

Technical Report

Oracle on Azure Deployment Best Practice Guide

Using Azure NetApp Files

Jeffrey Steiner, Jim Holl, Spencer Sells, Ebin Kadavy, NetApp
Geert van Teylingen, Andrew Chen, Juergen Thomas, Microsoft
May 2019 | TR-478/0

In partnership with



Abstract

This document provides deployment guidance and best practices for using Azure NetApp® Files for Oracle deployments. It describes various configuration and performance considerations for Oracle in Azure using Azure NetApp Files.

TABLE OF CONTENTS

1	Introduction	5
1.1	Oracle on Microsoft Azure	5
1.2	NetApp Values and Solutions on Microsoft Azure	5
1.3	Azure NetApp Files Key Value Proposition	8
1.4	Understanding Azure NetApp Files Performance and Capacity Tiering	9
2	Provisioning Azure NetApp Files Storage for Oracle Database	11
3	Network Configuration	14
3.1	TCP Parameters	14
4	NFS Configuration	15
4.1	Operating Systems	15
4.2	NFS Versions	15
4.3	TCP Slot Tables	15
4.4	actimeo=0, noac	15
4.5	Direct NFS	16
4.6	Direct NFS and Host File System Access	16
4.7	Mount Options	16
4.8	Nosharecache	16
4.9	Automatic Storage Management	17
5	Oracle Configuration	17
5.1	filesystemio_options	17
5.2	db_file_multiblock_read_count	18
5.3	Redo Block Size	18
6	Performance Optimization and Benchmarking	19
6.1	Oracle Automatic Workload Repository and Benchmarking	19
6.2	Oracle AWR and Troubleshooting	19
6.3	Scaling and Performance	20
	Appendix A: Stale NFS Locks	20
	Where to Find Additional Information	21
	Version History	21

LIST OF TABLES

Table 1)	Linux NFSv3 mount options—single instance	16
----------	---	----

LIST OF FIGURES

Figure 1) Cloud Volumes ONTAP.	6
Figure 2) Azure NetApp Files.	7
Figure 3) NetApp ONTAP offerings.	8
Figure 4) Azure NetApp Files - capacity versus performance.	9
Figure 5) Azure NetApp Files service levels.	10
Figure 6) Volumes added in capacity pool.	13
Figure 7) Shared volume layout.	14
Figure 8) Dedicated volume layout.	14

1 Introduction

Today, many customers are using Microsoft Azure to accelerate their Oracle deployments to reduce costs and provide increased agility for their business processes. These issues are of paramount importance for IT leaders who have a Cloud First strategy. Moreover, integrating Oracle with the Azure suite of platform as a service (PaaS) services such as Azure Data Factory, Azure Internet of Things (IoT) Hub, and Azure Machine Learning can create business value and support digitalization.

1.1 Oracle on Microsoft Azure

More large enterprises choose Azure as the cloud platform of choice for their enterprise applications including Oracle. Many customers embraced the DevOps paradigm by first moving their test and development systems. Recently, however, customers are choosing to migrate their complete Oracle infrastructures, including production, into the cloud.

Azure's vast compute offerings range from small to large virtual machines for the most demanding database workloads. Microsoft introduced the Azure M-Series virtual machines in 2018 with up to 6TB of memory and, recently, 12TB virtual machines were announced. These colossal virtual machines are targeted at specific workloads such as high-performance computing (HPC), SAP HANA, and Oracle.

1.2 NetApp Values and Solutions on Microsoft Azure

For many customers, the NetApp® storage and data management solutions, based on the ONTAP® software, have been the foundation for their enterprise workloads such as Oracle. For more than 20 years, the ONTAP system and NFS services have been used in many of the largest Oracle deployments to enable a secure and stable operation. This solution simplifies data management, accelerates projects, and reduces risk.

As a global Oracle technology partner, NetApp has a long history of providing excellent solutions and products with a deep integration into Oracle environments. This partnership enables customers to use NetApp Snapshot™ technology for fast, storage efficient, and reliable backup and recovery. It also provides fast and storage efficient cloning for quicker time to market while improving quality. The fully supported products help Oracle customers to automate a comprehensive backup and disaster recovery strategy while also considering other important workflows. You can focus on the complete Oracle application lifecycle management by using Snapshot-based cloning operations.

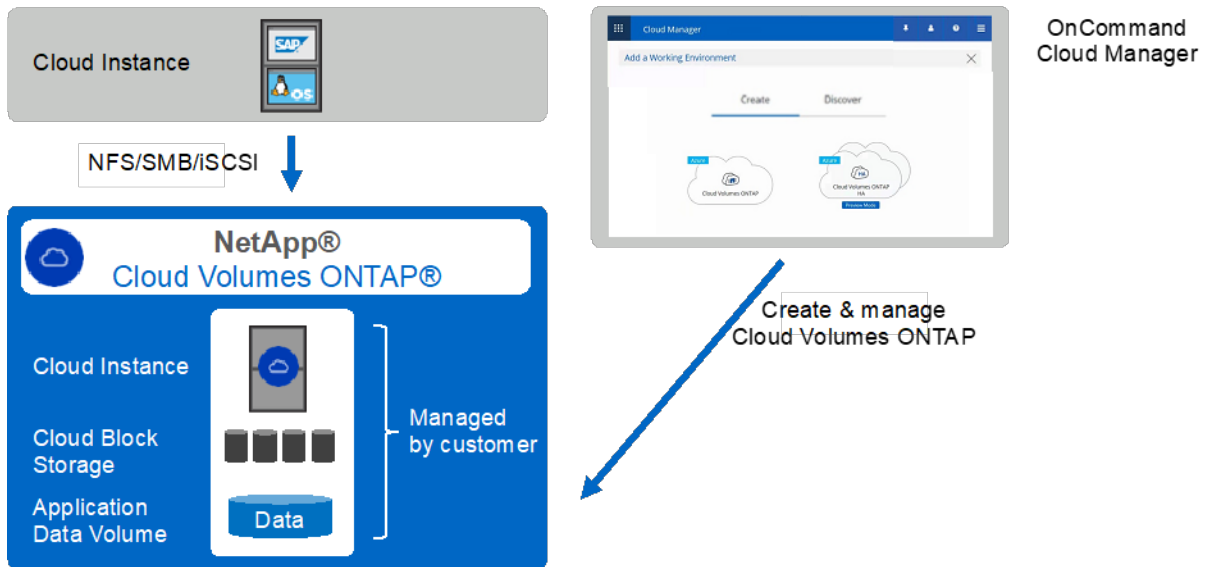
Many Oracle customers who want to move their systems to the cloud still want to use the NetApp storage benefits for their Oracle projects and operations. Customers do not want to give up on the performance, reliability, and enterprise data management capabilities when moving enterprise file-based workloads to the cloud. Not every cloud can offer a highly available, enterprise-grade, fast, reliable, feature-rich, but simple to manage shared file service based around NFS, which is required for those Oracle environments.

