



New York Oracle User Group

# ORACLE®

## HA/DR for Oracle Database on Azure

Tim Gorman

Principal CSA, Oracle SME


Customer Success, Microsoft

# Agenda

- Prerequisites – Redundancy in Azure
  - Storage redundancy levels
  - Fault isolation zones
  - Azure Site Recovery
  - Azure Backup
- 3 of the 6 Pillars of Oracle Architecture
  - ~~Compute~~
  - ~~Storage~~
  - ~~Sizing Assessment~~
  - Data protection (*Backups*)
  - Service protection (*High Availability*)
  - Business continuity (*Disaster Resiliency*)

# Prerequisites - Redundancy in Azure

- Prerequisites – Redundancy in Azure
  - Fault isolation in Azure
  - Storage redundancy levels
  - Azure Backup
  - Azure Site Recovery

Three circular maps of the world are arranged on the left side of the slide. Each map shows a network of nodes and connections, with nodes represented by colored dots (blue, orange, green, purple) and lines representing network paths. The maps focus on different geographical areas: the top map shows North America and the Atlantic, the middle map shows Europe and Africa, and the bottom map shows Asia and Australia.

# Fault isolation in Azure

- *a **VM** which resides within*
  - *zero, one, or more **fault domains (availability set)** which resides within*
    - *one or more **data centers** which reside within*
      - *zero or three **Availability Zones** which reside within*
        - *a **Region** which resides within*
          - *a **Geography***

# Storage redundancy levels

LRS	ZRS	GRS	GZRS	RA-GRS	RA-GZRS
Blob	Blob	Blob	Blob	Blob	Blob
Queue	Queue	Queue	Queue	Queue	Queue
Table	Table	Table	Table	Table	Table
Files	Files	Files <i>(standard, &lt;5 TB)</i>	Files <i>(standard, &lt;5 TB)</i>		
Managed disk	Managed disk				

## Managed Disk

- Standard HDD, Standard SSD
  - Not suitable for any Oracle Database files
- Premium SSD
  - Suitable for all Oracle Database file usages
- UltraDisk
  - Suitable mostly for Oracle online redo log files

## Blob

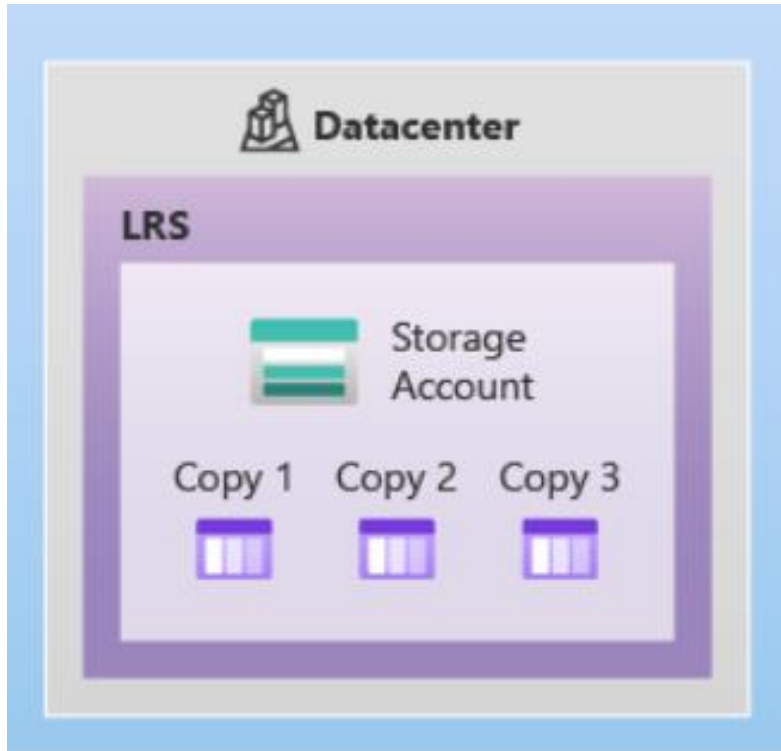
- **Standard GPv2** or **premium block-blob storage** with *hierarchical* storage namespace for fileshare mounted on NFS v3.0
  - Suitable for archived redo log files or RMAN backupset files

## Files

- Azure Files standard fileshare mounted on SMB (CIFS)
  - Not suitable for Oracle datafiles, tempfiles, controlfiles, online redo log files
  - Suitable for archived redo log files or backup set files
- Azure Files premium fileshare mounted on NFS v4.1
  - Suitable for all Oracle Database files
    - Not yet compatible with Oracle direct NFS (dNFS)

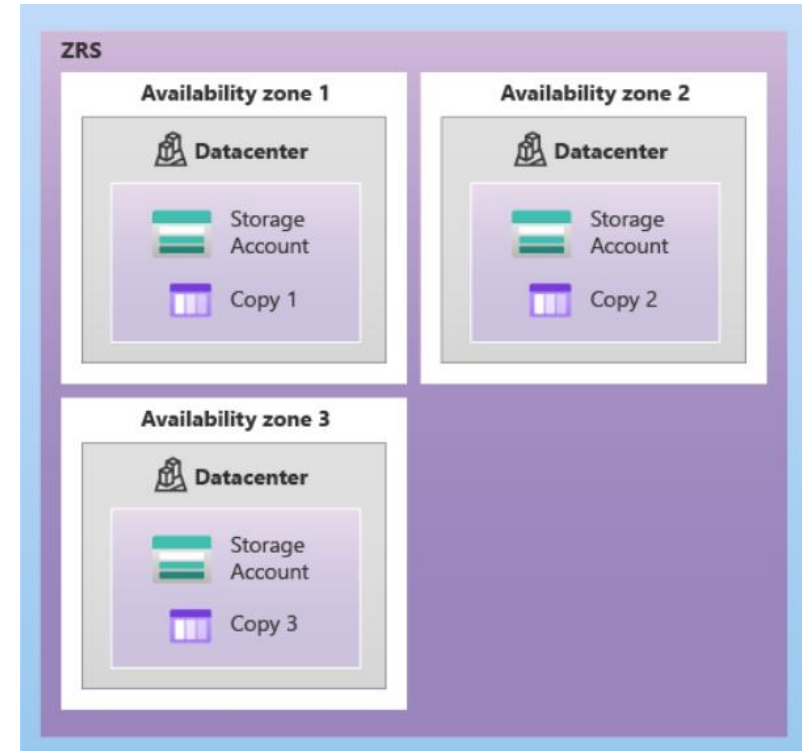
# Storage redundancy levels

- Locally Redundancy Storage (LRS)



