



# Oracle Database Deployment and Migration Best Practices for ANF

NetApp Solutions

NetApp  
February 12, 2025

# Table of Contents

- Oracle Database Deployment and Migration Best Practices for ANF ..... 1
  - TR-4954: Oracle Database Deployment and Migration Best Practices for ANF ..... 1
  - Solution Architecture ..... 2
  - Factors to consider for Oracle database deployment ..... 4
  - Step-by-Step Oracle deployment procedures on Azure VM and Azure NetApp Files ..... 9
  - Database migration from on-premises to Azure cloud ..... 29

# Oracle Database Deployment and Migration Best Practices for ANF

## TR-4954: Oracle Database Deployment and Migration Best Practices for ANF

This best practice guide provides details of a solution for deploying and migrating Oracle database on Azure NetApp file storage and Azure VM.

Author(s): Allen Cao, Niyaz Mohamed, NetApp

### Overview

Many mission-critical Oracle enterprise databases are still hosted on-premises, and many enterprises are looking to migrate these Oracle databases to a public cloud. Often, these Oracle databases are application centric and thus require user-specific configurations, a capability that is missing from many database-as-a-service public-cloud offerings. Therefore, the current database landscape calls for a public-cloud-based Oracle database solution built from a high-performance, scalable compute and storage service that can accommodate unique requirements. Azure virtual machine compute instances and the Azure NetApp Files storage service might be the missing pieces of this puzzle that you can leverage to build and migrate your mission-critical Oracle database workloads to a public cloud.

### Azure Virtual Machine

Azure virtual machines are one of several types of on-demand, scalable computing resources that Azure offers. Typically, you choose a virtual machine when you need more control over the computing environment than the other choices offer. Azure virtual machines offer a quick and easy way to create a computer with specific configurations required to run your Oracle database, whether it is for compute- or memory-intensive workloads. Virtual machines in an Azure virtual network can easily be connected to your organization's network, for example through a secured VPN tunnel.

### Azure NetApp Files (ANF)

Azure NetApp Files is a fully managed Microsoft service that will take your database workload to the cloud faster and more securely than ever before. It was designed to meet the core requirements of running high-performance workloads such as Oracle databases in the cloud, and it provides performance tiers that reflect the real-world range of IOPS demands, low latency, high availability, high durability, manageability at scale, and fast and efficient backup, recovery, and cloning. These capabilities are possible because Azure NetApp Files is based on physical all-flash NetApp ONTAP systems running within the Azure data center environment. Azure NetApp Files is completely integrated into the Azure DCs and portal, and customers can use the same comfortable graphical interface and APIs for creating and managing shared files as with any other Azure object. With Azure NetApp file, you can unlock the full capabilities of Azure without extra risk, cost, or time and trust the only enterprise file service native to Azure.

### Conclusion

This documentation describes in detail how to deploy, configure, and migrate an Oracle database with an Azure virtual machine and Azure NetApp Files storage service that delivers performance and durability similar to an on-premises system. For best-practices guidance, see TR-4780 [Oracle Databases on Microsoft Azure](#). More importantly, NetApp also provides automation toolkits that automate most of the tasks that are required


for the deployment, configuration, data protection, migration, and management of your Oracle database workload in the Azure public cloud. The automation toolkits are available for download at NetApp public GitHub site: [NetApp-Automation](https://github.com/NetApp/NetApp-Automation).

## Solution Architecture

The following architecture diagram illustrates a highly available Oracle database deployment on Azure VM instances and the Azure NetApp Files storage.

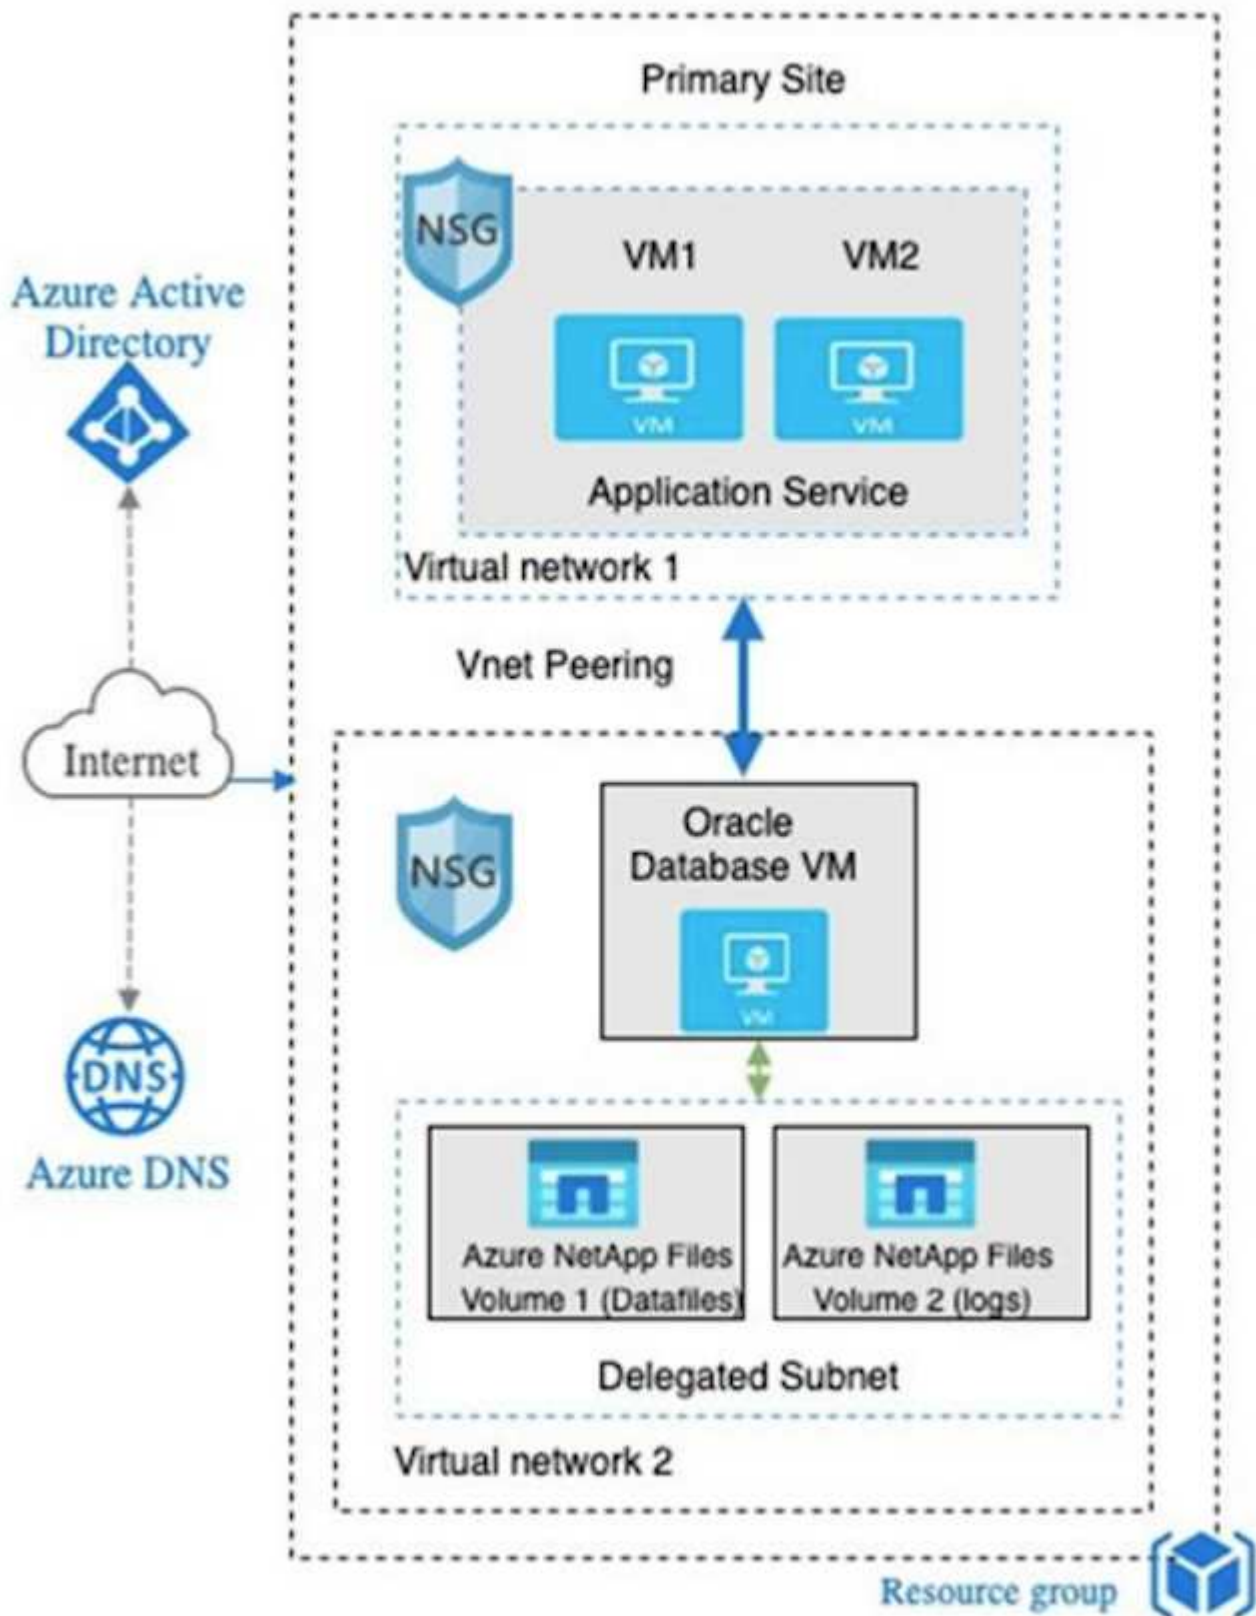
Within the environment, the Oracle compute instance is deployed via an Azure services VM console. There are multiple Azure instance types available from the console. NetApp recommends deploying a database-oriented Azure VM instance that meets your expected workload.

