



Openshift for on-premises

NetApp public and hybrid cloud solutions

NetApp

February 04, 2026

Table of Contents

Openshift for on-premises	1
NetApp Solution with Red Hat OpenShift Container platform workloads on VMware	1
Data protection and migration solution for OpenShift Container workloads using Trident Protect	1
Deploy and configure the Red Hat OpenShift Container platform on VMware	1
Data protection using Astra	4
Snapshot with ACC	4
Backup and Restore with ACC	5
Application specific execution hooks	5
Sample execution hook for pre-Snapshot of a redis application.....	6
Replication with ACC	6
Business Continuity with MetroCluster	7
Data migration using Trident Protect	8
Data Migration between different Kubernetes environment.....	8

Openshift for on-premises

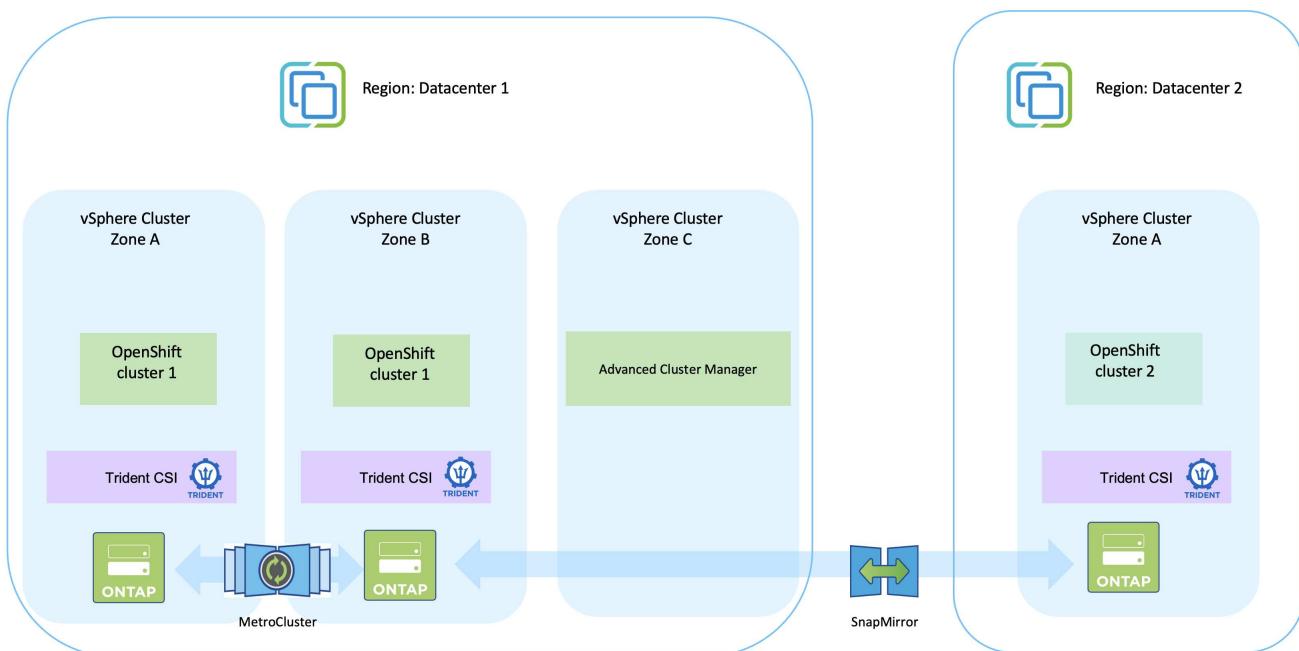
NetApp Solution with Red Hat OpenShift Container platform workloads on VMware

If customers have a need to run their modern containerized applications on infrastructure in their private data centers, they can do so. They should plan and deploy the Red Hat OpenShift container platform (OCP) for a successful production-ready environment for deploying their container workloads. Their OCP clusters can be deployed on VMware or bare metal.

NetApp ONTAP storage delivers data protection, reliability, and flexibility for container deployments. Trident serves as the dynamic storage provisioner to consume persistent ONTAP storage for customers' stateful applications. NetApp Trident Protect can be used for the many data management requirements of stateful applications such as data protection, migration, and business continuity.