On Azure, customers can now benefit from two distinct ONTAP-based offerings to build their Oracle systems on. Although the following sections provide a brief overview of both solutions, Cloud Volumes ONTAP and Azure NetApp Files, the remainder of this document will focus on Azure NetApp Files only. See [Oracle Databases on ONTAP Cloud with Microsoft Azure](#) for detailed description of Oracle on Microsoft Azure Using Cloud Volumes ONTAP.

Cloud Volumes ONTAP on Azure

[Cloud Volumes ONTAP](#) extends the trusted enterprise data management capabilities of ONTAP to leading cloud platforms such as Microsoft Azure. In Azure, it provides NFS/iSCSI-based services to hosted Oracle workloads. By using the underlying Azure storage and compute resources Cloud Volumes ONTAP adds storage efficiency features such as thin provisioning, deduplication, compression, and now tiered storage to Azure Blob Storage as well.

Figure 1) Cloud Volumes ONTAP.



Cloud Volumes ONTAP is NetApp's proven data management software running in a cloud instance using Cloud Storage. For the initial provisioning of Cloud Volume ONTAP and a cloud-like simplified but efficient management of all resources, customers need to install a small appliance known as OnCommand® Cloud Manager as shown in Figure 1. Cloud Manager can deploy and manage multiple Cloud Volumes ONTAP instances either as a single node or a highly available dual node configuration. Cloud Manager can manage Cloud Volumes ONTAP in Azure, or you can also use it to manage your on-premises ONTAP systems and Cloud Volumes ONTAP instances at other data centers or even other cloud providers. When provisioning Cloud Volumes ONTAP, you can select from different system classes and license types. This choice defines the maximum storage capacity and the possible throughput and performance.

Customers can provision their 'data volumes' and shared files to the Cloud Instance to run their Oracle application and databases. Customers new to NetApp can use Cloud Manager for this provisioning while customers used to NetApp can use all the NetApp tools and workflows they are using in their on-premises datacenters.

Azure NetApp Files

For the first time, Microsoft will deliver an Azure native, first-party service for enterprise NFS/SMB file services based on NetApp ONTAP technology. This new development is driven by a strategic partnership between NetApp and Microsoft and further extends the reach of NetApp's world-class data services to Azure.

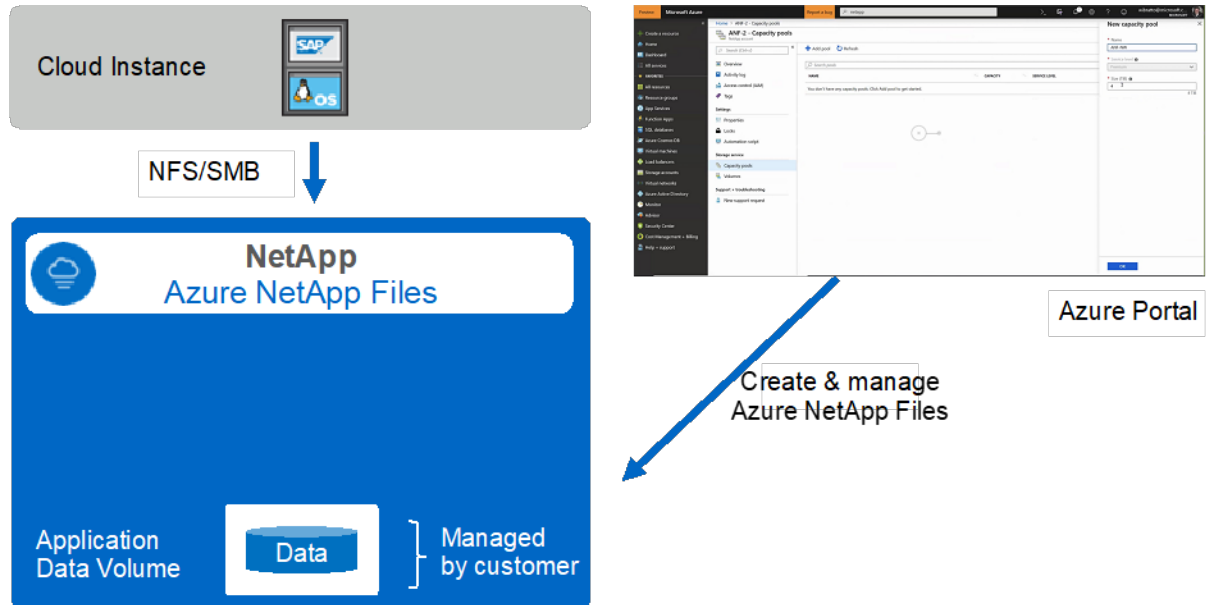
This Azure cloud-native data service delivers high performance, reliability, and enterprise data management and security for customers who are moving enterprise NFS/SMB workloads to Azure.

Azure NetApp Files is integrated into the Azure DCs, ARM architecture and portal, and you can use the same comfortable UI and APIs for creating and managing shared files as with any other Azure object (see Figure 2). Azure NetApp Files provides NetApp's enterprise class storage and delivers many of the data management capabilities such as easy creation and resizing of volumes, adapting capacity and

performance without downtime, creating space and time efficient storage Snapshot copies and clones¹ in seconds that are valuable to use to optimize SAP operations.

As a comparison to Cloud Volumes ONTAP, Azure NetApp Files is built on NetApp's proven ONTAP storage hardware hosted in Azure datacenters and operated and maintained by Microsoft directly. This results in high storage performance with low latency I/O.

Figure 2) Azure NetApp Files.

