Oracle on Azure IaaS

Kellyn Gorman

SME Oracle on Azure, (Principal CSA)

Cloud Architecture and Engineering Team at Microsoft

The Path to the Cloud- Most data estates are tightly coupled on-premises and must moved together

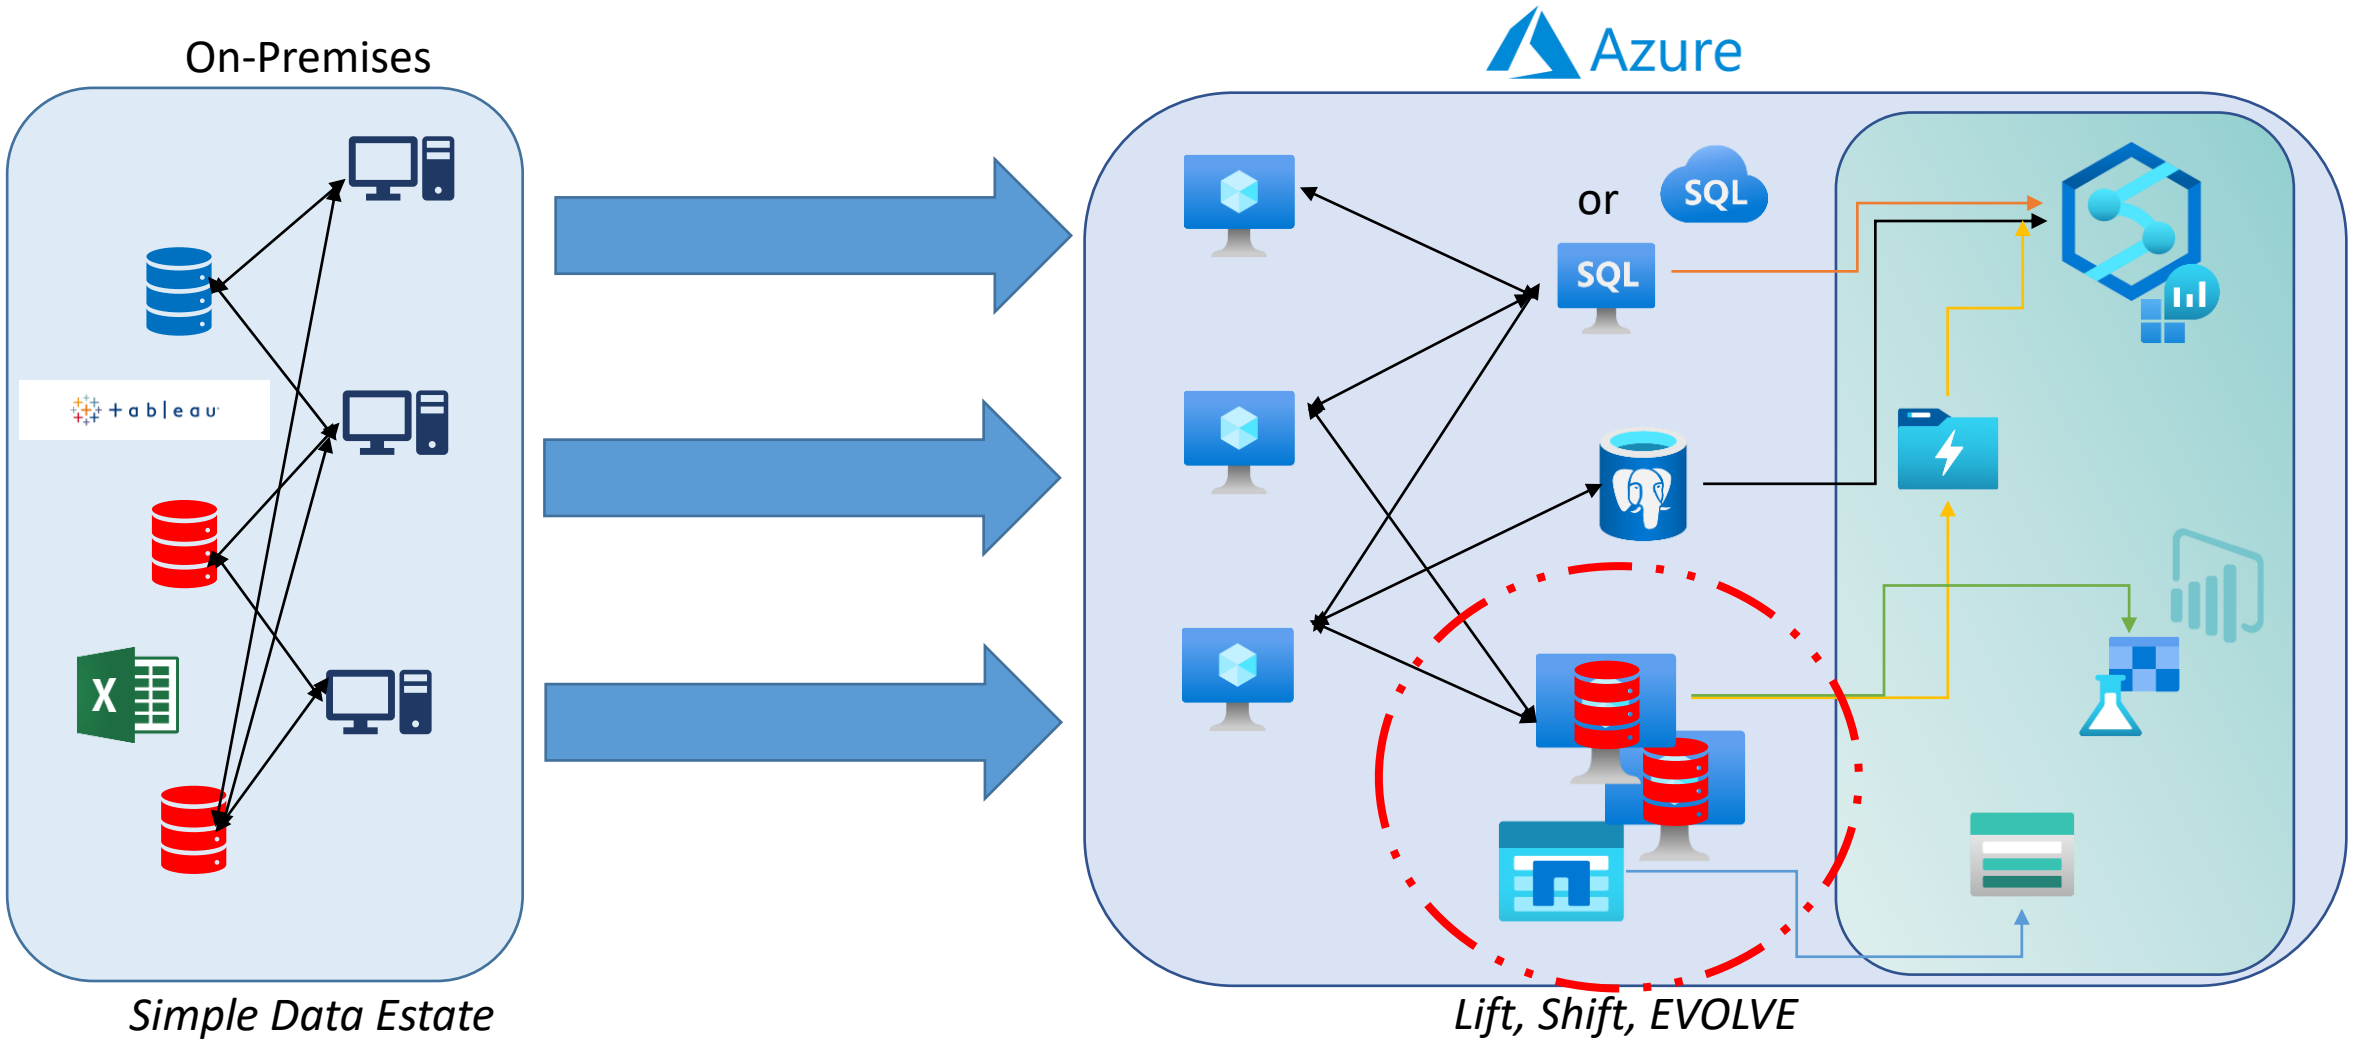

Simple Data Estate

Lift, Shift, EVOLVE

How to do Oracle on Azure and Not Waste Time

Oracle Migration Workflow

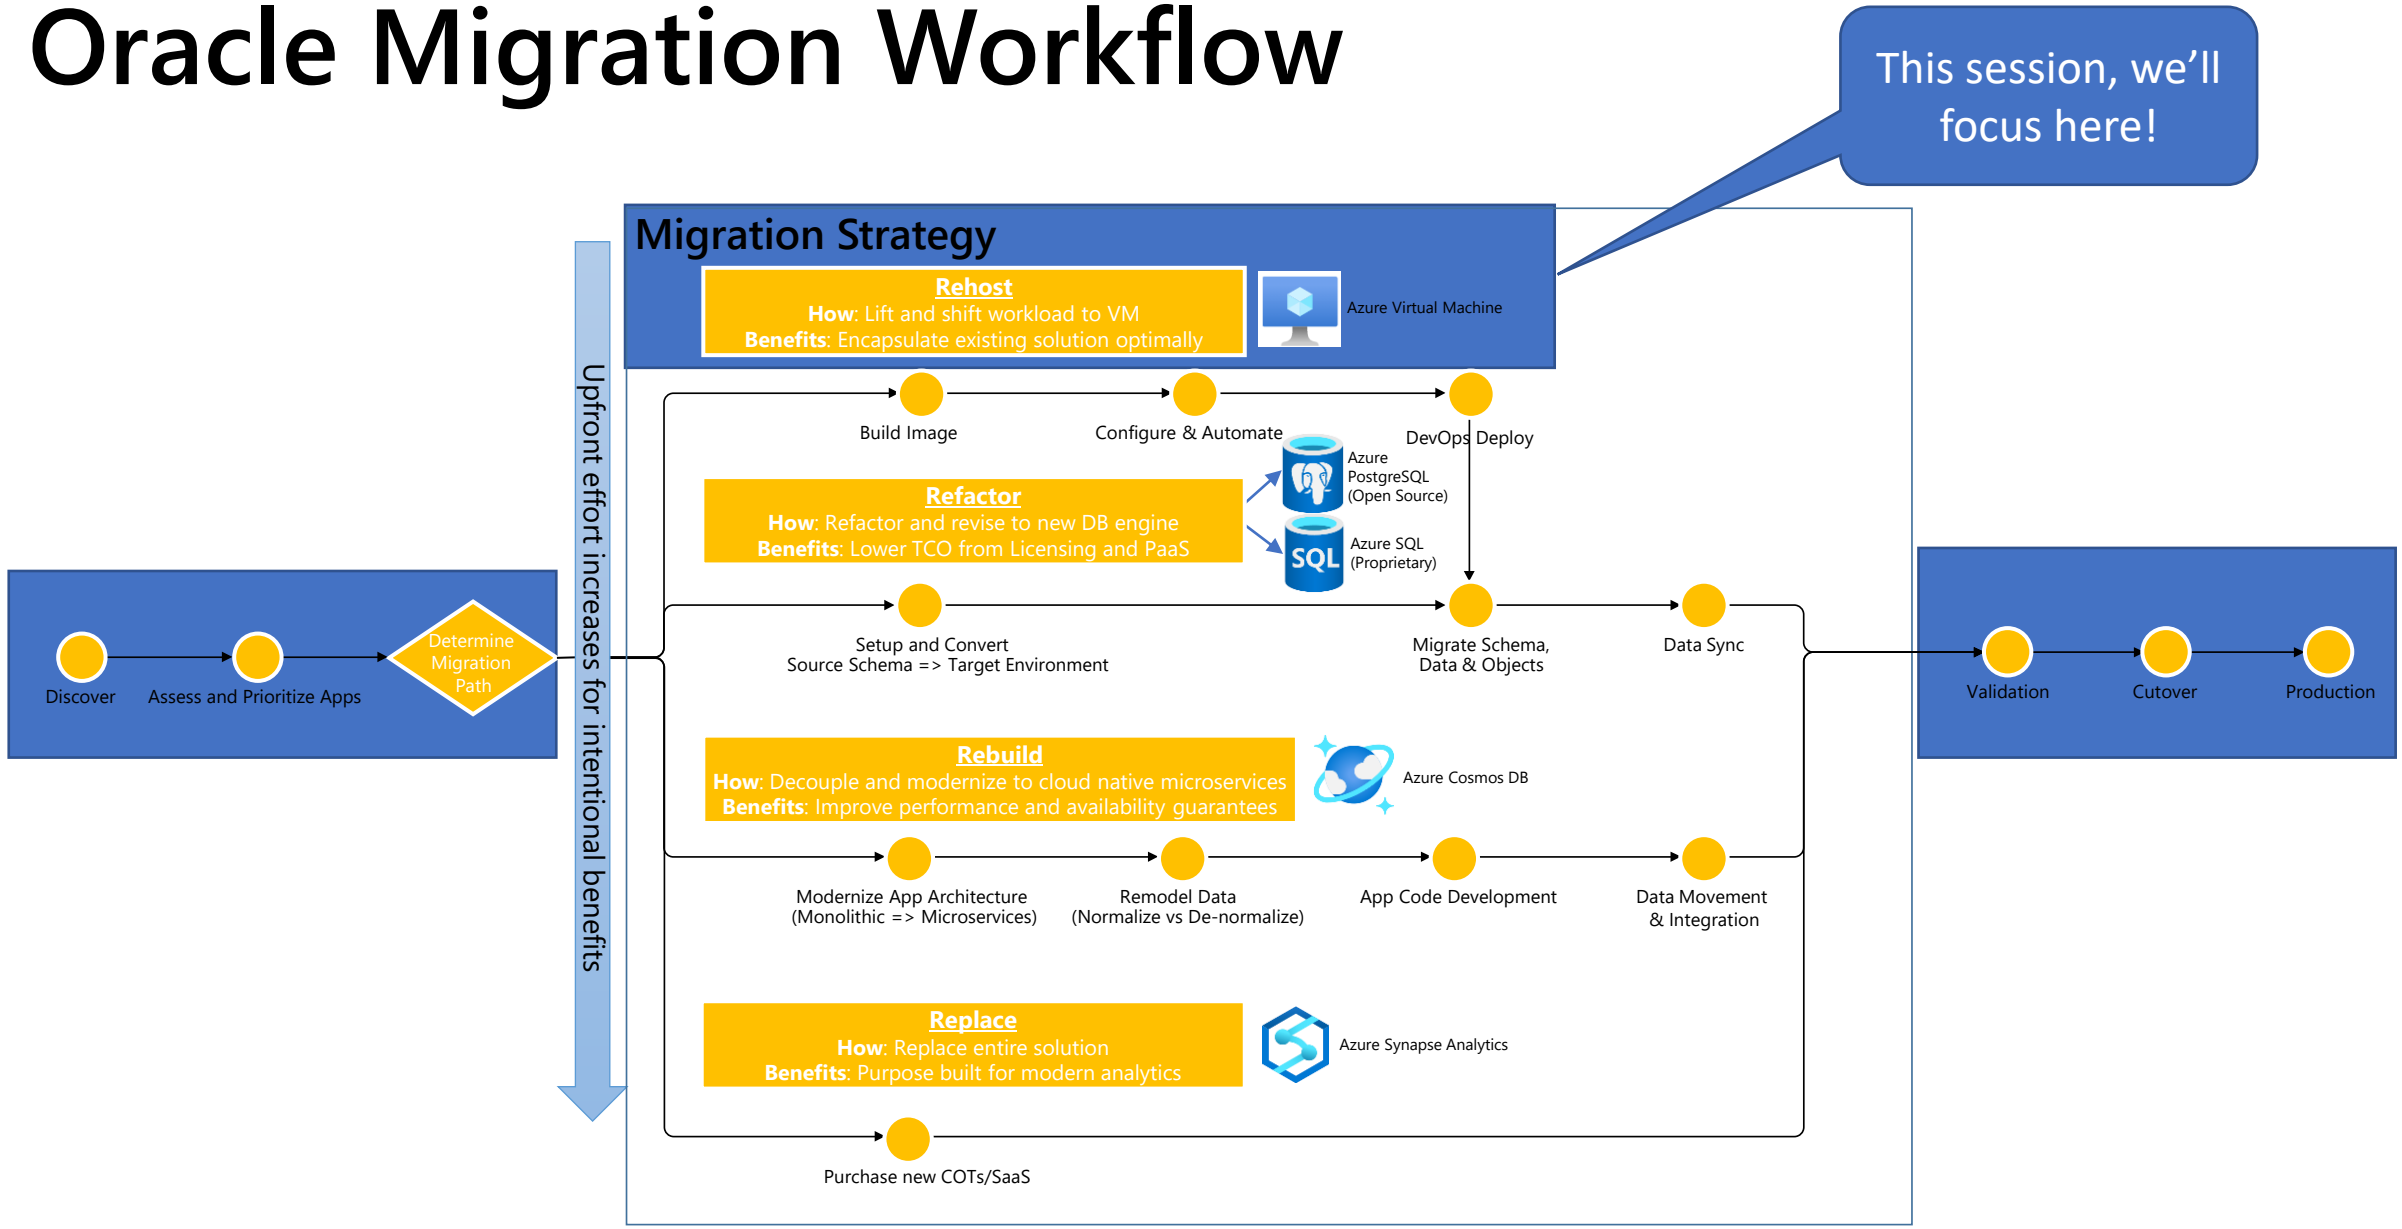

An Oracle on Azure IaaS Migration

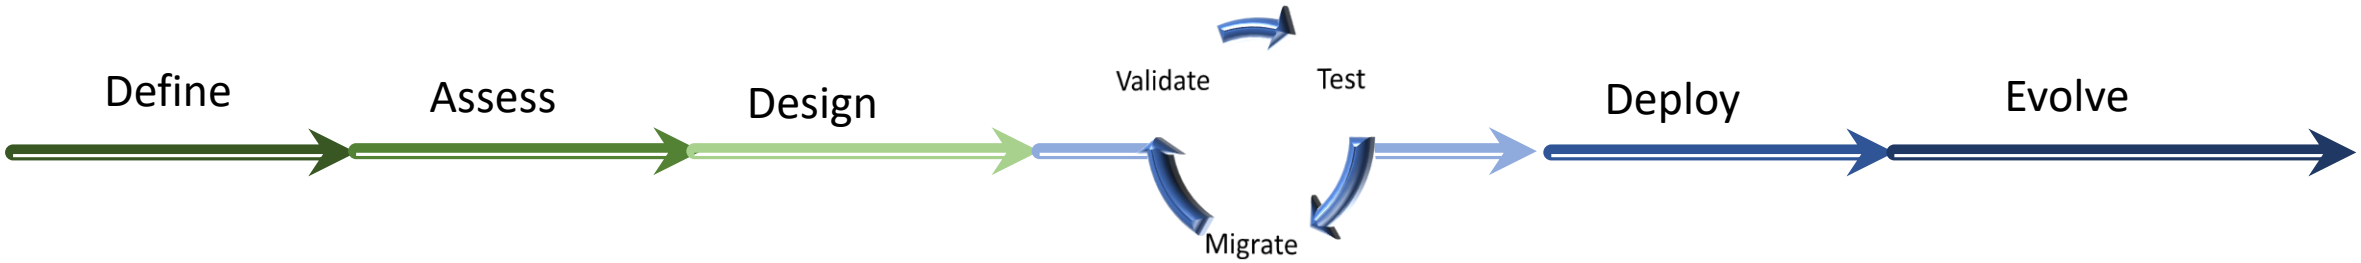

- Customer wants to stay on Oracle.
- Strong data estate strategy.

- Identify databases and applications going to Azure
- Assess Workload using AWR.

- Design architecture to meet workload requirements.
- Design backup, HA and DR on RPO/RTO and SLAs
- Build out deployment model

- Preform POC.
- Test backup, DR and HA process
- Test Migration process.
- Test, refine and improve.

- Endure Customer Readiness
- Verify recommended practices were followed.
- Lifecycle Support

- Optimize Existing IaaS with Roadmap opportunities.
- Identify potential refactoring at database (schema) opportunities
- Introduce Azure services to do more with Oracle data.

The First Three Pillars of Oracle on Azure IaaS

1. Compute
2. Storage
3. Sizing Assessments
4. RAC and Disaster Resilience
5. Everything Else

Pillars of Oracle on Azure IaaS

1. Compute
2. Storage
3. Sizing Assessments
4. RAC and Disaster Resilience
5. Everything Else

https://azure.microsoft.com/en-us/pricing/details/virtual-machines/series/

Understand IaaS VM Series

- A and B-series aren't optimal for Oracle workloads and are highly avoided.
- D-series are acceptable for small workloads, but reconsider if high demand on vCPU or memory.
- L and H-series are outliers for database workloads- Simply not designed for relational systems.

- Identify workload needs
- D-series is for general use
- **E-series and M-series are the most common VMs in the database industry**
  - E-series for average production databases
  - M-series for VLDB, (very large databases or heavy memory usage) and preferred with network attached storage.

Virtual machines for Oracle databases

Type	vCPUs	vRAM	Max throughput (MBps) for SCSI storage	Max IP egress rate (Mbps) for NFS, iSCSI storage
Ds_v4/v5	**2-64**	**8-256**	**125-4000**	**1000-30000**
Es_v4/v5	**2-80**	**16-504**	**125-4000**	**1000-30000**
Ms_v2	32-192	875-4096	500-2000	8000-30000
Mv2	208-416	2840-11400	1000-2000	16000-32000
Dds_v5 *	***2-96***	***8-384***	***125-4000***	***12500-35000***
Eds_v5 *	***2-104***	***16-672***	***125-4000***	***1000-100000***

**not available in all regions currently, fall back to v4 if unavailable.*

https://azure.microsoft.com/en-us/pricing/details/virtual-machines/series

When one VM is too Much- Constrained VMs

- Allows for isolation of vCPU to application licensing for database and app workloads
- Matched in existing series VMs in the Azure Pricing Calculator
- Share storage between databases or apps using network storage options like ANF