With VMware vSphere, NetApp ONTAP tools provides a vCenter Plugin which can be utilized to provision datastores. Apply tags and use it with OpenShift for storing the node configuration and data. NVMe based storage provides lower latency and high performance.

Data protection and migration solution for OpenShift Container workloads using Trident Protect



Deploy and configure the Red Hat OpenShift Container platform on VMware

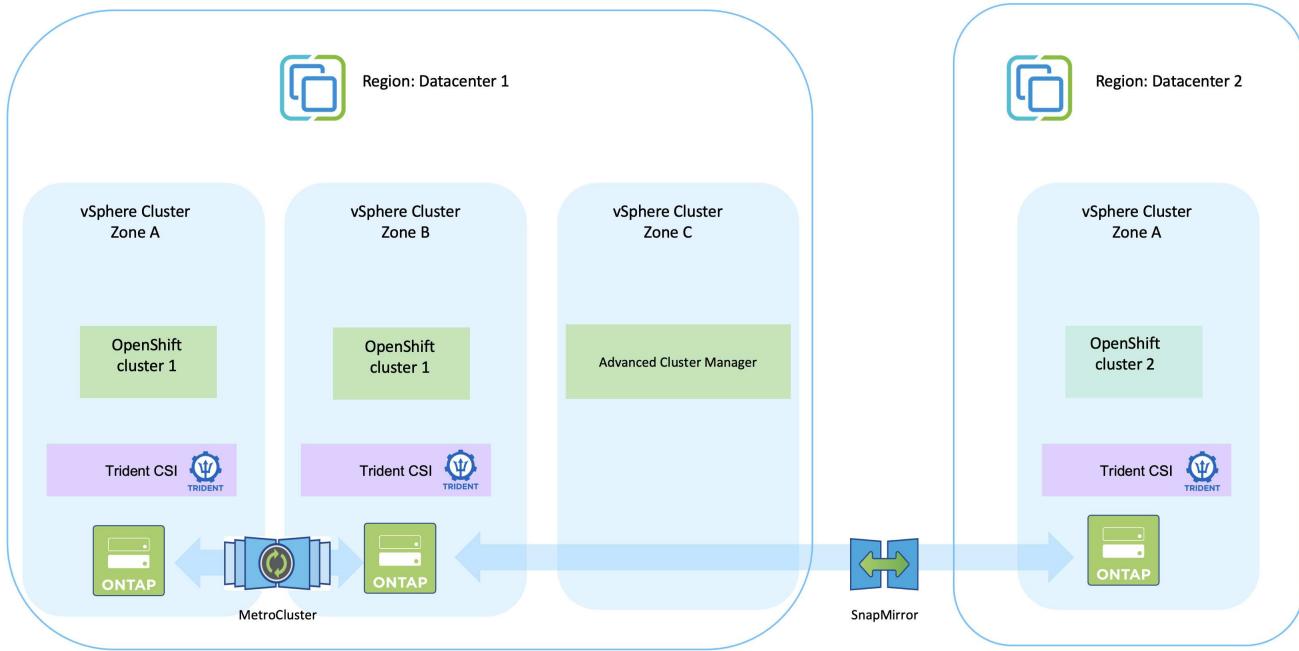
This section describes a high-level workflow of how to set up and manage OpenShift

clusters and manage stateful applications on them. It shows the use of NetApp ONTAP storage arrays with the help of Trident to provide persistent volumes.



There are several ways of deploying Red Hat OpenShift Container platform clusters. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the [resources section](#).

Here is a diagram that depicts the clusters deployed on VMware in a data center.



The setup process can be broken down into the following steps:

Deploy and configure a CentOS VM

- It is deployed in the VMware vSphere environment.
- This VM is used for deploying some components such as NetApp Trident and NetApp Trident Protect for the solution.
- A root user is configured on this VM during installation.

Deploy and configure an OpenShift Container Platform cluster on VMware vSphere (Hub Cluster)

Refer to the instructions for the [Assisted deployment](#) method to deploy an OCP cluster.