- All 3 copies updates synchronously
- Durability: 99.99999999999999% (16 nines)
- Availability: 99.9%

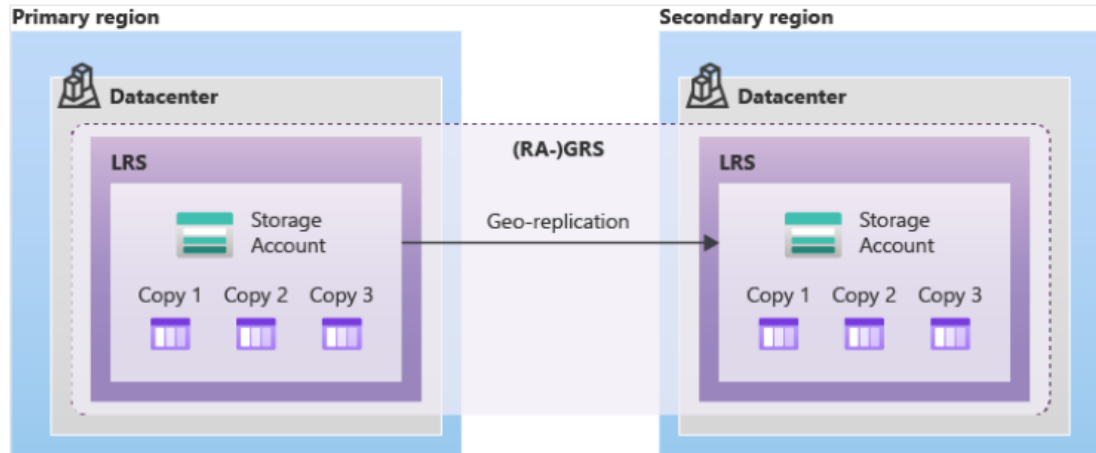
- Zone Redundancy Storage (ZRS)



- All 3 copies updates synchronously
- Durability: 99.99999999999999% (16 nines)
- Availability: 99.9%

# Storage redundancy levels

- Geo-Redundant Storage (LRS)

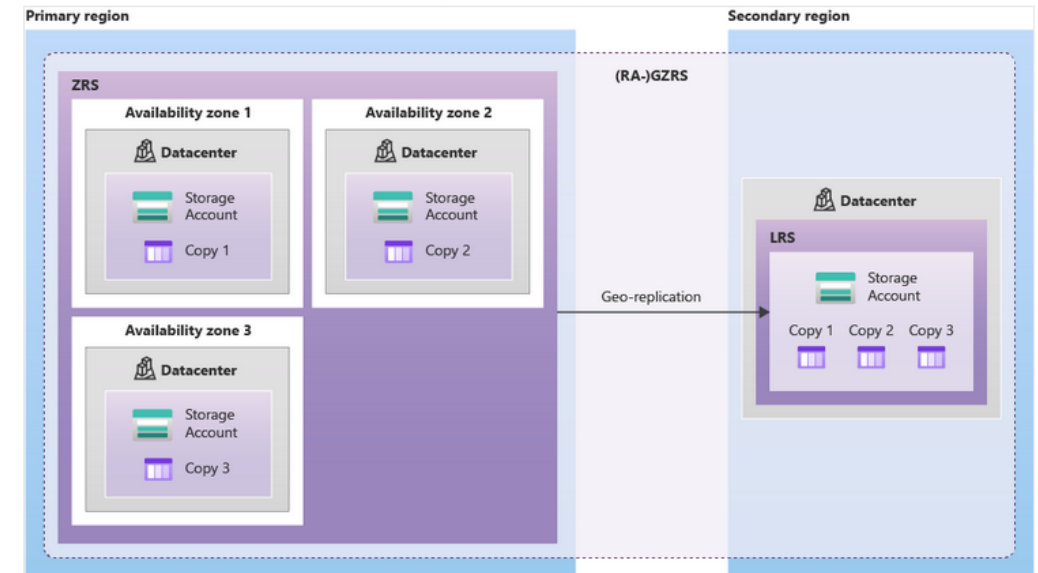


- Local LRS asynchronously replicated to remote LRS
  - 6 copies total
- Durability: 99.99999999999999% (16 nines)
- Availability: 99.9%

Data in the secondary region...

- **GRS** and **GZRS** storage is available for read- or write-access only after a **failover**
- **RA-GRS** and **RA-GZRS** storage is available for read-access at **all times**

- Geo Zone Redundancy Storage (GZRS)

