



NetApp storage reference architectures for Epic

FlexPod

NetApp
June 03, 2021

This PDF was generated from https://docs.netapp.com/us-en/flexpod/healthcare/ehr-epic-sizing_netapp_storage_reference_architectures_for_epic.html on September 12, 2021. Always check docs.netapp.com for the latest.

Table of Contents

- NetApp storage reference architectures for Epic 1
 - Storage layout and LUN configuration 1
 - Small configuration: four-node reference architecture for fewer than 5M global references (up to ~22K total IOPS) 2
 - Medium configuration: six-node reference architecture for greater than 5M global references (22K-50K total IOPS) 4
 - Large configuration: reference architecture for greater than 10M global references (more than 50K IOPS) . 6

NetApp storage reference architectures for Epic

An appropriate storage architecture can be determined by the overall database size and the total IOPS. Performance alone is not the only factor, and you might decide to use a larger node count based on additional customer requirements.

Given the storage requirements for Epic software environments, NetApp has three reference architectures based on the size of the environment. Epic requires the use of NetApp sizing methods to properly size a NetApp storage system for use in Epic environments. For quantitative performance requirements and sizing guidance, see NetApp [TR-3930i: NetApp Sizing Guidelines for Epic](#). NetApp Field Portal access is required to view this document.

The architectures listed here are a starting point for the design. The workloads must be validated in the SPM tool for the number of disks and controller utilization. Work with the NetApp Epic team to validate all designs.

All Epic production is deployed on all-flash arrays. In this report, the disk pools required for spinning disk have been consolidated to three disk pools for all-flash arrays. Before reading this section, you should review the Epic All-Flash Reference Architecture Strategy Handbook for the Epic storage layout requirements.

The three storage reference architectures are as follows:

- **Small.** Four-node architecture with two nodes in production and two nodes in DR (fewer than 5M global references)
- **Medium.** Six-node architecture with four nodes in production and two nodes in DR (more than 5M global references)
- **Large.** Twelve-or-more node architecture with six to ten nodes in production (5M-10M global references)



Global references = (Read IOPS + (Write Operations per 80-Second Write Burst / 45)) * 225.
These numbers are taken from the customer-specific Epic Hardware Configuration Guide.

Storage layout and LUN configuration

The first step in satisfying Epic's high-availability (HA_) and redundancy requirements is to design the storage layout specifically for the Epic software environment. The design considerations should include isolating disk pool 1 from disk pool 2 on dedicated high-performance storage. See the Epic All-Flash Reference Architecture Strategy Handbook for information about what workloads are in each disk pool.

Placing each disk pool on a separate node creates the fault domains required for the isolation of Epic's production and nonproduction workloads. Using one aggregate per node maximizes disk utilization and aggregate affinity to provide better performance. This design also maximizes storage efficiency with aggregate-level deduplication.

Because Epic allows storage resources to be shared for nonproduction needs, a storage system can often service both the Clarity server and production services storage needs, such as virtual desktop infrastructure (VDI), CIFS, and other enterprise functions.

The Epic Database Storage Layout Recommendations document provides recommendations for the size and number of LUNs for each database. These recommendations might need to be adjusted according to your environment. It is important to review these recommendations with Epic support and finalize the number of LUNs and LUN sizes.



NetApp recommends starting with larger size LUNs because the size of the LUNs themselves have no cost to storage. For ease of operation, make sure that the number of LUNs and initial size can grow well beyond expected requirements after 3 years. Growing LUNs is much easier to manage than adding LUNs while scaling. With thin-provisioned LUNs and volumes, the storage- used space shows up in the aggregate.

Epic requires database, journal, and application or system storage to be presented to database servers as LUNs through FC.

Use one LUN per volume for Epic production and for Clarity. For larger deployments, NetApp recommends 24 to 32 LUNs for the Epic database. Factors that determine the number of LUNs to use are:

- Overall size of the Epic DB after 3 years. For larger DBs, determine the maximum size of the LUN for that operating system (OS) and make sure that you have enough LUNs to scale. For example, if you need a 60TB Epic database and the OS LUNs have a 4TB maximum, you need 24 to 32 LUNs to provide scale and head room.

