



Review solution design

BeeGFS on NetApp with E-Series Storage

NetApp

January 31, 2023

Table of Contents

- Review solution design. 1
- Design overview. 1
- Hardware configuration 1
- Software configuration 2
- Design verification 6
- Sizing guidelines 9
- Performance tuning 10
- High capacity building block 12

Review solution design

Design overview

Specific equipment, cabling, and configurations are required to support the BeeGFS on NetApp solution, which combines the BeeGFS parallel file system with the NetApp EF600 storage systems.

Learn more:

- [Hardware configuration](#)
- [Software configuration](#)
- [Design verification](#)
- [Sizing guidelines](#)
- [Performance tuning](#)

Derivative architectures with variations in design and performance:

- [High Capacity Building Block](#)

Hardware configuration

The hardware configuration for BeeGFS on NetApp includes file nodes and network cabling.

File node configuration

File nodes have two CPU sockets configured as separate NUMA zones, which include local access to an equal number of PCIe slots and memory.

InfiniBand adapters must be populated in the appropriate PCI risers or slots, so the workload is balanced over the available PCIe lanes and memory channels. You balance the workload by fully isolating work for individual BeeGFS services to a particular NUMA node. The goal is to achieve similar performance from each file node as if it were two independent single socket servers.

The following figure shows the file node NUMA configuration.

[beegfs design image5 small]

The BeeGFS processes are pinned to a particular NUMA zone to ensure that the interfaces used are in the same zone. This configuration avoids the need for remote access over the inter-socket connection. The inter-socket connection is sometimes known as the QPI or GMI2 link; even in modern processor architectures, they can be a bottleneck when using high-speed networking like HDR InfiniBand.

Network cabling configuration

Within a building block, each file node is connected to two block nodes using a total of four redundant InfiniBand connections. In addition, each file node has four redundant connections to the InfiniBand storage network.

In the following figure, notice that:

- All file node ports outlined in green are used to connect to the storage fabric; all other file node ports are the direct connects to the block nodes.
- Two InfiniBand ports in a specific NUMA zone connect to the A and B controllers of the same block node.
- Ports in NUMA node 0 always connect to the first block node.
- Ports in NUMA node 1 connect to the second block node.

[beegfs design image6]



For storage networks with redundant switches, ports outlined in light green should connect to one switch, and ports in dark green to another switch.

The cabling configuration depicted in the figure allows each BeeGFS service to:

- Run in the same NUMA zone regardless of which file node is running the BeeGFS service.
- Have secondary optimal paths to the front-end storage network and to the back-end block nodes regardless of where a failure occurs.
- Minimize performance effects if a file node or controller in a block node requires maintenance.

Cabling to leverage bandwidth

To leverage the full PCIe bidirectional bandwidth, make sure one port on each InfiniBand adapter connects to the storage fabric, and the other port connects to a block node. The theoretical maximum speed of an HDR InfiniBand port is 25GBps (not accounting for signaling and other overhead). The maximum single direction bandwidth of a PCIe 4.0 x16 slot is 32GBps, creating a potential bottleneck when implementing file nodes that incorporate dual port InfiniBand adapters that can theoretically handle 50GBps of bandwidth.

The following figure shows the cabling design used to leverage the full PCIe bidirectional bandwidth.

[beegfs design image7]

For each BeeGFS service, use the same adapter to connect the preferred port used for client traffic with the path to the block nodes controller that is the primary owner of that services volumes. For more information, see [Software configuration](#).

Software configuration

The software configuration for BeeGFS on NetApp includes BeeGFS network components, EF600 block nodes, BeeGFS file nodes, resource groups, and BeeGFS services.

BeeGFS network configuration

The BeeGFS network configuration consists of the following components.

- **Floating IPs**

Floating IPs are a kind of virtual IP address that can be dynamically routed to any server in the same network. Multiple servers can own the same Floating IP address, but it can only be active on one server at any given time.

Each BeeGFS server service has its own IP address that can move between file nodes depending on the run location of the BeeGFS server service. This floating IP configuration allows each service to fail over independently to the other file node. The client simply needs to know the IP address for a particular BeeGFS service; it does not need to know which file node is currently running that service.

- **BeeGFS server multi-homing configuration**

To increase the density of the solution, each file node has multiple storage interfaces with IPs configured in the same IP subnet.