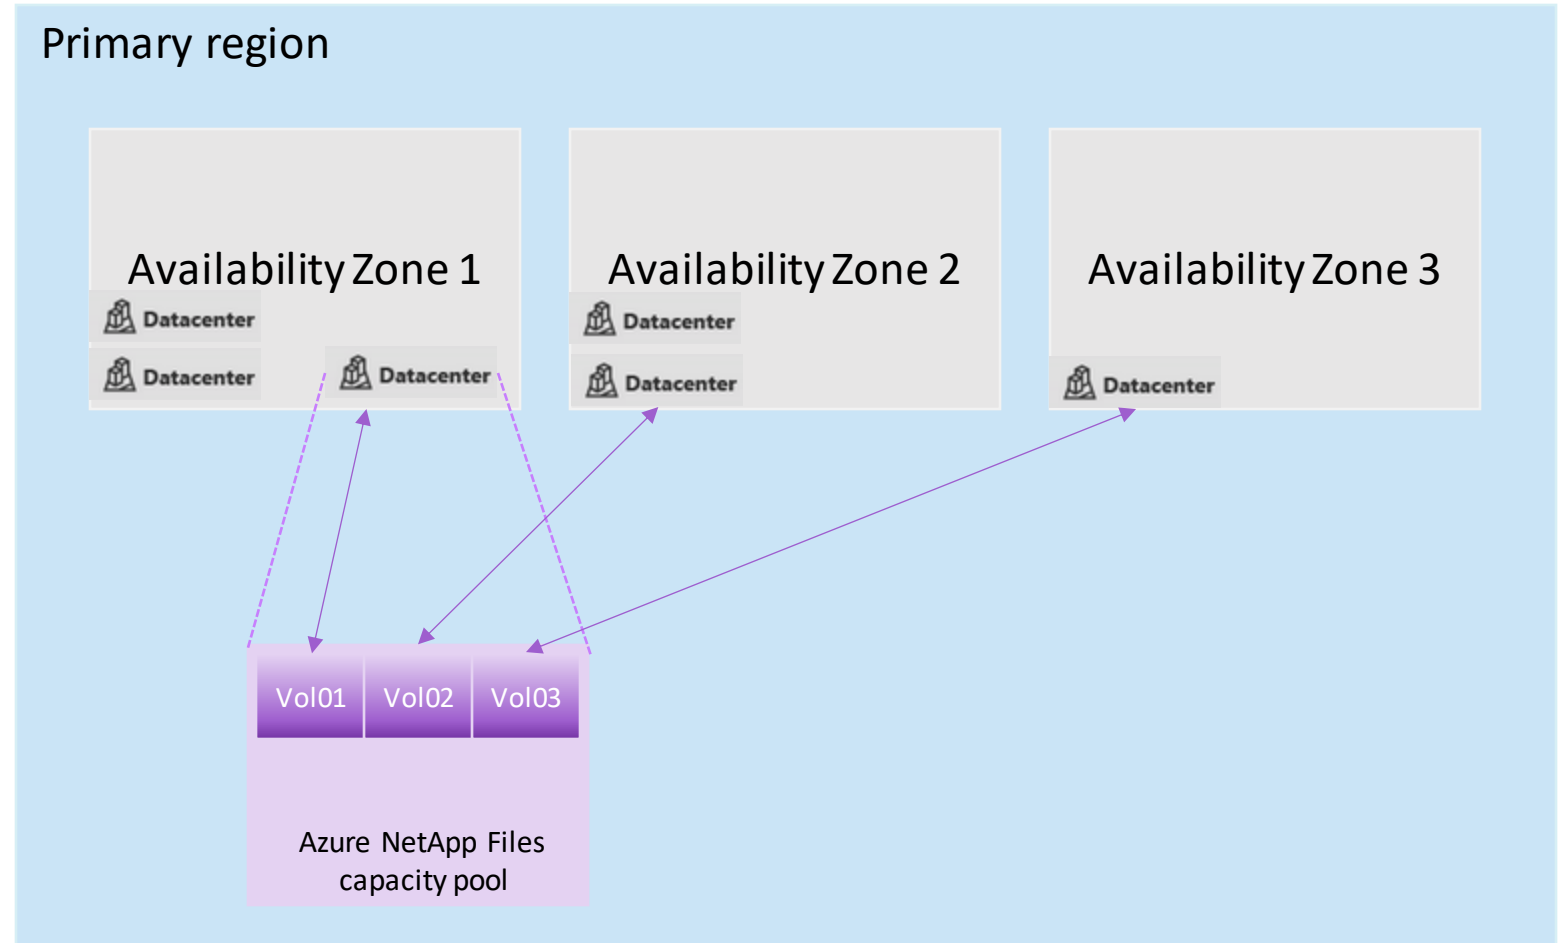


- Local ZRS asynchronously replicated to remote LRS
  - 6 copies total
- Durability: 99.99999999999999% (16 nines)
- Availability: 99.9%

# Storage redundancy levels

## Azure NetApp Files

- Licensed on NetApp ONTAP technology
- *Native* Azure storage engineered and supported by Microsoft
- Based on specialized hardware within one (or more) data centers within a region
  - No awareness of availability zones....
    - Not yet, but on the roadmap