Oracle database storage on the other hand is deployed with the Azure NetApp Files service available from Azure console. The Oracle binary, data, or log volumes are subsequently presented and mounted on an Azure VM instance Linux host.



In many respects, the implementation of Azure NetApp Files in Azure cloud is very similar to an on-premises ONTAP data storage architecture with many built-in redundancies, such as RAID and dual controllers. For disaster recovery, a standby site can be setup in different regions and database can be synced up with the primary site using application-level replication (for example, Oracle Data Guard).

In our test validation for Oracle database deployment and data protection, the Oracle database is deployed on a single Azure VM as illustrated in the following diagram:



The Azure Oracle environment can be managed with an Ansible controller node for automation using tool kits provided by NetApp for database deployment, backup, recovery, and database migration. Any updates to the Oracle Azure VM instance operating-system kernel or Oracle patching can be performed in parallel to keep the primary and standby in sync. In fact, the initial toolkits can be easily expanded to perform daily Oracle tasks if

needed. If you need help to set up a CLI Ansible controller, see [NetApp Solution Automation](#) to get started.

## Factors to consider for Oracle database deployment

A public cloud provides many choices for compute and storage, and using the correct type of compute instance and storage engine is a good place to start for database deployment. You should also select compute and storage configurations that are optimized for Oracle databases.

The following sections describe the key considerations when deploying an Oracle database in the Azure public cloud on an Azure virtual machine instance with Azure NetApp Files storage.

### VM type and sizing

Selecting the right VM type and size is important for optimal performance of a relational database in a public cloud. An Azure virtual machine provides a variety of compute instances that can be used to host Oracle database workloads. See the Microsoft documentation [Sizes for virtual machines in Azure](#) for different types of Azure virtual machines and their sizing. In general, NetApp recommends using a general-purpose Azure virtual machine for the deployment of small- and medium-sized Oracle databases. For the deployment of larger Oracle databases, a memory-optimized Azure VM is appropriate. With more available RAM, a larger Oracle SGA or smart flash cache can be configured to reduce the physical I/O, which in turn improves database performance.

Azure NetApp Files works as an NFS mount attached to an Azure virtual machine, which offers higher throughput and overcomes the storage-optimized VM throughput limit with local storage. Therefore, running Oracle on Azure NetApp Files could reduce the licensable Oracle CPU core count and licensing costs. See [TR-4780: Oracle Databases on Microsoft Azure](#), Section 7 - How Does Oracle Licensing Work?

Other factors to consider include the following:

- Choose the correct vCPU and RAM combination based on workload characteristics. As the RAM size increases on the VM, so does the number of vCPU cores. There should be a balance at some point as the Oracle license fees are charged on the number of vCPU cores.
- Add swap space to a VM. The default Azure VM deployment does not create a swap space, which is not optimal for a database.

### Azure NetApp Files performance

Azure NetApp Files volumes are allocated from a capacity pool the customer must provision in their Azure NetApp Files storage account. Each capacity pool is assigned as follows:

- To a service level that defines the overall performance capability.
- The initially provisioned storage capacity or tiering for that capacity pool. A quality of service (QoS) level that defines the overall maximum throughput per provisioned space.

The service level and initially provisioned storage capacity determines the performance level for a particular Oracle database volume.

#### 1. Service Levels for Azure NetApp Files

Azure NetApp Files supports three service levels: Ultra, Premium, and Standard.

- **Ultra storage.** This tier provides up to 128MiBps of throughput per 1TiB of volume quota assigned.
- **Premium storage.** This tier provides up to 64MiBps of throughput per 1TiB of volume quota assigned.
- **Standard storage.** This tier provides up to 16MiBps of throughput per 1TiB of volume quota assigned.

## 2. Capacity pool and quality of service

Each of the desired service levels has an associated cost for provisioned capacity and includes a quality-of-service (QoS) level that defines the overall maximum throughput for provisioned space.

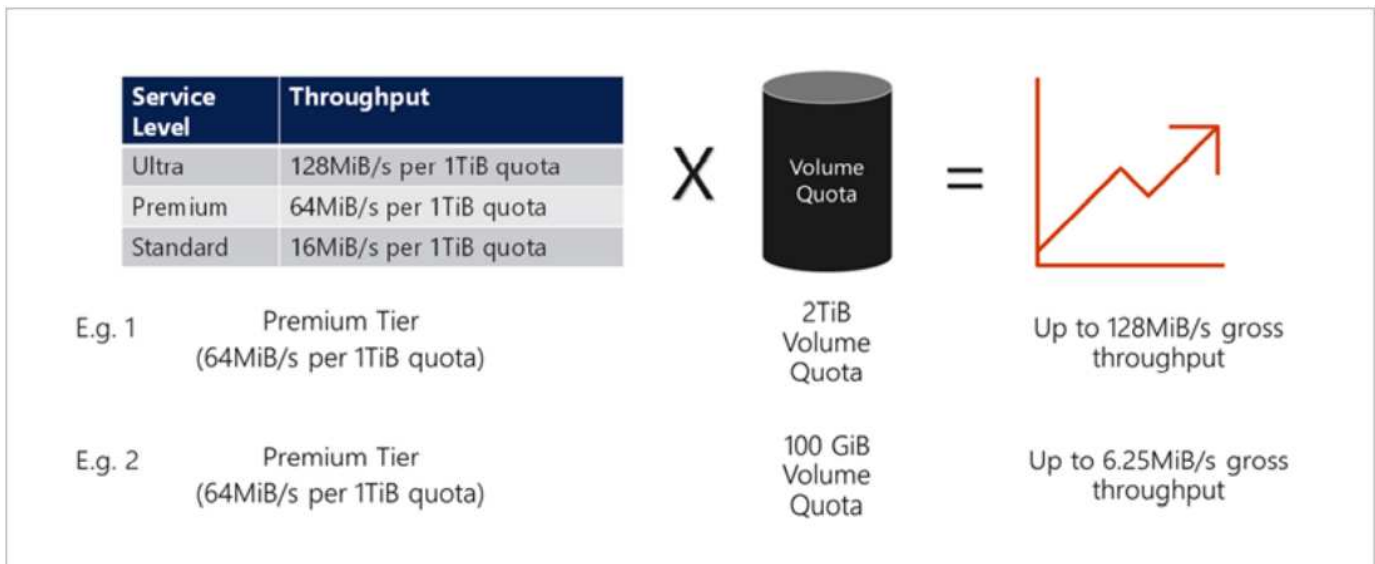
For example, a 10TiB-provisioned single-capacity pool with the premium service level provides an overall available throughput for all volumes in this capacity pool of 10x 64MBps, so 640MBps with 40,000 (16K) IOPs or 80,000 (8K) IOPs.