Remember the following:

- Create ssh public and private key to provide to the installer. These keys will be used to login to the master and worker nodes if needed.
- Download the installer program from the assisted installer. This program is used to boot the VMs that you create in the VMware vSphere environment for the master and worker nodes.
- VMs should have the minimum CPU, memory, and hard disk requirement. (Refer to the `vm create` commands on [this](#) page for the master and the worker nodes which provide this information)
- The `diskUUID` should be enabled on all VMs.
- Create a minimum of 3 nodes for master and 3 nodes for worker.
- Once they are discovered by the installer, turn on the VMware vSphere integration toggle button.



Install Advanced Cluster Management on the Hub cluster

This is installed using the Advanced Cluster Management Operator on the Hub Cluster. Refer to the instructions [here](#).

Install two additional OCP clusters (Source and Destination)

- The additional clusters can be deployed using the ACM on the Hub Cluster.
- Refer to the instructions [here](#).

Configure NetApp ONTAP storage

- Install an ONTAP cluster with connectivity to the OCP VMs in VMWare environment.
- Create an SVM.
- Configure NAS data lif to access the storage in SVM.

Install NetApp Trident on the OCP clusters

- Install NetApp Trident on all three clusters: Hub, source, and destination clusters
- Refer to the instructions [here](#).
- Create a storage backend for `ontap-nas` .
- Create a storage class for `ontap-nas`.
- Refer to instructions [here](#).

Deploy an Application on Source Cluster

Use OpenShift GitOps to deploy an application. (eg. Postgres, Ghost)

The next step is to use the Trident Protect for Data protection and Data migration from the source to the destination cluster.

Refer [here](#) for instructions.

Data protection using Astra

This page shows the data protection options for Red Hat OpenShift Container based applications running on VMware vSphere using Trident Protect (ACC).

As users take their journey of modernizing their applications with Red Hat OpenShift, a data protection strategy should be in place to protect them from accidental deletion or any other human errors. Often a protection strategy is also required for regulatory or compliance purposes to protect their data from a disaster.

The requirements of data protection varies from reverting back to a point in time copy to automatically failing over to a different fault domain without any human intervention. Many customers pick ONTAP as their preferred storage platform for their Kubernetes applications because of its rich features like multitenancy, multi-protocol, high performance and capacity offerings, replication and caching for multi-site locations, security and flexibility.

Data protection in ONTAP can be achieved using ad-hoc or policy controlled

- **Snapshot**
- **backup and restore**

Both Snapshot copies and backups protect the following types of data:

- **The application metadata that represents the state of the application**
- **Any persistent data volumes associated with the application**
- **Any resource artifacts belonging to the application**

Snapshot with ACC

A point in time copy of data can be captured using Snapshot with ACC. Protection policy defines the number of Snapshot copies to keep. Minimum schedule option available is hourly. Manual, on-demand Snapshot copies can be taken at any time and at shorter intervals than scheduled Snapshot copies. Snapshot copies are stored on the same provisioned volume as the app.

Configuring Snapshot with ACC

Backup and Restore with ACC

A backup is based on a Snapshot. Trident Protect can take Snapshot copies using CSI and perform backup using the point in time Snapshot copy. The backup is stored in an external object store (any s3 compatible including ONTAP S3 at a different location). Protection policy can be configured for scheduled backups and the number of backup versions to keep. The minimum RPO is one hour.

Restoring an application from a backup using ACC

ACC restores application from the S3 bucket where the backups are stored.

Application specific execution hooks

In addition, execution hooks can be configured to run in conjunction with a data protection operation of a managed app. Even though storage array level data protection features are available, often additional steps are needed to make backups and restores, application consistent. The app-specific additional steps could be:

- before or after a Snapshot copy is created.
- before or after a backup is created.
- after restoring from a Snapshot copy or backup.

Astra Control can execute these app-specific steps coded as custom scripts called execution hooks.

[NetApp Verda GitHub project](#) provides execution hooks for popular cloud-native applications to make protecting applications straightforward, robust, and easy to orchestrate. Feel free to contribute to that project if you have enough information for an application that is not in the repository.

Sample execution hook for pre-Snapshot of a redis application.