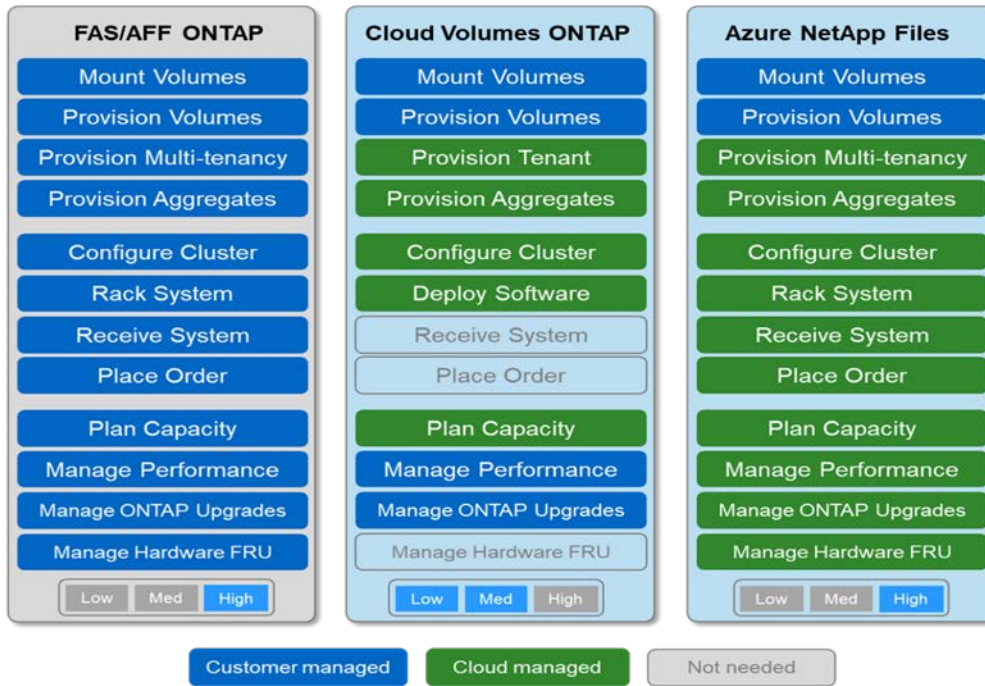


Comparison

Figure 3 shows the difference in level of management for the various options for ONTAP-based systems, ranging from On-prem (FAS/AFF ONTAP) to infrastructure as a service (IaaS) (Cloud Volumes ONTAP) to PaaS (Azure NetApp Files). Clearly, Azure NetApp Files provides the optimal “on-demand cloud service” experience while providing “on-premises-like” performance. If you require “cloud-based storage” with a higher level of management control and require access to all ONTAP features, Cloud Volumes ONTAP might be a better option.

¹ While Azure NetApp Files offers a cloning feature today, its use with Oracle volumes is not recommended. In the future, it will be.

Figure 3) NetApp ONTAP offerings.



1.3 Azure NetApp Files Key Value Proposition

Azure NetApp Files was designed to meet the core requirements of running high-performance workloads like databases in the cloud, and provides:

- Performance tiers that reflect the real-world range of IOPS demands
- Low latency
- High availability
- High durability
- Manageability at scale
- Fast and efficient backup and recovery and cloning²

These capabilities are possible because Azure NetApp Files is based on physical all-flash NetApp® ONTAP® systems running within Azure data center environment. The result is an ideal database storage technology that can be provisioned and consumed just like other native cloud storage options.

This document addresses the requirements for operating an Oracle database on Azure NetApp Files in two ways. First, when a clear best practice exists, it is called out specifically. Second, this document reviews the many design considerations that must be addressed by architects of Oracle storage solutions based on their specific business requirements.

² While Azure NetApp Files offers a cloning feature today, its use with Oracle volumes is not recommended. In the future, it will be.

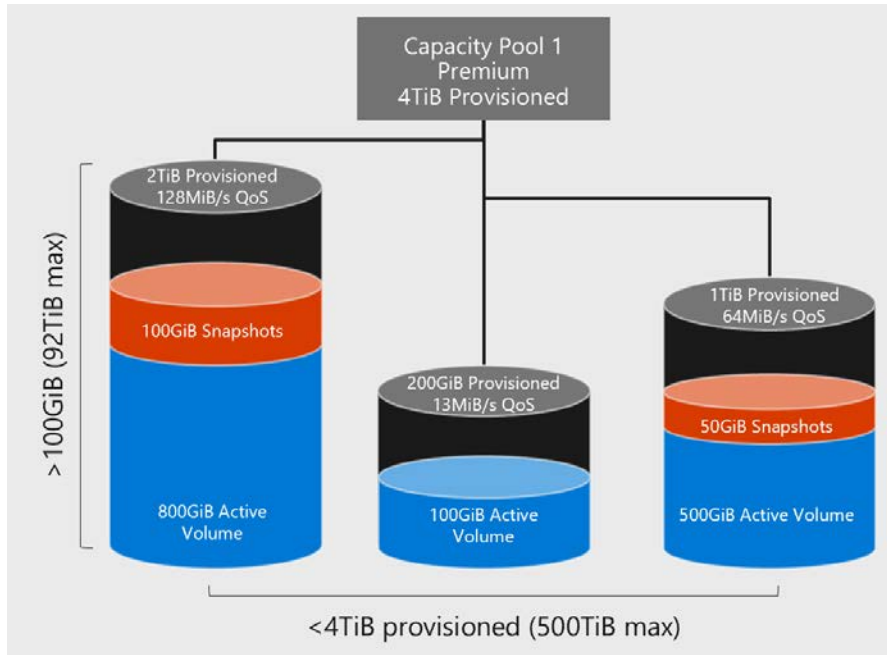
1.4 Understanding Azure NetApp Files Performance and Capacity Tiering

To understand how to optimize Azure NetApp Files' capacity in relation to performance and costs a closer look needs to be taken on how Azure NetApp Files is provisioned.

Azure NetApp Files volumes are allocated from a capacity pool that you must provision in your Azure NetApp Files storage account. Each capacity pool is assigned the following:

- A service level that defines the overall performance capability,
- The initially provisioned storage capacity for that capacity pool.

Figure 4) Azure NetApp Files - capacity versus performance.




The performance of a volume is based on the capacity pool service level, with the number of TiBs provisioned for that volume. Customers can dynamically grow and shrink the volumes and pool capacity, to manage performance, capacity or both. Billing occurs for the provisioned **capacity pool**, on an hourly base.

