prerequisites and steps, check [create a failover cluster](#).

Here's a step-by-step guide to setting up a Windows Stretch Cluster:

1. Install Windows 2022 on all four servers hyperv1, hyperv2, hyperv3 and hyperv4
2. Join all four servers to the same Active Directory domain: hyperv.local.
3. Install the Windows features failover-clustering, Hyper-V, Hyper-V_Powershell and MPIO on each server.

```
Install-WindowsFeature -Name "Failover-Clustering", "Hyper-V", "Hyper-V-Powershell", "MPIO" -IncludeManagementTools
```

4. Configure MPIO, add support for iSCSI devices.



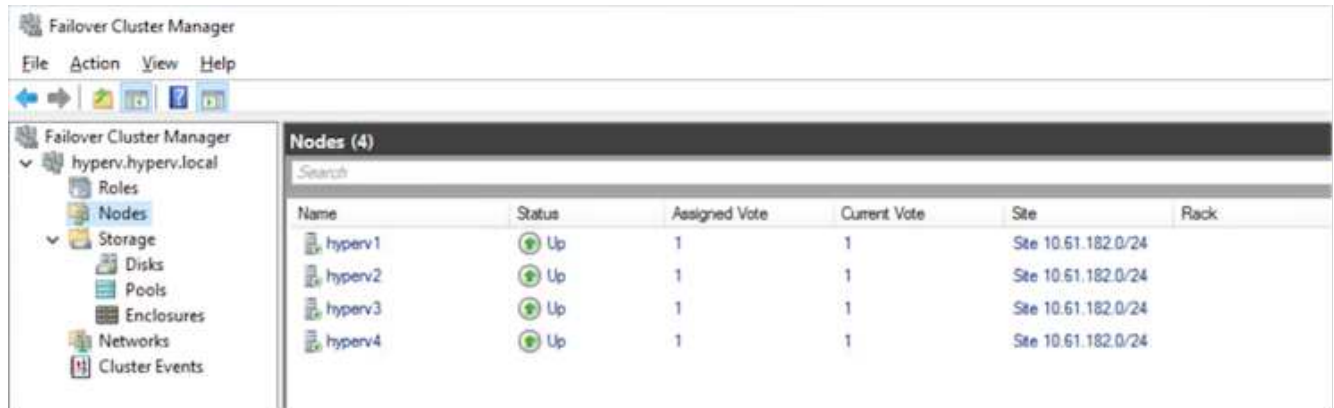
5. On site 1 and site 2 ONTAP storage, create two iSCSI LUNs (SQLdata and SQLlog) and map to the windows servers iqn group. Use Microsoft iSCSI software initiator to connect the LUNs. For more details, check [iSCSI configuration for Windows](#).

- Run the Cluster Validation report for any errors or warnings.

```
Test-Cluster -Node hyperv1, hyperv2, hyperv3, hyperv4
```


- Create a failover cluster, assign a static IP address,

```
New-Cluster -Name <clustername> -Node hyperv1, hyperv2, hyperv3, hyperv4, StaticAddress <IPaddress>
```



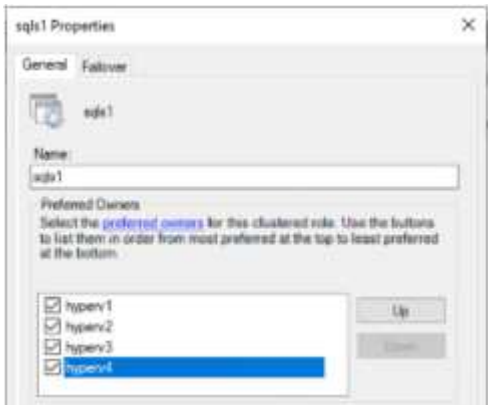
- Add the mapped iSCSI storages to the failover cluster.
- Configure a witness for quorum, right-click the cluster → More Actions → Configure Cluster Quorum Settings, choose disk witness.

The diagram below shows four clustered shared LUNs – two sites sqldata and sqllog and one disk witness in quorum.



Always On Failover Cluster Instance


An Always On Failover Cluster Instance (FCI) is a SQL Server instance that is installed across nodes with SAN shared disk storage in a WSFC. During a failover, the WSFC service transfers ownership of instance's resources to a designated failover node. The SQL Server instance is then re-started on the failover node, and databases are recovered as usual. For more details on setup check Windows Failover Clustering with SQL. Create two Hyper-V SQL FCI VMs on each site and set priority. Use hyperv1 and hyperv2 as the preferred owners for the site 1 VMs and hyperv3 and hyperv4 as the preferred owners for site 2 VMs.