The minimum capacity pool size is 4TiB. You can change the size of a capacity pool in 1TiB increments in response to changes in your workload requirements to manage storage needs and costs.

## 3. Calculate the service level at a database volume

The throughput limit for an Oracle database volume is determined by a combination of the following factors: The service level of the capacity pool to which the volume belongs and The quota assigned to the volume.

The following diagram shows how the throughput limit for an Oracle database volume is calculated.



In example 1, a volume from a capacity pool with the Premium storage tier that is assigned 2TiB of quota is assigned a throughput limit of 128MiBps (2TiB \* 64MiBps). This scenario applies regardless of the capacity pool size or the actual volume consumption.


In example 2, a volume from a capacity pool with the Premium storage tier that is assigned 100GiB of quota is assigned a throughput limit of 6.25MiBps (0.09765625TiB \* 64MiBps). This scenario applies regardless of the capacity pool size or the actual volume consumption.

Please note that the minimum volume size is 100GiB.

## Storage layout and settings

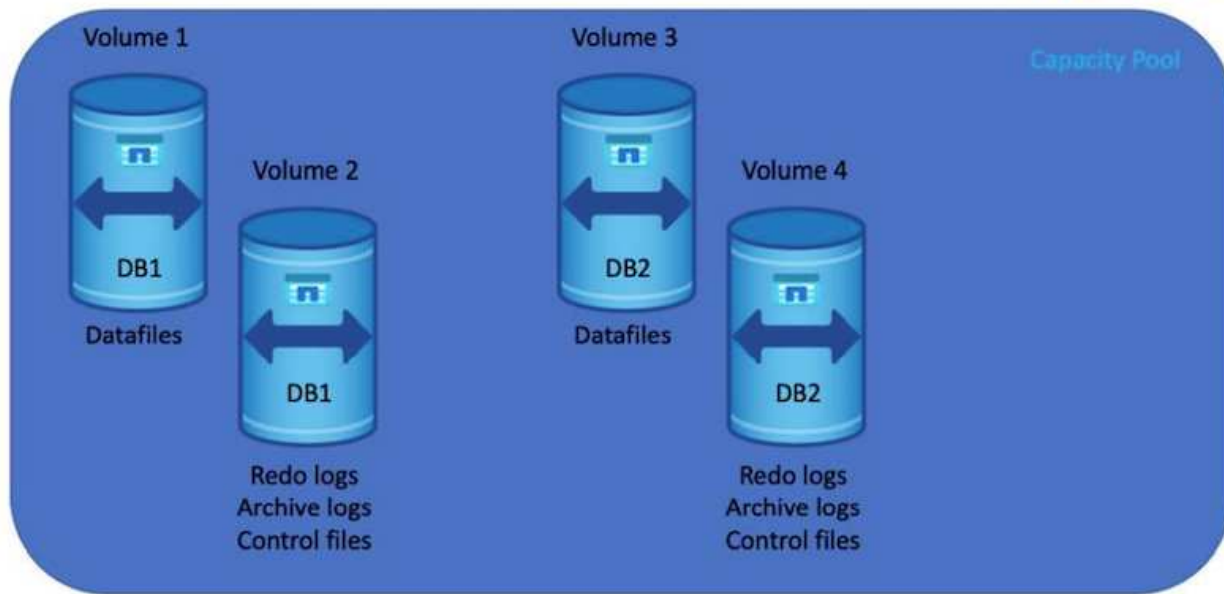
NetApp recommends the following storage layout:

- For small databases, using single volume layout for all Oracle files.



- For large databases, the recommended volume layout is multiple volumes: one for Oracle data and a duplicate control file and one for the Oracle active log, archived log, and control file. NetApp highly recommends allocating a volume for the Oracle binary instead of the local drive so that the database can be relocated to a new host and quickly restored.






## NFS configuration

Linux, the most common operating system, includes native NFS capabilities. Oracle offers a direct NFS (dNFS) client natively integrated into Oracle. Oracle dNFS bypasses the OS cache and enables parallel processing to improve database performance. Oracle has supported NFSv3 for over 20 years, and NFSv4 is supported with Oracle 12.1.0.2 and later.

By using dNFS (available since Oracle 11g), an Oracle database running on an Azure Virtual Machine can drive significantly more I/O than the native NFS client. Automated Oracle deployment using the NetApp automation toolkit automatically configures dNFS on NFSv3.

The following diagram demonstrates the SLOB benchmark on Azure NetApp Files with Oracle dNFS.

## Oracle and Azure NetApp Files Comparing dNFS and Native NFS (Kernel NFS) 75% Read, 25% Update SLOB2 Workload



Other factors to consider:

- 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 TCP slot tables to 128.

```
sysctl -a | grep tcp.*.slot_table
```

- The following table provides recommended NFS mount options for a single instance of Linux NFSv3.

File Type	Mount Options
<ul style="list-style-type: none"> <li>• Control files</li> <li>• Data files</li> <li>• Redo logs</li> </ul>	rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsiz=65536
<ul style="list-style-type: none"> <li>• ORACLE_HOME</li> <li>• ORACLE_BASE</li> </ul>	rw,bg,hard,vers=3,proto=tcp,timeo=600,rsize=65536,wsiz=65536



Before using dNFS, verify that the patches described in Oracle Doc 1495104.1 are installed. The NetApp Support matrix for NFSv3 and NFSv4 do not include specific operating systems. All OSs that obey the RFC are supported. When searching the online IMT for NFSv3 or NFSv4 support, do not select a specific OS because no matches will be displayed. All OSs are implicitly supported by the general policy.

## Step-by-Step Oracle deployment procedures on Azure VM and Azure NetApp Files

This section describes the deployment procedures of deploying Oracle RDS custom database with FSx storage.

### Deploy an Azure VM with ANF for Oracle via Azure portal console

If you are new to Azure, you first need to set up an Azure account environment. This includes signing up your organization to use Azure Active Directory. The following section is a summary of these steps. For details, see the linked Azure-specific documentation.

#### Create and consume Azure resources

After your Azure environment is set up and an account is created and associated with a subscription, you can log into Azure portal with the account to create the necessary resources to run Oracle.

##### 1. Create a virtual network or VNet

Azure Virtual Network (VNet) is the fundamental building block for your private network in Azure. VNet enables many types of Azure resources, such as Azure Virtual Machines (VMs), to securely communicate with each other, the internet, and on-premises networks. Before provisioning an Azure VM, a VNet (where a VM is deployed) must first be configured.

See [Create a virtual network using the Azure portal](#) to create a VNet.

##### 2. Create a NetApp storage account and capacity pool for ANF

In this deployment scenario, an Azure VM OS is provisioned using regular Azure storage, but ANF volumes are provisioned to run Oracle database via NFS. First, you need to create a NetApp storage account and a capacity pool to host the storage volumes.

See [Set up Azure NetApp Files and create an NFS volume](#) to set up an ANF capacity pool.

##### 3. Provision Azure VM for Oracle

Based on your workload, determine what type of Azure VM you need and the size of the VM vCPU and RAM to deploy for Oracle. Then, from the Azure console, click the VM icon to launch the VM deployment workflow.