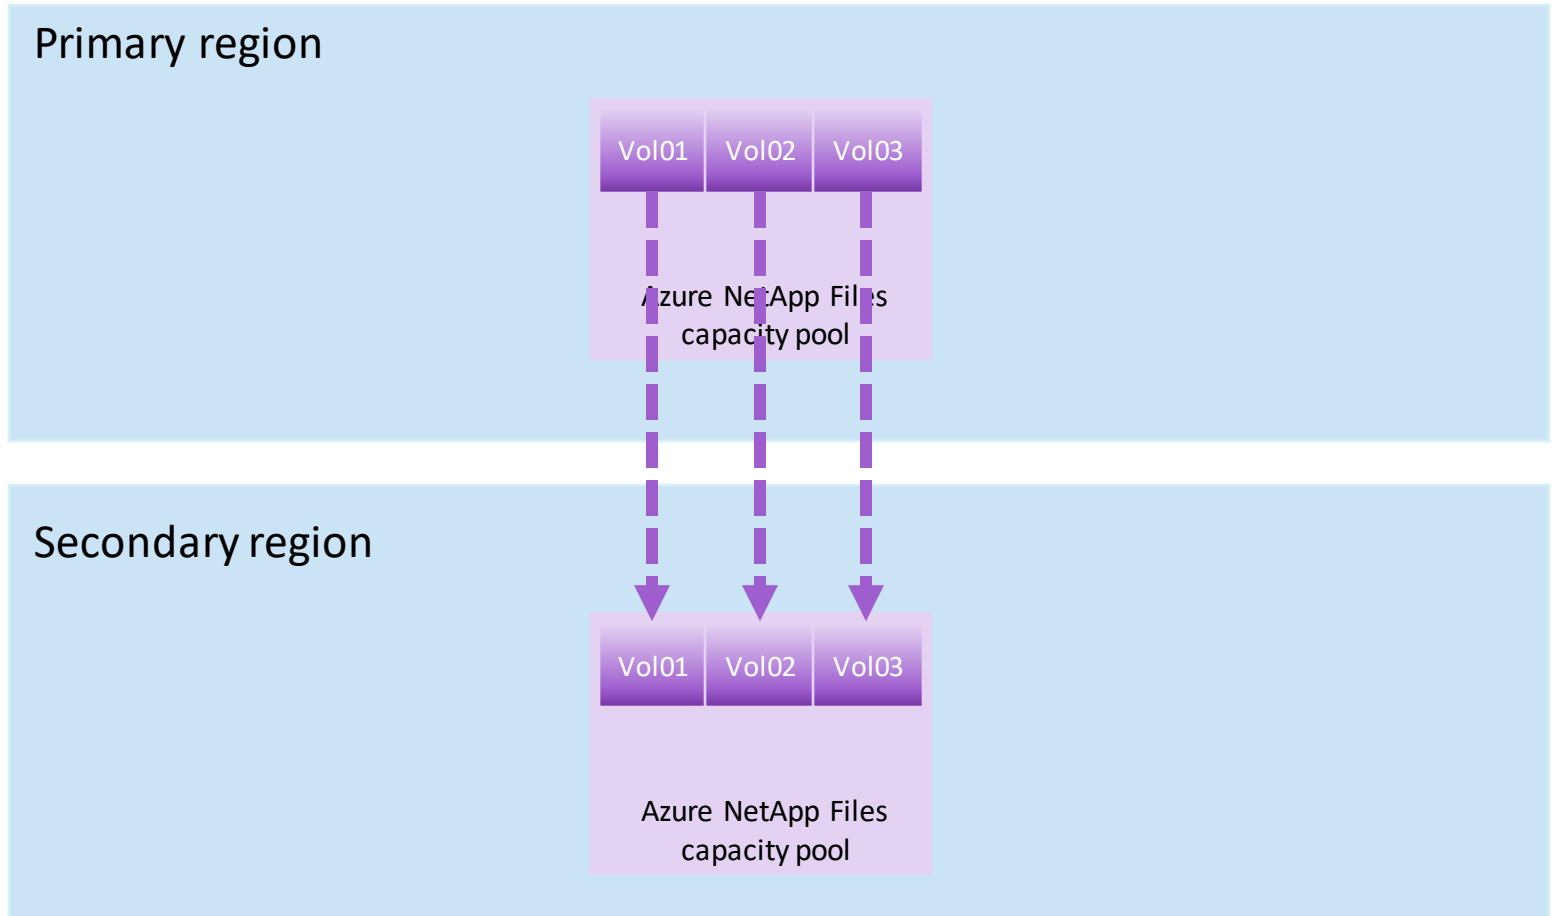




# Storage redundancy levels


## Azure NetApp Files

- Cross-region replication
- After zonal awareness, also cross-zonal replication will be available



# Azure Backup

- Daily snapshots of all attached managed disk
  - Images are incremental
  - Retained locally for configured period (*default: 2 days*)
  - UltraDisk not included
- Snapshot images are copied to an immutable Recovery Services Vault storage account
  - Retained in vault for configured period (*default: 30 days*)
- Restore actions...
  1. Restore to VM (original VM or new VM)
  2. Restore individual disk images





# Azure Site Recovery

---

- Near real-time continuous asynchronous copy of all storage attached to VM, to a geo-redundant storage account
- Upon site failure, VMs are instantiated in secondary region from the latest image in geo-redundant storage account
- The current limit for per virtual machine data churn is 54 MB/s, regardless of size, per [HERE](#)
  - Roadmap includes raising limit to 100 MB/s

<https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-overview>

# 3 of the 6 Pillars of Oracle architecture

- ~~1. Compute~~ *(covered in previous presentation)*
- ~~2. Storage~~ *(covered in previous presentation)*
- ~~3. Sizing Assessment~~ *(covered in previous presentation)*
4. Data protection (*Backups*)
5. Service protection (*High Availability*)
6. Business continuity (*Disaster Resiliency*)

# Data Protection (Backups)

Two possible types of database backup...

1. **Streaming backups** using Oracle RMAN (Recovery Manager)
  - No dependency on type of storage
  - Based on the concept of backups being streamed to sequential tape media
  - Two destinations for RMAN backupsets...
    - Virtual tape libraries in Azure Marketplace (CommVault, Veeam, Veritas, etc)
    - Fileshares (Azure Blob NFS, Azure Files, Azure NetApp Files)
2. **Storage-level snapshots**
  - Dependent on the type of storage used for database files
    - Managed disk (*premium SSD*): Azure Backup integrated with Oracle Database
    - Azure NetApp Files: ANF Backup, AzAcSnap, and cross-region replication
    - SILK: scripting using SILK APIs

# Data Protection – streaming backups

- RMAN backs up datafiles to backup media
  - sending backupset files or datafile copies to either...
    - virtual tape library managers like *CommVault*, *Veeam*, *EMC*, or *Veritas*
      - Syntax: DEVICE TYPE SBT
    - Filesystem files (local-attach or remote fileshare)
      - Syntax: DEVICE TYPE DISK
- Three types of backups are available
  - FULL or LEVEL-0 INCREMENTAL backup the entire database
  - LEVEL-1 INCREMENTAL backup just changes since last LEVEL-0 or LEVEL-1
  - ARCHIVELOG backup the archived redo log files
  - Each type can also include control files and parameter files in the backupset

# Data Protection – streaming backups


Backup media	Attach protocol	Write Throughput limit (MBps)	Doc link detailing throughput expectations	Comments
Managed disk	SCSI	4000*	<a href="#">LINK</a>	Use LVM striping for highest throughput
Blob NFS	NFS v3.0	1000*	<a href="#">LINK</a>	
Azure Files standard	SMB 3.1 (CIFS)	60 (300)	<a href="#">LINK</a>	<a href="#">Automated backups</a>
Azure Files premium	NFS v4.1	1200	<a href="#">LINK</a>	
Azure NetApp Files	NFS v3.0 or v4.1	1733	<a href="#">LINK</a>	<a href="#">Automated backups</a>
blobfuse	FUSE libraries	20-50*		Open-source on GitHub supported by Azure Storage team