Create Intercluster Peering

You must create peer relationships between source and destination clusters before you can replicate Snapshot copies using SnapMirror.

1. Add intercluster network interfaces on both clusters



2. You can use the cluster peer create command to create a peer relationship between a local and remote cluster. After the peer relationship has been created, you can run cluster peer create on the remote cluster to authenticate it to the local cluster.



Configure Mediator with ONTAP

The ONTAP Mediator receives health information about peered ONTAP clusters and nodes, orchestrating between the two and determining if each node/cluster is healthy and running. SM-as allows data to be replicated to the target as soon as it is written to the source volume. The mediator must be deployed at the third failure domain.

Prerequisites

- HW Specs: 8GB RAM, 2x2GHz CPU, 1Gb Network (<125ms RTT)
- Installed Red Hat 8.7 OS, check [ONTAP Mediator version and supported Linux version](#).
- Configure the Mediator Linux host: network setup and firewall ports 31784 and 3260
- Install the yum-utils package
- [Register a security key when UEFI Secure Boot is enabled](#)

Steps

1. Download the Mediator installation package from the [ONTAP Mediator download page](#).
2. Verify the ONTAP Mediator code signature.
3. Run the installer and respond to the prompts as required:

```
./ontap-mediator-1.8.0/ontap-mediator-1.8.0 -y
```

4. When Secure Boot is enabled, you must take additional steps to register the security key after installation:
 - a. Follow the instructions in the README file to sign the SCST kernel module:


```
/opt/netapp/lib/ontap_mediator/ontap_mediator/SCST_mod_keys/README.module-signing
```

- b. Locate the required keys:

```
/opt/netapp/lib/ontap_mediator/ontap_mediator/SCST_mod_keys
```

5. Verify the installation

- a. Confirm the processes:

```
systemctl status ontap_mediator mediator-scst
```

```
[root@mediator2 home]# systemctl status mediator-scst ontap_mediator
● mediator-scst.service
Loaded: loaded (/etc/systemd/system/mediator-scst.service; enabled; vendor preset: disabled)
Active: active (running) since Thu 2024-04-04 11:41:57 EDT; 2 weeks 5 days ago
Process: 41042 ExecStartPost=/usr/sbin/multiplexd scst_vdisk 0000000000, status=0/SUCCESS
Process: 41042 ExecStart=/usr/sbin/ontap_mediator-scst start 0000000000, status=0/SUCCESS
Main PID: 41042 (multiplexd)
Tasks: 2 (limit: 19904)
Memory: 2.0M
CGroup: /system.slice/mediator-scst.service
└─ 41042 /usr/sbin/multiplexd --scst /usr/sbin/ontap_mediator-scst

Jun 06 11:41:54 mediator2 hyperv.local systemd[1]: Starting mediator-scst.service...
Jun 06 11:41:54 mediator2 hyperv.local systemd[1]: Starting and mounting SCST.
Jun 06 11:41:54 mediator2 hyperv.local multiplexd[41042]: mounting on
Jun 06 11:41:57 mediator2 hyperv.local systemd[1]: Started mediator-scst.service.

● ontap_mediator.service: ONTAP Mediator
Loaded: loaded (/etc/systemd/system/ontap_mediator.service; disabled; vendor preset: disabled)
Active: active (running) since Thu 2024-04-04 11:41:57 EDT; 2 weeks 5 days ago
Process: 41042 ExecStartPost=/opt/netapp/lib/ontap_mediator/ontap_mediator/ontap_mediator --scst /usr/sbin/ontap_mediator-scst --scst /usr/sbin/ontap_mediator-scst --scst /usr/sbin/ontap_mediator-scst
Main PID: 41042 (multiplexd)
Tasks: 2 (limit: 19904)
Memory: 2.0M
CGroup: /system.slice/ontap_mediator.service
└─ 41042 /opt/netapp/lib/ontap_mediator/ontap_mediator/ontap_mediator --scst /usr/sbin/ontap_mediator-scst --scst /usr/sbin/ontap_mediator-scst --scst /usr/sbin/ontap_mediator-scst

Jun 06 11:41:56 mediator2 hyperv.local multiplexd[41042]: Creating FibreChannel with 409600 kb blocks and 409600 blocks
Jun 06 11:41:56 mediator2 hyperv.local multiplexd[41042]: FibreChannel GUID: 559d1331-99ed-4d30-9050-32c1212261d1
Jun 06 11:41:56 mediator2 hyperv.local multiplexd[41042]: Registering backlogs started on blocks:
Jun 06 11:41:56 mediator2 hyperv.local multiplexd[41042]: 22748, 92204, 102816
Jun 06 11:41:56 mediator2 hyperv.local multiplexd[41042]: (428 Block data)
Jun 06 11:41:56 mediator2 hyperv.local multiplexd[41042]: (1988 Block data)
Jun 06 11:41:56 mediator2 hyperv.local multiplexd[41042]: Creating Jovosak (1000 blocks): done
Jun 06 11:41:56 mediator2 hyperv.local multiplexd[41042]: (1758 Block data)
Jun 06 11:41:56 mediator2 hyperv.local multiplexd[41042]: (1758 Block data)
Jun 06 11:41:56 mediator2 hyperv.local multiplexd[41042]: (49500) getting FBI metadata from /opt/netapp/lib/ontap_mediator/
```