- Poorly named, vCPU constrained counts are the # of vCPUs on VM.
- Carefully match workloads on IO and memory, not just vCPU usage when sizing.

https://docs.microsoft.com/en-us/azure/virtual-machines/windows/constrained-vcpu

Specialized Constrained vCPU VMs

Name	vCPU	Specs
Standard_E4-2s_v5	2	Same as E4s_v5
Standard_E8-4s_v5	4	Same as E8s_v5
Standard_E8-2s_v5	2	Same as E8s_v5
Standard_E16-8s_v5	8	Same as E16s_v5
Standard_E16-4s_v5	4	Same as E16s_v5
Standard_E32-16s_v5	16	Same as E32s_v5
Standard_E32-8s_v5	8	Same as E32s_v5
Standard_E64-32s_v5	32	Same as E64s_v5
Standard_E4-2s_v5	2	Same as E4s_v5
Standard_E8-4s_v5	4	Same as E8s_v5
Standard_E8-2s_v5	2	Same as E8s_v5
Standard_E16-8s_v5	8	Same as E16s_v5
Standard_E16-4s_v5	4	Same as E16s_v5
Standard_E32-16s_v5	16	Same as E32s_v5

https://docs.microsoft.com/en-us/azure/virtual-machines/windows/constrained-vcpu

Marketplace Images- OS and Oracle

Most Common

- Customer Leverages OS Image without Oracle installed
- RedHat or Oracle Linux
- Builds out OS Image
- Installs Oracle version(s)
- Installs ASM (50% use ASM)
- Sets up directories
- Builds Image for Image Gallery

Less Common

- Uses Oracle image from marketplace.
- Uses Windows with Oracle.
- Brings over on-premises image and turns into Azure image.
- About 25% use some type of automation for deployment, (Terraform, Ansible or even Bash scripts)

Overview of Oracle DB on Azure

BYOL license for Oracle Database **on Azure Infrastructure-as-a-Service**

VM images with Oracle Database preloaded, available in Azure Marketplace, which also includes VM images for OEL, RHEL, and SLES Linux

Oracle Database 12.1, 12.2, 18.3, 19c,
21c **Enterprise Edition**

Oracle Database 12.1, 12.2, 18.3
Standard Edition

Customers can bring other Oracle Database images as well

Oracle provides license mobility from on-premises to Azure

Locating Azure Images

Oracle Linux Images Are Not Visible from the Portal

Using the Azure Command Line:

```
az vm image list \
--offer Oracle-Linux \
--all \
--publisher Oracle \
--output table
```

RedHat, (RHEL) Images from Command Line

```
az vm image list \
--offer RHEL-byos\
--all \
--publisher RedHat \
--output table
```

Locating Oracle Enterprise Linux Images for Azure (dbakevlar.com)

Oracle Apps

Bring your Oracle Apps to Azure

- E-Business Suite
- JD Edwards EnterpriseOne
- PeopleSoft
- Oracle Retail applications
- Oracle Hyperion Financial Management

Oracle Apps are supported in Azure due to the Oracle Partnership with Azure

Pillars of Oracle on Azure IaaS

1. Compute
2. Storage
3. Sizing Assessments
4. RAC and Disaster Resilience
5. Everything Else

Storage is SEPARATE and Important

- Ensure you know the IO workload for your database going to the cloud- DO NOT GUESS!
- Understand both the MBPs and the IOPS, (Requests) for the database.
- Oracle has demonstrated, on average, much higher demands for IO than MSSQL, MySQL or PostgreSQL.