\* = observed, not documented - YMMV

# Data Protection – streaming backups


- The problem with blobfuse

- Fuse protocol uses TMP as cache storage
  - Reads are cached, writes cannot be cached
  - Writes must be written through



- Compared to *normal* fileshares

- NFS or SMB (CIFS on Linux) protocols





# Data Protection – streaming backups

## **Recovery Point Objective (RPO)**

- RPO is the tolerance for data loss recovering from a failure
- Restored backups can be recovered to the point-in-time of the latest backed-up archived redo log file
  - So RPO=0 might be achievable in certain rare situations, but can't be promised
  - Setting of Oracle parameter ARCHIVE\_LAG\_TARGET represents a way to choose the number of seconds of data lost during restore/recovery

## **Recovery Time Objective (RTO)**

- RTO is the expectation for time to return-to-service after a failure
- When the time to fully restore a database using RMAN exceeds RTO
  - Then it is time to consider replacing RMAN as the method of data protection...

# Data Protection – storage-level snapshots

Database stored on managed disk...

- [Azure Backup](#)

Database stored on Azure NetApp Files...

- [Azure NetApp Files Backup](#)

Database stored on SILK...

- Volume snapshot and replication using REST API for SILK

# Data Protection – Azure Backup

## Azure Backup is integrated with Oracle Database

- Azure VM Agent runs within Linux VM
  - Runs *pre-script* before snapshot
    - Flushes transaction logs to archive, and then enters database into BACKUP mode
  - Runs *post-script* after snapshot
    - Reverts database from BACKUP mode, and then flushes transaction logs to archive
- Pricing for integrated Oracle Database backups [HERE](#)
  - Use pricing in the tab **Azure VMs and on-premises servers**
- Step-by-step documentation for setup [HERE](#)
  - Brief videos explaining setup available from GitHub [HERE](#)
    - Third of three videos depicts database restore in a new VM
  - Diagnostic verification script available from GitHub [HERE](#)

# Data Protection – Azure NetApp Files

## [Understand Azure NetApp Files Backup](#)

- Currently in preview – [sign-up form](#)
- Snapshots of volumes are captured via manual or scheduled backups
  - [How Azure NetApp Files snapshots work](#)
  - Backups can be scheduled according to Azure policies
- Snapshot images are copied to immutable Azure storage accounts using ZRS storage
  - Snapshots are retained locally (*default: 5 most-recent images*)
  - Snapshots images are retained in storage account for longer-term restore (*default: 30 most-recent images*)

# Data Protection - Backups

- While a database is in “backup mode” after ALTER DATABASE BEGIN BACKUP has executed, normal database operations proceed as usual, with two exceptions...
  1. A six-byte field in the header of each datafile, holding an SCN (system change number) value for the “latest change number” is frozen until END BACKUP is executed
  2. The entire block is written to the redo stream on the first change after BEGIN BACKUP, instead of the just the changed bytes

# Data Protection - Backups

- As a result, the volume of redo generated during the time the database is in “backup mode” can greatly increase
  - But with storage-level snapshots, the time spent in “backup mode” is seconds
- The purpose of the LAST\_CHANGE# field in the datafile headers affects roll-forward recovery after datafile restore
  - Recovery from archived redo begins in each datafile for all redo generated after LAST\_CHANGE#
  - While the database is in “backup mode”, that start point is the same for all datafiles, which is the moment BEGIN BACKUP ran

# Data Protection – Azure Backup

- Fully-automated *accelerator* scripts in bash using Azure CLI
  - cr\_oravm.sh in Github ([HERE](#))
    - Script to automate the creation of a single VM with a configured database running Oracle Enterprise Linux (OEL)
      - specify version of Oracle Database including 12.1, 12.2, 18.3, and 19.3 from the Azure Marketplace with sample schemas enabled
      - specify either *premium SSD* storage or Azure NetApp Files *premium* service-level storage
      - Azure Backup integrated with Oracle Database already configured
      - Azure Files standard fileshare using CIFS protocol configured for archived redo log files


# Service protection (High Availability)

- Architectural redundancy for the purpose of minimizing downtime from service outage
  - Planned outages
    - Patching, updates, upgrades
      - Rolling upgrades
  - Unplanned outage
    - Failures