- b. Confirm the ports that are used by the ONTAP Mediator service:

```
[root@mediator2 server_config]# netstat -antl | grep -E '3260|31784'
tcp        0      0 0.0.0.0:3260          0.0.0.0:*            LISTEN
tcp        0      0 0.0.0.0:31784        0.0.0.0:*            LISTEN
tcp        0      0 10.61.182.163:31784 10.61.182.148:26429  ESTABLISHED
tcp        0      0 10.61.182.163:31784 10.61.182.148:24546  FIN_WAIT2
tcp6       0      0 :::3260              :::*                  LISTEN
```

6. Initialize the ONTAP Mediator for SnapMirror active sync using self-signed certificates

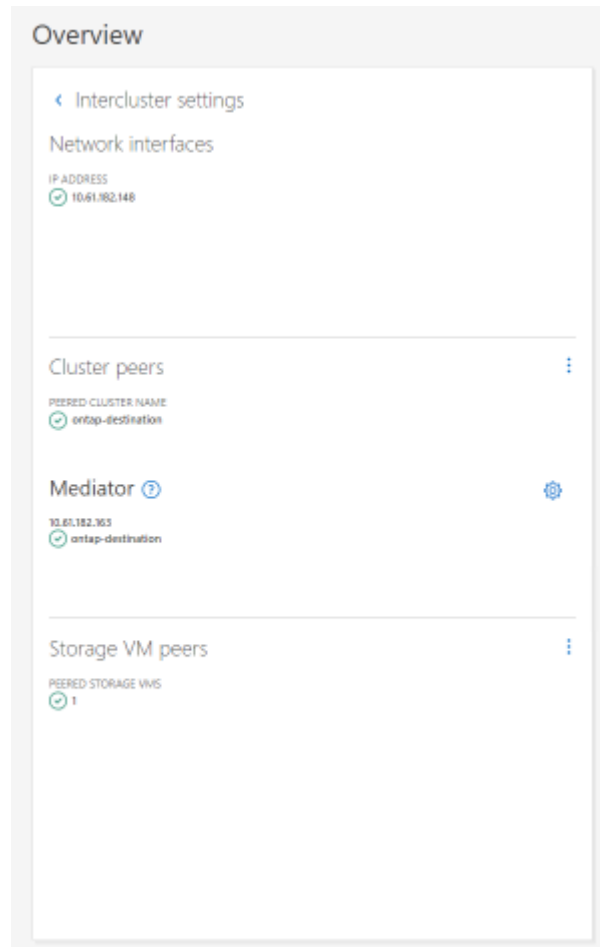
- a. Find the ONTAP Mediator CA certificate from the ONTAP Mediator Linux VM/host software installation location `cd /opt/netapp/lib/ontap_mediator/ontap_mediator/server_config`.
- b. Add the ONTAP Mediator CA certificate to an ONTAP cluster.

```
security certificate install -type server-ca -vserver <vserver_name>
```

- 7. Add the mediator, go to System Manager, protect>Overview>mediator, enter the mediator's IP address,

username (API User default is mediatoradmin), password and the port 31784.

The following diagram shows the intercluster network interface, cluster peers, mediator, and SVM peer are all setup.




Configure Symmetric active/active protection

Consistency groups facilitate application workload management, providing easily configured local and remote protection policies and simultaneous crash-consistent or application-consistent Snapshot copies of a collection of volumes at a point in time. For more details refer to [consistency group overview](#). We use a uniform configuration for this setup.

Steps for a uniform configuration

1. When creating the consistency group, specify host initiators to create igroups.
2. Select the checkbox to Enable SnapMirror then choose the AutomatedFailoverDuplex policy.
3. In the dialog box that appears, select the Replicate initiator groups checkbox to replicate igroups. In Edit proximal settings, set proximal SVMs for your hosts.



4. Select Save

The protection relationship is established between the source and destination.



Perform Cluster Failover Validation Test

We recommend you perform planned failover tests to do a cluster validation check, the SQL databases or any clustered software on both sites – primary or mirrored site should continue to be accessible during tests.