- *Separate storage offers us the opportunity to create the right combination for success in IaaS.*

Database Savvy Storage

	Protocol	Max Throughput (MBps)	Min Latency (ms)	Pricing	Notes
Premium SSD	SCSI	900/device, 2000/VM cumulatively	0.7 w/ host-caching, 2 w/o host-caching	$	Snapshot capable, bursting capable, LRS/ZRS redundancy
UltraDisk	SCSI	2000/device, 2000/VM cumulatively	1	$$$	No snapshots, LRS redundancy
Azure Files premium	NFS v4.1	100 + (0.1 * GiB-provisioned)	1	$$	No snapshots, LRS/ZRS redundancy
Azure NetApp Files	NFS v3.0 NFS v4.1 SMB	4500	0.25	$$$	Snapshot capable, LRS/GRS redundancy
SILK	iSCSI	12000	0.5	$$	Snapshot capable, thin cloning, LRS/ZRS redundancy
FlashGrid	iSCSI	25000	0.5	$$	Oracle Only, Requires Grid and supported only through FlashGrid

Storage Considerations

What is the storage to be used for?

Data- OLTP, DSS, OLAP, Big Data?

Logging

Backup

Ensure that backups and data refresh requirements are calculated into the IO demands for the database.

Host Caching vs. Bursting

Name	Capacity (GiB)	IOPS per disk	Max burstable IOPS	Throughput per disk (MB/s)	Max burstable throughput per disk (MB/s)	Cache limit per disk (GiB)
P1	4	120	3,500	25	170	4
P2	8	120	3,500	25	170	8
P3	16	120	3,500	25	170	16
P4	32	120	3,500	25	170	32
P6	64	240	3,500	50	170	64
P10	128	500	3,500	100	170	128
P15	256	1,100	3,500	125	170	256
P20	512	2,300	3,500	150	170	512
P30	1,024	5,000	30000	200	1000	500 MBPs/1,024
P40	2,048	7,500	30000	250	1000	750 MBPs/2,048
P50	4,096	7,500	30000	250	1000	750 MBPs/4,095
P60	8,192	16,000	30000	500	1000	750 MBPs/4,095
P70	16,384	18,000	30000	750	1000	750 MBPs/4,095
P80	32,727	20,000	30000	900	1000	750 MBPs/4,095

Source: Managed disks pricing

Ultra Disk

Ultra Disk Offerings									
Disk Size (GiB)	4	8	16	32	64	128	256	512	1,024-65,536 (in increments of 1 TiB)
IOPS Range	1,200	2,400	4,800	9,600	19,200	38,400	76,800	160,000	160,000
Throughput Range (MB/s)	300	600	1,200	2,000	2,000	2,000	2,000	2,000	2,000

Ultra Disk Configuration	Unit	Hourly Price	Monthly Price**
Disk Capacity (GiB)	GiB	$0.000164	$0.11972
Provisioned IOPS	IOPS	$0.000068	$0.04964
Provisioned Throughput (MB/s)	MBps	$0.000479	$0.34967
Provisioned vcpu reservation charge*	vCPU	$0.006	$4.38

Ultra Disk Pricing

GiB * .000164, MBPs * .34967, IOPs * ..04964, vCPU * 4.38

https://docs.microsoft.com/en-us/azure/virtual-machines/disks-enable-ultra-ssd#ga-scope-and-limitations

Ultradisks

- Often the first recommendation by Infra
- Be aware of the limitations before recommending for database workloads:
  - Oracle 12.2 later is supported.
  - Only supports un-cached reads and un-cached writes
  - Doesn't support disk snapshots, must use RMAN
  - VM images, OS Disk, availability sets, Azure Dedicated Hosts, or Azure disk encryption
  - No integration with Azure Backup or Azure Site Recovery
- Offers up to 16 TiB per region per subscription unless upped via support.
- Isn’t available in all regions.
- Must use a minimum of single AZ.

Disk Size (GiB)	IOPS Cap	Throughput Cap (MBps)
4	1,200	300
8	2,400	600
16	4,800	1,200
32	9,600	2,400
64	19,200	4,000
128	38,400	4,000
256	76,800	4,000
512	153,600	4,000
1,024-65,536 (sizes in this range increasing in increments of 1 TiB)	160,000	4,000

Redo Logs, Yeah!

Types of cache Settings

- Available to Premium Storage
  - A Multi-tier caching technology, aka BlobCache
  - OS Disk- ReadWrite is fine, which is the default, but not for datafiles.
  - ReadOnly Cache is, as it caches reads, while letting writes pass through to disk.
  - Limit of 4095Gib on per individual premium disk
    - Results in any disk above a P40 for entirety will silently disable read caching.
    - Larger disks are preferably used without caching, otherwise additional space is wasted. P50, just allocate 4095 of the 4096 size.
    - Use smaller disks and choose to stripe and mirror.
  - Verify SKU Availability

IO Throttling

- Why does it happen?
  - No, you can't have all the resources for yourself.
- What all can be involved?
  - It's not just the database.
- How to identify it?
  - What do to when it is identified?

Just because you have a big VM, doesn't mean you have high throughput!

Size	vCPU	Memory: GiB	Temp storage (SSD) GiB	Max data disks	Max cached and temp storage throughput: IOPS/MBps (cache size in GiB)
Standard_M8ms	8	218.75	256	8	10000/100 (793)
Standard_M16ms	16	437.5	512	16	20000/200 (1587)
Standard_M32ts	32	192	1024	32	40000/400 (3174)

https://docs.microsoft.com/en-us/azure/virtual-machines/sizes-memory?toc=/azure/virtual-machines/linux/toc.json&bc=/azure/virtual-machines/linux/breadcrumb/toc.json

Bring in Additional Solutions

- High IOPS-
  - MBPs: Azure NetApp Files
  - Higher IO throughput: Consider Silk, Flashgrid Storage, Pure Storage
  - Consider disk striping of smaller disks and parallel processing at the database level.
- Backups, batch loading and other challenges:
  - Offload backups with secondary backup solutions.
  - Refactor batch processing with other services, (Azure Data Factory, Azure Analysis Services, Databricks, etc.)

Oracle on Azure – Improve Agility with ANF

Simple to Manage

- Native Azure experience for easy deployment & scale

Highly Performant

- On-prem like performance
- On-demand scalability

Ent Data Management

- Space / Time efficient snapshot and cloning

Hybrid

- Works with Dataguard for hybrid cloud deployments

https://docs.microsoft.com/en-us/azure/azure-netapp-files/azure-netapp-files-solution-architectures#oracle

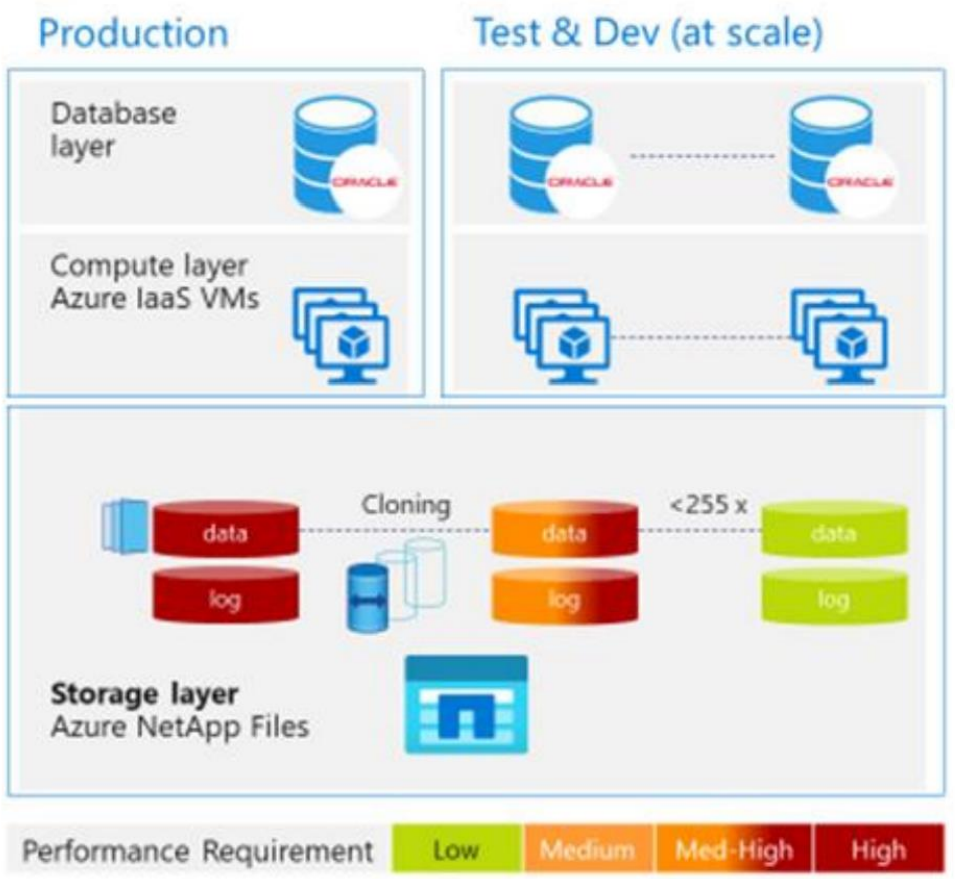

Azure NetApp Files

- Fully Managed, PaaS, Microsoft Azure Storage Service
- All Flash Bare metal Storage
- Only dependent on Nic, not VM.
- *Available in Standard, Premium, (common) and Ultra, (optimal)
- ANF is native to Azure

*Be aware of pricing with scaling to meet IO

	Azure Files	Premium Files	Azure NetApp Files	Premium Disk
Performance	1K IOPs	100K IOPs	**320K IOPs**	20K IOPs
Capacity Pool	5TB	100TB	**500TB**	32TB
AD Integration	Azure AD	N/A	**Bring Your Own AD / Azure AD**	N/A
Protocol	SMB	SMB	**NFS & SMB**	Disk
Data Protection		LRS Only	**Snapshots Back Up Tools**	Snapshots

FAQs About Azure NetApp Files | Microsoft Docs

Silk Platform Architecture

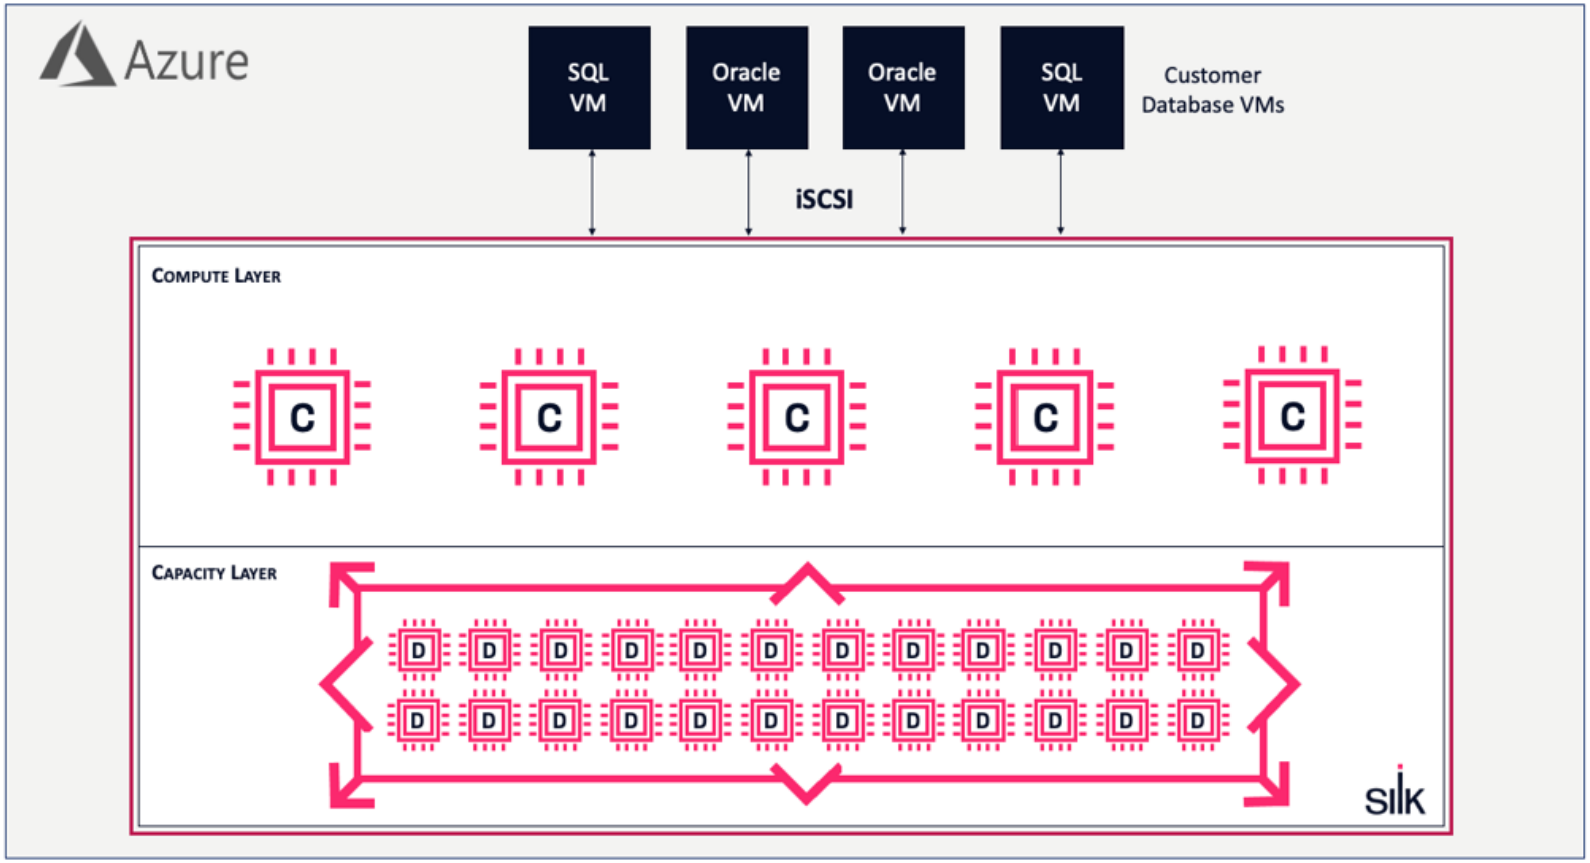

A Silk data pod, deployed in a single zone, RG

Silk Platform components

C

Compute:
2 - 8 c.nodes

D

Capacity:
1 m.node =
24 d.nodes

Silk Performance

Why Silk on Azure ?