Each of the service levels available³ has an associated cost per provisioned capacity and includes a Quality of Service (QoS) level that defines the overall maximum throughput per provisioned space. For example, a 10TiB provisioned single capacity pool with premium service level provides an overall available throughput for all volumes in this capacity pool of 10 x 64MBps, so 640MBps with 40,000 (16K) IOPs or 80,000 (8K) IOPS.

Within a capacity pool, each volume is provisioned with a specific quota in the range between 100GB up to the maximum volume size. This quota defines the maximum throughput or IOPS for this volume.

³ At or around time of general availability of the Azure NetApp Files service, all three tiers are expected to be available.

Figure 5) Azure NetApp Files service levels.



More details can be found here:
[Service levels for Azure NetApp Files.](#)

The following chapters will further dive into the specific considerations and best practices around networking, NFS protocol, Oracle configurations, performance, and sizing.

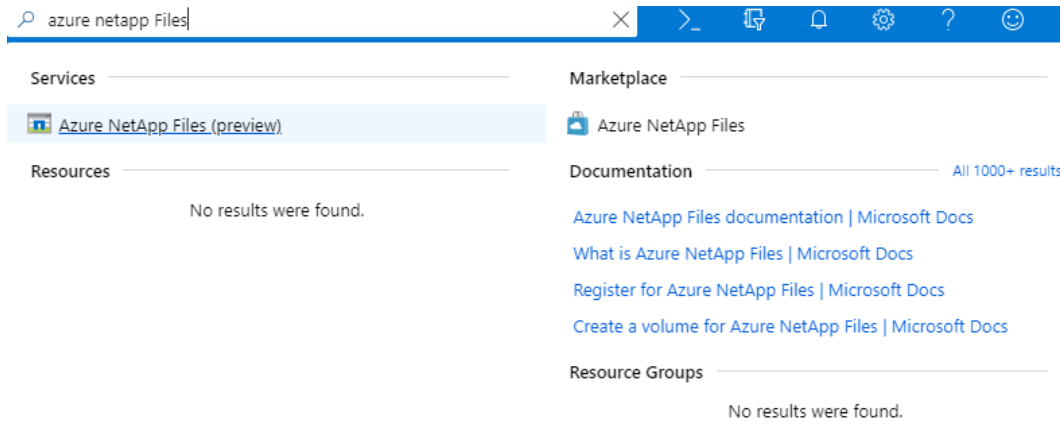
2 Provisioning Azure NetApp Files Storage for Oracle Database

Provisioning Azure NetApp Files storage for Oracle database covers two parts:

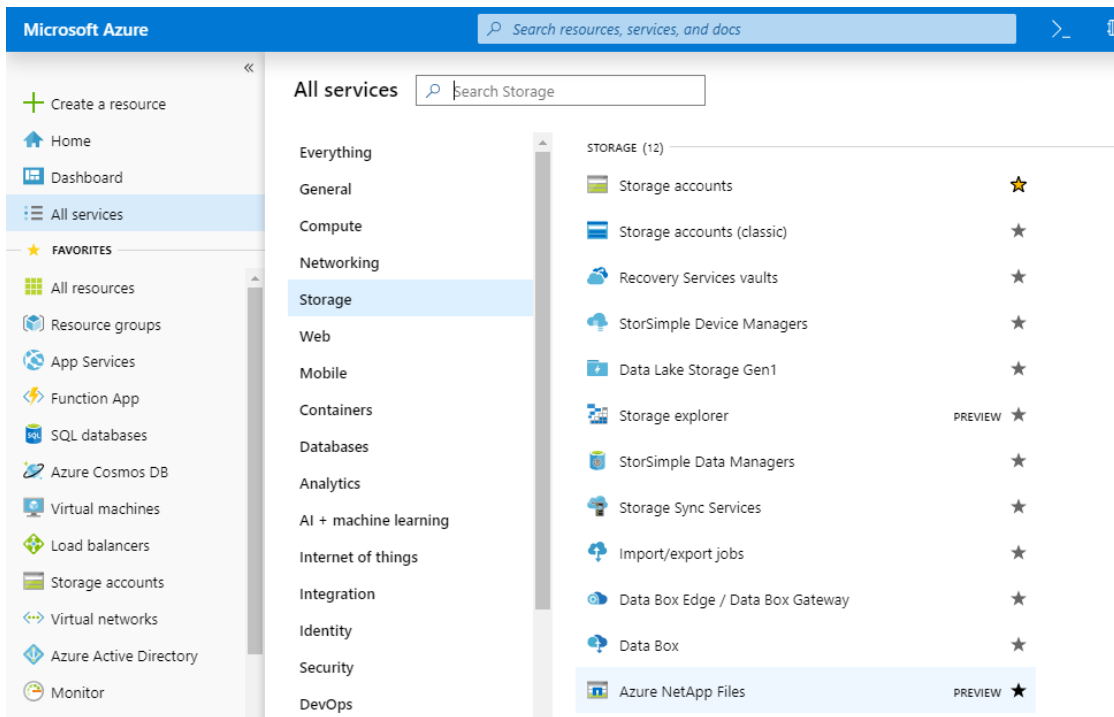
1. Creating a NetApp account in Microsoft Azure portal.
2. Setting up capacity pools and volumes

Creating a NetApp account enables you to set up a capacity pool and then create a volume. You can log in to Azure portal and use any of the approaches to open a NetApp account.

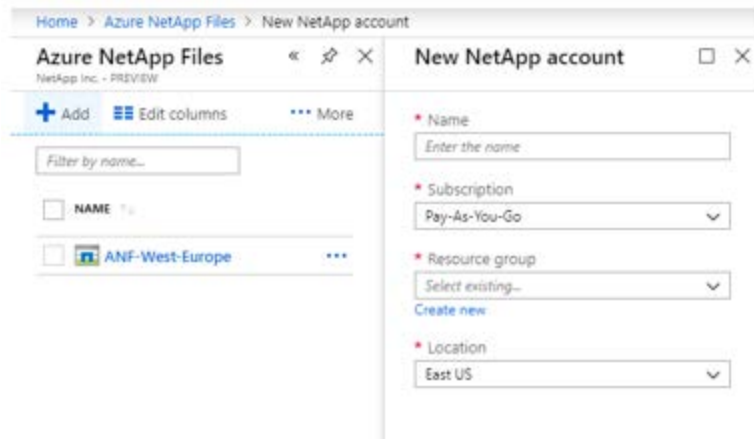
1. Search for **Azure NetApp Files** in the Azure portal search box.