Additional configuration is required to make sure that this configuration works as expected with the Linux networking stack, because by default, requests to one interface can be responded to on a different interface if their IPs are in the same subnet. In addition to other drawbacks, this default behavior makes it impossible to properly establish or maintain RDMA connections.

The Ansible-based deployment handles tightening of the reverse path (RP) and address resolution protocol (ARP) behavior, along with ensuring when floating IPs are started and stopped; corresponding IP routes and rules are dynamically created to allow the multihomed network configuration to work properly.

- **BeeGFS client multi-rail configuration**

Multi-rail refers to the ability of an application to use multiple independent network “rails” to increase performance.

Although BeeGFS can use RDMA for connectivity, BeeGFS uses IPoIB to simplify discovering and establishing RDMA connections. To allow BeeGFS clients to use multiple InfiniBand interfaces, you can configure each client with an IP address located in a different subnet and then configure the preferred interfaces for half of the BeeGFS server services in each subnet.

In the following diagram, interfaces highlighted in light green are located in one IP subnet (for example, 100.127.0.0/16) and the dark green interfaces are located in another subnet (for example, 100.128.0.0/16).

The following figure shows the balancing of traffic across multiple BeeGFS client interfaces.

[beegfs design image8]

Because each file in BeeGFS is typically striped across multiple storage services, the multi-rail configuration allows the client to achieve more throughput than is possible with a single InfiniBand port. For example, the following code sample shows a common file-striping configuration that allows the client to balance traffic across both interfaces:

```

root@ictad21h01:/mnt/beegfs# beegfs-ctl --getentryinfo myfile
Entry type: file
EntryID: 11D-624759A9-65
Metadata node: meta_01_tgt_0101 [ID: 101]
Stripe pattern details:
+ Type: RAID0
+ Chunksize: 1M
+ Number of storage targets: desired: 4; actual: 4
+ Storage targets:
  + 101 @ stor_01_tgt_0101 [ID: 101]
  + 102 @ stor_01_tgt_0101 [ID: 101]
  + 201 @ stor_02_tgt_0201 [ID: 201]
  + 202 @ stor_02_tgt_0201 [ID: 201]

```

Using two IPoIB subnets is a logical distinction. You can use a single physical InfiniBand subnet (storage network), if desired.



Multi-rail support was added in BeeGFS 7.3.0 to allow the use of multiple IB interfaces in a single IPoIB subnet. The design for the BeeGFS on NetApp solution was developed before the general availability of BeeGFS 7.3.0, and thus demonstrates the use of two IP subnets to use two IB interfaces on the BeeGFS clients. One advantage of the multiple IP subnet approach is eliminating the need to configure multihoming on BeeGFS client nodes (for more information, see [BeeGFS RDMA support](#)).

EF600 block node configuration

Block nodes are comprised of two active/active RAID controllers with shared access to the same set of drives. Typically, each controller owns half the volumes configured on the system, but can take over for the other controller as needed.

Multipathing software on the file nodes determines the active and optimized path to each volume and automatically moves to the alternate path in the event of a cable, adapter, or controller failure.

The following diagram shows the controller layout in EF600 block nodes.

[beegfs design image9]

To facilitate the shared-disk HA solution, volumes are mapped to both file nodes so that they can take over for each other as needed. The following diagram shows an example of how the BeeGFS service and preferred volume ownership is configured for maximum performance. The interface to the left of each BeeGFS service indicates the preferred interface that the clients and other services use to contact it.

[beegfs design image10]

In the previous example, clients and server services prefer to communicate with storage service 1 using interface i1b. Storage service 1 uses interface i1a as the preferred path to communicate with its volumes (storage_tgt_101, 102) on controller A of the first block node. This configuration makes use of the full bidirectional PCIe bandwidth available to the InfiniBand adapter and achieves better performance from a dual-port HDR InfiniBand adapter than would otherwise be possible with PCIe 4.0.

BeeGFS file node configuration

The BeeGFS file nodes are configured into a High-Availability (HA) cluster to facilitate failover of BeeGFS services between multiple file nodes.

The HA cluster design is based on two widely used Linux HA projects: Corosync for cluster membership and Pacemaker for cluster resource management. For more information, see [Red Hat training for high-availability add-ons](#).

NetApp authored and extended several open cluster framework (OCF) resource agents to allow the cluster to intelligently start and monitor the BeeGFS resources.