1. From the Azure VM page, click **Create** and then choose **Azure virtual machine**.

Microsoft Azure Search resources, services, and docs (G+)

Home > Virtual machines Hybrid Cloud TME

Create Switch to classic Reservations Manage view Refresh Export to CSV Open query Assign tags Start Restart Stop Delete Services Maintenance

Filter for any field... Subscription equals all Type equals all Resource group equals all Location equals all Add filter

No grouping List view

Name	Type	Subscription	Resource group	Location	Status	Operating system	Size	Public IP address	Disks
acao-ora01	Virtual machine	Hybrid Cloud TME Onprem	TMEtstres	South Central US	Stopped (deallocated)	Linux	Standard_B4ms	13.65.63.157	1
ANFAV5val2JH	Virtual machine	Hybrid Cloud TME Onprem	ANFAV5VAL2	West Europe	Running	Windows	Standard_DS2_v2	20.229.80.88	1
ANFAV5f001	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Stopped (deallocated)	Linux	Standard_DS2ds_v4	-	1
ANFAV5f0AZ1	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Running	Linux	Standard_E32as_v4	40.124.74.246	1
ANFAV5f0AZ2	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Stopped (deallocated)	Linux	Standard_E32as_v4	40.124.178.111	1
ANFAV5f0AZ3	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Stopped (deallocated)	Linux	Standard_E32as_v4	40.124.194.32	1
ANFAV5valDC	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Stopped (deallocated)	Windows	Standard_B4ms	-	1
ANFAV5valIH	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Running	Windows	Standard_B2ms	70.37.66.218	1
ANFAV5valIH2	Virtual machine	Hybrid Cloud TME Onprem	anfavsrg	South Central US	Running	Windows	Standard_B2s	20.225.210.195	1
ANFCVOCM	Virtual machine	Hybrid Cloud TME Onprem	anfavsval2	West Europe	Running	Linux	Standard_DS3_v2	-	1
ANFCVODRDC2	Virtual machine	Hybrid Cloud TME Onprem	anfavsval2	West Europe	Running	Windows	Standard_B2s	-	1
ANFCVODRDemo	Virtual machine	Hybrid Cloud TME Onprem	anfvcodrdemo-rg	West Europe	Running	Linux	Standard_E4s_v3	-	5
AVSCVOPerfinguest	Virtual machine	Hybrid Cloud TME Onprem	avscvoperfinguest-rg	West Europe	Stopped (deallocated)	Linux	Standard_DS15_v2	-	5

2. Choose the subscription ID for the deployment, and then choose the resource group, region, host name, VM image, size, and authentication method. Go to the Disk page.

[Home](#) > [Virtual machines](#) >

# Create a virtual machine ...

[Basics](#) [Disks](#) [Networking](#) [Management](#) [Advanced](#) [Tags](#) [Review + create](#)

Create a virtual machine that runs Linux or Windows. Select an image from Azure marketplace or use your own customized image. Complete the Basics tab then Review + create to provision a virtual machine with default parameters or review each tab for full customization. [Learn more](#) ↗

## Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription * ⓘ	<input type="text" value="Hybrid Cloud TME Onprem"/>
Resource group * ⓘ	<input type="text" value="ANFAVSRG"/> <a href="#">Create new</a>

## Instance details

Virtual machine name * ⓘ	<input type="text" value="acao-ora01"/>
Region * ⓘ	<input type="text" value="(US) South Central US"/>
Availability options ⓘ	<input type="text" value="No infrastructure redundancy required"/>
Security type ⓘ	<input type="text" value="Standard"/>
Image * ⓘ	<input type="text" value="Red Hat Enterprise Linux 8.0 (LVM) - Gen2"/> <a href="#">See all images</a>   <a href="#">Configure VM generation</a>
Run with Azure Spot discount ⓘ	<input type="checkbox"/>
Size * ⓘ	<input type="text" value="Standard_D8s_v3 - 8 vcpus, 32 GiB memory (\$273.02/month)"/> <a href="#">See all sizes</a>

## Administrator account

Authentication type ⓘ	<input type="radio"/> SSH public key
	<input checked="" type="radio"/> Password

[Review + create](#)[< Previous](#)[Next : Disks >](#)

[Home](#) > [Virtual machines](#) >

## Create a virtual machine ...

Size \* ⓘ  See all sizes

### Administrator account

Authentication type ⓘ  SSH public key  
 Password

Username \* ⓘ  ✓

Password \* ⓘ  ✓

Confirm password \* ⓘ  ✓

### Inbound port rules

Select which virtual machine network ports are accessible from the public internet. You can specify more limited or granular network access on the Networking tab.

Public inbound ports \* ⓘ  None  
 Allow selected ports

Select inbound ports \*

**⚠ This will allow all IP addresses to access your virtual machine.** This is only recommended for testing. Use the Advanced controls in the Networking tab to create rules to limit inbound traffic to known IP addresses.

### Licensing

If you have eligible Red Hat Enterprise Linux subscriptions that are enabled for Red Hat Cloud Access, you can use Azure Hybrid Benefit to attach your Red Hat subscriptions to this VM and save money on compute costs [Learn more](#)

Your Azure subscription is currently not a part of Red Hat Cloud Access. In order to enable AHB for this VM, you must add this Azure subscription to Cloud Access. [Learn more](#)

[Review + create](#)[< Previous](#)[Next : Disks >](#)

3. Choose **premium SSD** for OS local redundancy and leave the data disk blank because the data disks are mounted from ANF storage. Go to the Networking page.

[Home](#) > [Virtual machines](#) >

# Create a virtual machine

[Basics](#) [Disks](#) [Networking](#) [Management](#) [Advanced](#) [Tags](#) [Review + create](#)

Azure VMs have one operating system disk and a temporary disk for short-term storage. You can attach additional data disks. The size of the VM determines the type of storage you can use and the number of data disks allowed. [Learn more](#)

## Disk options

OS disk type \* Delete with VM Enable encryption at host 

**i** Encryption at host is not registered for the selected subscription. [Learn more about enabling this feature](#)

Encryption type \* Enable Ultra Disk compatibility 

## Data disks for acao-ora01

You can add and configure additional data disks for your virtual machine or attach existing disks. This VM also comes with a temporary disk.

LUN	Name	Size (GiB)	Disk type	Host caching	Delete with VM
<a href="#">Create and attach a new disk</a>	<a href="#">Attach an existing disk</a>				

[Advanced](#)[Review + create](#)[< Previous](#)[Next : Networking >](#)

4. Choose the VNet and subnet. Allocate a public IP for external VM access. Then go to the Management page.

[Home](#) > [Virtual machines](#) >

## Create a virtual machine

### Network interface

When creating a virtual machine, a network interface will be created for you.

Virtual network *	<input type="text" value="ANFAVSVal"/>
	<a href="#">Create new</a>
Subnet *	<input type="text" value="VM_Sub (172.30.137.128/25)"/>
	<a href="#">Manage subnet configuration</a>
Public IP	<input type="text" value="(new) acao-ora01-ip"/>
	<a href="#">Create new</a>
NIC network security group	<input type="radio"/> None <input checked="" type="radio"/> Basic <input type="radio"/> Advanced
Public inbound ports *	<input type="radio"/> None <input checked="" type="radio"/> Allow selected ports
Select inbound ports *	<input type="text" value="SSH (22)"/>

**⚠ This will allow all IP addresses to access your virtual machine.** This is only recommended for testing. Use the Advanced controls in the Networking tab to create rules to limit inbound traffic to known IP addresses.

Delete public IP and NIC when VM is deleted

Enable accelerated networking

### Load balancing

You can place this virtual machine in the backend pool of an existing Azure load balancing solution. [Learn more](#)

Place this virtual machine behind an existing load balancing solution?

[Review + create](#)[< Previous](#)[Next : Management >](#)

5. Keep all defaults for Management and move to the Advanced page.



[Home](#) > [Virtual machines](#) >

# Create a virtual machine

[Basics](#) [Disks](#) [Networking](#) **[Management](#)** [Advanced](#) [Tags](#) [Review + create](#)

Configure monitoring and management options for your VM.