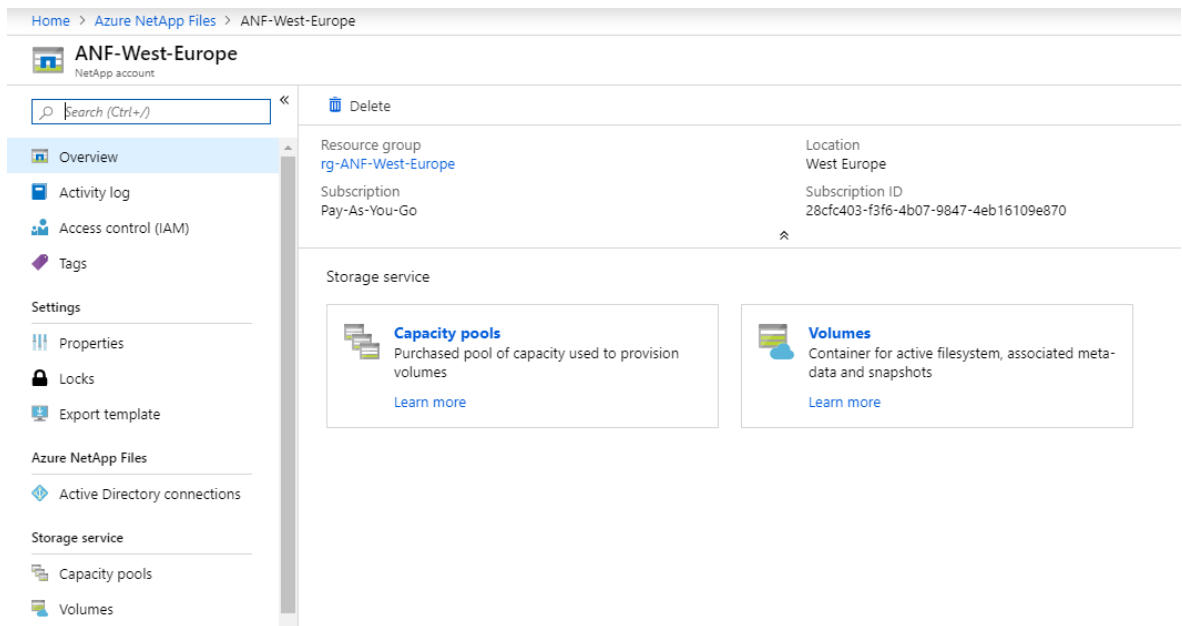
2. Click **All services** in the navigation, and then filter to Azure NetApp Files.



3. After launching the Azure NetApp files wizard, you can create an account by pressing the Add button



4. After the NetApp account is created, it should look similar to the following snippet.



5. The last step is to add a capacity pool and tag the appropriate service levels followed by creating the required volumes.

New capacity pool ✕

* Name

* Service level ⓘ

* Size (TiB) ⓘ

4 TiB

Here is a snippet of volumes that are added in a given capacity pool. The volume layouts for any given workload should be based on best practices discussed in the next session.

Figure 6) Volumes added in capacity pool.

Home > Azure NetApp Files > ANF-West-Europe - Capacity pools > ANF-WE-01 (ANF-West-Europe/ANF-WE-01) - Volumes

ANF-WE-01 (ANF-West-Europe/ANF-WE-01) - Volumes

Capacity pool

Search (Ctrl+/)

- Overview
- Activity log
- Access control (IAM)
- Tags
- Settings
 - Properties
 - Locks
 - Export template
- Storage service
 - Volumes**
- Monitoring
 - Metrics

+ Add volume Refresh

Search volumes

NAME	QUOTA	PROTO...
D01-data-mnt00001	500 GiB	NFS
D01-log-mnt00001	500 GiB	NFS
D01-shared	496 GiB	NFS
P01-data-mnt00001	500 GiB	NFS
P01-log-mnt00001	200 GiB	NFS
P01-shared	200 GiB	NFS
Q01-data-mnt00001	200 GiB	NFS
Q01-data-mnt00001-...	200 GiB	NFS
Q01-log-mnt00001	500 GiB	NFS

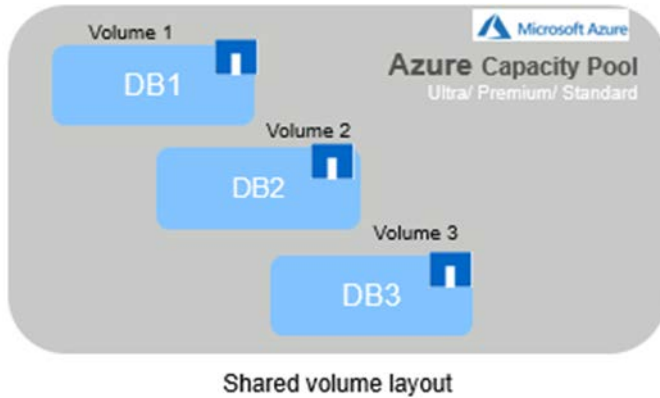
Volume Layout Best Practices

The volume layout for Oracle database can be designed based on the criticality of the workload. You can choose between any of these approaches:

- Dedicated volume layout
- Shared volume layout

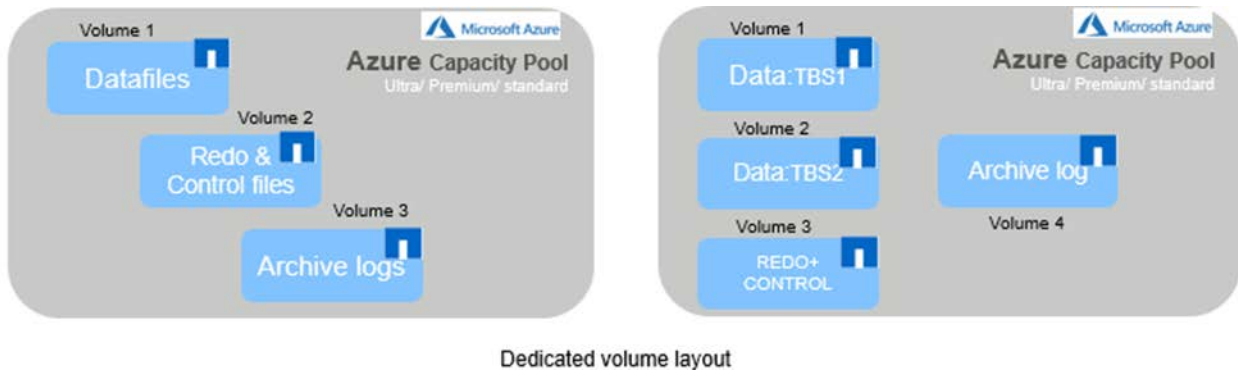
The simplest possible layout is the shared volume layout where all the data files, redo logs, controlfile, and archive logs are kept in the same volume (As shown in below snippet).

