



Exadata Cloud at Customer

Who we are




The Company

- > Founded in 2010
- > More than 80 specialists
- > Specialized in the Middleware Infrastructure
 - > The invisible part of IT
- > Customers in Switzerland and all over Europe

Our Offer

- > Consulting
- > Service Level Agreements (SLA)
- > Trainings
- > License Management



About me

Clemens Bleile

Principal Consultant & Technology Leader Oracle


+41 78 677 51 09

clemens.bleile[at]dbi-services.com




Agenda

- 1.ExaCC - What is it and why?
- 2.Customer Objectives
- 3.Things to consider
- 4.Migration Approach / Planning
- 5.Backup & Recovery
- 6.Summary



ExaCC – What is it and why?

- > Deployment models
- > Data flow
- > Why?



ExaCC – What is it?

Deployment Models

On-Premises

Exadata Database Machine



Customer Data Center
Purchased
Customer Managed

Cloud at Customer

Exadata Cloud @Customer



Customer Data Center
Subscription
Oracle Managed

Public Cloud

Exadata Cloud Service



Oracle Cloud
Subscription
Oracle Managed

ExaCC – What is it?

Responsibilities

The Database

- Customer managed
- Provisioned through the Oracle Cloud
- Customer patching
- All data is encrypted and **not** visible to Oracle
- Cost is based on consumption