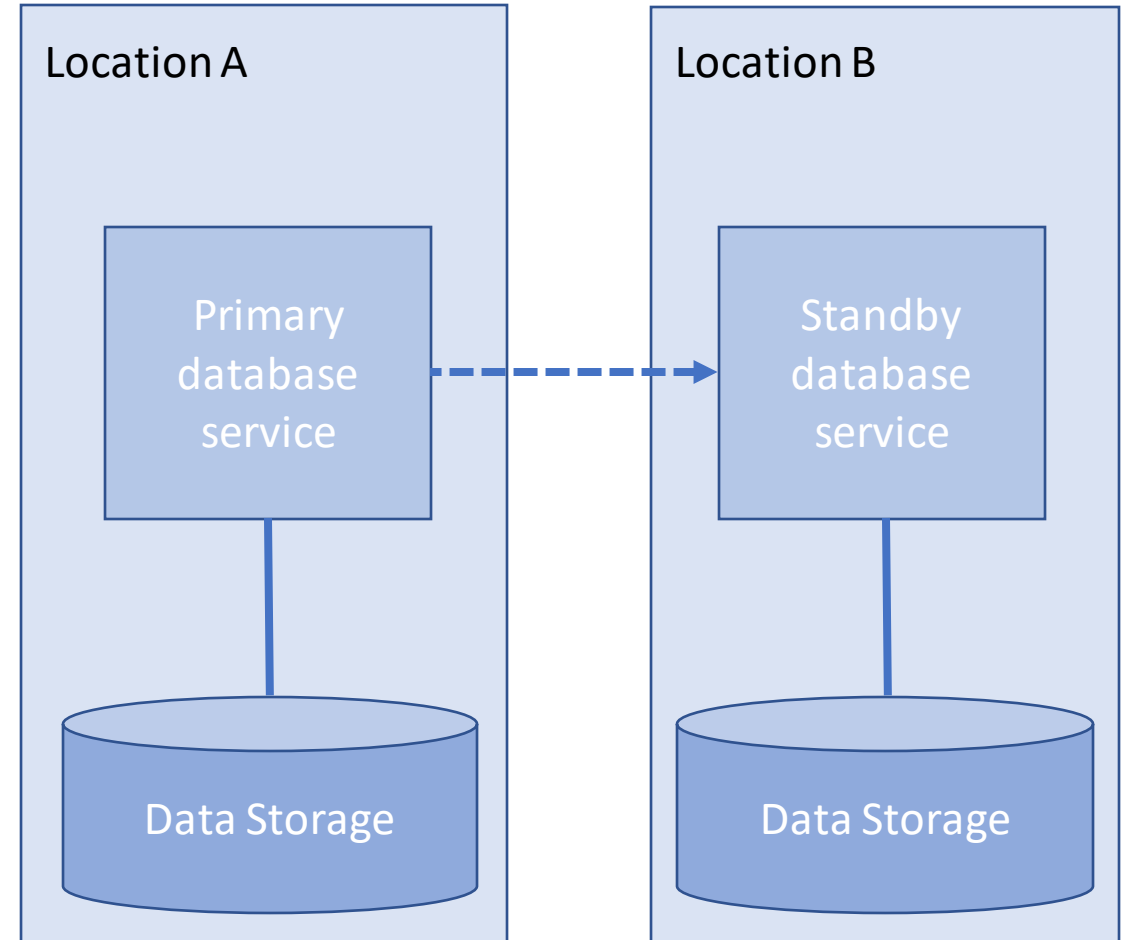
# Service protection (High Availability)

Clusters



*Two possible architectures*

Replication



# Service protection (High Availability)

- **Clusters**

- *Active/passive*
  - Linux Pacemaker/Corosync (PCS)
  - Oracle RAC OneNode
- *Active/active*
  - Oracle RAC

- **Replication**

- *Active/passive*
  - Oracle DataGuard
  - Single-master
    - Oracle GoldenGate, Quest SharePlex, HVR, Striim, etc
  - Azure Site Recovery
- *Active/active*
  - Multi-master
    - Oracle GoldenGate, Quest SharePlex, etc

# Service protection (High Availability)

- **Clusters**

- *Active/passive*
  - Linux Pacemaker/Corosync (PCS)
  - Oracle RAC OneNode
- *Active/active*
  - Oracle RAC

- **Replication**

- *Active/passive*
  - Oracle DataGuard
  - Single-master
    - Oracle GoldenGate, Quest SharePlex, HVR, Striim, etc
  - Azure Site Recovery
- *Active/active*
  - Multi-master
    - Oracle GoldenGate, Quest SharePlex, etc

# Service protection (High Availability)

Unplanned outage scenario	RMAN	RAC	Linux PCS	DataGuard	GoldenGate
Database software failure	0	5	4	5	4
Datafile or database corruption	1	0	0	3	3
Host failure or OS panic	0	5	4	5	4
TOR (top of rack) network failure	0	5	4	5	4
Data center failure	1	0	0	5	4
<b>TOTAL</b>	<b>2</b>	<b>15</b>	<b>12</b>	<b>23</b>	<b>18</b>

Value	Description
0	Cannot protect at all
1	Restore service within 4 hours
2	Restore service within 1 hour
3	Restore service within 10 minutes, and all database sessions will need to reconnect
4	Restore service within 2 minutes, and all database sessions will need to reconnect
5	Restore service within 2 minutes, including no interruption of connected database sessions

# Service protection (High Availability)

- Most Azure customers use DataGuard for service protection
  - Non-mission-critical workloads, MAX PERFORMANCE is sufficient
    - *No guarantees, just best effort, and least impact to the primary database*
    - RPO: Data-loss less than 2-5 minutes is very possible, RPO=2-5 mins
    - Performance of the primary database is given the highest priority
      - above data protection and service availability
  - Mission-critical workloads, MAX AVAILABILITY is most common
    - RPO: Data-loss is possible in a few unlikely corner situations
    - Availability of the database service is given the highest priority
      - above data protection and performance
  - MAX PROTECTION does not permit data-loss under any circumstances
    - Protection of committed transactions is given the highest priority
      - above availability and performance


# Service protection (High Availability)

- Fully-automated *accelerator* scripts in bash using Azure CLI
  - cr\_oradg.sh in GitHub ([HERE](#))
    - Specify version of Oracle Database including 12.1, 12.2, 18.3, and 19.3 from the Azure Marketplace with sample schemas enabled on three VMs...
      - one VM in AZ1 with a configured primary database running Oracle Enterprise Linux (OEL)
      - one VM in AZ2 with a configured standby database on OEL
      - one VM in AZ3 with Oracle DataGuard Broker running as an “observer” for FSFO
  - cr\_orapcs.sh in GitHub ([HERE](#))
    - Specify version of Oracle Database including 12.1, 12.2, 18.3, and 19.3 from the Azure Marketplace with sample schemas enabled on three VMs...
      - two VMs within an availability set, also within a proximity placement group, attached to shared premium SSD, running Oracle Enterprise Linux (OEL)
      - Pacemaker/Corosync managing cluster resources for fencing, network, storage, database, and database listener
      - one additional VM running OEL intended as a “jump box” for launching tests