Figure 7) Shared volume layout.



This layout might not be a good fit for workloads that needs best performance. For critical workloads that demands higher IOPS and lower latency can follow dedicated volume layout by spreading the data across different volumes and isolating the other non-data files on different volumes.

Figure 8) Dedicated volume layout.



3 Network Configuration

The TCP/IP settings required for Oracle database software installation are sufficient to provide good performance for Azure NetApp Files. There are a few exceptions.

3.1 TCP Parameters

Three settings are frequently misconfigured: TCP timestamps, selective acknowledgment (SACK), and TCP window scaling. Many out-of-date documents on the internet recommend disabling one or more of these parameters to improve performance. There was some merit to this recommendation many years

ago when CPU capabilities were much lower and there was a benefit to reducing the overhead on TCP processing whenever possible.

However, with modern operating systems (OSs), disabling any of these TCP features usually results in no detectable benefit or might result in performance damage. Performance damage is especially likely in virtualized networking environments because these features are required for efficient handling of packet loss and changes in network quality.

The following settings are recommended:

- Enable TCP timestamps, SACK, and TCP window scaling on the host. Check the Host Operating system's network performance tuning guide on how to enable these parameters.

Note: NetApp does not recommend enabling TCP timestamps for SAP on Oracle database deployments.

4 NFS Configuration

4.1 Operating Systems

The most common database platforms in cloud environments are Linux and Microsoft Windows. The Linux operating system includes native NFS capabilities and is the most common operating system.

For customers that prefer Windows, Oracle offers the direct NFS (dNFS) client, natively integrated into Oracle. See section, 3.5, "Direct NFS" for additional information. The dNFS feature offers a path to the management benefits of NFS, including the ability to view files across environments, dynamically resize volumes, and use a less expensive IP protocol. See the official Oracle documentation for information about installing and configuring a database on Microsoft Windows using dNFS. No special best practices exist.

4.2 NFS Versions

Oracle has supported NFSv3 for over 20 years, and NFSv4 is supported with Oracle 12.1.0.2 and higher.

4.3 TCP Slot Tables

TCP slot tables are the NFS equivalent of host bus adapter (HBA) queue depth. These tables control the number of NFS operations that can be outstanding at any one time. The default value is usually 16, which is far too low for optimum performance. The opposite problem occurs on newer Linux kernels, which can automatically increase the TCP slot table limit to a level that saturates the NFS server with requests.

For optimum performance and to prevent performance problems, adjust the kernel parameters that control the TCP slot tables.

Run `sysctl -a | grep tcp.*.slot_table`, and check the following parameters:

```
# sysctl -a | grep tcp.*.slot_table
sunrpc.tcp_max_slot_table_entries = 128
sunrpc.tcp_slot_table_entries = 128
```

All Linux systems should include `sunrpc.tcp_slot_table_entries`, but only some will include `sunrpc.tcp_max_slot_table_entries`. They should all be set to 128.

4.4 `actimeo=0, noac`

The presence of the following mount options in `ORACLE_HOME` causes host caching to be disabled, which badly damages performance for many workloads, especially installation and patching:

```
actimeo=0, noac
```

These mount options are unavoidable with Oracle RAC, but as RAC is not supported in Azure currently, these options are unnecessary for Azure NetApp Files.

4.5 Direct NFS

Oracle's dNFS client is critical for optimum performance in a cloud environment.

The dNFS client provides multiple benefits. It is designed to bypass the host NFS client and perform NFS file operations directly on an NFS server. Enabling it only requires changing the Oracle Disk Manager (odm) library. Instructions for this process are provided in the Oracle documentation.

Using dNFS results in a general improvement in I/O performance and decreases the load on the host and the storage system because I/O is performed in the most efficient way possible.

When dNFS is used, it is critical that all patches described in Oracle Doc 1495104.1 are installed. If a patch cannot be installed, the environment must be evaluated to make sure that the bugs described in that document do not cause problems. Sometimes, an inability to install the required patches prevents the use of dNFS.

Caution

- Before using dNFS, verify that the patches described in Oracle Doc 1495104.1 are installed.
- Starting with Oracle 12c, DNFS includes support for NFSv3, NFSv4, and NFSv4.1. NetApp Support policies cover v3 and v4 for all clients, but at the time of writing NFSv4.1 is not supported for use with Oracle dNFS.

4.6 Direct NFS and Host File System Access

Using dNFS can occasionally cause problems for applications or user activities that rely on the visible file systems mounted on the host. These problems occur because the DNFS client accesses the file system out of band from the host OS. The dNFS client can create, delete, and modify files without the knowledge of the OS. There might be a lag at times before a file change becomes visible to OS users.

4.7 Mount Options

Table 3 lists Linux NFSv3 mount options.

Table 1) Linux NFSv3 mount options—single instance.

File Type	Mount Options
Control files Data files Redo logs	<code>rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsiz=65536</code>
ORACLE_HOME ORACLE_BASE	<code>rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsiz=65536</code>

Generally, non-database files should be mounted with the same options used for single-instance data files, although specific applications might have different requirements. Avoid the mount options `noac` and `actimeo=0` if possible because these options disable filesystem-level read ahead and buffering. This can cause severe performance problems for processes such as extract, transform, load (ETL).

4.8 Nosharecache

One more mount option, called `nosharecache`, is required in the following circumstances:

- dNFS is enabled

- A source volume is mounted more than once on a single server
- The volume mounts are nested

This `nosharecache` configuration is seen primarily in environments supporting SAP applications.