BeeGFS HA clusters

Typically, when you start a BeeGFS service (with or without HA), a few resources must be in place:

- IP addresses where the service is reachable, typically configured by Network Manager.
- Underlying file systems used as the targets for BeeGFS to store data.

These are typically defined in `/etc/fstab` and mounted by Systemd.

- A Systemd service responsible for starting BeeGFS processes when the other resources are ready.

Without additional software, these resources start only on a single file node. Therefore, if the file node goes offline, a portion of the BeeGFS file system is inaccessible.

Because multiple nodes can start each BeeGFS service, Pacemaker must make sure each service and dependent resources are only running on one node at a time. For example, if two nodes try to start the same BeeGFS service, there is a risk of data corruption if they both try to write to the same files on the underlying target. To avoid this scenario, Pacemaker relies on Corosync to reliably keep the state of the overall cluster in sync across all nodes and establish quorum.

If a failure occurs in the cluster, Pacemaker reacts and restarts BeeGFS resources on another node. In some scenarios, Pacemaker might not be able to communicate with the original faulty node to confirm the resources are stopped. To verify that the node is down before restarting BeeGFS resources elsewhere, Pacemaker fences off the faulty node, ideally by removing power.

Many open-source fencing agents are available that enable Pacemaker to fence a node with a power distribution unit (PDU) or by using the server baseboard management controller (BMC) with APIs such as Redfish.

When BeeGFS is running in an HA cluster, all BeeGFS services and underlying resources are managed by Pacemaker in resource groups. Each BeeGFS service and the resources it depends on, are configured into a resource group, which ensures resources are started and stopped in the correct order and collocated on the same node.

For each BeeGFS resource group, Pacemaker runs a custom BeeGFS monitoring resource that is responsible for detecting failure conditions and intelligently triggering failovers when a BeeGFS service is no longer accessible on a particular node.

The following figure shows the Pacemaker-controlled BeeGFS services and dependencies.

[beegfs design image11]



So that multiple BeeGFS services of the same type are started on the same node, Pacemaker is configured to start BeeGFS services using the Multi Mode configuration method. For more information, see the [BeeGFS documentation on Multi Mode](#).

Because BeeGFS services must be able to start on multiple nodes, the configuration file for each service (normally located at `/etc/beegfs`) is stored on one of the E-Series volumes used as the BeeGFS target for that service. This makes the configuration along with the data for a particular BeeGFS service accessible to all nodes that might need to run the service.

```
# tree stor_01_tgt_0101/ -L 2
stor_01_tgt_0101/
├── data
│   ├── benchmark
│   ├── buddymir
│   ├── chunks
│   ├── format.conf
│   ├── lock.pid
│   ├── nodeID
│   ├── nodeNumID
│   ├── originalNodeID
│   ├── targetID
│   └── targetNumID
├── storage_config
│   ├── beegfs-storage.conf
│   ├── connInterfacesFile.conf
│   └── connNetFilterFile.conf
```

Design verification

The second-generation design for the BeeGFS on NetApp solution was verified using three building block configuration profiles.

The configuration profiles include the following:

- A single base building block, including BeeGFS management, metadata, and storage services.
- A BeeGFS metadata plus a storage building block.
- A BeeGFS storage-only building block.

The building blocks were attached to two Mellanox Quantum InfiniBand (MQM8700) switches. Ten BeeGFS clients were also attached to the InfiniBand switches and used to run synthetic benchmark utilities.

The following figure shows the BeeGFS configuration used to validate the BeeGFS on NetApp solution.

[beegfs design image12]

BeeGFS file striping

A benefit of parallel file systems is the ability to stripe individual files across multiple storage targets, which could represent volumes on the same or different underlying storage systems.

In BeeGFS, you can configure striping on a per-directory and per-file basis to control the number of targets used for each file and to control the chunksize (or block size) used for each file stripe. This configuration allows the file system to support different types of workloads and I/O profiles without the need for reconfiguring or restarting services. You can apply stripe settings using the `beegfs-ctl` command line tool or with applications that use the striping API. For more information, see the BeeGFS documentation for [Striping](#) and [Striping API](#).

To achieve the best performance, stripe patterns were adjusted throughout testing, and the parameters used for each test are noted.