Metric	Azure + Silk	Azure Ultra SSD	Gain vs. Ultra SSD	Azure Premium SSD	Gain vs. Premium SSD
Read IOPS	1m+	160K	6.3x	20K	50x
Read BW	15 GB/s	2 GB/s	7.5x	0.9 GB/s	15x
Write IOPs	800K	160K	5x	20K	40x
Write BW	8.5 GB/s	2 GB/s	4x	0.9 GB/s	9x

All performance numbers achieved at 1.5ms consistent latency or lower, with data services enabled

- Designed for mixed workloads – OLTP & BI
- Elastically scale performance up and down
- Patented algorithms for high parallelism
- Automatic tuning for optimal CX
- Shared performance for all applications

SILK IN THE PUBLIC CLOUD

c.nodes	2 c.nodes	3 c.nodes	4 c.nodes	5 c.nodes	6 c.nodes	7 c.nodes	8 c.nodes
Usable Capacity*	12TB-0.5PB	12TB-1PB	12TB-1.5PB	12TB-2PB	12TB-2.5PB	12TB-3PB	12TB-3.5PB
IOPS	Up to 220K	Up to 330K	Up to 440K	Up to 550K	Up to 660K	Up to 770K	Up to 880K
Throughput	Up to 3.5GB/s	Up to 5.25GB/s	Up to 7GB/s	Up to 8.75GB/s	Up to 10.5GB/s	Up to 12.25GB/s	Up to 14GB/s
Latency	0.25ms						

FlashGrid Storage Fabric

- Upwards of 30K MBPs in Azure
- Requires a grid installation
- Single instance or failover cluster provided as part of deployment

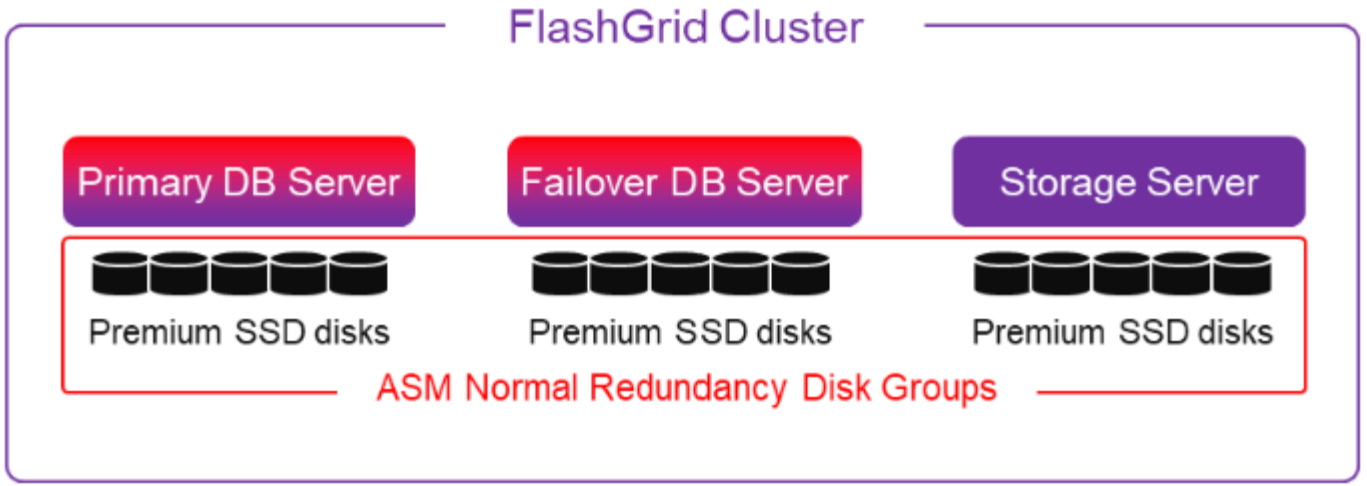

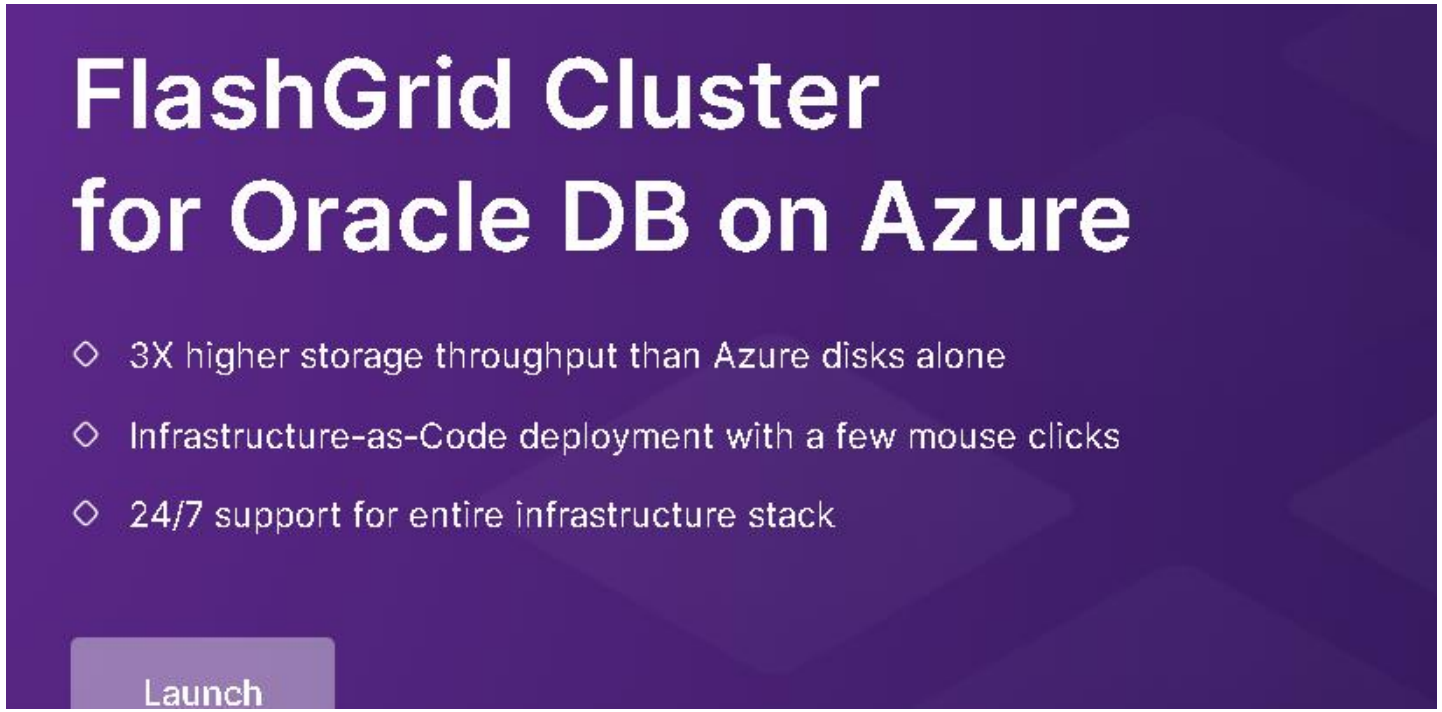

When To Go Old-School

- Depending on the combination of storage, striping and RAID, performance can vary greatly.
  - Verify that disk is striped correctly, (log creation commands and document.)
  - Consider smaller disk size and stripe vs. larger, single drive to offer better performance.
  - In Linux, consider huge pages and use LVM, (Linux Volume Manager) or Oracle ASM, (Automatic Storage Management) to provide advanced features for diskgroup layout.
  - Keep an eye on disk sector size, (there's a bug requiring 512 byte sector size in Oracle 12.1)

Storage Decision Matrix

Workload	Premium SSD	Ultra Disk	Azure NetApp(ANF)	Silk	Flashgrid
IO up to 750MBPs	X	* Rarely used, not cost effective	X	X	X
IO up to 2000MBPs		X	X	X	X
IO up to 5000MBPs			X	X	X
IO 10K MBPs +				X	*with RAC
Exadata Offloading with HCC				X Silk Compression	* With RAC can compensate
Non-Platform specific	X	X	X	X	Oracle Only
RAC Capable			* On RAC Bare Metal	X	X
Check Region Availability		X	X	X	X
Non-AZ Deployment			X	X	
Limited Use Cases		* Mostly used with Redo Logs			
Thin Snapshots	X	Private Preview	X	X	
Thin Clones				X	
Scaling Storage Options	X	X	X		X
Enterprise Cost	$	$$$	$$$$	$$$	$$$

Datafile Optimization through Storage

Distribution of Data Across Storage- Level 1

Managed Disk Premium with combination of Datafiles and Redo Datafiles. If ASM is used, the alias' are listed above the volume:

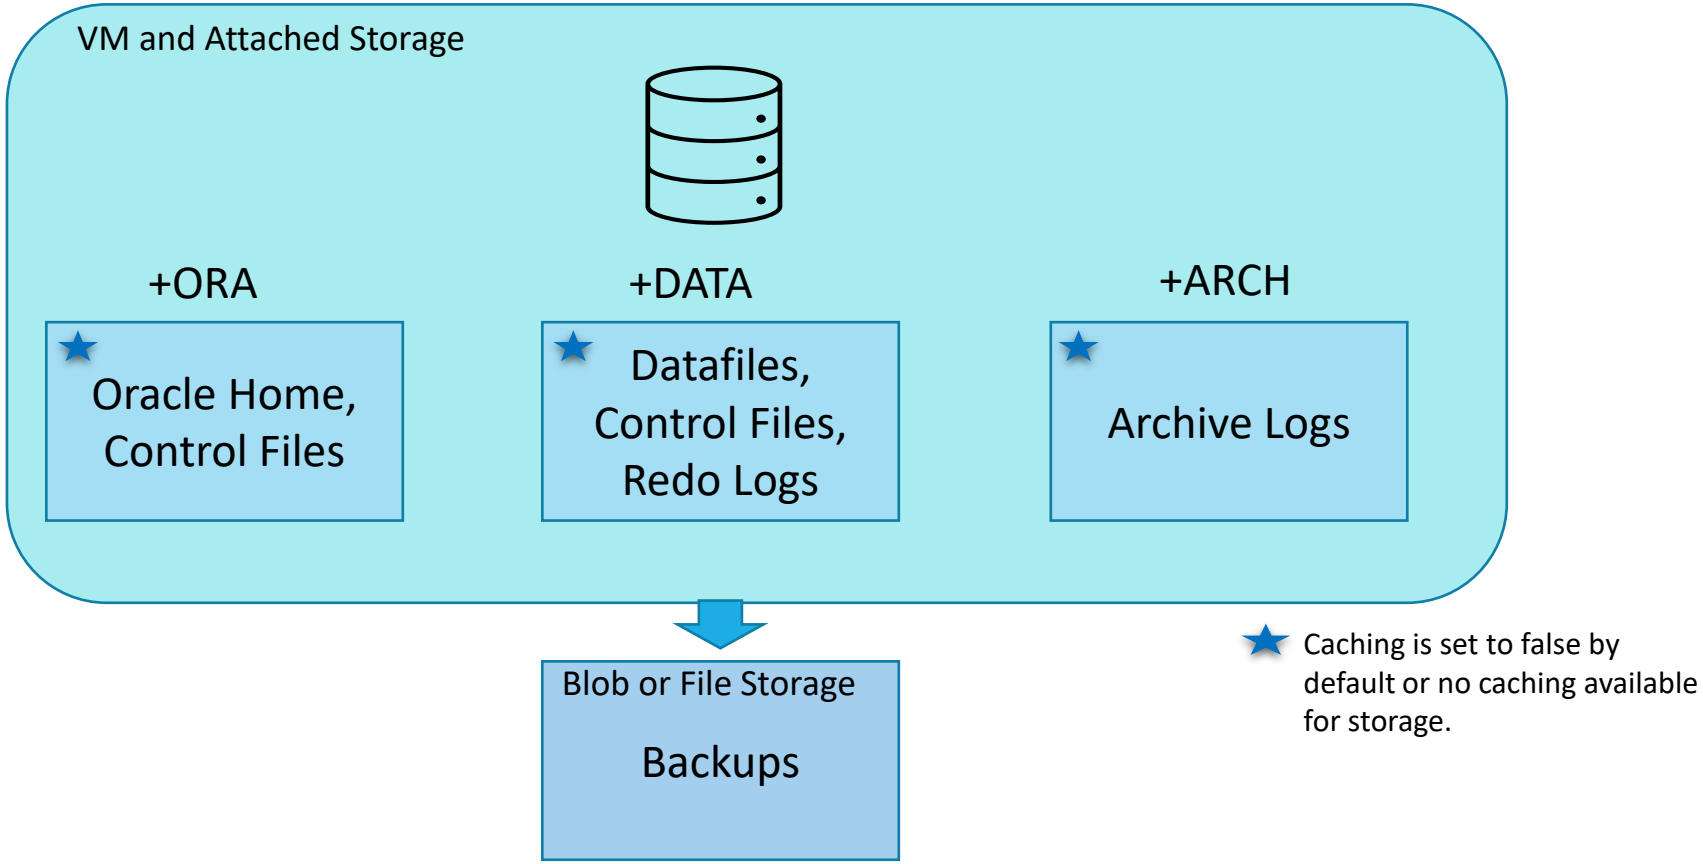

Distribution of Data Across Storage- Level 2

Managed Disk Premium with separation of Datafiles and Redo Datafiles.