In general, a single volume is only mounted once but this practice is not mandatory. For example, an Azure NetApp Files volume might have a path at `/vol1234/base` and `/vol1234/home`. If `/vol1234/base` is mounted at `/oracle` and `/vol1234/home` is mounted at `/oracle/home`, the result is nested NFS mounts that originate on the same source volume.

The OS can detect the fact that `/oracle` and `/oracle/home` reside on the same volume, which is the same source file system. The OS then uses the same device handle for accessing the data. Doing so improves the use of OS caching and certain other operations, but it interferes with dNFS. If dNFS must access a file on `/oracle/home`, it might erroneously attempt to use the wrong path to the data. The result is a failed I/O operation. In these configurations, add the `nosharecache` mount option to any NFS file system that shares a source volume with another NFS file system on that host. Doing so forces the Linux OS to allocate an independent device handle for that file system.

4.9 Automatic Storage Management

Automatic Storage Management (ASM) is supported with NFS. ASM uses the space inside one or more files and presents it to the database as a single pool of storage. Normally, these files are LUN devices at paths such as `/dev/sdab` or `/dev/mapper/oracleasm/disk0`, but they can also be paths to an NFS file system, including Azure NetApp Files volumes.

Re-virtualizing files on an NFS file system does not offer any significant benefits over placing Oracle files directly on NFS file systems, but if management practices make ASM desirable, then ASM over NFS might be used. Commands work as usual, and the `odm` library used by ASM also includes the ability to use dNFS. Performance should be nearly identical.

5 Oracle Configuration

5.1 `filesystemio_options`

The Oracle initialization parameter `filesystemio_options` controls the use of asynchronous and direct I/O. Contrary to common belief, asynchronous and direct I/O are not mutually exclusive. NetApp has observed that this parameter is frequently misconfigured in customer environments, and this misconfiguration is directly responsible for many performance problems.

Asynchronous I/O means that Oracle I/O operations can be parallelized. Before the availability of asynchronous I/O on various operating systems, users configured numerous dbwriter processes and changed the server process configuration. With asynchronous I/O, the OS itself performs I/O on behalf of the database software in a highly efficient and parallel manner. This process does not place data at risk, and critical operations, such as Oracle redo logging, are still performed synchronously.

Direct I/O bypasses the OS buffer cache. I/O on a UNIX system ordinarily flows through the OS buffer cache. This is useful for applications that do not maintain an internal cache, but Oracle has its own buffer cache within the System Global Area (SGA). In almost all cases, it is better to enable direct I/O and allocate server RAM to the SGA rather than to rely on the OS buffer cache. The Oracle SGA uses the memory more efficiently. In addition, when I/O flows through the OS buffer, it is subject to extra processing, which increases latencies. The increased latencies are especially noticeable with heavy write I/O when low latency is a critical requirement.

The options for `filesystemio_options` are:

- `async`. Oracle submits I/O requests to the OS for processing. This process allows Oracle to perform other work rather than waiting for I/O completion and thus increases I/O parallelization.

- `directio`. Oracle performs I/O directly against physical files rather than routing I/O through the host OS cache.
- `none`. Oracle uses synchronous and buffered I/O. In this configuration, the choice between shared and dedicated server processes and the number of dbwriters are more important.
- `setall`. Oracle uses both asynchronous and direct I/O.

In almost all cases, the use of `setall` is optimal, but consider the following issues:

- If a database has been using buffered I/O, a switch to direct I/O might also warrant a change in the SGA size. Disabling buffered I/O eliminates the performance benefit that the host OS cache provides for the database. Adding RAM back to the SGA repairs this problem. The net result should be an improvement in I/O performance.
- Although it is almost always better to use RAM for the Oracle SGA than for OS buffer caching, it might be impossible to determine the best value. For example, it might be preferable to use buffered I/O with small SGA sizes on a database server with many intermittently active Oracle instances. This arrangement allows the flexible use of the remaining free RAM on the OS by all running database instances. This is a highly unusual situation, but it has been observed at some customer sites.

Note: The `filesystemio_options` parameter has no effect in DNFS and ASM environments. The use of dNFS or ASM automatically results in the use of both asynchronous and direct I/O.

NetApp recommends the following:

- Set `filesystemio_options` to `setall`, but sometimes the loss of the host buffer cache might require an increase in the Oracle SGA.

5.2 db_file_multiblock_read_count

The `db_file_multiblock_read_count` parameter controls the maximum number of Oracle database blocks that Oracle reads as a single operation during sequential I/O. This parameter does not, however, affect the number of blocks that Oracle reads during all read operations, nor does it affect random I/O. Only sequential I/O is affected.

Oracle recommends that the user leave this parameter unset. Doing so allows the database software to automatically set the optimum value. This generally means that this parameter is set to a value that yields an I/O size of 1MB. For example, a 1MB read of 8KB blocks would require 128 blocks to be read, and the default value for this parameter would therefore be 128.

Most database performance problems observed by NetApp at customer sites involve an incorrect setting for this parameter. There were valid reasons to change this value with Oracle versions 8 and 9. As a result, the parameter might be unknowingly present in `init.ora` files because the database was upgraded in place to Oracle 10 and later. A legacy setting of 8 or 16, compared to a default value of 128, significantly damages sequential I/O performance.

NetApp recommends the following:

- The `db_file_multiblock_read_count` parameter should not be present in the `init.ora` file. NetApp has never encountered a situation in which changing this parameter improved performance, but there are many cases in which it damaged sequential I/O throughput.

5.3 Redo Block Size

Oracle supports either a 512-byte or 4KB redo block size. The default is 512 bytes. The best option is expected to be 512 bytes because this size minimizes the amount of data written during redo operations. However, it is possible that the 4KB size could offer a performance benefit at high logging rates. For example, a single database with 50MBps of redo logging might be more efficient if the redo block size is larger. A storage system supporting many databases with a large total amount of redo logging might

benefit from a 4KB redo block size. This is because this setting would eliminate inefficient partial I/O processing when only a part of a 4KB block must be updated.

It is not correct that all I/O operations are performed in single units of the redo log block size. At high logging rates, the database generally performs large I/O operations composed of multiple redo blocks. The actual size of those redo blocks does not generally affect the efficiency of logging.

NetApp recommends the following:

- Only change the default block size for cause, such as a documented requirement for a particular application or because of a recommendation made by NetApp or Oracle customer support.