HOOK DETAILS

Operation: Pre-snapshot

Hook arguments (optional): `pre`

Hook name: redis-pre-snapshot

CONTAINER IMAGES

Apply to all container images

Container image names to match: redis

SCRIPT

+ Add

Name: mariadb_mysql.sh

Name: postgresql.sh

Name: redis_hook.sh

Save

Replication with ACC

For regional protection or for a low RPO and RTO solution, an application can be replicated to another Kubernetes instance running at a different site, preferably in another region. Trident Protect utilizes ONTAP async SnapMirror with RPO as low as 5 minutes. Replication is done by replicating to ONTAP and then a fail over creates the Kubernetes resources in the destination cluster.

i Note that replication is different from the backup and restore where the backup goes to S3 and restore is performed from S3. Refer xref:./openshift/ [here](#) to get additional details about the differences between the two types of data protection.

Refer [here](#) for SnapMirror setup instructions.

SnapMirror with ACC

The screenshot shows the Astra application protection interface. On the left, a sidebar lists navigation options: Dashboard, Applications, Client, Cloud Instances, Backends, Buckets, Account, Activity, and Support. The main area is titled 'ghost' and contains two tabs: 'APPLICATION STATUS' and 'APPLICATION PROTECTION'. The 'APPLICATION STATUS' tab shows 'ghost' as 'Healthy' and 'ghost-clone' as 'Replicated?'. The 'APPLICATION PROTECTION' tab shows 'Fully protected' and 'Protection policy configured'. Below these tabs are tabs for 'Data protection', 'Storage', 'Resources', 'Execution hooks', 'Activity', and 'Tasks'. A 'Configure' button is also present. On the right, an 'Actions' dropdown menu includes options like Snapshot, Back up, Clone, Restore, Fail-over, Reverse replication, and Unmanage. Below the 'Actions' menu are buttons for Snapshots, Backups, and Replication. The 'Replication relationship' section on the right details the status as 'Established', the schedule as 'Replicate snapshot every 5 minutes to ghost-clone', and the last sync as '2023/04/26 19:16 UTC' with a duration of '30 seconds'.



san-economy and nas-economy storage drivers do not support replication feature. Refer [here](#) for additional details.

Demo video:

[Demonstration video of disaster recovery with Trident Protect](#)

[Data protection with Trident Protect](#)