Regardless of whether the architecture is small, medium, or large:

- ONTAP allows easy nondisruptive scale up and scale out. Disks and nodes can be upgraded, added, or removed by using ONTAP nondisruptive operations. Customers can start with four nodes and move to six nodes or upgrade to larger controllers nondisruptively.
- NetApp OnCommand Workflow Automation workflows can back up and refresh Epic full-copy test environments. This solution simplifies the architecture and saves on storage capacity with integrated efficiencies.
- The DR shadow database server is part of a customer's business continuity strategy (used to support SRO functionality and potentially configured to be an SRW instance). Therefore, the placement and sizing of the third storage system are usually the same as in the production database storage system.
- Database consistency requires some consideration. If NetApp SnapMirror backup copies are used in relation to business continuity, see the Epic document Business Continuity Technical Solutions Guide. For information about the use of SnapMirror technologies, see [TR-3446: SnapMirror Async Overview and Best Practices Guide](#).
- Isolation of production from potential bully workloads is a key design objective of Epic. A storage pool is a fault domain in which workload performance must be isolated and protected. Each node in an ONTAP cluster is a fault domain and can be considered as a pool of storage.

All platforms in the ONTAP family can run the full host of feature sets for Epic workloads.

Small configuration: four-node reference architecture for fewer than 5M global references (up to ~22K total IOPS)

The small reference architecture is a four-node architecture with two nodes in production and two nodes in DR, with fewer than 5M global references. This architecture can be used by customers with fewer than 5M global references. At this size, the separation of Report and Clarity is not required.

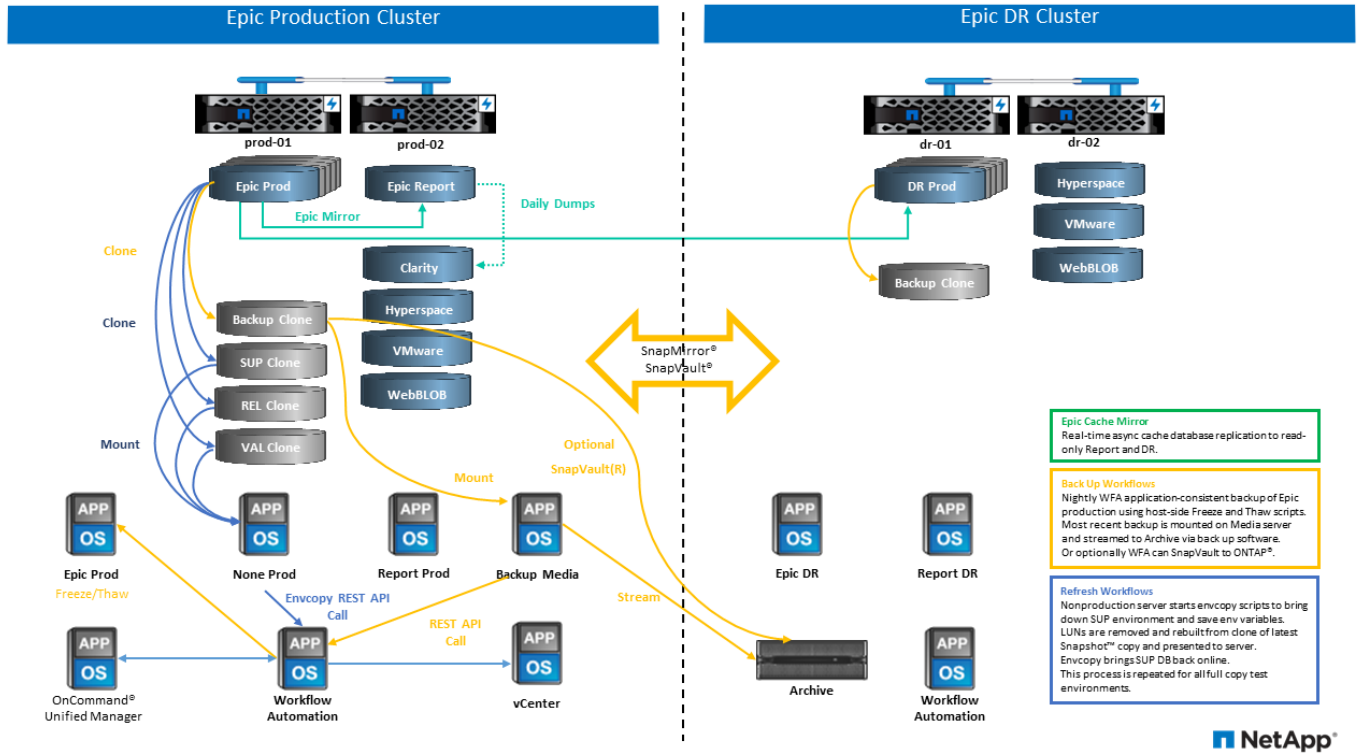
With unique multiprotocol support from NetApp, QoS, and the ability to create fault domains in the same cluster, you can run all the production workload for disk pool1 and disk pool2 on a single HA pair and meet all of NetApp best practices and Epic's High Comfort rating requirements. All of disk pool1 would run on node1 and all of disk pool 2 would run on pool2.