Customer



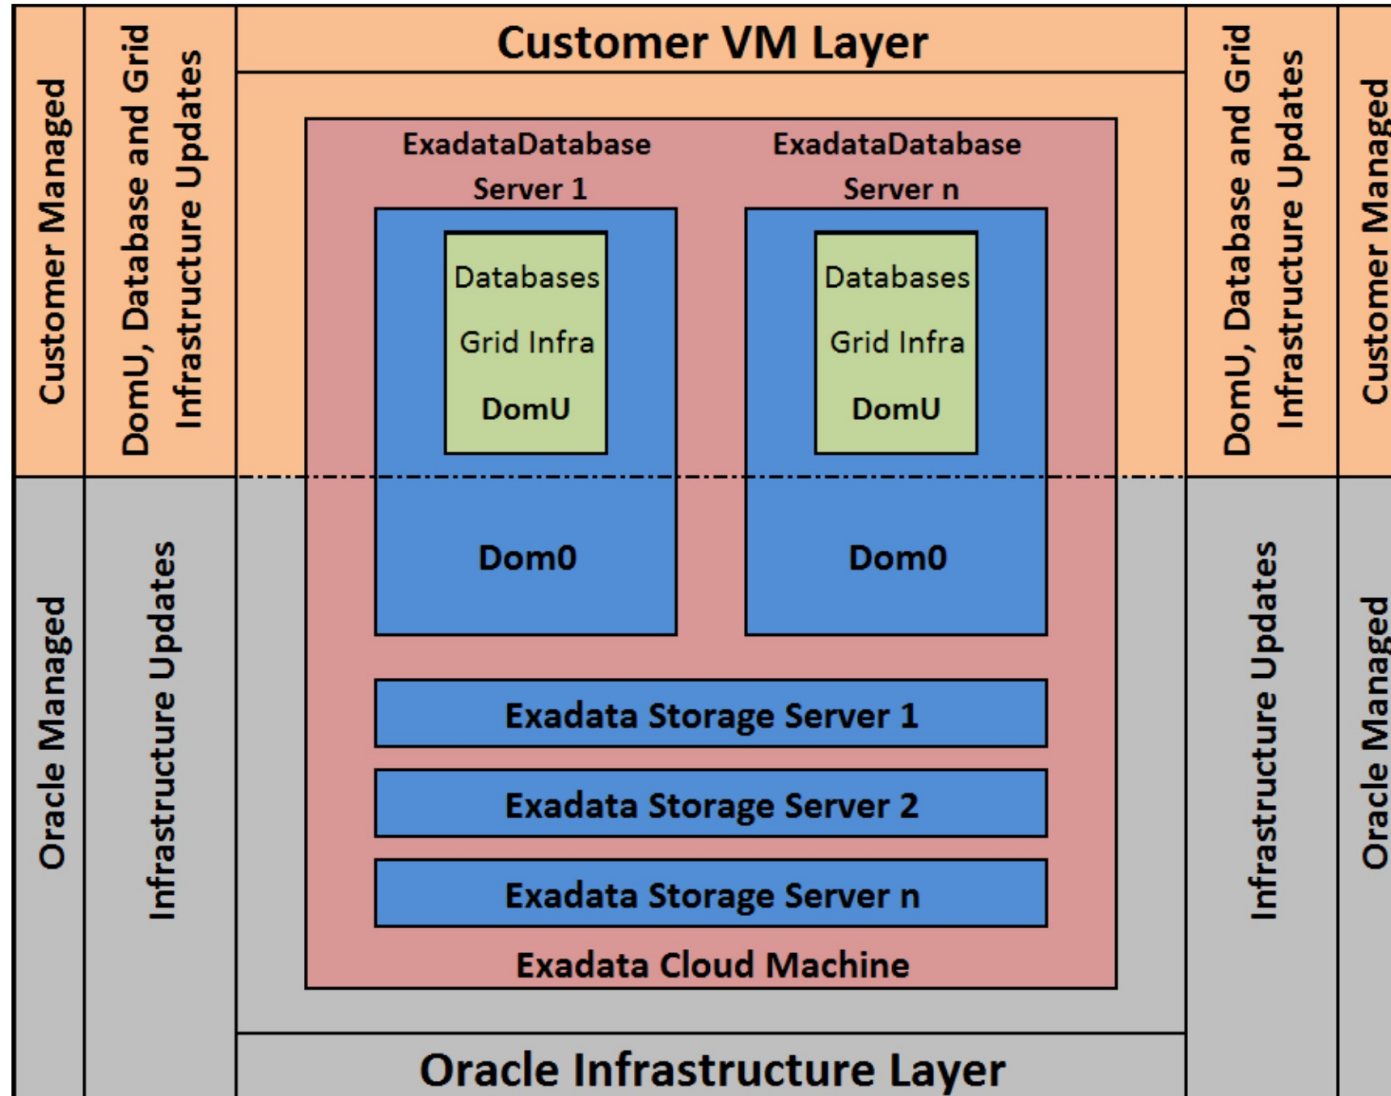
Exadata Infrastructure

- Oracle managed
- Transparently patched via Oracle Cloud
- Isolated from the customer
- Cost is based on the “shape” sizing

ORACLE

ExaCC – What is it?