# Service protection (High Availability)


- Database *connect-strings* can contain instructions to reconnect after a failover
  - [Retrying a connection using FAILOVER MODE syntax in connect-string](#)
- Most often, the important thing is getting the application to disconnect first
  - Even when the network connection drops, database connections may not time out for minutes
- Seamless application failover requires configuring database clients to use ONS
  - Oracle Notifications Services (ONS) which is part of Oracle Grid Infrastructure (GI)
    - GI is normally part of RAC, but a standalone version called **Oracle Restart** is available with non-RAC DataGuard
  - Most Oracle clients (i.e. SQL\*Plus, DataPump, SQL\*Loader, etc) are already instrumented for ONS
    - Most JDBC applications are not
  - Both RAC and DataGuard use the same tools and methods to keep *database sessions* alive across a failover
    - [Oracle OpenWorld 2010 presentation: Seamless Application Failover with Oracle DataGuard](#)
    - [Oracle Support note #1429223.1: How to Configure Client Failover For Data Guard Connections Using Database Services](#)

# Service protection (High Availability)






# Service protection (High Availability)




# Service protection (High Availability)




3. Network connections also remain, waiting for timeout

# Service protection (High Availability)



4. After database failover or switchover completes, ONS on new primary notifies the client-side FAN event handler to implement reconnect

# Service protection (High Availability)



5. Oracle Net uses `FAILOVER_MODE` clause in TNS connect-string to reconnect

# Service protection (High Availability)

*“Zero downtime on maintenance operations with RAC”*

An **overstatement** for RAC and an **understatement** for DataGuard


- Oracle RAC is **incapable** of facilitating zero downtime using the technique of rolling upgrades on **all** maintenance operations
  - Rolling upgrade is possible for patches to HW, OS, and some RDBMS patches
  - But rolling upgrade is not possible for all RDBMS patches or any application patches

...in contrast...

- Since Oracle12c, Oracle DataGuard has been capable of facilitating zero downtime using the technique of rolling upgrades on **all** maintenance operations
  - Including HW, OS, RDBMS, and application patches

# Rolling software patching with RAC


Clustered solutions



## Rolling software-only patch in RAC


1. Stop one node
2. Apply patch to stopped node
3. Restart node
4. Stop other node
5. Apply patch to stopped node
6. Restart node

Replicated solutions



# Rolling software patching with RAC


Clustered solutions



## Rolling software-only patch in RAC


1. Stop one node
2. Apply patch to stopped node
3. Restart node
4. Stop other node
5. Apply patch to stopped node
6. Restart node

Replicated solutions



# Rolling software patching with RAC


Clustered solutions



## Rolling software-only patch in RAC

1. Stop one node
2. Apply patch to stopped node
3. Restart node
4. Stop other node
5. Apply patch to stopped node
6. Restart node


Replicated solutions





# Rolling software patching with RAC


Clustered solutions




## Rolling software-only patch in RAC

1. Stop one node
2. Apply patch to stopped node
3. Restart node
4. Stop other node
5. Apply patch to stopped node
6. Restart node

Replicated solutions




g with RAC



# Rolling software patching with RAC


Clustered solutions



Rolling software-only patch in RAC


1. Stop one node
2. Apply patch to stopped node
3. Restart node
4. Stop other node
5. Apply patch to stopped node
6. Restart node

Replicated solutions



# Rolling software patching with RAC


Clustered solutions



## Rolling software-only patch in RAC


1. Stop one node
2. Apply patch to stopped node
3. Restart node
4. Stop other node
5. Apply patch to stopped node
6. Restart node

Replicated solutions



# Rolling software patching with RAC


Clustered solutions



## Rolling software-only patch in RAC

1. Stop one node
2. Apply patch to stopped node
3. Restart node
4. Stop other node
5. Apply patch to stopped node
6. Restart node

Replicated solutions




# Rolling software patching with RAC

- RAC is fine when patching or upgrading non-shared resources...
  - OS software
  - Oracle Grid Infrastructure software
  - Oracle RDBMS software
  - Application software or middleware
- RAC is not OK when patching or upgrading shared resources...
  - Oracle RDBMS database patches or upgrades
  - Application data patches or upgrades

...as demonstrated next...

# Rolling database patches or upgrades with RAC


Clustered solutions



Rolling software and database patch/upgrade in RAC


1. Stop one node
2. Apply patch to stopped node
3. Restart node
4. Stop other node
5. Apply patch to stopped node
6. Restart node

Replicated solutions



# Rolling database patches or upgrades with RAC


Clustered solutions



Rolling software and database patch/upgrade in RAC


1. Stop one node
2. Apply patch to stopped node
3. Restart node
4. Stop other node
5. Apply patch to stopped node
6. Restart node

Replicated solutions



# Rolling database patches or upgrades with RAC


Clustered solutions



Rolling software and database patch/upgrade in RAC

1. Stop one node
2. Apply patch to stopped node
3. Restart node
4. Stop other node
5. Apply patch to stopped node
6. Restart node


Replicated solutions





# Rolling database patches or upgrades with RAC

Clustered solutions




Rolling software and database patch/upgrade in RAC

1. Stop one node
2. Apply patch to stopped node
3. Restart node
4. Stop other node
5. Apply patch to stopped node
6. Restart node

**The unpatched/non-upgraded software on node 2 is no longer able to work with the patched/updated database contents**

Replicated solutions




# Rolling patches or upgrades with DataGuard

- Rolling patches or upgrades with DataGuard employ a technique called *transient logical standby*
  - Temporarily converts a *physical standby database* into a *logical standby database*
    - Not possible to open a physical standby database for read-write activity
      - But a logical standby database is inherently open for read-write activity
  - Apply patch or upgrade on the open logical standby database while it is also being updated with transactions from the primary database

# Rolling patches or upgrades with DataGuard


Clustered solutions



## Rolling software and database patch/upgrade in DataGuard


1. Convert physical standby to logical standby
2. Apply patch to logical standby
3. Revert logical standby to physical standby
4. Switchover
5. Convert physical standby to logical standby
6. Apply patch to logical standby
7. Revert logical standby to physical standby
8. Switchover

Replicated solutions



# Rolling patches or upgrades with DataGuard


Clustered solutions



## Rolling software and database patch/upgrade in DataGuard


1. Convert physical standby to logical standby
2. Apply patch to logical standby
3. Revert logical standby to physical standby
4. Switchover
5. Convert physical standby to logical standby
6. Apply patch to logical standby
7. Revert logical standby to physical standby
8. Switchover