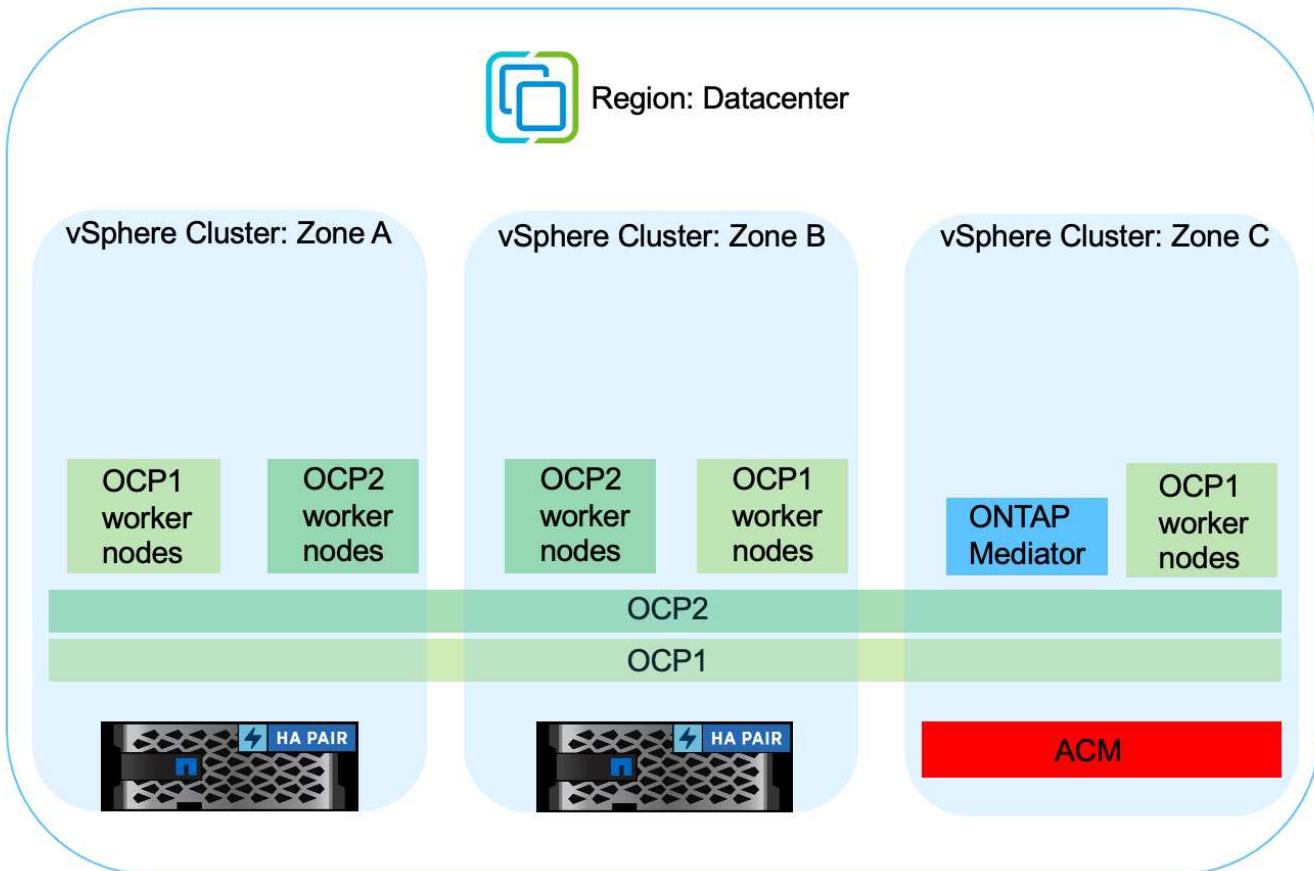
Business Continuity with MetroCluster

Most of our hardware platform for ONTAP has high availability features to protect from device failures avoiding the need to perform disaster recovery. But to protect from fire or any other disaster and to continue the business with zero RPO and low RTO, often a MetroCluster solution is used.

Customers who currently have an ONTAP system can extend to MetroCluster by adding supported ONTAP systems within the distance limitations for providing zone level disaster recovery.

Trident, the CSI (Container Storage Interface) supports NetApp ONTAP including MetroCluster configuration as well as other options like Cloud Volumes ONTAP, Azure NetApp Files, AWS FSx ONTAP, etc. Trident provides five storage driver options for ONTAP and all are supported for MetroCluster configuration. Refer [here](#) for additional details about ONTAP storage drivers supported by Trident.

The MetroCluster solution requires layer 2 network extension or capability to access the same network address from both fault domains. Once MetroCluster configuration is in place, the solution is transparent to application owners as all the volumes in the MetroCluster svm are protected and get the benefits of SyncMirror (zero RPO).



For Trident Backend Configuration (TBC), do not specify the dataLIF and SVM when using MetroCluster configuration. Specify SVM management IP for managementLIF and use vsadmin role credentials.

Details on Trident Protect Data Protection features are available [here](#)

Data migration using Trident Protect

This page shows the data migration options for container workloads on Red Hat OpenShift clusters with Trident Protect.

Kubernetes Applications are often required to be moved from one environment to another. To migrate an application along with its persistent data, NetApp Trident Protect can be utilized.

Data Migration between different Kubernetes environment

ACC supports various Kubernetes flavors including Google Anthos, Red Hat OpenShift, Tanzu Kubernetes Grid, Rancher Kubernetes Engine, Upstream Kubernetes, etc. For additional details, refer [here](#).