Responsibilities

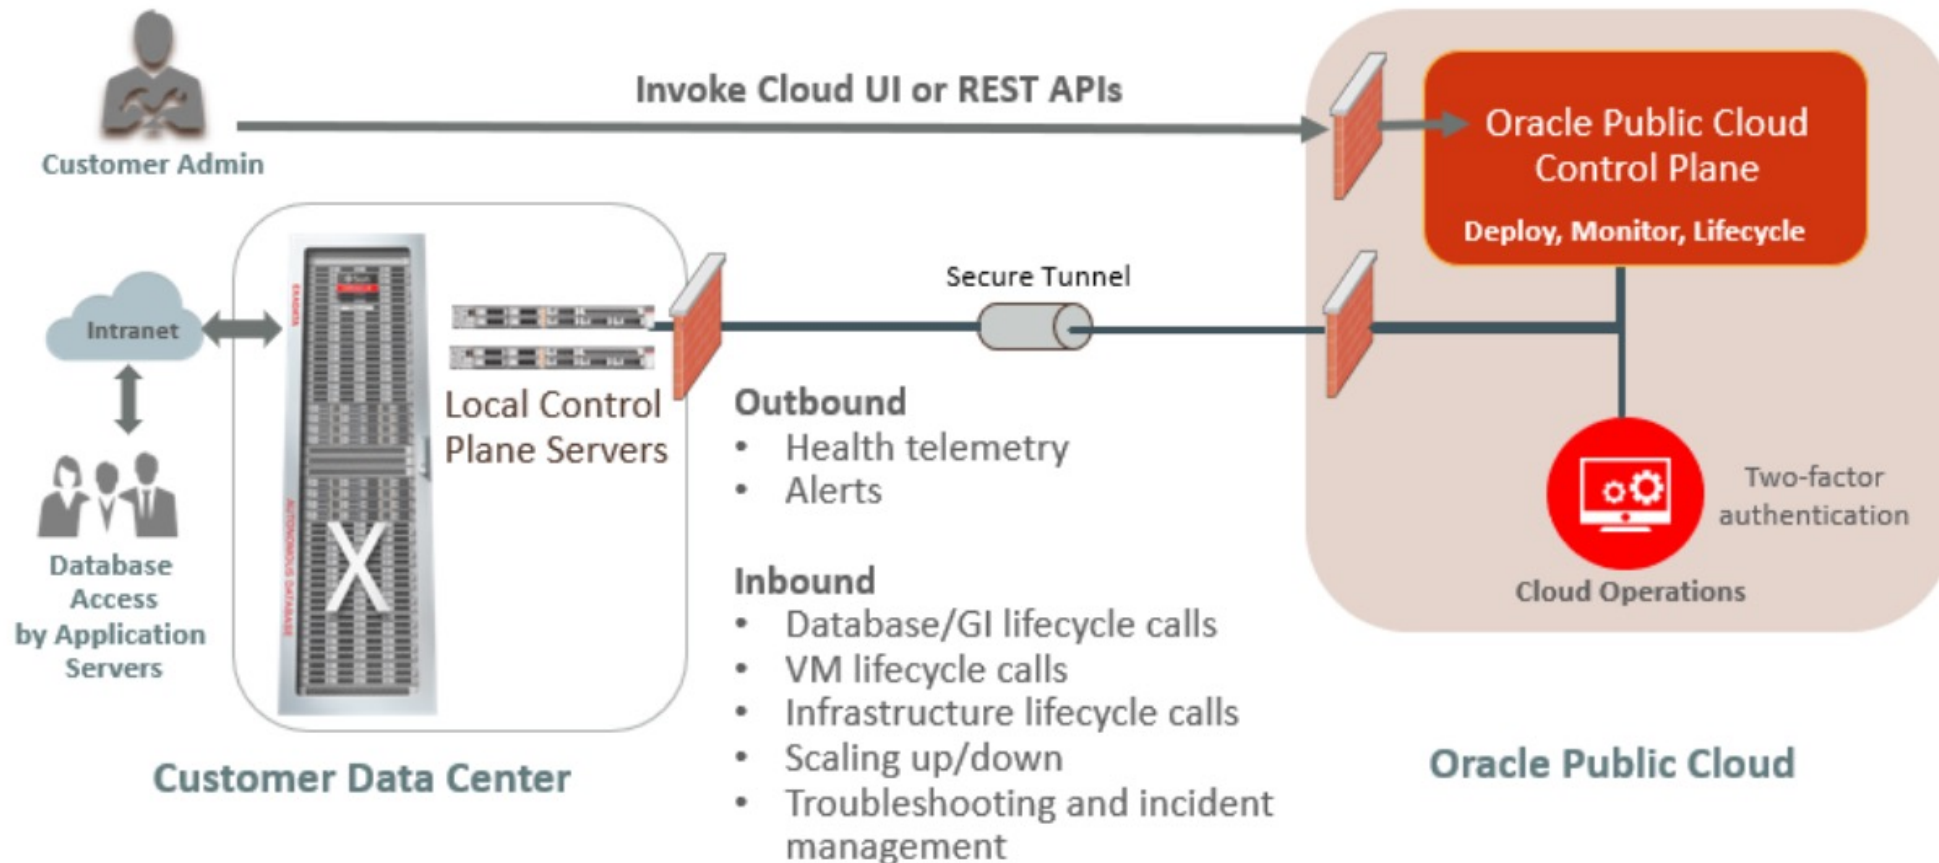


Original source: Oracle

ExaCC – What is it?

Management flow

Gen 2 Exadata Cloud at Customer— Management Flow



ExaCC – Why?