Replicated solutions



# Rolling patches or upgrades with DataGuard


Clustered solutions



Rolling software and database patch/upgrade in DataGuard


1. Convert physical standby to logical standby
2. Apply patch to logical standby
3. Revert logical standby to physical standby
4. Switchover
5. Convert physical standby to logical standby
6. Apply patch to logical standby
7. Revert logical standby to physical standby
8. Switchover

Replicated solutions



# Rolling patches or upgrades with DataGuard


Clustered solutions



## Rolling software and database patch/upgrade in DataGuard


1. Convert physical standby to logical standby
2. Apply patch to logical standby
3. Revert logical standby to physical standby
4. Switchover
5. Convert physical standby to logical standby
6. Apply patch to logical standby
7. Revert logical standby to physical standby
8. Switchover

Replicated solutions



# Rolling patches or upgrades with DataGuard


Clustered solutions



## Rolling software and database patch/upgrade in DataGuard


1. Convert physical standby to logical standby
2. Apply patch to logical standby
3. Revert logical standby to physical standby
4. Switchover
5. Convert physical standby to logical standby
6. Apply patch to logical standby
7. Revert logical standby to physical standby
8. Switchover

Replicated solutions



# Rolling patches or upgrades with DataGuard


Clustered solutions



## Rolling software and database patch/upgrade in DataGuard

1. Convert physical standby to logical standby
2. Apply patch to logical standby
3. Revert logical standby to physical standby
4. Switchover
5. Convert physical standby to logical standby
6. Apply patch to logical standby
7. Revert logical standby to physical standby
8. Switchover


Replicated solutions





# Rolling patches or upgrades with DataGuard


Clustered solutions



Rolling software and database patch/upgrade in DataGuard


1. Convert physical standby to logical standby
2. Apply patch to logical standby
3. Revert logical standby to physical standby
4. Switchover
5. Convert physical standby to logical standby
6. Apply patch to logical standby
7. Revert logical standby to physical standby
8. Switchover

Replicated solutions



# Rolling patches or upgrades with DataGuard


Clustered solutions



Rolling software and database patch/upgrade in DataGuard


1. Convert physical standby to logical standby
2. Apply patch to logical standby
3. Revert logical standby to physical standby
4. Switchover
5. Convert physical standby to logical standby
6. Apply patch to logical standby
7. Revert logical standby to physical standby
8. Switchover

Replicated solutions



# Rolling patches or upgrades with DataGuard


Clustered solutions



Rolling software and database patch/upgrade in DataGuard

1. Convert physical standby to logical standby
2. Apply patch to stopped node
3. Revert logical standby to physical standby
4. Switchover
5. Convert physical standby to logical standby
6. Apply patch to logical standby
7. Revert logical standby to physical standby
8. Switchover

Replicated solutions



# Service protection (High Availability)

- Oracle has forced customers running Oracle on public clouds to the best database service protection, which is DataGuard

# Business continuity (Disaster resiliency)

- Architectural redundancy for the purpose of minimizing downtime from multiple component or from geography-wide failures

# Business continuity (Disaster resiliency)

- Oracle DataGuard
  - Most common option chosen
    - MAX PERFORMANCE mode
    - MAX AVAILABILITY mode with Far Sync enabled
- Oracle GoldenGate
  - Next most common DR option chosen
- Azure Site Recovery
  - Due to 54 MBps limit on cumulative VM data changes, it is unlikely that ASR can be used for any database workloads
    - Not just Oracle, but SQL Server, PostgreSQL, MySQL, etc
  - Application tier VMs fit under the limit well

# Business continuity (Disaster resiliency)

- Fully-automated *accelerator* script in bash using Azure CLI
  - cr\_oradg.sh in GitHub ([HERE](#))
    - Specify version of Oracle Database including 12.1, 12.2, 18.3, and 19.3 from the Azure Marketplace with sample schemas enabled on three VMs...
      - one VM in AZ1 with a configured primary database running Oracle Enterprise Linux (OEL)
      - one VM in AZ2 with a configured standby database on OEL
      - one VM in AZ3 with Oracle DataGuard Broker running as an “observer” for FSFO



# References referenced

[Azure Backup overview](#)

[Azure Site Recovery overview](#)

[Oracle RMAN – Getting Started](#)

[Oracle DataGuard – Getting Started](#)

[Oracle GoldenGate – Data Sheet](#)

[Oracle RAC architecture](#)

[Azure Evs\\_v5 instance types](#)

[Azure Blob NFS comparison](#)

[Azure Files – performance scale targets](#)

[Azure Files share backups - Overview](#)

[Azure NetApp Files Calculator](#)

[Azure NetApp Files Backups - Introduction](#)

[Azure NetApp Files Backup preview signup](#)

[How Azure NetApp Files snapshots work - Introduction](#)

[Azure Backups pricing](#)

[Azure Backups integrated with Oracle – Step By Step](#)

[Azure Backups integrated with Oracle – training videos](#)

[Azure Backups integrated with Oracle – diagnostic script](#)

[GitHub – accelerator script to create a VM with Oracle database and Oracle Backup integrated](#)

[GitHub – accelerator script to create a DataGuard FSFO cluster with DataGuard Broker observer](#)

[GitHub – accelerator script to create a Pacemaker/Corosync HA cluster with Oracle database](#)

[TNS syntax to retry session reconnect after failover](#)

[Oracle OpenWorld 2010 presentation: Seamless Application Failover with Oracle DataGuard](#)

[Oracle Support note #1429223.1: \*How to Configure Client Failover For Data Guard Connections Using Database Services\*](#)





# New York Oracle User Group

## Q & A

Tim Gorman  
Principal CSA – Azure Core  
Customer Success Architecture and Engineering

[tigorman@microsoft.com](mailto:tigorman@microsoft.com)   
[linkedin.com/in/timgorman](https://www.linkedin.com/in/timgorman)   
[@TimGormanTech](https://twitter.com/TimGormanTech) 