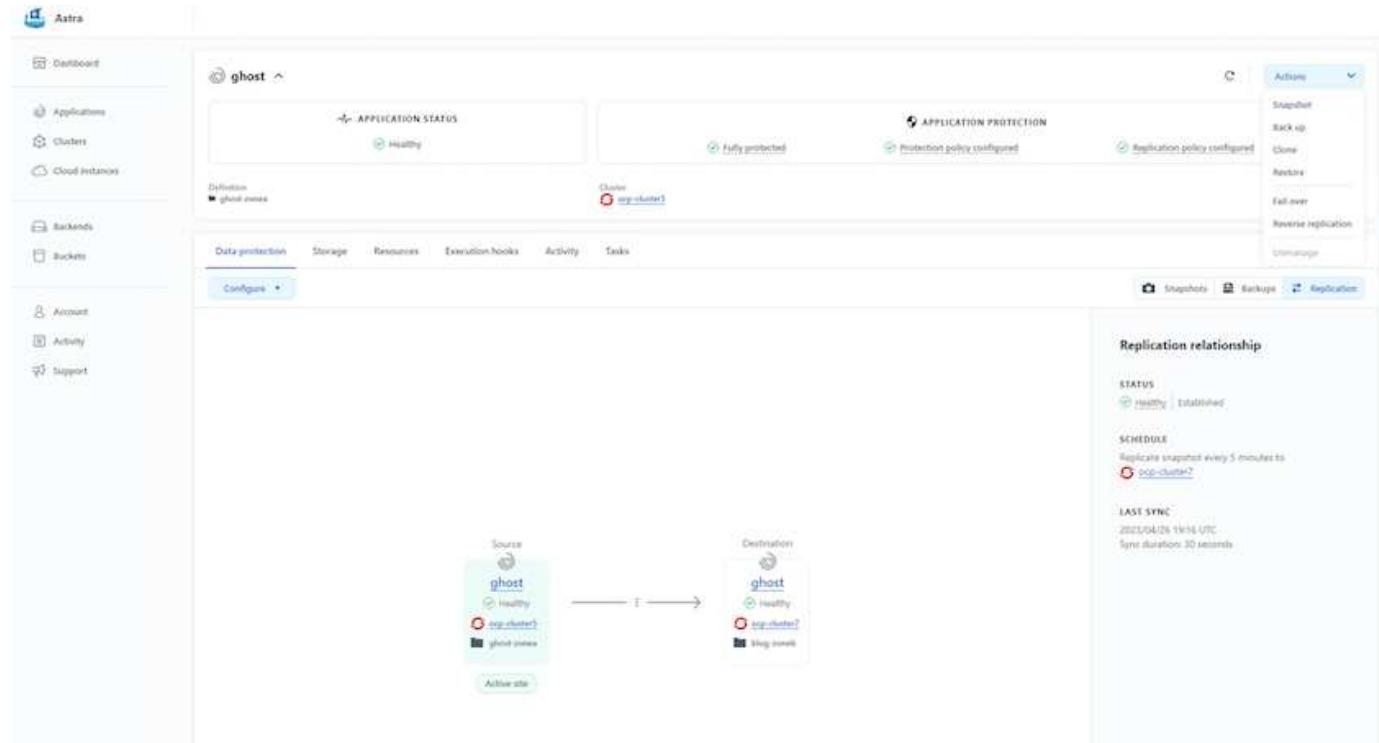
To migrate application from one cluster to another, you can use one of the following features of ACC:

- **replication**
- **backup and restore**
- **clone**

Refer to the [data protection section](#) for the **replication and backup and restore** options.

Refer [here](#) for additional details about **cloning**.

Performing data replication using ACC



The screenshot shows the Astra Cloud Control Center (ACC) interface. On the left, a sidebar navigation includes: Dashboard, Applications (ghost), Cluster, Cloud Instances, Backends, Buckets, Account, Activity, and Support. The main content area is for the application 'ghost'. It shows the following details:

- APPLICATION STATUS:** The application is marked as **Healthy**.
- DATA PROTECTION:** The status is **Fully protected**. The 'Cluster' dropdown is set to **src-cluster?**.
- APPLICATION PROTECTION:** Protection policy is **configured**.
- Replication relationship:**
 - STATUS:** **Established**.
 - SCHEDULE:** Replicate snapshot every 5 minutes to **dst-cluster?**
 - LAST SYNC:** 2023/04/26 19:16 UTC, sync duration: 30 seconds.

The interface also includes tabs for Storage, Resources, Execution hooks, Activity, and Tasks, and a 'Configure' button. A 'Actions' dropdown menu on the right provides options for Snapshot, Back up, Clone, Restore, Fail-over, Reverse replication, and Unmanage.

Copyright information

Copyright © 2026 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.