## Microsoft Defender for Cloud

Microsoft Defender for Cloud provides unified security management and advanced threat protection across hybrid cloud workloads. [Learn more](#)

Your subscription is protected by Microsoft Defender for Cloud basic plan.

## Monitoring

Boot diagnostics

- Enable with managed storage account (recommended)  
 Enable with custom storage account  
 Disable

Enable OS guest diagnostics

## Identity

Enable system assigned managed identity

## Azure AD

Login with Azure AD

RBAC role assignment of Virtual Machine Administrator Login or Virtual Machine User Login is required when using Azure AD login. [Learn more](#)

Azure AD login now uses SSH certificate-based authentication. You will need to use an SSH client that supports OpenSSH certificates. You can use Azure CLI or Cloud Shell from the Azure Portal. [Learn more](#)

## Auto-shutdown

Enable auto-shutdown

## Backup

[Review + create](#)[< Previous](#)[Next : Advanced >](#)

6. Keep all defaults for the Advanced page unless you need to customize a VM after deployment with custom scripts. Then go to Tags page.

[Home](#) > [Virtual machines](#) >

# Create a virtual machine

[Basics](#) [Disks](#) [Networking](#) [Management](#) **[Advanced](#)** [Tags](#) [Review + create](#)


Add additional configuration, agents, scripts or applications via virtual machine extensions or cloud-init.

## Extensions

Extensions provide post-deployment configuration and automation.


Extensions  [Select an extension to install](#)

## VM applications



VM applications contain application files that are securely and reliably downloaded on your VM after deployment. In addition to the application files, an install and uninstall script are included in the application. You can easily add or remove applications on your VM after create. [Learn more](#) 

[Select a VM application to install](#)


## Custom data

Pass a script, configuration file, or other data into the virtual machine **while it is being provisioned**. The data will be saved on the VM in a known location. [Learn more about custom data for VMs](#) 

Custom data

 Your image must have a code to support consumption of custom data. If your image supports cloud-init, custom-data will be processed by cloud-init. [Learn more about custom data for VMs](#) 

## User data

Pass a script, configuration file, or other data that will be accessible to your applications **throughout the lifetime of the virtual machine**. Don't use user data for storing your secrets or passwords. [Learn more about user data for VMs](#) 

Enable user data

[Review + create](#)[< Previous](#)[Next : Tags >](#)

7. Add a tag for the VM if desired. Then, go to the Review + create page.


[Home](#) > [Virtual machines](#) >

## Create a virtual machine ...

Basics Disks Networking Management Advanced **Tags** Review + create

Tags are name/value pairs that enable you to categorize resources and view consolidated billing by applying the same tag to multiple resources and resource groups. [Learn more about tags](#)

Note that if you create tags and then change resource settings on other tabs, your tags will be automatically updated.

Name ⓘ	Value ⓘ	Resource
<input type="text" value="database"/>	<input type="text" value="oracle"/>	12 selected  
<input type="text"/>	<input type="text"/>	12 selected 

Review + create

< Previous

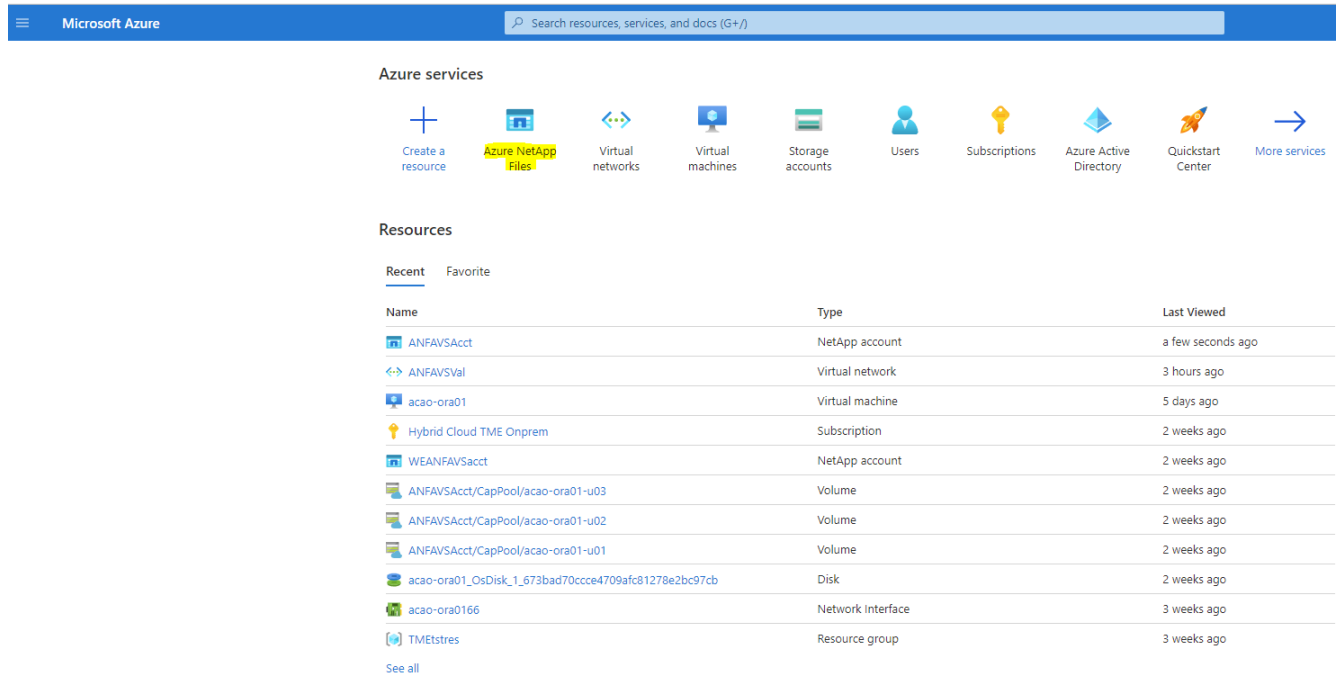
Next : Review + create >

- The deployment workflow runs a validation on the configuration, and, if the validation passes, click **Create** to create the VM.

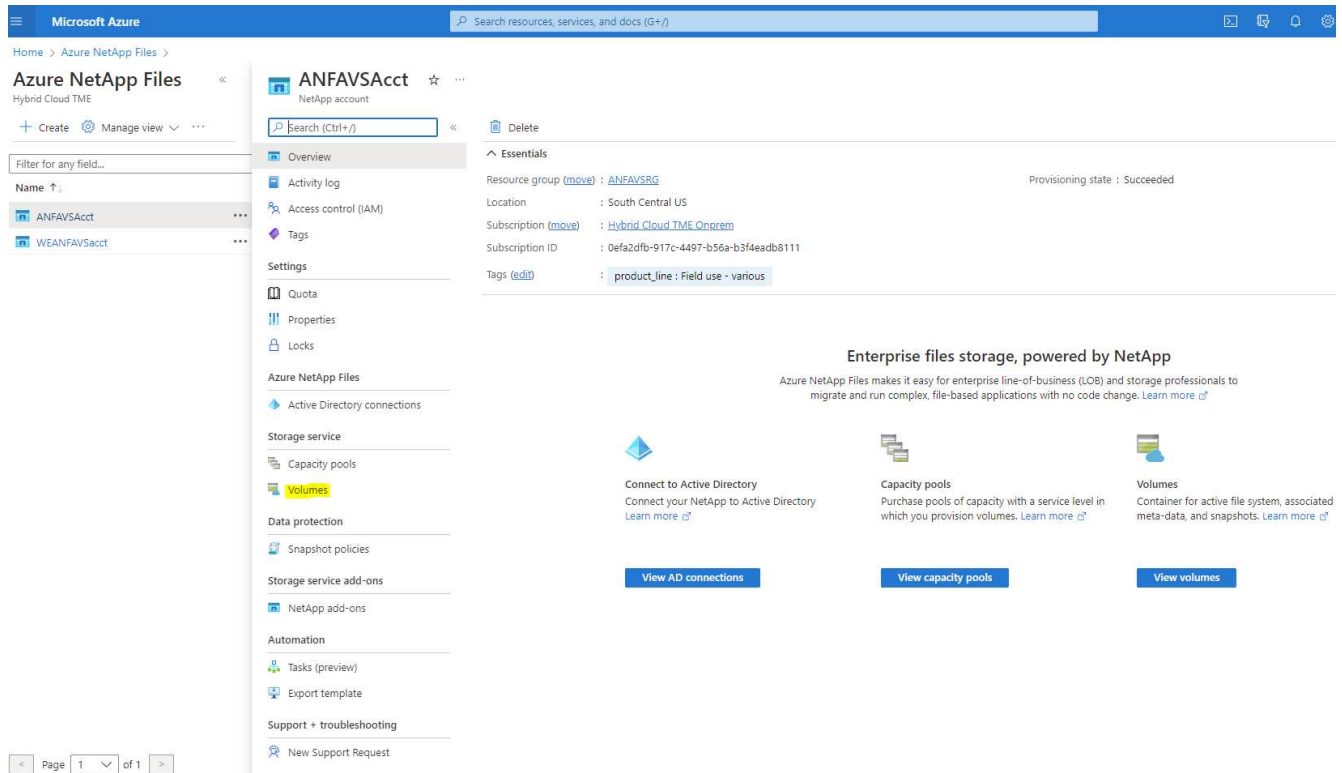
#### 4. Provision ANF database volumes for Oracle