With the ability of ONTAP to segregate workloads in the same cluster, and ONTAP multiprotocol support, all

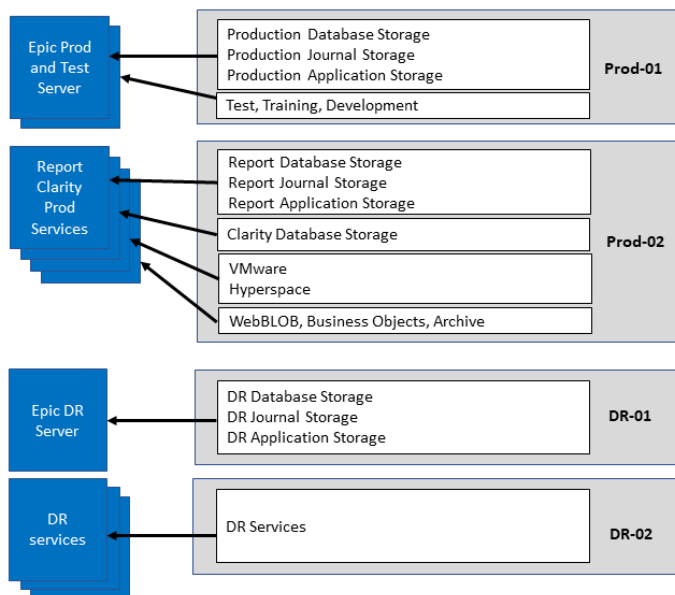
the production Epic workloads (Production, Report, Clarity, VMware, Citrix, CIFS, and Epic- related workloads) can be run on a single HA pair in a single cluster. This capability enables you to meet all of Epic's requirements (documented in the Epic All-Flash Reference Architecture Strategy Handbook) and all the NetApp best practices. Basically, pool1 runs on node prod-01 and pool2 runs on prod-02, as shown in the figure below. The NAS 1 workload can be placed on node 2 with NetApp multiprotocol NAS and SAN capabilities.

For disaster recovery, Epic DR pool 3 is split between the two nodes in the HA pair. Epic DR runs on node dr-01 and DR services run on dr-02.

NetApp SnapMirror or SnapVault replication can be set up as needed for workloads.



From a storage design and layout perspective, the following figure shows a high-level storage layout for the production database and the other constructs that comprise the Epic workload.



Medium configuration: six-node reference architecture for greater than 5M global references (22K-50K total IOPS)

The medium reference architecture is a six-node architecture with four nodes in production and two nodes in DR, with 5M-10M global references.

For this size, the All-Flash Reference Architecture Strategy Handbook states that you need to separate Epic Report workloads from Clarity, and that you need at least four nodes in production.