6 Performance Optimization and Benchmarking

Accurate testing of database storage performance is a complicated subject. It requires not just an understanding of IOPS and throughput, but also understand of the following:

- Differences between foreground and background I/O operations
- Affect of latency upon the database
- Multiple OS and network settings that affect storage performance
- Non-storage database tasks

Note: There is a point where optimizing storage performance yields no useful benefits because storage performance is no longer a limiting factor for performance.

Any evaluation of database performance in a cloud environment must focus on real-world needs. There will nearly always be significant differences at maximum loading levels. The cloud environment will almost inevitably be slower at the top end, especially since nearly all cloud resources have multiple types of quality of service (QoS) controls, both visible and invisible.

Rather than testing maximums, identify the actual requirements and see whether they can be met in the Cloud.

6.1 Oracle Automatic Workload Repository and Benchmarking

The gold standard for Oracle performance comparison is an Oracle Automatic Workload Repository (AWR) report.

There are multiple types of AWR reports. From a storage point of view, a report generated by running the `awrrpt.sql` command is the most comprehensive and valuable because it targets a specific database instance and includes some detailed histograms that break down storage I/O events based on latency.

Comparing two performance arrays ideally involves running the same workload on each array and producing an AWR report that precisely targets the workload. If there is a long-running workload, you can use a single AWR report with an elapsed time that encompasses the start and stop time, but it is preferable to break out the AWR data as multiple reports. For example, if a batch job ran from midnight to 6 a.m., create a series of one-hour AWR reports from midnight–1 a.m., 1 a.m.–2 a.m., and so on.

In other cases, a short query should be optimized. The best option is an AWR report based on an AWR snapshot copy that is created when the query begins and a second AWR snapshot copy created when the query ends. The database server should be otherwise quiet to minimize the background activity that would obscure the activity of the query under analysis.

Note: Where AWR reports are not available, Oracle stats pack reports are a good alternative. They contain most of the same I/O statistics as an AWR report.

6.2 Oracle AWR and Troubleshooting

An AWR report is also the most important tool for analyzing a performance problem.

As with benchmarking, performance troubleshooting requires that you precisely measure a particular workload. When possible, provide AWR data when reporting a performance problem to the NetApp or Microsoft support center or when working with a NetApp, Microsoft, or partner account team about a new solution.

When providing AWR data, consider the following requirements:

- Run the `awrrpt.sql` command to generate the report. The output can be either text or HTML.
- Target the specific time the problem existed. The maximum acceptable elapsed time of an AWR report is generally one hour. If a problem persists for multiple hours or involves a multi-hour operation such as a batch job, provide multiple one-hour AWR reports that cover the entire period to be analyzed.
- If possible, adjust the AWR snapshot interval to 15 minutes. This setting allows a more detailed analysis to be performed. This also requires more executions of `awrrpt.sql` to provide a report for each 15-minute interval.
- If the problem is a very short running query, provide an AWR report based on an AWR snapshot created when the operation begins and a second AWR snapshot created when the operation ends. The database server should be otherwise quiet to minimize the background activity that would obscure the activity of the operation under analysis.
- If a performance problem is reported at certain times but not others, provide more AWR data that demonstrates good performance for comparison.

6.3 Scaling and Performance

One of the innovative features of Azure NetApp Files is the capability to scale capacity and performance dynamically. Currently the performance of any given volume can be increased and decreased dynamically by resizing the volume quota, thus resizing the performance limits of the volume. This enables short-term (burst) performance increase of the volume and the databases hosted on it. This is especially useful if high-performance is only needed for short periods of time, for example, during specific data warehouse or batch workloads.

Another way to increase performance continuously would be to distribute the database layout across multiple volumes, spreading the load and aggregating performance.

More detailed information about performance scalability of Azure NetApp Files can be found here:

- [Performance benchmarks for Azure NetApp Files.](#)
- [Azure NetApp Files: Benchmarks](#) (select “Oracle” tab)

Appendix A: Stale NFS Locks

If an Oracle database server crashes, it might have problems with stale NFS locks upon restart. This problem is avoidable by paying careful attention to the configuration of name resolution on the server.

This problem arises because creating a lock and clearing a lock use two slightly different methods of name resolution. Two processes are involved, the Network Lock Manager (NLM) and the NFS client. The NLM uses `uname -n` to determine the host name, while the `rpc.statd` process uses `gethostbyname()`. These host names must match for the OS to properly clear stale locks. For example, the host might be looking for locks owned by `dbserver5`, but the locks were registered by the host as `dbserver5.mydomain.org`. If `gethostbyname()` does not return the same value as `uname -a`, then the lock release process did not succeed.

The following sample script verifies whether name resolution is fully consistent:

```
#!/usr/bin/perl
```

```
$uname=`uname -n`;  
chomp($uname);  
($name, $aliases, $addrtype, $length, @addrs) = gethostbyname $uname;  
print "uname -n yields: $uname\n";  
print "gethostbyname yields: $name\n";
```

If `gethostbyname` does not match `uname`, stale locks are likely. For example, this result reveals a potential problem:

```
uname -n yields: dbserver5  
gethostbyname yields: dbserver5.mydomain.org
```

The solution is found by changing the order in which hosts appear in `/etc/hosts`. For example, assume that the hosts file includes this entry:

```
10.156.110.201 dbserver5.mydomain.org dbserver5 loghost
```

To resolve this issue, change the order in which the fully qualified domain name and the short host name appear:

```
10.156.110.201 dbserver5 dbserver5.mydomain.org loghost
```

`gethostbyname()` now returns the short `dbserver5` host name, which matches the output of `uname`. Locks are thus cleared automatically after a server crash.

Where to Find Additional Information

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

- NetApp Documentation Center
<https://www.netapp.com/us/documentation/index.aspx>
- NetApp Product Documentation
<https://docs.netapp.com>
- Cloud Volumes ONTAP
<https://cloud.netapp.com/ontap-cloud>
- Azure NetApp Files Documentation
<https://docs.microsoft.com/en-us/azure/azure-netapp-files/>
- TR-4691 ONTAP Databases on ONTAP Cloud with Microsoft Azure
<https://www.netapp.com/us/media/tr-4691.pdf>

Version History

Version	Date	Document Version History
Version 1.0	May 2019	Initial release.

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 © 2019 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

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.