IOR bandwidth tests: Multiple clients

The IOR bandwidth tests used OpenMPI to run parallel jobs of the synthetic I/O generator tool IOR (available from [HPC GitHub](#)) across all 10 client nodes to one or more BeeGFS building blocks. Unless otherwise noted:

- All tests used direct I/O with a 1MiB transfer size.
- BeeGFS file striping was set to a 1MB chunksize and one target per file.

The following parameters were used for IOR with the segment count adjusted to keep the aggregate file size to 5TiB for one building block and 40TiB for three building blocks.

```
mpirun --allow-run-as-root --mca btl tcp -np 48 -map-by node -hostfile  
10xnodes ior -b 1024k --posix.odirect -e -t 1024k -s 54613 -z -C -F -E -k
```

One BeeGFS base (management, metadata, and storage) building block

The following figure shows the IOR test results with a single BeeGFS base (management, metadata, and storage) building block.

[beegfs design image13]

BeeGFS metadata + storage building block

The following figure shows the IOR test results with a single BeeGFS metadata + storage building block.

[beegfs design image14]

BeeGFS storage-only building block

The following figure shows the IOR test results with a single BeeGFS storage-only building block.

[beegfs design image15]

Three BeeGFS building blocks

The following figure shows the IOR test results with three BeeGFS building blocks.

[beegfs design image16]

As expected, the performance difference between the base building block and the subsequent metadata + storage building block is negligible. Comparing the metadata + storage building block and a storage-only

building block shows a slight increase in read performance due to the additional drives used as storage targets. However, there is no significant difference in write performance. To achieve higher performance, you can add multiple building blocks together to scale performance in a linear fashion.

IOR bandwidth tests: Single client

The IOR bandwidth test used OpenMPI to run multiple IOR processes using a single high-performance GPU server to explore the performance achievable to a single client.

This test also compares the reread behavior and performance of BeeGFS when the client is configured to use the Linux kernel page-cache (`tuneFileCacheType = native`) versus the default `buffered` setting.

The native caching mode uses the Linux kernel page-cache on the client, allowing reread operations to come from local memory instead of being retransmitted over the network.

The following diagram shows the IOR test results with three BeeGFS building blocks and a single client.

[beegfs design image17]



BeeGFS striping for these tests was set to a 1MB chunksize with eight targets per file.

Although write and initial read performance is higher using the default buffered mode, for workloads that reread the same data multiple times, a significant performance boost is seen with the native caching mode. This improved reread performance is important for workloads like deep learning that reread the same dataset multiple times across many epochs.

Metadata performance test

The Metadata performance tests used the MDTest tool (included as part of IOR) to measure the metadata performance of BeeGFS. The tests utilized OpenMPI to run parallel jobs across all ten client nodes.

The following parameters were used to run the benchmark test with the total number of processes scaled from 10 to 320 in step of 2x and with a file size of 4k.

```
mpirun -h 10xnodes -map-by node np $processes mdtest -e 4k -w 4k -i 3 -I  
16 -z 3 -b 8 -u
```

Metadata performance was measured first with one then two metadata + storage building blocks to show how performance scales by adding additional building blocks.

One BeeGFS metadata + storage building block

The following diagram shows the MDTest results with one BeeGFS metadata + storage building blocks.

[beegfs design image18]

Two BeeGFS metadata + storage building blocks

The following diagram shows the MDTest results with two BeeGFS metadata + storage building blocks.

[beegfs design image19]

Functional validation

As part of validating this architecture, NetApp executed several functional tests including the following:

- Failing a single client InfiniBand port by disabling the switch port.
- Failing a single server InfiniBand port by disabling the switch port.
- Triggering an immediate server power off using the BMC.
- Gracefully placing a node in standby and failing over service to another node.
- Gracefully placing a node back online and failing back services to the original node.
- Powering off one of the InfiniBand switches using the PDU. All tests were performed while stress testing was in progress with the `sysSessionChecksEnabled: false` parameter set on the BeeGFS clients. No errors or disruption to I/O was observed.



There is a known issue (see the [Changelog](#)) when BeeGFS client/server RDMA connections are disrupted unexpectedly, either through loss of the primary interface (as defined in `connInterfacesFile`) or a BeeGFS server failing; active client I/O can hang for up to ten minutes before resuming. This issue does not occur when BeeGFS nodes are gracefully placed in and out of standby for planned maintenance or if TCP is in use.