The six-node architecture is the most commonly deployed architecture in Epic environments. Customers with more than 5,000,000 global references are required to place Report and Clarity in separate fault domains. See the Epic All-Flash Reference Architecture Strategy Handbook.

Customers with fewer than 5,000,000 global references can opt to go with six nodes rather than four nodes for the following key advantages:

- Offload backup archive process from production
- Offload all test environments from production

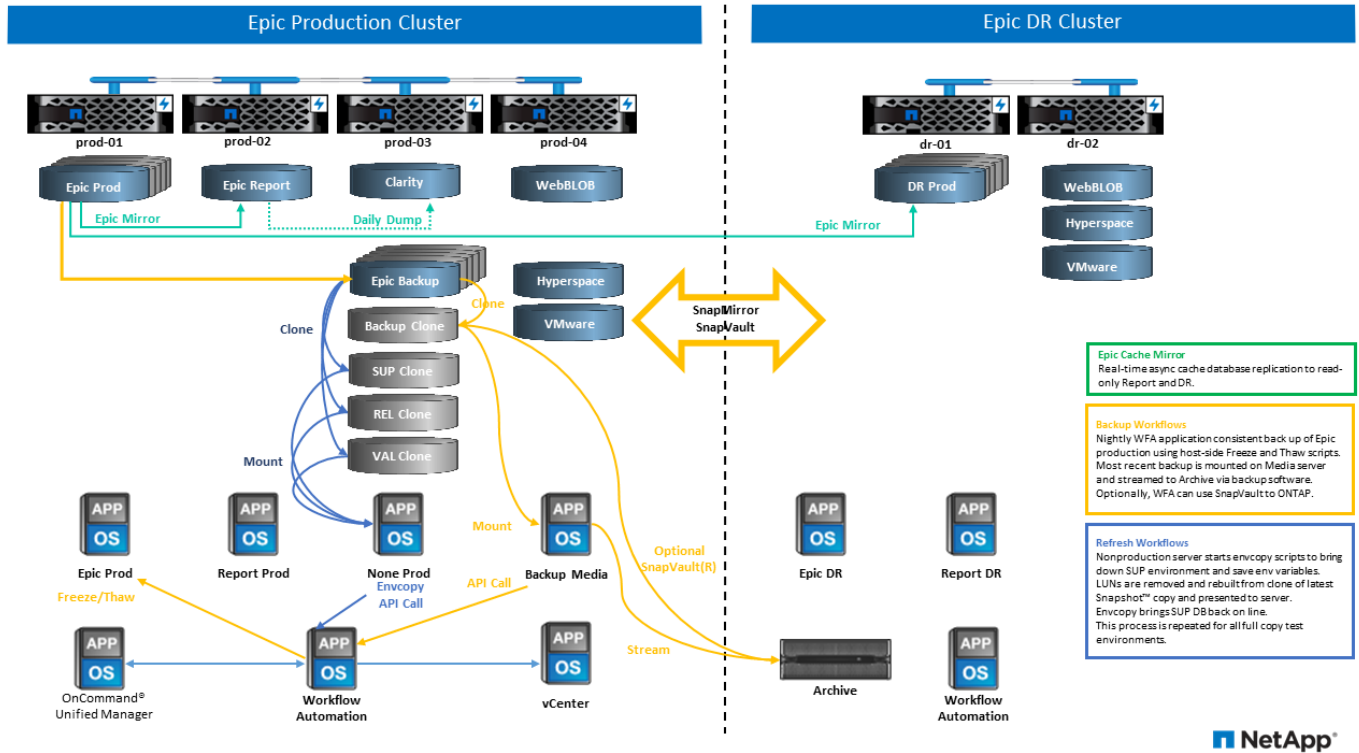
Production runs on node prod-01. Report runs on node prod-02, which is an up-to-the-minute Epic mirror copy of production. Test environments like support, release, and release validation can be cloned from either Epic production, Report, or DR. The figure below shows clones made from production for full-copy test environments.