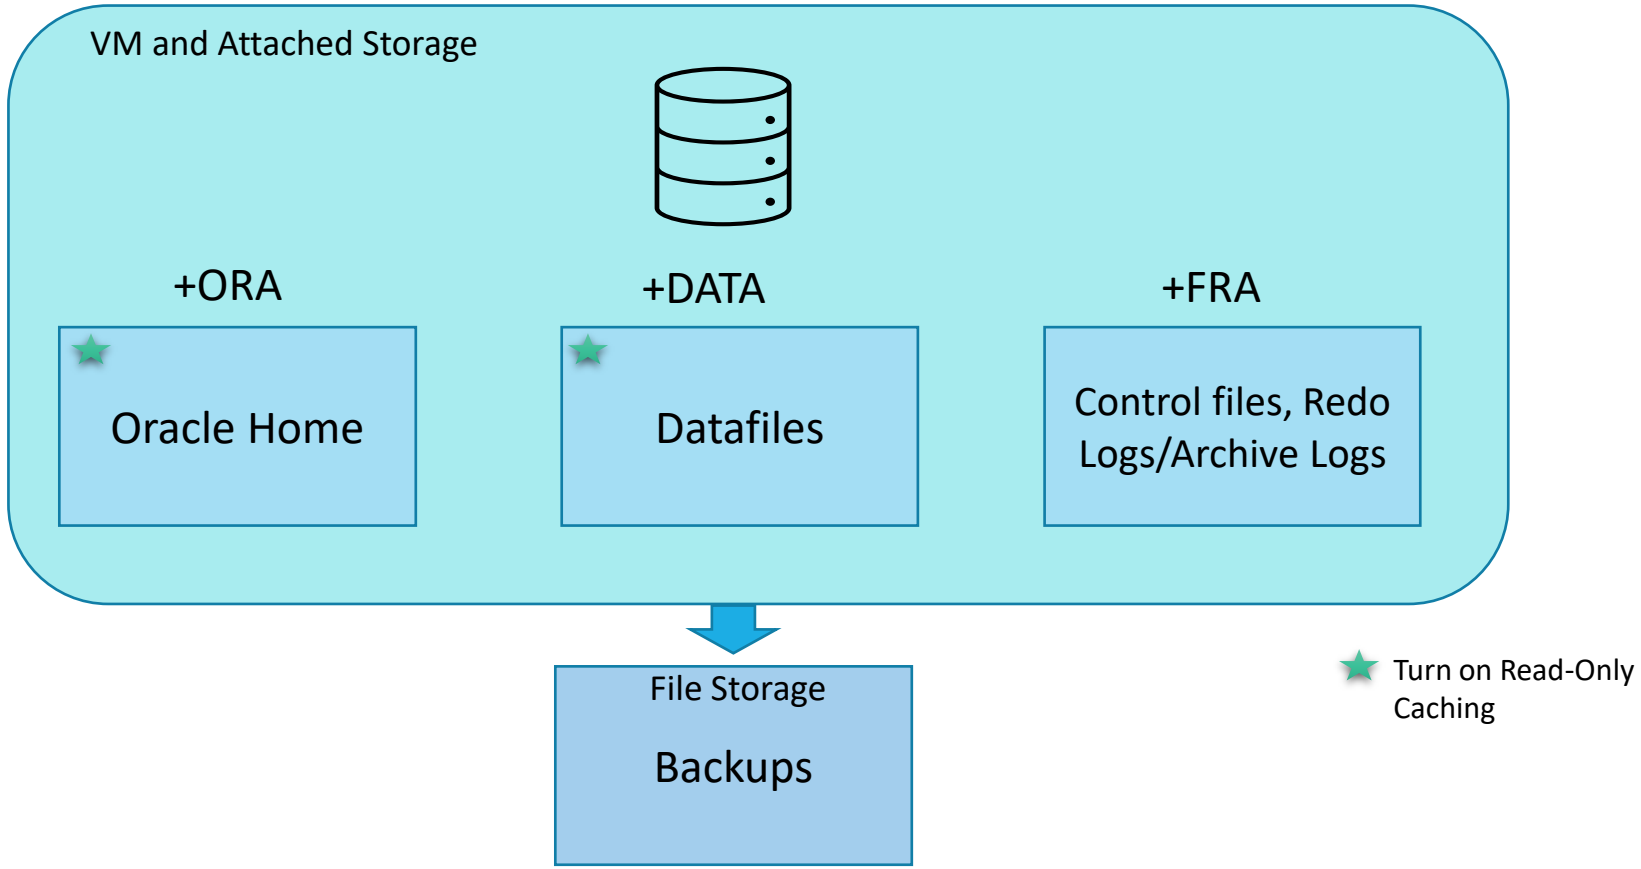

Distribution of Data Across Storage- Level 3

Managed Disk Premium with separation of Datafiles, Redo and Archive logs-

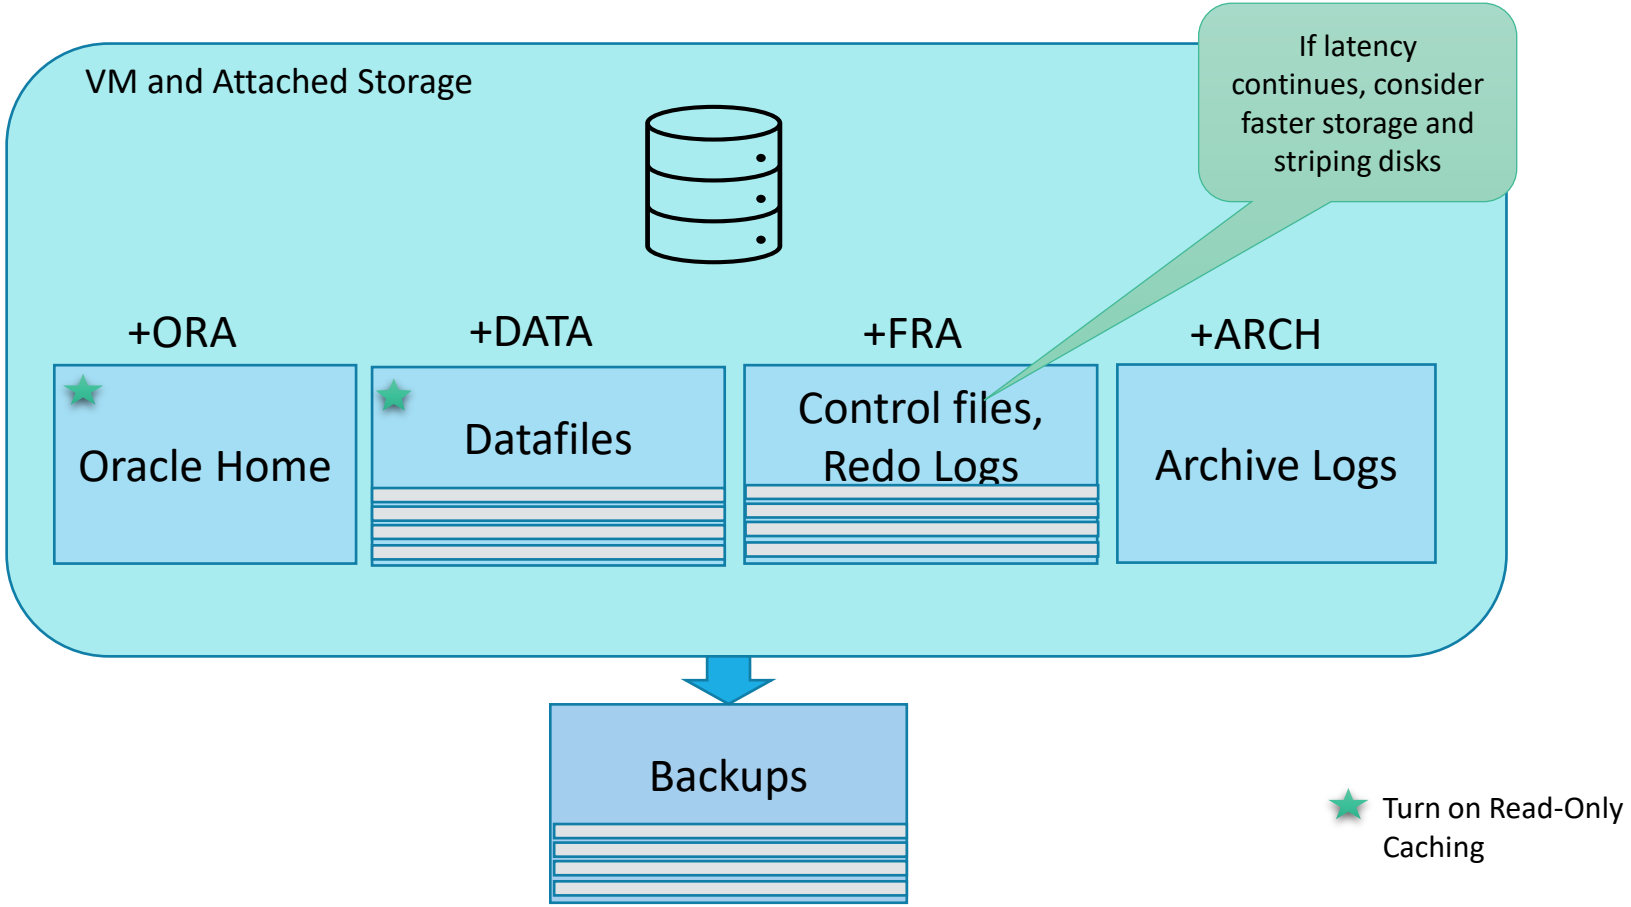

Distribution of Data Across Storage- Level 4

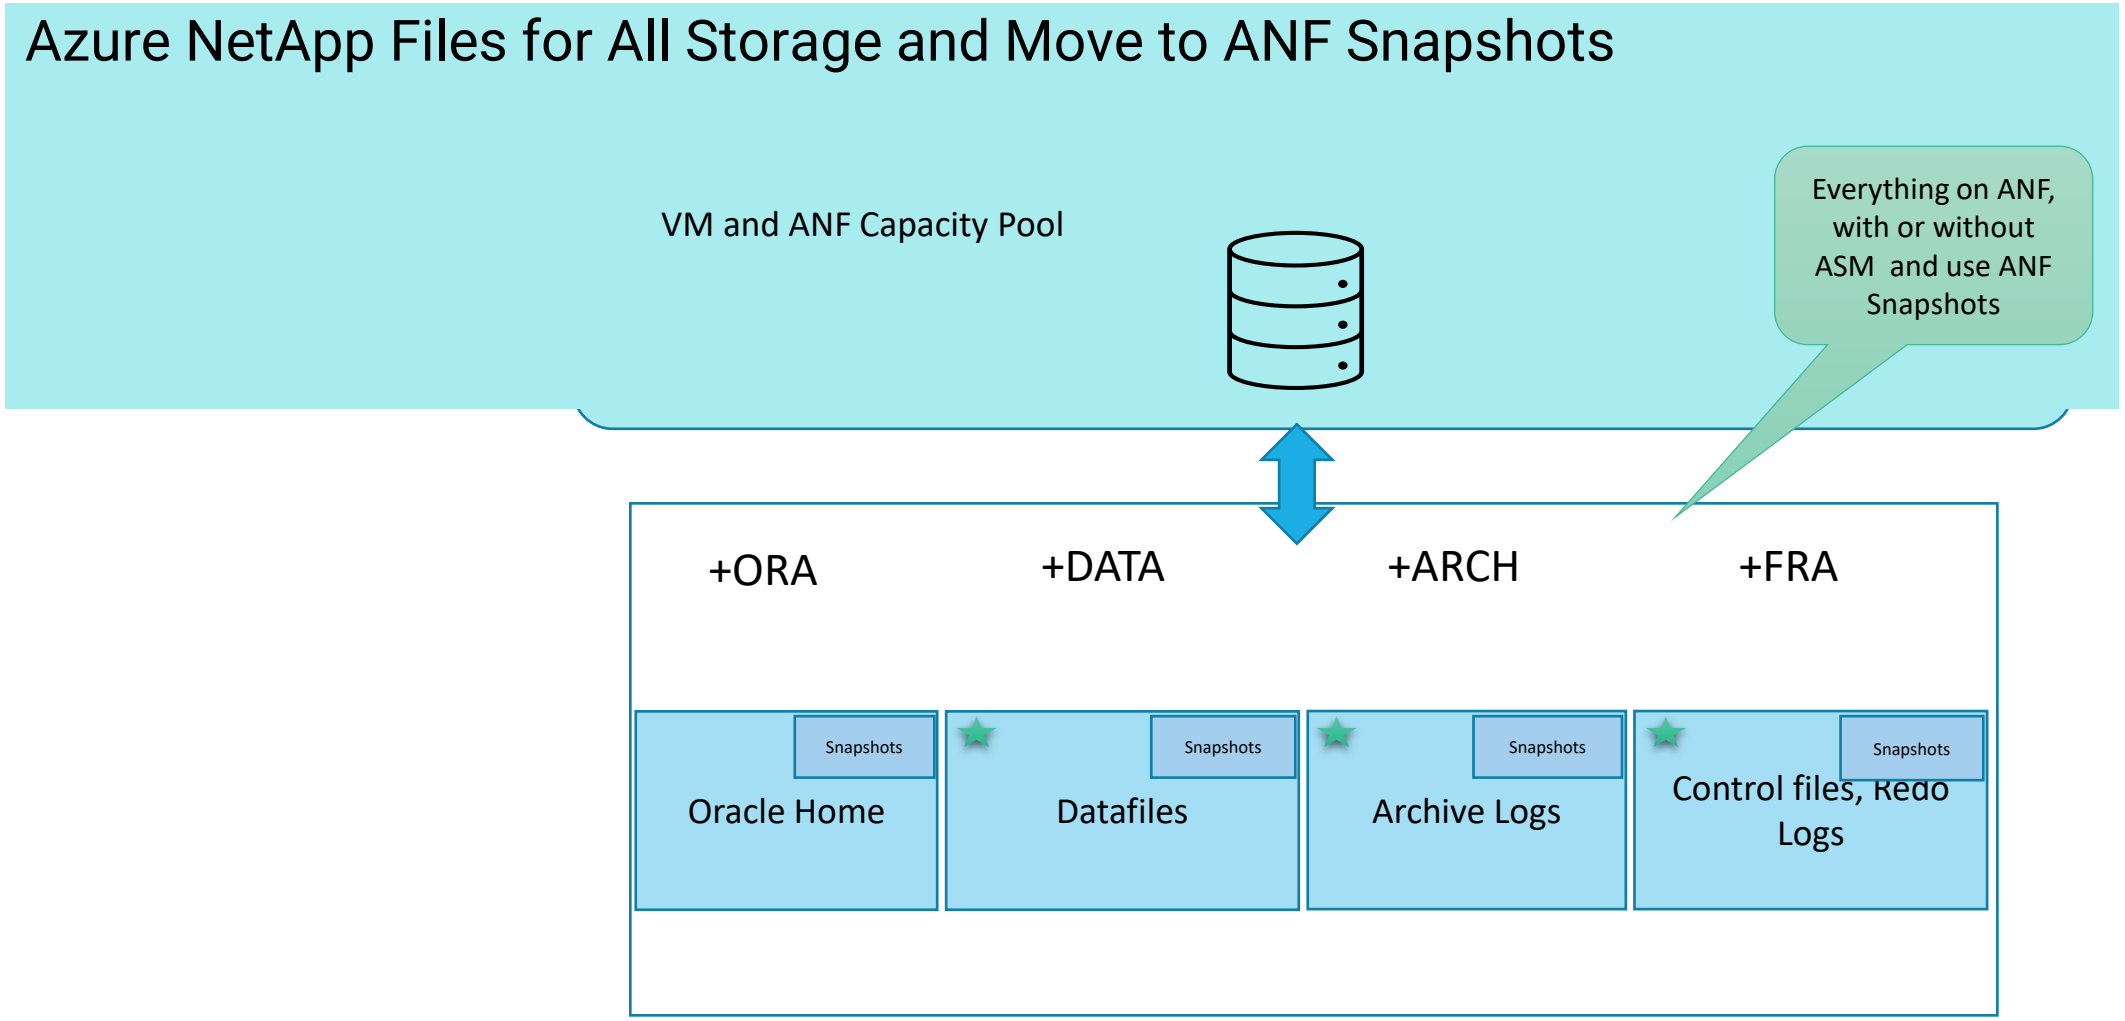

Pillars of Oracle on Azure

1. Compute
2. Storage
3. Sizing Assessments
4. RAC and Disaster Resilience
5. Everything Else

Licensing Oracle products within Azure

Do NOT accept statements about licensing or supporting Oracle products on Azure from Oracle account teams

Insist that Oracle's corporate License Management Services (LMS) team be consulted when encountering any suspicious or unreasonable assertions

Licensing Oracle products in Azure

According to Oracle

"When counting Oracle Processor license requirements in Authorized Cloud Environments, the Oracle Processor Core Factor Table is not applicable. For the purposes of licensing Oracle programs in an Authorized Cloud Environment (Azure), customers are required to count as follows - count two vCPUs as equivalent to one Oracle Processor license if hyperthreading is enabled, and one vCPU as equivalent to one Oracle Processor license if hyperthreading is not enabled."

http://www.oracle.com/us/corporate/pricing/cloud-licensing-070579.pdf

For example: The intel core factor is 0.5, so an 8 core physical box requires 4 cores of licensing. In Azure an 8 core VM (16 vCPUs Hyperthreaded or 8 vCPUs non-hyperthreaded) requires 8 cores of licensing.

*Oracle licensing folks as third-party cloud specialists try to lift and shift on-premises Oracle environments into Azure.

Migrate the Workload, not the Hardware

Servers may not have been sized appropriately for the workload.

Workload of database may have changed over time.

May cost you more in licensing than what the workload actually requires.

For different databases, there are different tools to assist:

SQL Server: DMVs, PerfMon, Scripting, (Randal, Klee, etc) Redgate SQL Monitor

Oracle: AWR, OEM, ASH, SASH, Statspack, Tracing

MySQL: Solarwinds DPA, Instrumental, Panopta

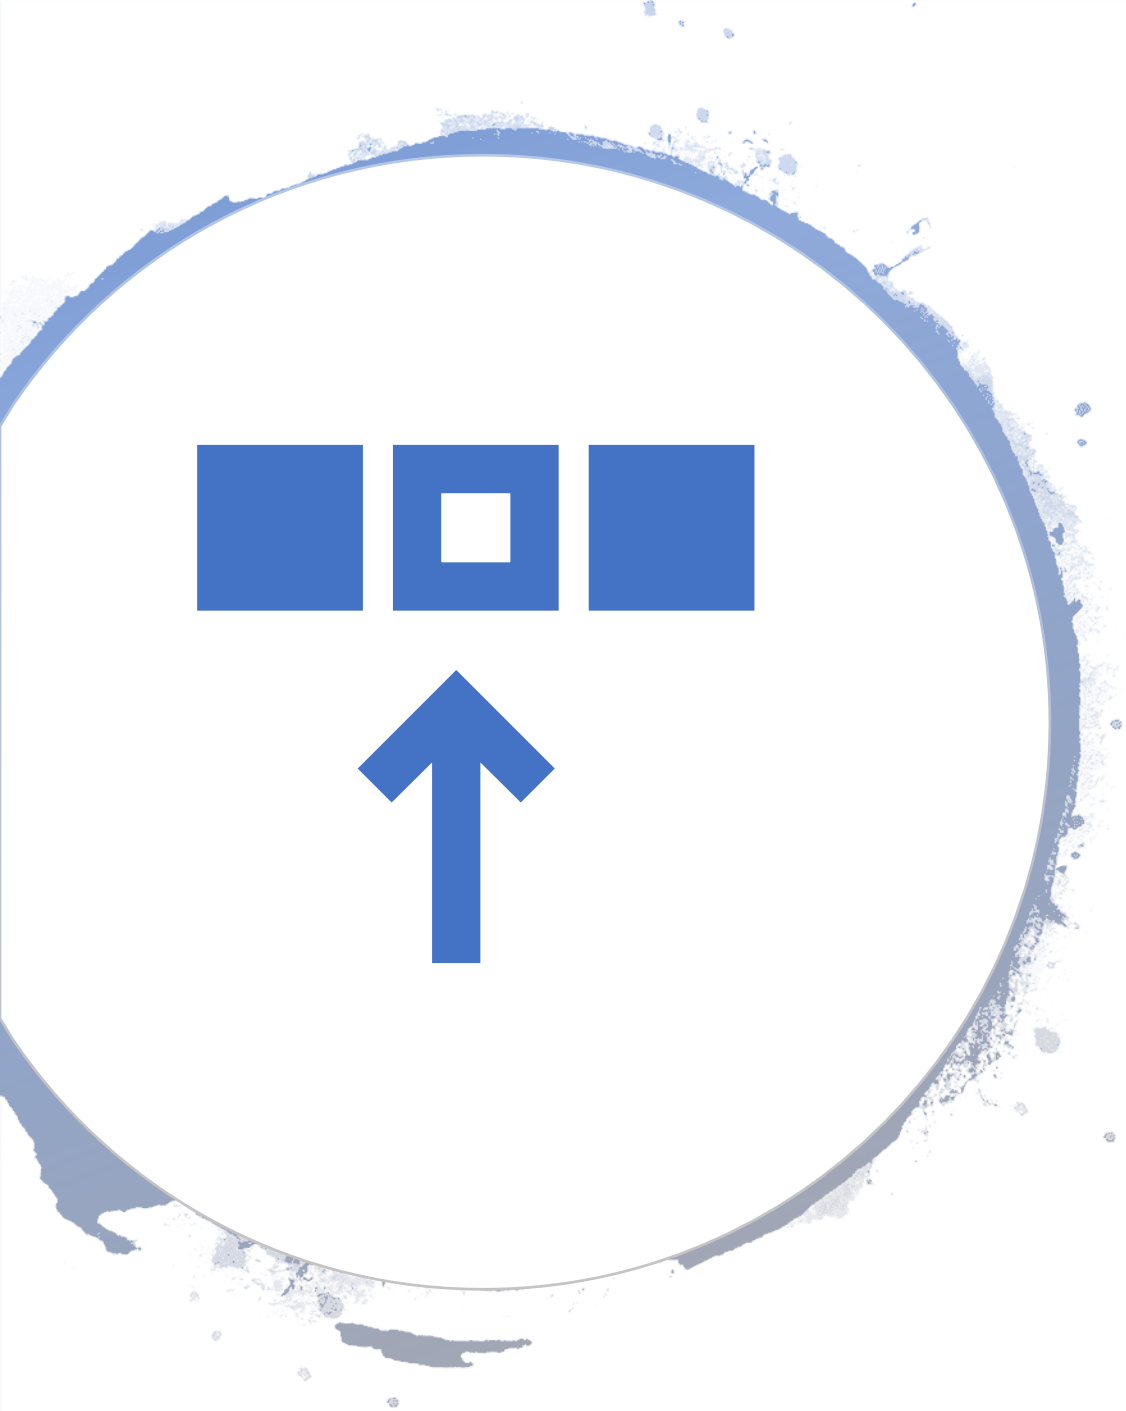

Can Exadata be Built in Azure?

- As Exadata is an engineered system, this solution can't be duplicated in Azure, but it can be migrated:
  - With Oracle Dataguard as a single Oracle instance with the correctly scaled VM, especially with advanced storage solutions such as Silk or ANF.
  - With Oracle Dataguard, clustering solutions such as SIOS, Pacemaker or Flashgrid.
  - Using Goldengate or another migration solution to Azure DW, which can simulate some of the workload features:
    - MPP, (massive parallel processing) across secondary nodes
    - materialized views
    - columnar features
    - compression

How to size (and price) resources for Oracle database

Going to stress this again…

Move the Oracle workload, not the server to the cloud.