NVIDIA DGX A100 SuperPOD and BasePOD validation

NetApp validated a storage solution for NVIDIA's DGX A100 SuperPOD using a similar BeeGFS file system consisting of three building blocks with the metadata plus storage configuration profile applied. The qualification effort involved testing the solution described by this NVA with twenty DGX A100 GPU servers running a variety of storage, machine learning, and deep learning benchmarks. All storage certified for use in NVIDIA's DGX A100 SuperPOD is automatically certified for use in NVIDIA BasePOD architectures as well.

For more information, see [NVIDIA DGX SuperPOD with NetApp](#) and [NVIDIA DGX BasePOD](#).

Sizing guidelines

The BeeGFS solution includes recommendations for performance and capacity sizing that were based on verification tests.

The objective with a building-block architecture is to create a solution that is simple to size by adding multiple building blocks to meet the requirements for a particular BeeGFS system. Using the guidelines below, you can estimate the quantity and types of BeeGFS building blocks that are needed to meet the requirements of your environment.

Keep in mind that these estimates are best-case performance. Synthetic benchmarking applications are written and utilized to optimize the use of underlying file systems in ways that real-world applications might not.

Performance sizing

The following table provides recommended performance sizing.

Configuration profile	1MiB reads	1MiB writes
Metadata + storage	62GiBps	21GiBps

Configuration profile	1MiB reads	1MiB writes
Storage only	64GiBps	21GiBps

Metadata capacity sizing estimates are based on the "rule of thumb" that 500GB of capacity is sufficient for roughly 150 million files in BeeGFS. (For more information, see the BeeGFS documentation for [System Requirements](#).)

The use of features like access control lists and the number of directories and files per directory also affect how quickly metadata space is consumed. Storage capacity estimates do account for usable drive capacity along with RAID 6 and XFS overhead.

Capacity sizing for metadata + storage building blocks

The following table provides recommended capacity sizing for metadata plus storage building blocks.

Drive size (2+2 RAID 1) metadata volume groups	Metadata capacity (number of files)	Drive size (8+2 RAID 6) storage volume groups	Storage capacity (file content)
1.92TB	1,938,577,200	1.92TB	51.77TB
3.84TB	3,880,388,400	3.84TB	103.55TB
7.68TB	8,125,278,000	7.68TB	216.74TB
15.3TB	17,269,854,000	15.3TB	460.60TB



When sizing metadata plus storage building blocks, you can reduce costs by using smaller drives for metadata volume groups versus storage volume groups.

Capacity sizing for storage-only building blocks

The following table provides rule-of-thumb capacity sizing for storage-only building blocks.

Drive size (10+2 RAID 6) storage volume groups	Storage capacity (file content)
1.92TB	59.89TB
3.84TB	119.80TB
7.68TB	251.89TB
15.3TB	538.55TB



The performance and capacity overhead of including the management service in the base (first) building block are minimal, unless global file locking is enabled.

Performance tuning

The BeeGFS solution includes recommendations for performance tuning that were based on verification tests.

Although BeeGFS provides reasonable performance out of the box, NetApp has developed a set of

recommended tuning parameters to maximize performance. These parameters take into account the capabilities of the underlying E-Series block nodes and any special requirements needed to run BeeGFS in a shared-disk HA architecture.

Performance tuning for file nodes

The available tuning parameters that you can configure include the following:

1. System settings in the UEFI/BIOS of file nodes.

To maximize performance, we recommend configuring the system settings on the server model you use as your file nodes. You configure the system settings when you set up your file nodes by using either the system setup (UEFI/BIOS) or the Redfish APIs provided by the baseboard management controller (BMC).

The system settings vary depending on the server model you use as your file node. The settings must be manually configured based on the server model in use. To learn how to configure the system settings for the validated Lenovo SR665 file nodes, see [Tune file node system settings for performance](#).

2. Default settings for required configuration parameters.

The required configuration parameters affect how BeeGFS services are configured and how E-Series volumes (block devices) are formatted and mounted by Pacemaker. These required configuration parameters include the following:

- BeeGFS Service configuration parameters

You can override the default settings for the configuration parameters as needed. For the parameters that you can adjust for your specific workloads or use cases, see the [BeeGFS service configuration parameters](#).