The second HA pair is used for production services storage requirements. These workloads include storage for Clarity database servers (SQL or Oracle), VMware, Hyperspace, and CIFS. Customers might have non-Epic workloads that could be added to nodes 3 and node 4 in this architecture, or preferably added to a separate HA pair in the same cluster.

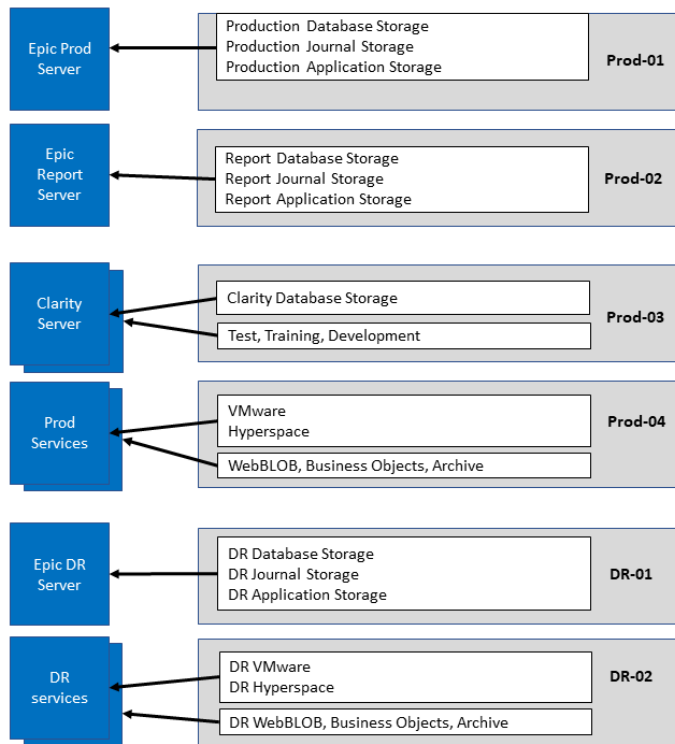
SnapMirror technology is used for storage-level replication of the production database to the second HA pair.

SnapMirror backup copies can be used to create NetApp FlexClone volumes on the second storage system for nonproduction environments such as support, release, and release validation. Storage-level replicas of the production database can also support customers' implementation of their DR strategy.

Optionally, to be more storage efficient, full-test clones can be created from the Report NetApp Snapshot copy backup and run directly on node 2. In this design, a SnapMirror destination copy is not required to be saved on disk.



The following figure shows the storage layout for a six-node architecture.