Data Residency Law / Data Sovereignty

- > Keep my data local in my DC

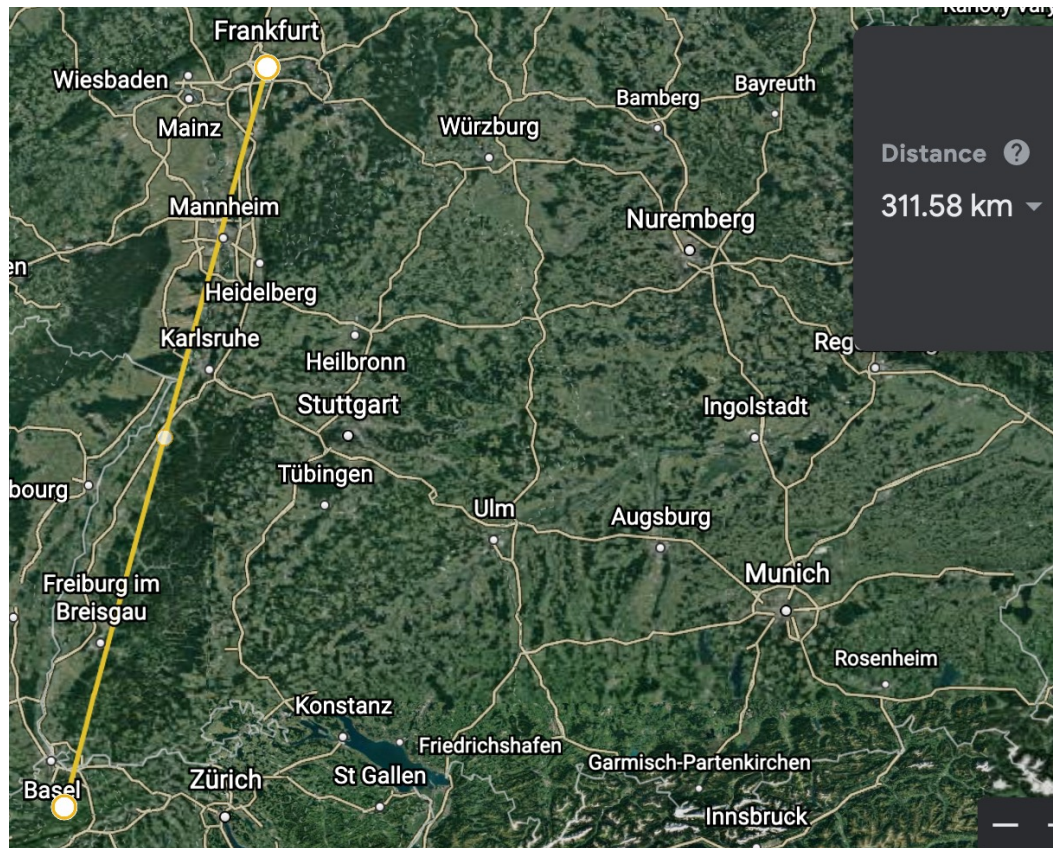


Original source: Oracle

ExaCC – Why?

Latency


- > Keep application and DB close together



WAN Latency Estimator

Enter a number in either field, then click outside of the text box.

Distance in:	<input type="radio"/> Miles (mi)	<input checked="" type="radio"/> Kilometres (KM)
Distance between sites:	<input type="text" value="312"/>	KM
Fibre path adjustment:	<input type="text" value="10"/>	%
Metro fibre & local loop length:	<input type="text" value="80"/>	KM
Speed of light in fibre:	<input type="text" value="200"/>	KM / millisecond
Equipment latency:	<input type="text" value="1"/>	millisecond
Round-trip time (RTT):	6.23	millisecond



Latency

Oracle Cloud Infrastructure Network Test

</> Javascript Tag

Results for downlink and latency tests from your connection to Oracle Cloud Infrastructure. For test result details, place the mouse cursor over result bars in the table below.

Service	Downlink [1 - 128KB / 4 threads]						Downlink [256KB - 10MB / 2 threads]						Latency								
	Mb/s	0	11	22	33	44	55	Mb/s	0	50	100	150	200	250	ms	0	10	20	30	40	50
Oracle Cloud Compute EU-FRANKFURT-AD-1	27.72							214.67							17						
Oracle Cloud Compute EU-ZURICH-1-AD-1	28.44							214.01							14						

<http://cloudharmony.com/speedtest-for-oracle:compute-eu-zurich-1-and-oracle:compute-eu-frankfurt-1>

Latency

Oracle Cloud Infrastructure Network Test

</> Javascript Tag

Results for downlink and latency tests from your connection to Oracle Cloud Infrastructure. For test result details, place the mouse cursor over result bars in the table below.

Service	Downlink [256KB - 10MB / 2 threads]							Downlink [1 - 128KB / 4 threads]							Latency						
	Mb/s	0	53	106	159	212	265	Mb/s	0	14	28	42	56	70	ms	0	91	182	273	364	455
Oracle Cloud Compute EU-ZURICH-1-AD-1	228.45							40.51							13						
Oracle Cloud Compute AP-SYDNEY-1-AD-1	2.09							2.12							408.5						

<http://cloudharmony.com/speedtest-for-oracle:compute-eu-zurich-1-and-oracle:compute-ap-sydney-1>

ExaCC – Why?