You must create three NFS volumes for an ANF capacity pool for the Oracle binary, data, and log volumes respectively.

- From the Azure console, under the list of Azure services, click Azure NetApp Files to open a volume creation workflow. If you have more than one ANF storage account, click the account that you would like to provision volumes from.




- Under your NetApp storage account, click **Volumes**, and then **Add volume** to create new Oracle volumes.





- As a good practice, identify Oracle volumes with the VM hostname as a prefix and then followed by the mount point on the host, such as u01 for Oracle binary, u02 for Oracle data, and u03 for Oracle log. Choose the same VNet for the volume as for the VM. Click **Next: Protocol**.



- Choose the NFS protocol, add the Oracle host IP address to the allowed client, and remove the default policy that allows all IP addresses 0.0.0.0/0. Then click **Next: Tags**.

Microsoft Azure

Search resources, services, and docs (G+)

Home > Azure NetApp Files > ANFAVSAcct | Volumes >

### ANFAVSAcct | Volumes

NetApp account

Search (Ctrl+/)

+ Add volume

Search volumes

Name	Quota
anf2-z1-stdds01	200 GiB
anf2-z1-stdds02	200 GiB
anf2-z1-stdds03	100 GiB
anf2-z1-stdds04	100 GiB
anf2-z1-stdds05	100 GiB
anf2-z1-stdds06	100 GiB
anf2-z1-stdds07	100 GiB
anf2-z1-stdds08	100 GiB
anf-z1-stdds01	6 TiB
anf-z1-stdds02	200 GiB
anf-z1-stdds03	1 TiB
anf-z1-stdds04	200 GiB
anf-z1-stdds06	200 GiB
anf-z1-stdds07	200 GiB
anf-z1-stdds08	200 GiB
anf-zq-stdds05	1 TiB
vol1	1 TiB
vol3basic	100 GiB
volnfsbasic	100 GiB
volnfsstd	100 GiB
volnfsstdnew	100 GiB
zone1basic	6 TiB
zone2basic	100 GiB

### Create a volume

Basics Protocol Tags Review + create

Configure access to your volume.

**Access**

Protocol type  NFS  SMB  Dual-protocol

**Configuration**

File path \*

Versions \*

Kerberos  Enabled  Disabled

LDAP  Enabled  Disabled

Azure VMware Solution DataStore

**Export policy**

Configure the volume's export policy. This can be edited later. [Learn more](#)

↑ Move up ↓ Move down ↕ Move to top ↓ Move to bottom 🗑 Delete

<input type="checkbox"/>	Index	Allowed clients	Access	Root Access	...
<input type="checkbox"/>	1	<del>0.0.0.0</del>	Read & Write	On	...
<input type="checkbox"/>	2	172.30.137.142 ✓	Read & Write	On	...

**Review + create** < Previous Next: Tags >

5. Add a volume tag if desired. Then click **Review + Create**.

Microsoft Azure Search resources, services, and docs (G+)

Home > Azure NetApp Files > ANFAVSAcct | Volumes >

## ANFAVSAcct | Volumes

NetApp account

Search (Ctrl+/) Add volume

Search volumes

Name	Quota
anf2-z1-stdds01	200 GiB
anf2-z1-stdds02	200 GiB
anf2-z1-stdds03	100 GiB
anf2-z1-stdds04	100 GiB
anf2-z1-stdds05	100 GiB
anf2-z1-stdds06	100 GiB
anf2-z1-stdds07	100 GiB
anf2-z1-stdds08	100 GiB
anf-z1-stdds01	6 TiB
anf-z1-stdds02	200 GiB
anf-z1-stdds03	1 TiB
anf-z1-stdds04	200 GiB
anf-z1-stdds06	200 GiB
anf-z1-stdds07	200 GiB
anf-z1-stdds08	200 GiB
anf-zq-stdds05	1 TiB
vol1	1 TiB
vol3basic	100 GiB
volnfsbasic	100 GiB
volnfsstd	100 GiB
volnfsstdnew	100 GiB
zone1basic	6 TiB
zone2basic	100 GiB

### Create a volume

Basics Protocol **Tags** Review + create

Tags are name/value pairs that enable you to categorize resources and view consolidated billing by applying the same tag to multiple resources and resource groups. [Learn more about tags](#)

Note that if you create tags and then change resource settings on other tabs, your tags will be automatically updated.

Name Value

database : oracle

Review + create < Previous Next : Review + create >

6. If the validation passes, click **Create** to create the volume.

## Install and configure Oracle on Azure VM with ANF

The NetApp solutions team has created many Ansible-based automation toolkits to help you deploy Oracle in Azure smoothly. Follow these steps to deploy Oracle on an Azure VM.

### Set up an Ansible controller

If you have not set up an Ansible controller, see [NetApp Solution Automation](#), which has detailed instructions on how to setup an Ansible controller.

### Obtain Oracle deployment automation toolkit

Clone a copy of the Oracle deployment toolkit in your home directory under the user ID that you use to log into the Ansible controller.

```
git clone https://github.com/NetApp-Automation/na_oracle19c_deploy.git
```

### Execute the toolkit with your configuration

See the [CLI deployment Oracle 19c Database](#) to execute the playbook with the CLI. You can ignore the ONTAP portion of the variables configuration in the global VARS file when you create database volumes from



the Azure console rather than the CLI.



The toolkit default deploys Oracle 19c with RU 19.8. It can be easily adapted for any other patch level with minor default configuration changes. Also default seed-database active log files are deployed into the data volume. If you need active log files on the log volume, it should be relocated after initial deployment. Reach out to the NetApp Solution team for help if needed.

## Set up AzAcSnap backup tool for app-consistent snapshots for Oracle

The Azure Application-Consistent Snapshot tool (AzAcSnap) is a command-line tool that enables data protection for third-party databases by handling all the orchestration required to put them into an application-consistent state before taking a storage snapshot. It then returns these databases to an operational state. NetApp recommends installing the tool on the database server host. See the following installation and configuration procedures.

### Install AzAcSnap tool

1. Get the most recent version of the [the AzArcSnap Installer](#).
2. Copy the downloaded self-installer to the target system.
3. Execute the self-installer as the root user with the default installation option. If necessary, make the file executable using the `chmod +x *.run` command.

```
./azacsnap_installer_v5.0.run -I
```

### Configure Oracle connectivity

The snapshot tools communicate with the Oracle database and need a database user with appropriate permissions to enable or disable backup mode.