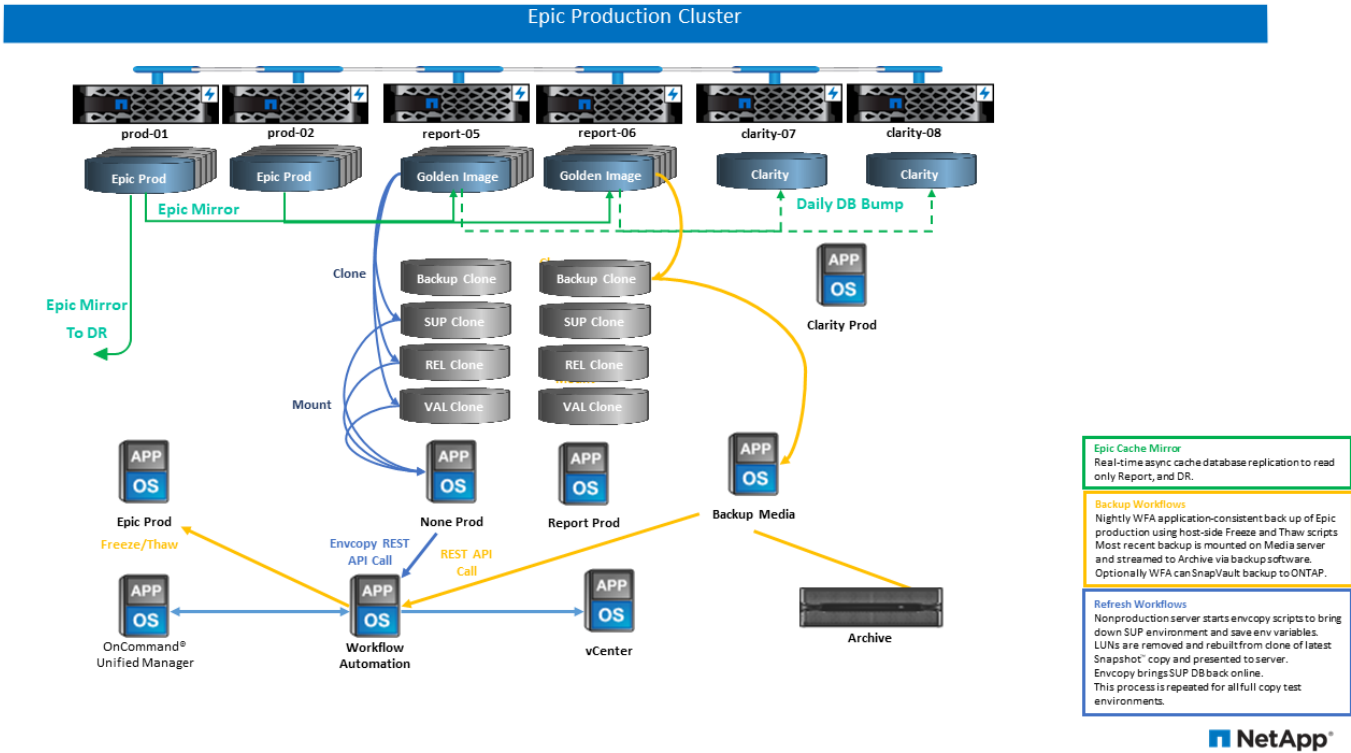
Large configuration: reference architecture for greater than 10M global references (more than 50K IOPS)

The large architecture is typically a twelve-or-more-node architecture with six to ten nodes in production, with more than 10M global references. For large Epic deployments, Epic Production, Epic Report, and Clarity can be placed on a dedicated HA pair with storage evenly balanced among the nodes, as shown in the figure below.

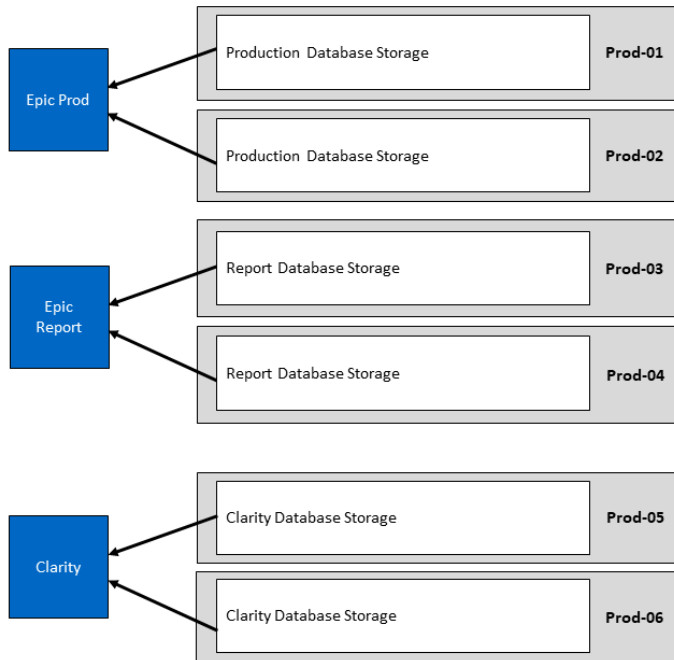
Larger customers have two options:

- Retain the six-node architecture and use AFF A700 controllers.
- Run Epic production, report, and DR on a dedicated AFF A300 HA pair.

You must use the SPM to compare controller utilization. Also, consider rack space and power when selecting controllers.



The following figure shows the storage layout for a large reference architecture.



Copyright Information

Copyright © 2021 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.