- The hardware is often quite old or was sized to last years.
- The database workload has grown and/or the workload has changed from the time the hardware was purchased.
- The hardware was incorrectly sized to begin with.

Moving the on-premises hardware is ripe with failure…

- The Automatic Workload Repository, (AWR) report is the best data to collect and size a workload from Oracle
- AWR is always on in Oracle Enterprise Edition, (SE can use Statspack similarly.)

Capture peak workload in the AWR report, if possible

- For simplicity, we usually request the most recent 7-days in a single AWR report
- But ideally, we'd like to see a snapshot of the **peak workload**
  - SQL*Plus script `busiest_awr.sql` can assist in finding peak workloads recorded within AWR repository
- AWR reports don't capture all the information we need for sizing (i.e. *database size*, *average daily redo generated*, *expected annual growth rate*) so…
  - SQL*Plus script `dbspace.sql` can assist in capturing this information from the data dictionary

Tips

Ensure customer resources are onboard and understand why a substantial workload is required for an accurate sizing.

Advisors for SGA and PGA may recommend more memory than allocated.

Although rare, there are some bugs discovered in the AWR:

- RAC global report % busy for second node is often less than utilized
- 19c Exadata IOPS is miscalculated in IOSTAT Function Summary, use IOSTAT File Type Summary
- 10.2.0.3 missing totals and must be calculated for IO, separated by tablespace
- Statspack, (AWR predecessor requires manual calculations)
- 12.1 bug will show error at top of AWR report

AWR Architecture

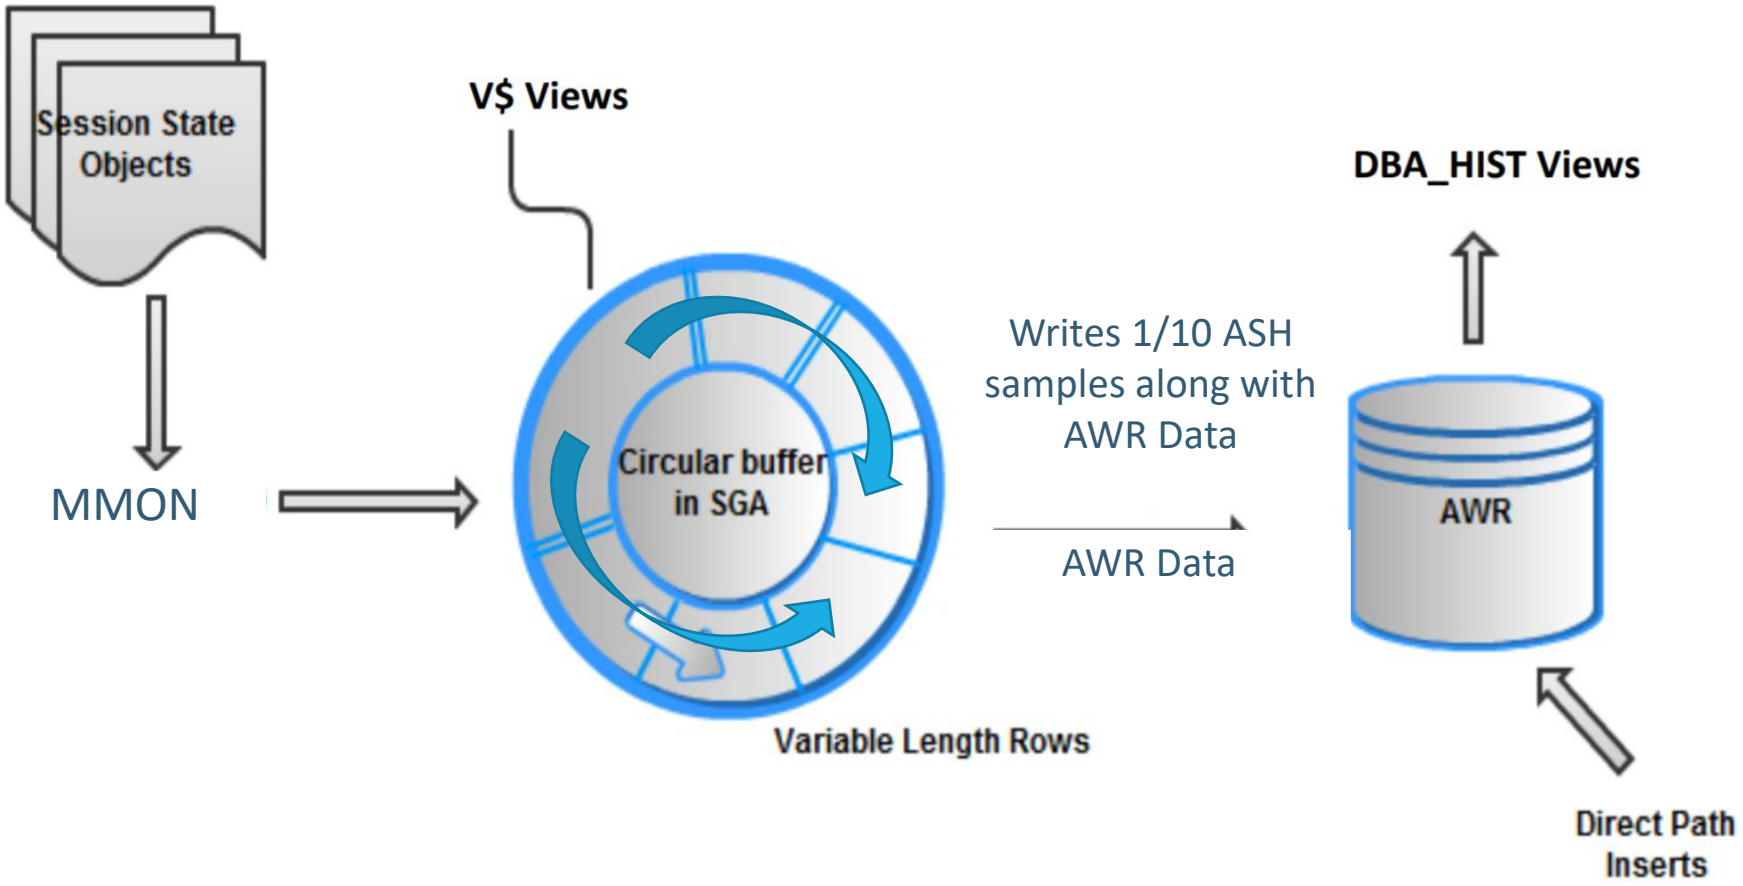

AWR Cont.

- Canned reports housed in $ORACLE_HOME/rdbms/admin directory
  - Single instance report: awrrpt.sql
  - RAC global instance report: awrgrpt.sql
  - ASH report: ashrpt.sql

- Due to hourly snapshots, a report for a single day would be a variance of 24 snapshot IDs between the beginning snapshot ID and the ending snapshot ID.
- ASH reports go by time, (day, hour, minute, seconds) and samples can overlap

The AWR Report

- AWR Reports can be between 10-30 pages- Don't get overwhelmed. We only need 12 values. ☺
- Always use a clean worksheet template
- Take care pasting values into worksheet. Can corrupt calculations on second sheet.
- Focus on production workloads.
- Dev and test sizing this way can be wrought with pitfalls, (boiling the ocean).

WORKLOAD REPOSITORY report for

DB Name	DB Id	Unique Name	Role	Edition	Release	RAC	CDB
DB01	423509001	RCP01	PRIMARY	EE	19.0.0.0.0	NO	NO

Instance	Inst Num	Startup Time	User Name	System Data Visible
DB01	1	11-Jul-21 00:09	SYSTEM	YES

Host Name	Platform	CPUs	Cores	Sockets	Memory (GB)
ex06db05.bus01.com	Linux x86 64-bit	56	28	2	754.16

	Snap Id	Snap Time	Sessions	Cursors/Session
Begin Snap:	82951	08-Aug-21 14:00:53	1116	16.0
End Snap:	82953	08-Aug-21 16:00:05	1123	16.0
Elapsed:		119.21 (mins)		
DB Time:		1,531.25 (mins)		

Rarely Change “Fudge Factors”

- These are the numbers that are used to take missing workload data, address averages, additional immediate growth and calculation issues into consideration.
- Leave to the default values unless a specific type of workload, (Exadata, Sparc, etc.)

Peak CPU factor	2.00
Est'd RAM factor	2.00
vCPU HT factor	2.00
%Busy CPU threshold	0.75
%Busy CPU multiplier	1.25
IO metrics (IOPS & MB/s) fudge factor	2.00

Capturing observed CPU utilization from AWR

Database Summary

Database				Snapshot Ids		Number of Instances		Number of Hosts		Report Total (minutes)	
Id	Name	RAC	Block Size	Begin	End	In Report	Total	In Report	Total	DB time	Elapsed time
3172330132		YES	8192	152699	152700	4	4	8	4	5,731.96	29.80

Database Instances Included In Report

- Listed in order of instance number, I#

I#	Instance	Host	Startup	Begin Snap Time	End Snap Time	Release	Elapsed Time(min)	DB time(min)	Up Time(hrs)	Avg Active Sessions	Platform
1	c		21-May-21 23:33	03-Jun-21 22:30	03-Jun-21 23:00	12.1.0.2.0	29.78	1,514.75	311.45	50.86	Linux x86 64-bit
2	c		21-May-21 23:33	03-Jun-21 22:30	03-Jun-21 23:00	12.1.0.2.0	29.78	1,460.85	311.45	49.05	Linux x86 64-bit
3	c		21-May-21 23:33	03-Jun-21 22:30	03-Jun-21 23:00	12.1.0.2.0	29.80	1,423.54	311.45	47.77	Linux x86 64-bit
4	c		21-May-21 23:33	03-Jun-21 22:30	03-Jun-21 23:00	12.1.0.2.0	29.78	1,332.82	311.45	44.75	Linux x86 64-bit

WORKLOAD REPOSITORY report for

DB Name	DB Id	Unique Name	Role	Edition	Release	RAC	CDB
PRD1	641091274	prd1	PRIMARY	EE	19.0.0.0.0	NO	NO

Instance	Inst Num	Startup Time	User Name	System Data Visible
prd1	1	17-Nov-21 03:40	SYSMAN	YES

Host Name	Platform	CPUs	Cores	Sockets	Memory (GB)
	AIX-Based Systems (64-bit)	24	3		50.00

	Snap Id	Snap Time	Sessions	Cursors/Session
Begin Snap:	126167	17-Nov-21 05:00:47	114	1.3
End Snap:	126177	17-Nov-21 15:00:32	144	2.0
Elapsed:		599.75 (mins)		
DB Time:		1,202.93 (mins)		

Capturing observed memory (RAM) utilization from AWR

Report Summary

Cache Sizes

- All values are in Megabytes
- Listed in order of instance number, I#
- End values displayed only if different from Begin values

	Memory Target		Sga Target		DB Cache		Shared Pool		Large Pool		Java Pool		Streams Pool		PGA Target		
I#	Begin	End	Begin	End	Begin	End	Begin	End	Begin	End	Begin	End	Begin	End	Begin	End	Log Buffer
1			81,920		49,920		20,480		1,024		1,792		8,192		30,720		128.00
2			81,920		54,784		20,224		1,024		256		5,120		30,720		128.00
3			81,920		54,528		20,480		1,024		256		5,120		30,720		128.00
4			81,920		54,528		20,480		1,024		256		5,120		30,720		128.00
Avg			81,920		53,440		20,416		1,024		640		5,888		30,720		128.00
Min			81,920		49,920		20,224		1,024		256		5,120		30,720		128.00
Max			81,920		54,784		20,480		1,024		1,792		8,192		30,720		128.00

Memory Statistics

- PGA Aggregate Target Statistics
- Process Memory Summary

Back to Top

PGA Aggregate Target Statistics

- all stats are reported in MegaBytes

	PGA Aggr Target		Auto PGA Target		PGA Mem Alloc		Auto Workareas		Manual Workarea		Global Mem Bound	
I#	Begin	End	Begin	End	Begin	End	Begin	End	Begin	End	Begin	End
1	30,720.00	30,720.00	13,407.92	21,126.97	23,616.78	16,561.72	1,212.94	2,699.75	8,862.85	126.98	1,024.00	1,024.00
2	30,720.00	30,720.00	18,349.12	18,994.61	20,009.45	21,692.76	1,586.95	3,608.22	235.14	250.40	1,024.00	1,024.00
3	30,720.00	30,720.00	19,368.50	19,365.81	23,210.45	18,993.92	4,903.25	2,395.80	38.92	87.45	1,024.00	1,024.00
4	30,720.00	30,720.00	19,539.50	19,553.51	18,810.11	16,849.74	1,195.16	706.38	173.14	102.75	1,024.00	1,024.00

Memory Statistics

	Begin	End
Host Mem (MB):	51,200.0	51,200.0
SGA use (MB):	21,440.0	21,440.0
PGA use (MB):	1,392.7	1,529.9
% Host Mem used for SGA+PGA:	44.60	44.86

Use the larger of the values displayed for each Begin and End capture for PGA.