#### 1. Set up AzAcSnap database user

The following examples show the setup of the Oracle database user and the use of sqlplus for communication to the Oracle database. The example commands set up a user (AZACSNAP) in the Oracle database and change the IP address, usernames, and passwords as appropriate.

1. From the Oracle database installation, launch sqlplus to log into the database.

```
su - oracle  
sqlplus / AS SYSDBA
```

2. Create the user.

```
CREATE USER azacsnap IDENTIFIED BY password;
```

3. Grant the user permissions. This example sets the permission for the AZACSNAP user to enable putting the database into backup mode.

```
GRANT CREATE SESSION TO azacsnap;  
GRANT SYSBACKUP TO azacsnap;
```

4. Change the default user's password expiration to unlimited.

```
ALTER PROFILE default LIMIT PASSWORD_LIFE_TIME unlimited;
```

5. Validate azacsnap connectivity for the database.

```
connect azacsnap/password  
quit;
```

## 2. Configure Linux-user azacsnap for DB access with Oracle wallet

The AzAcSnap default installation creates an azacsnap OS user. It's Bash shell environment must be configured for Oracle database access with the password stored in an Oracle wallet.

1. As root user, run the `cat /etc/oratab` command to identify the `ORACLE_HOME` and `ORACLE_SID` variables on the host.

```
cat /etc/oratab
```

2. Add `ORACLE_HOME`, `ORACLE_SID`, `TNS_ADMIN`, and `PATH` variables to the azacsnap user bash profile. Change the variables as needed.

```
echo "export ORACLE_SID=ORATEST" >> /home/azacsnap/.bash_profile  
echo "export ORACLE_HOME=/u01/app/oracle/product/19800/ORATST" >>  
/home/azacsnap/.bash_profile  
echo "export TNS_ADMIN=/home/azacsnap" >> /home/azacsnap/.bash_profile  
echo "export PATH=\$PATH:\$ORACLE_HOME/bin" >>  
/home/azacsnap/.bash_profile
```

3. As the Linux user azacsnap, create the wallet. You are prompted for the wallet password.

```
sudo su - azacsnap  
  
mkstore -wrl $TNS_ADMIN/.oracle_wallet/ -create
```

4. Add the connect string credentials to the Oracle Wallet. In the following example command, `AZACSNAP` is the `ConnectString` to be used by `AzAcSnap`, `azacsnap` is the Oracle Database User, and `AzPasswd1` is the Oracle User's database password. You are again prompted for the wallet password.

```
mkstore -wrl $TNS_ADMIN/.oracle_wallet/ -createCredential AZACSNAP
azacsnap AzPasswdl
```

5. Create the `tnsnames-ora` file. In the following example command, `HOST` should be set to the IP address of the Oracle Database and the `Server SID` should be set to the Oracle Database SID.

```
echo "# Connection string
AZACSNAP=\"(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (HOST=172.30.137.142) (PORT=1521)) (CONNECT_DATA=(SID=ORATST)))\"
" > $TNS_ADMIN/tnsnames.ora
```

6. Create the `sqlnet.ora` file.

```
echo "SQLNET.WALLET_OVERRIDE = TRUE
WALLET_LOCATION=(
    SOURCE=(METHOD=FILE)
    (METHOD_DATA=(DIRECTORY=\$TNS_ADMIN/.oracle_wallet))
) " > $TNS_ADMIN/sqlnet.ora
```

7. Test Oracle access using the wallet.

```
sqlplus /@AZACSNAP as SYSBACKUP
```

The expected output from the command:

```
[azacsnap@acao-ora01 ~]$ sqlplus /@AZACSNAP as SYSBACKUP

SQL*Plus: Release 19.0.0.0.0 - Production on Thu Sep 8 18:02:07 2022
Version 19.8.0.0.0

Copyright (c) 1982, 2019, Oracle. All rights reserved.

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.8.0.0.0

SQL>
```

## Configure ANF connectivity

This section explains how to enable communication with Azure NetApp Files (with a VM).

1. Within an Azure Cloud Shell session, make sure that you are logged into the subscription that you want to be associated with the service principal by default.

```
az account show
```

2. If the subscription isn't correct, use the following command:

```
az account set -s <subscription name or id>
```

3. Create a service principal using the Azure CLI as in the following example:

```
az ad sp create-for-rbac --name "AzAcSnap" --role Contributor --scopes /subscriptions/{subscription-id} --sdk-auth
```

The expected output:

```
{
  "clientId": "00aa000a-aaaa-0000-00a0-00aa000aaa0a",
  "clientSecret": "00aa000a-aaaa-0000-00a0-00aa000aaa0a",
  "subscriptionId": "00aa000a-aaaa-0000-00a0-00aa000aaa0a",
  "tenantId": "00aa000a-aaaa-0000-00a0-00aa000aaa0a",
  "activeDirectoryEndpointUrl": "https://login.microsoftonline.com",
  "resourceManagerEndpointUrl": "https://management.azure.com/",
  "activeDirectoryGraphResourceId": "https://graph.windows.net/",
  "sqlManagementEndpointUrl":
"https://management.core.windows.net:8443/",
  "galleryEndpointUrl": "https://gallery.azure.com/",
  "managementEndpointUrl": "https://management.core.windows.net/"
}
```

4. Cut and paste the output content into a file called `oracle.json` stored in the Linux user `azacsnap` user `bin` directory and secure the file with the appropriate system permissions.



Make sure the format of the JSON file is exactly as described above, especially with the URLs enclosed in double quotes (").

### Complete the setup of AzAcSnap tool

Follow these steps to configure and test the snapshot tools. After successful testing, you can perform the first database-consistent storage snapshot.

1. Change into the snapshot user account.

```
su - azacsnap
```

2. Change the location of commands.

```
cd /home/azacsnap/bin/
```

3. Configure a storage backup detail file. This creates an azacsnap.json configuration file.

```
azacsnap -c configure --configuration new
```

The expected output with three Oracle volumes:

```
[azacsnap@acao-ora01 bin]$ azacsnap -c configure --configuration new
Building new config file
Add comment to config file (blank entry to exit adding comments): Oracle
snapshot bkup
Add comment to config file (blank entry to exit adding comments):
Enter the database type to add, 'hana', 'oracle', or 'exit' (for no
database): oracle

=== Add Oracle Database details ===
Oracle Database SID (e.g. CDB1): ORATST
Database Server's Address (hostname or IP address): 172.30.137.142
Oracle connect string (e.g. /@AZACSNAP): /@AZACSNAP

=== Azure NetApp Files Storage details ===
Are you using Azure NetApp Files for the database? (y/n) [n]: y
--- DATA Volumes have the Application put into a consistent state before
they are snapshot ---
Add Azure NetApp Files resource to DATA Volume section of Database
configuration? (y/n) [n]: y
Full Azure NetApp Files Storage Volume Resource ID (e.g.
/subscriptions/.../resourceGroups/.../providers/Microsoft.NetApp/netAppA
ccounts/.../capacityPools/Premium/volumes/...): /subscriptions/0efa2dfb-
917c-4497-b56a-
b3f4eadb8111/resourceGroups/ANFAVSRG/providers/Microsoft.NetApp/netAppAc
counts/ANFAVSAacct/capacityPools/CapPool/volumes/acao-ora01-u01
Service Principal Authentication filename or Azure Key Vault Resource ID
(e.g. auth-file.json or https://...): oracle.json
Add Azure NetApp Files resource to DATA Volume section of Database
configuration? (y/n) [n]: y
Full Azure NetApp Files Storage Volume Resource ID (e.g.
/subscriptions/.../resourceGroups/.../providers/Microsoft.NetApp/netAppA
```

```

ccounts/.../capacityPools/Premium/volumes/...): /subscriptions/0efa2dfb-
917c-4497-b56a-
b3f4eadb8111/resourceGroups/ANFAVSRG/providers/Microsoft.NetApp/netAppAc
counts/ANFAVSAcct/capacityPools/CapPool/volumes/acao-ora01-u02
Service Principal Authentication filename or Azure Key Vault Resource ID
(e.g. auth-file.json or https://...): oracle.json
Add Azure NetApp Files resource to DATA Volume section of Database
configuration? (y/n) [n]: n
--- OTHER Volumes are snapshot immediately without preparing any
application for snapshot ---
Add Azure NetApp Files resource to OTHER Volume section of Database
configuration? (y/n) [n]: y
Full Azure NetApp Files Storage Volume Resource ID (e.g.
/subscriptions/.../resourceGroups/.../providers/Microsoft.NetApp/netAppA
ccounts/.../capacityPools/Premium/volumes/...): /subscriptions/0efa2dfb-
917c-4497-b56a-
b3f4eadb8111/resourceGroups/ANFAVSRG/providers/Microsoft.NetApp/netAppAc
counts/ANFAVSAcct/capacityPools/CapPool/volumes/acao-ora01-u03
Service Principal Authentication filename or Azure Key Vault Resource ID
(e.g. auth-file.json or https://...): oracle.json
Add Azure NetApp Files resource to OTHER Volume section of Database
configuration? (y/n) [n]: n

=== Azure Managed Disk details ===
Are you using Azure Managed Disks for the database? (y/n) [n]: n

=== Azure Large Instance (Bare Metal) Storage details ===
Are you using Azure Large Instance (Bare Metal) for the database? (y/n)
[n]: n

Enter the database type to add, 'hana', 'oracle', or 'exit' (for no
database): exit

Editing configuration complete, writing output to 'azacsnap.json'.