- Volume formatting and mounting parameters are set to recommended defaults, and should only be adjusted for advanced use cases. The default values will do the following:
 - Optimize initial volume formatting based on the target type (such as management, metadata, or storage), along with the RAID configuration and segment size of the underlying volume.
 - Adjust how Pacemaker mounts each volume to ensure that changes are immediately flushed to E-series block nodes. This prevents data loss when file nodes fail with active writes in progress.

For the parameters that you can adjust for your specific workloads or use cases, see the [volume formatting and mounting configuration parameters](#).

3. System settings in the Linux OS installed on the file nodes.

You can override the default Linux OS system settings when you create the Ansible inventory in step 4 of [Create the Ansible inventory](#).

The default settings were used to validate the BeeGFS on NetApp solution, but you can change them to adjust for your specific workloads or use cases. Some examples of the Linux OS system settings that you can change include the following:

- I/O queues on E-Series block devices.

You can configure I/O queues on the E-Series block devices used as BeeGFS targets to:

- Adjust the scheduling algorithm based on the device type (NVMe, HDD, and so on).
- Increase the number of outstanding requests.
- Adjust request sizes.

- Optimize read ahead behavior.
- Virtual memory settings.

You can adjust virtual memory settings for optimal sustained streaming performance.

- CPU settings.

You can adjust the CPU frequency governor and other CPU configurations for maximum performance.

- Read request size.

You can increase the maximum read request size for Mellanox HCAs.

Performance tuning for block nodes

Based on the configuration profiles applied to a particular BeeGFS building block, the volume groups configured on the block nodes change slightly. For example, with a 24-drive EF600 block node:

- For the single base building block, including BeeGFS management, metadata, and storage services:
 - 1x 2+2 RAID 10 volume group for BeeGFS management and metadata services
 - 2x 8+2 RAID 6 volume groups for BeeGFS storage services
- For a BeeGFS metadata + storage building block:
 - 1x 2+2 RAID 10 volume group for BeeGFS metadata services
 - 2x 8+2 RAID 6 volume groups for BeeGFS storage services
- For BeeGFS storage only building block:
 - 2x 10+2 RAID 6 volume groups for BeeGFS storage services



As BeeGFS needs significantly less storage space for management and metadata versus storage, one option is to use smaller drives for the RAID 10 volume groups. Smaller drives should be populated in the outermost drive slots. For more information, see the [deployment instructions](#).

These are all configured by the Ansible-based deployment, along with several other settings generally recommended to optimize performance/behavior including:

- Adjusting the global cache block size to 32KiB and adjusting demand-based cache flushing to 80%.
- Disabling autoload balancing (ensuring controller volume assignments stay as intended).
- Enabling read caching and disabling read-ahead caching.
- Enabling write caching with mirroring and requiring battery backup, so that caches persist through failure of a block node controller.
- Specifying the order that drives are assigned to volume groups, balancing I/O across available drive channels.

High capacity building block

The standard BeeGFS solution design is built with high performance workloads in mind. Customers looking for a high capacity use cases should observe the variations in design

and performance characteristics outlined here.

Hardware and software configuration

Hardware and software configuration for the high capacity building block is standard except that the EF600 controllers should be replaced with a EF300 controllers with an option to attach between 1 and 7 IOM expansion trays with 60 drives each for each storage array, totaling 2 to 14 expansions trays per building block.

Customers deploying a high capacity building block design are likely to use only the base building block style configuration consisting of BeeGFS management, metadata, and storage services for each node. For cost efficiency, high capacity storage nodes should provision metadata volumes on the NVMe drives in the EF300 controller enclosure and should provision storage volumes to the NL-SAS drives in the expansion trays.

[high capacity rack diagram]

Sizing guidelines

These sizing guidelines assume high capacity building blocks are configured with one 2+2 NVMe SSD volume group for metadata in the base EF300 enclosure and 6x 8+2 NL-SAS volume groups per IOM expansion tray for storage.

Drive size (capacity HDDs)	Capacity per BB (1 Tray)	Capacity per BB (2 Trays)	Capacity per BB (3 Trays)	Capacity per BB (4 Trays)
4TB	439TB	878 TB	1317 TB	1756 TB
8 TB	878 TB	1756 TB	2634 TB	3512 TB
10 TB	1097 TB	2195 TB	3292 TB	4390 TB
12 TB	1317 TB	2634 TB	3951 TB	5268 TB
16 TB	1756 TB	3512 TB	5268 TB	7024 TB
18 TB	1975 TB	3951 TB	5927 TB	7902 TB

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—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.