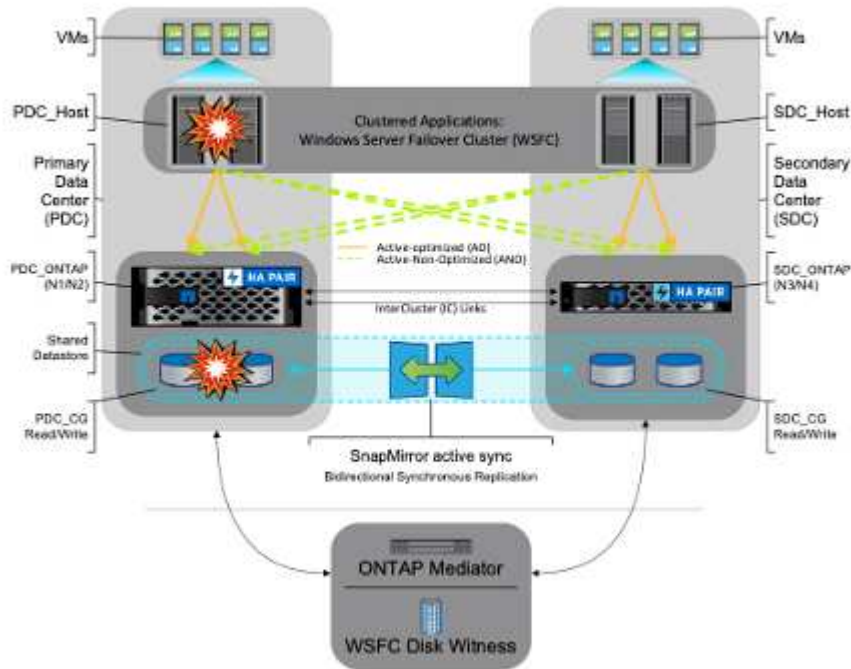
Hyper-V failover cluster requirements include:

- The SnapMirror active sync relationship must be in sync.
- You cannot initiate a planned failover when a nondisruptive operation is in process. Nondisruptive operations include volume moves, aggregate relocations, and storage failovers.
- The ONTAP Mediator must be configured, connected, and in quorum.
- At least two Hyper-V cluster nodes on each site with the CPU processors belongs to the same CPU family to optimize the process of VM migration. CPUs should be CPUs with support for hardware-assisted virtualization and hardware-based Data Execution Prevention (DEP).

- Hyper-V cluster nodes should be the same Active Directory Domain members to ensure resiliency.
- Hyper-V Cluster nodes and NetApp Storage Nodes should be connected by redundant networks to avoid a single point of failure.
- Shared storage, which can be accessed by all cluster nodes via iSCSI, Fibre Channel, or SMB 3.0 protocol.

Test Scenarios

There are many ways that trigger a failover on a host, storage or network.



Hyper-V failed node or a site

- Node failure
A failover cluster node can take over the workload of a failed node, a process known as failover.
Action: Power off a hyper-V node
Expect result: The other node in the cluster will take over the workload. VMs will be migrated to the other node.
- One site failure
We can also fail the entire site and trigger the primary site failover to the mirror site:
Action: Turn off both Hyper-V nodes on one site.
Expect result: VMs on the primary site will migrate to the mirror site Hyper-V cluster because SnapMirror active sync symmetric active/active serves IO locally with bidirectional replication, no workload impact with zero RPO and zero RTO.

Storage failure on one site

- Offline volumes
Action: cluster1::> volume offline vol1
Expected results: ONTAP will detect the primary site volume offline, the cluster will communicate with the mediator and detect the state of the storage. Primary site hyper-V communicate with mirror site storage

volume to achieve zero RPO and zero RTO.

- Stop a SVM on primary site

Action: Stop the iSCSI SVM

Expected results: Hyper-v primary cluster has already connected to the mirrored site and with SnapMirror active sync symmetric active/active no workload impact with zero RPO and zero RTO.

Success criteria

During the tests, observe the following:

- Observe the cluster's behavior and ensure that services are transferred to the remaining nodes.
- Check for any errors or service interruptions.
- Ensure that the cluster can handle storage failures and continue operating.
- Verify that database data remains accessible and that services continue to operate.
- Verify that database data integrity is maintained.
- Validate that specific applications can fail over to another node without user impact.
- Verify that the cluster can balance load and maintain performance during and after a failover.

Summary

SnapMirror active sync can help multisite application data, for example, MSSQL and Oracle to be actively accessible and in sync across both sites. If a failure occurs, applications are immediately redirected to the remaining active site, with no loss of data and no loss of access.

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.