```

4. As the azacsnap Linux user, run the azacsnap test command for an Oracle backup.

```

cd ~/bin
azacsnap -c test --test oracle --configfile azacsnap.json

```

The expected output:

```
[azacsnap@acao-ora01 bin]$ azacsnap -c test --test oracle --configfile
azacsnap.json
BEGIN : Test process started for 'oracle'
BEGIN : Oracle DB tests
PASSED: Successful connectivity to Oracle DB version 1908000000
END   : Test process complete for 'oracle'
[azacsnap@acao-ora01 bin]$
```

5. Run your first snapshot backup.

```
azacsnap -c backup --volume data --prefix ora_test --retention=1
```

## Database migration from on-premises to Azure cloud

As a result of the Oracle decision to phase out single-instance databases, many organizations have converted single-instance Oracle databases to multitenant container databases. This enables the easy relocation of a subset of container databases called PDB to cloud with the maximum availability option, which minimize downtime during migration.

However, if you still have a single instance of a Oracle database, it can first be converted into a multitenant container database in place before attempting PDB relocation.

The following sections provide details for the migration of on-premises Oracle databases to Azure cloud in either scenarios.

### Converting a single instance non-CDB to a PDB in a multitenant CDB

If you still have a single-instance Oracle database, it must be converted into a multitenant container database whether you wish to migrate it to the cloud or not, because Oracle will stop supporting single-instance databases some time soon.

The following procedures plug a single instance database into a container database as a pluggable database or PDB.

1. Build a shell container database on the same host as the single-instance database in a separate `ORACLE_HOME`.
2. Shut down the single instance database and restart it in read-only mode.
3. Run the `DBMS_PDB.DESCRIBE` procedure to generate the database metadata.

```

BEGIN
  DBMS_PDB.DESCRIBE(
    pdb_descr_file => '/home/oracle/ncdb.xml');
END;
/

```

4. Shut down the single-instance database.
5. Start up the container database.
6. Run the `DBMS_PDB.CHECK_PLUG_COMPATIBILITY` function to determine whether the non-CDB is compatible with the CDB.

```

SET SERVEROUTPUT ON
DECLARE
  compatible CONSTANT VARCHAR2(3) :=
    CASE DBMS_PDB.CHECK_PLUG_COMPATIBILITY(
      pdb_descr_file => '/disk1/oracle/ncdb.xml',
      pdb_name       => 'NCDB')
    WHEN TRUE THEN 'YES'
    ELSE 'NO'
END;
BEGIN
  DBMS_OUTPUT.PUT_LINE(compatible);
END;
/

```

If the output is YES, then the non-CDB is compatible, and you can continue with the next step.

If the output is NO, then the non-CDB is not compatible, and you can check the `PDB_PLUG_IN_VIOLATIONS` view to see why it is not compatible. All violations must be corrected before you continue. For example, any version or patch mismatches should be resolved by running an upgrade or the `opatch` utility. After correcting the violations, run `DBMS_PDB.CHECK_PLUG_COMPATIBILITY` again to ensure that the non-CDB is compatible with the CDB.

7. Plug in the single instance non-CDB.

```

CREATE PLUGGABLE DATABASE ncdb USING '/home/oracle/ncdb.xml'
  COPY
  FILE_NAME_CONVERT = ('/disk1/oracle/dbs/', '/disk2/oracle/ncdb/')
;

```



If there is not sufficient space on the host, the `NOCOPY` option can be used to create the PDB. In that case, a single-instance non-CDB is not useable after plug in as a PDB because the original data files has been used for the PDB. Make sure to create a backup before the conversion so that there is something to fall back on if anything goes wrong.



8. Start with PDB upgrade after conversion if the version between the source single-instance non-CDB and the target CDB are different. For the same-version conversion, this step can be skipped.

```
sqlplus / as sysdba;
alter session set container=ncdb
alter pluggable database open upgrade;
exit;
dbupgrade -c ncdb -l /home/oracle
```

Review the upgrade log file in the `/home/oracle` directory.

9. Open the pluggable database, check for pdb plug-in violations, and recompile the invalid objects.

```
alter pluggable database ncdb open;
alter session set container=ncdb;
select message from pdb_plug_in_violations where type like '%ERR%' and
status <> 'RESOLVED';
$ORACLE_HOME/perl/bin/perl $ORACLE_HOME/rdbms/admin/catcon.pl -n 1 -c
'ncdb' -e -b utlrp -d $ORACLE_HOME/rdbms/admin utlrp.sql
```

10. Execute `noncdb_to_pdb.sql` to update the data dictionary.

```
sqlplus / as sysdba
alter session set container=ncdb;
@$ORACLE_HOME/rdbms/admin/noncdb_to_pdb.sql;
```

Shut down and restart the container DB. The ncdb is taken out of restricted mode.

## Migrate on-premises Oracle databases to Azure with PDB relocation

Oracle PDB relocation with the maximum-availability option employs PDB hot-clone technology, which allows source PDB availability while the PDB is copying over to the target. At switchover, user connections are redirected to the target PDB automatically. Thus, downtime is minimized independent of the size of the PDB. NetApp provides an Ansible-based toolkit that automates the migration procedure.

1. Create a CDB in the Azure public cloud on an Azure VM with the same version and patch level.
2. From the Ansible controller, clone a copy of the automation toolkit.

```
git clone https://github.com/NetApp-Automation/na_ora_aws_migration.git
```

3. Read the instruction in the README file.
4. Configure the Ansible host variable files for both the source and target Oracle servers and the DB server host's configuration file for name resolution.

5. Install the Ansible controller prerequisites on Ansible controller.

```
ansible-playbook -i hosts requirements.yml
ansible-galaxy collection install -r collections/requirements.yml
--force
```

6. Execute any pre-migration tasks against the on-premises server.

```
ansible-playbook -i hosts ora_pdb_relocate.yml -u admin -k -K -t
ora_pdb_relo_onprem
```



The admin user is the management user on the on-premises Oracle server host with sudo privileges. The admin user is authenticated with a password.

7. Execute Oracle PDB relocation from on-premises to the target Azure Oracle host.

```
ansible-playbook -i hosts ora_pdb_relocate.yml -u azureuser --private
-key db1.pem -t ora_pdb_relo_primary
```



The Ansible controller can be located either on-premises or in the Azure cloud. The controller needs connectivity to the on-premises Oracle server host and the Azure Oracle VM host. The Oracle database port (such as 1521) is open between the on-premises Oracle server host and the Azure Oracle VM host.

## Additional Oracle database migration options

Please see the Microsoft documentation for additional migration options: [Oracle database migration decision process](#).

## Copyright information

Copyright © 2025 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.