Capturing observed I/O utilization from AWR

IOStat by File Type (per Second)

- Total Reads includes all Filetypes: Data File, Temp File, Archive Log, Backups, Control File, Data Pump Dump File, Flashback Log, Log File, Other, etc
- Total Writes includes all Filetypes: Data File, Temp File, Log File, Archive Log, Backup, Control File, Data Pump Dump File, Flashback Log, Log File, Other, etc

	Reads MB/sec			Writes MB/sec				Reads requests/sec			Writes requests/sec			
I#	Total	Data File	Temp File	Total	Data File	Temp File	Log File	Total	Data File	Temp File	Total	Data File	Temp File	Log File
1	57,475.37	57,446.01	25.16	44.93	6.50	18.14	10.23	76,566.57	76,436.76	116.93	1,396.94	513.49	40.59	814.40
2	80,123.45	80,040.80	12.06	40.67	6.66	12.14	12.03	102,279.97	102,075.71	160.22	1,387.56	504.47	37.77	814.41
3	28,944.44	28,873.44	12.89	34.70	6.07	10.76	8.94	49,829.96	49,629.32	165.20	1,350.46	493.71	28.92	802.85
4	22,839.97	22,784.28	17.72	56.37	9.38	23.46	11.46	42,212.09	42,034.23	150.76	1,621.09	780.39	57.37	744.95
Sum	189,383.23	189,144.53	67.84	176.67	28.62	64.50	42.66	270,888.58	270,176.02	593.11	5,756.05	2,292.07	164.65	3,176.61
Avg	47,345.81	47,286.13	16.96	44.17	7.15	16.13	10.66	67,722.15	67,544.00	148.28	1,439.01	573.02	41.16	794.15

IOStat by Filetype summary

- 'Data' columns suffixed with M,G,T,P are in multiples of 1024 other columns suffixed with K,M,G,T,P are in multiples of 1000
- Small Read and Large Read are average service times
- Ordered by (Data Read + Write) desc

Filetype Name	Reads: Data	Reqs per sec	Data per sec	Writes: Data	Reqs per sec	Data per sec	Small Read	Large Read
Data File	1.1T	2864.78	31.364M	10.2G	14.79	.29M	1.05ms	1.26ms
Temp File	38.6G	7.25	1.099M	42.4G	7.07	1.207M	9.73ms	30.08ms
Control File	3.3G	6.04	.094M	558M	0.99	.016M	113.14us	
Log File	0M	0.00	0M	2.5G	16.09	.071M	125.00us	
TOTAL:	1.1T	2878.07	32.558M	55.6G	38.94	1.583M	1.06ms	30.04ms

Fill in the Twelve Data Points to the AWR Worksheet

AWR detail collected by database instance

DB Name	Instance Name	Host Name	Elapsed Time (mins)	DB Time (mins)	DB CPU (s)	CPUs	Cores	Memory (GB)	%busy CPU	SGA use (MB)	PGA use (MB)	Read Throughput (MB/s)	Write Throughput (MB/s)	Read IOPS	Write IOPS
DBPROD1	dbprod1	host1	360.36	12,092.44	598,197	64	32	32.78	67.90%	15,200	16,384	8,421.00	496.90	3,574.22	200.93
DBPROD2	dbprod2	host2	10,080.34	82,662.54	2,641,050	16	8	377.87	67.60%	128,000	5,898	19.00	4.10	677.20	165.70
DBPROD3	dbprod3	host3	2,200.94	61,276.52	2,599,622	64	32	32.78	52.00%	13,440	4,000	68.28	2.62	4,223.16	184.55
DBRAC	dbrac11	host1	10,139.38	38,675.37	1,883,035	64	32	32.78	38.00%	12,000	3,200	686.81	89.74	250.71	39.60
DBRAC	dbrac12	host2	10,200.54	51,790.93	2,444,168	64	32	32.78	27.00%	7,500	3,200	451.19	24.27	281.83	21.89
DBPROD4	dbprod4	host4	10,079.25	89,952.03	3,955,400	64	32	32.78	16.00%	4,272	2,400	3,399.70	222.54	1,221.39	114.71

1. Copy the first two columns from the AWR sheet to the “Calculated detail by database instance” on the **Calculations** sheet.
2. Copy the third column, (Host Name) from the AWR sheet to the “Aggregated calculations by host” on the **Calculations** sheet.

Please add rows by specifying DB Name, Instance Name, or Host Name as appropriate.
The calculated cells already have formulas to aggregate the AWR information added on the other worksheet appropriately.

Calculated detail by database instance

DB Name	Instance Name	%DB Time of Elapsed Time	%DB CPU of server capacity	Total ORA (GB)	Total IOPS	Total Throughput (MB/s)	Est'd Azure vCPUs
DBPROD1	dbprod1	3355.655%	43.229%	31	3,775.15	8,917.90	25.17
DBPROD2	dbprod2	820.037%	27.292%	131	842.90	23.10	6.15
DBPROD3	dbprod3	2784.107%	30.759%	17	4,407.71	70.90	20.88
DBRAC	dbrac11	381.437%	4.836%	15	290.31	776.55	2.86
DBRAC	dbrac12	507.727%	6.240%	10	303.72	475.46	3.81
DBPROD4	dbprod4	892.448%	10.220%	7	1,336.10	3,622.24	6.69
		8741.412%	122.575%	210	10,955.89	13,886.15	65.56

Aggregated calculations by host

Host Name	%DB Time of Elapsed Time	%DB CPU of server capacity	Total ORA (GB)	Total IOPS	Total Throughput (MB/	Est'd Azure vCPUs
host1	489.861%	6.236%	58.42	13294.20	10014.53	1.22
host2	664.648%	10.443%	153.94	2300.62	538.05	1.66
host3	2784.107%	30.759%	17.03	4407.71	70.90	20.88
host1	489.861%	6.236%	58.42	13294.20	10014.53	1.22
host2	664.648%	10.443%	153.94	2300.62	538.05	1.66
host4	892.448%	10.220%	6.52	1336.10	3622.24	6.69
Total	5985.572%		448.28	36933.45	24798.30	33.35

Aggregated calculations by database

DB Name	%DB Time of Elapsed Time	Total vRAM (GiB) consumed only by Oracle	Est'd Azure vRAM for server	Total IOPS	Total Throughput (MB/s)	Est'd Azure IOPS for peak load	Est'd Azure Throughput (MB/s) for peak load	Est'd Azure vCPUs for avg load	Est'd Azure vCPUs for peak load	DB Size in TB	Bkup Info
DBPROD1	3355.655%	31	46	3,775.15	8,917.90	7,550.30	17,835.80	26	39		
DBPROD2	820.037%	131	196	842.90	23.10	1,685.80	46.20	7	11		
DBPROD3	2784.107%	17	26	4,407.71	70.90	8,815.42	141.80	21	32		
DBRAC	444.772%	25	38	594.03	1,252.01	1,188.06	2,504.02	7	11		
DBPROD4	892.448%	7	10	1,336.10	3,622.24	2,672.20	7,244.48	7	11		

Final Steps

- Paste each UNIQUE DB Name into the “Aggregated Calculations by Database”
- Rest of calculations will fill in for Azure sizing
- Manually add the size of the database and the backup storage required.

Assessment Tips

- Don’t try to boil the ocean- remember, you’re finding the “bucket” the workload fits into.
- Be careful not to confuse different DBCPU values and IO data.
- Take your time filling out the AWR worksheet. If you take the time to fill this out correctly, all the rest of the data will populate.
- Always size from the values for peak, (in RED) but also make intelligent decisions of where you can consolidate, working with the customer to identify opportunities.
- Add notes about versioning, SLAS for uptime, etc. in the notes section on the worksheet.
- Try not to change the fudge factors unless necessary.
- Get assistance when trying to migrate Exadata workloads. Extra guidance is required to have success with these engineered systems.

Simplify the Shift to the Cloud

- Migrate your tools that you already use to monitor and manage the database on-prem into the cloud whenever possible.
  - For Oracle, we implement Oracle Enterprise Manager, (Cloud Control) to ensure the cloud environment looks just like their onprem one.
  - Redgate SQL Monitor, Solarwinds SQL Sentry, Dynatrace, Idera Uptime Infrastructure Monitor, etc.
  - Use features to automate OS patching using Azure Linux/Windows automated patching service.
- Incorporate DevOps automation to the cloud changes FIRST.

Simulate PaaS in IaaS

- Use Azure Managed Instance for SQL Server
- Use Lifecycle Management Pack with Oracle Enterprise Manager to automate monitoring, management and database patching.
- Use Linux Automated Patching, (preview) to automate OS patching of VMs.
- Introduce Azure services to simplify the current products used onprem
- Automate using DevOps, including deployment builds with Terraform, Ansible, etc.

Modernization is the Future

- Identify, once moved, what can be modernized
  - OBIEE and other older BI tools can be replaced with Power BI. Great opportunities to do this with Cognos and Tableau with Power BI connected to Oracle.
  - Hyperion and Essbase are heavy Excel users- engage them with Power BI, a natural fit and evolve to analytics, AI and ML.
  - Most were Tableau, but they're Power BI due to the similarity with their favorite tool, Excel. very intrigued by
- New projects with a focus on PostgreSQL PaaS from schemas inside large Oracle databases, allowing for easier learning curve for developers.
- Then propose PostgreSQL with PaaS for current IaaS Oracle.
- Heavier data warehouse loads to Azure data lakes or Synapse
- New projects , (greenfield) that will lean on Azure PaaS solutions, not Oracle ones.

In Summary…

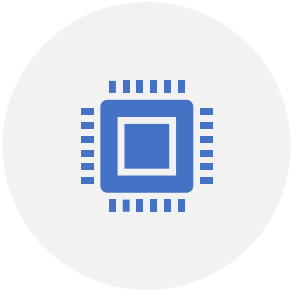

Demonstrate offerings for Linux images that could be built into a customer supported library of images to be reused.

E and M series deployments that could support the resource heavy Oracle workloads. Newest V5 series has doubled some IO workload max limits vs. a few months ago.

Lift and shift the workload, not the hardware, which was a new approach. Focused on the Automatic Workload Repository, (AWR) supports 85% of databases and creates almost 1:1 licensing parity for customers going to Azure.

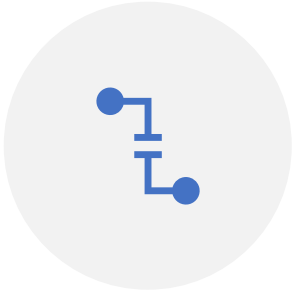

Heavy dependence on ANF and 3rd party storage solutions to bypass the VM level throttling on IOPs/MBPs for large Oracle workloads. Work closely with PM teams to help guide support for high IO workloads.

Your Go-To for Everything Oracle on Azure IaaS

- Oracle on Azure Azure/Oracle-Workloads-for-Azure: Oracle workloads for Azure infrastructure as a service (Oracle Azure IaaS) (github.com)
- Partners FOCUSED on All Oracle Solutions for Azure
- Data Architecture Blog - Microsoft Tech Community

- SQL Server Performance Guidelines on Azure: Checklist: Best practices & guidelines - SQL Server on Azure VM | Microsoft Docs
- Oracle on Azure: Oracle solutions on Microsoft Azure - Azure Virtual Machines | Microsoft Docs
- Understanding AZ and AS: Availability options for Azure Virtual Machines - Azure Virtual Machines | Microsoft Docs
- Virtual Machine and Disk Performance: Virtual machine and disk performance - Azure Virtual Machines | Microsoft Docs
- Azure Premium Storage: Azure Premium Storage: Design for high performance - Azure Virtual Machines | Microsoft Docs
- Azure Network Performance for IaaS: Optimize VM network throughput | Microsoft Docs
- Infrastructure Automation: Use infrastructure automation tools - Azure Virtual Machines | Microsoft Docs

Azure IaaS References

Oracle Specific Reference Reads

- Azure IO Performance for the RDBMS DBA- Part I (dbakevlar.com)
- Why a One-Week Report for AWR Sizing in Azure (dbakevlar.com)
- Prepping an Oracle Database for a Cloud Migration (dbakevlar.com)
- Oracle Storage Snapshots with Azure Backup - Microsoft Tech Community
- Backup Scenarios for Oracle on Azure IaaS - Microsoft Tech Community
- Estimate Tool for Sizing Oracle Workloads to Azure IaaS VMs - Microsoft Tech Community
- Script to create a simple Oracle VM configuration: GitHub - tigormanmsft/oravm: Azure CLI (bash) script to fully automate the creation of an Azure VM to run Oracle database

Thank you