Latency

Oracle Cloud Compute - EU-FRANKFURT	
Latency	
Status.....	Success
Tests Performed.....	12
Tests Successful.....	12
Median.....	18 ms
Mean.....	17.92 ms
Fastest.....	17 ms
Slowest.....	19 ms
90th Percentile.....	17 ms
75th Percentile.....	17 ms
25th Percentile.....	18.5 ms
10th Percentile.....	19 ms
Standard Deviation.....	0.76
Data Transferred.....	72 B

Amazon EC2 - eu-central-1	
Latency	
Status.....	Success
Tests Performed.....	12
Tests Successful.....	12
Median.....	19 ms
Mean.....	19.25 ms
Fastest.....	18 ms
Slowest.....	22 ms
90th Percentile.....	18.5 ms
75th Percentile.....	19 ms
25th Percentile.....	19.5 ms
10th Percentile.....	20 ms
Standard Deviation.....	1.01
Data Transferred.....	72 B


Microsoft Azure Virtual Machines - eu-west-1	
Latency	
Status.....	Success
Tests Performed.....	12
Tests Successful.....	12
Median.....	26 ms
Mean.....	26.08 ms
Fastest.....	25 ms
Slowest.....	27 ms
90th Percentile.....	25.5 ms
75th Percentile.....	26 ms
25th Percentile.....	26.5 ms
10th Percentile.....	27 ms
Standard Deviation.....	0.64
Data Transferred.....	72 B

<http://cloudharmony.com/speedtest-for-aws:ec2-eu-central-1-and-oracle:compute-eu-frankfurt-1-and-azure:compute-eu-west>

ExaCC – Why?

ExaCC versus Public Cloud: Network Latency

- Speed of light in vacuum: 300m/microsecond
- Switches/Hubs/Firewalls add latency time
- E.g. Local latency: 0.25ms
- Public cloud latency: 3.5ms



Factor 14



Low Latency is important if

- lots of data is transported between DB and App
- lots of fetches happen
- statements are called often (millions of times)

ExaCC may make the difference

ExaCC – Why?

Test the Latency between your client and database:



(Elapsed Time - DB Time) / network round trips

```
SQL> exec dbms_output.put_line(to_char(:roundtrips_end - :roundtrips_begin)||' network round trips.');
```

4953 network round trips.

```
SQL> exec dbms_output.put_line(to_char((:time_end - :time_begin)*10)||' ms elapsed time.');
```

18870 ms elapsed time.

```
SQL> exec dbms_output.put_line(to_char((:db_time_end - :db_time_start)/1000)||' ms DB time.');
```

876.625 ms DB time.


```
SQL> exec dbms_output.put_line(to_char(round((((:time_end - :time_begin)*10)-((:db_time_end - :db_time_start)/1000))/(:roundtrips_end - :roundtrips_begin),3)||' ms latency per round trip.');
```

3.633 ms latency per round trip.

<https://blog.dbi-services.com/script-to-calculate-the-network-latency-between-the-application-and-the-oracle-db-server>

Customer Objectives

- > Current state
- > Alternatives
- > Why ExaCC?
- > Target architecture



Objectives

Current state

Hardware

- > DBs running on Solaris servers
- > Capacity limit reached
- > Recurring performance problems
- > Reliability
- > Hardware maintenance: +15% per year

Software

- > DB license/support: +4% per year
- > No flexibility in terms of license with the fix number of CPUs
- > Need for additional licenses (Multitenant, Tuning Pack)



Objectives

Alternatives

- > Linux bare metal
- > VMWare
- > Oracle Database Appliance
- > Exadata Cloud@Customer

Criteria

- > Support
- > Patching
- > Price
- > Flexibility in terms of licensing
- > Availability
- > Isolation
- > Reuse of existing storage
- > DB options



Objectives

Why ExaCC?

Technical reasons

- > Flexibility and automation
- > Simplified management via Cloud portal and APIs
- > DB options included (multitenant, advanced security, compression, in-memory, etc.)
- > Hardware/firmware/hypervisor maintenance delegated to Oracle
- > High Availability (RAC cluster)
- > Engineered system (optimized for Oracle DBs)



Original source: Oracle

Objectives

Why ExaCC?

Security

- > Easy patching with bundles and rolling upgrades
- > Encrypted DB (customer has key)



Commercial reasons

- > Cloud model (flexibility)
- > Excellent conditions
 - > Hardware + Licenses included
 - > Pure Opex model
 - > Services offered



Original source: Oracle

Objectives

Why ExaCC?

ExaCC X8M specs

	Base*	Quarter Rack	Half Rack	Full Rack
DB servers	2	2	4	8
Max OCPU	48	100	200	400
Total Memory	656 GB	2,780 GB	5,560 GB	11,120 GB
Storage servers	3	3	6	12
Usable disk storage	74.8 TB	149.7 TB	299.4 TB	598.7 TB

Original source: APACOUK

Objectives

Why ExaCC?

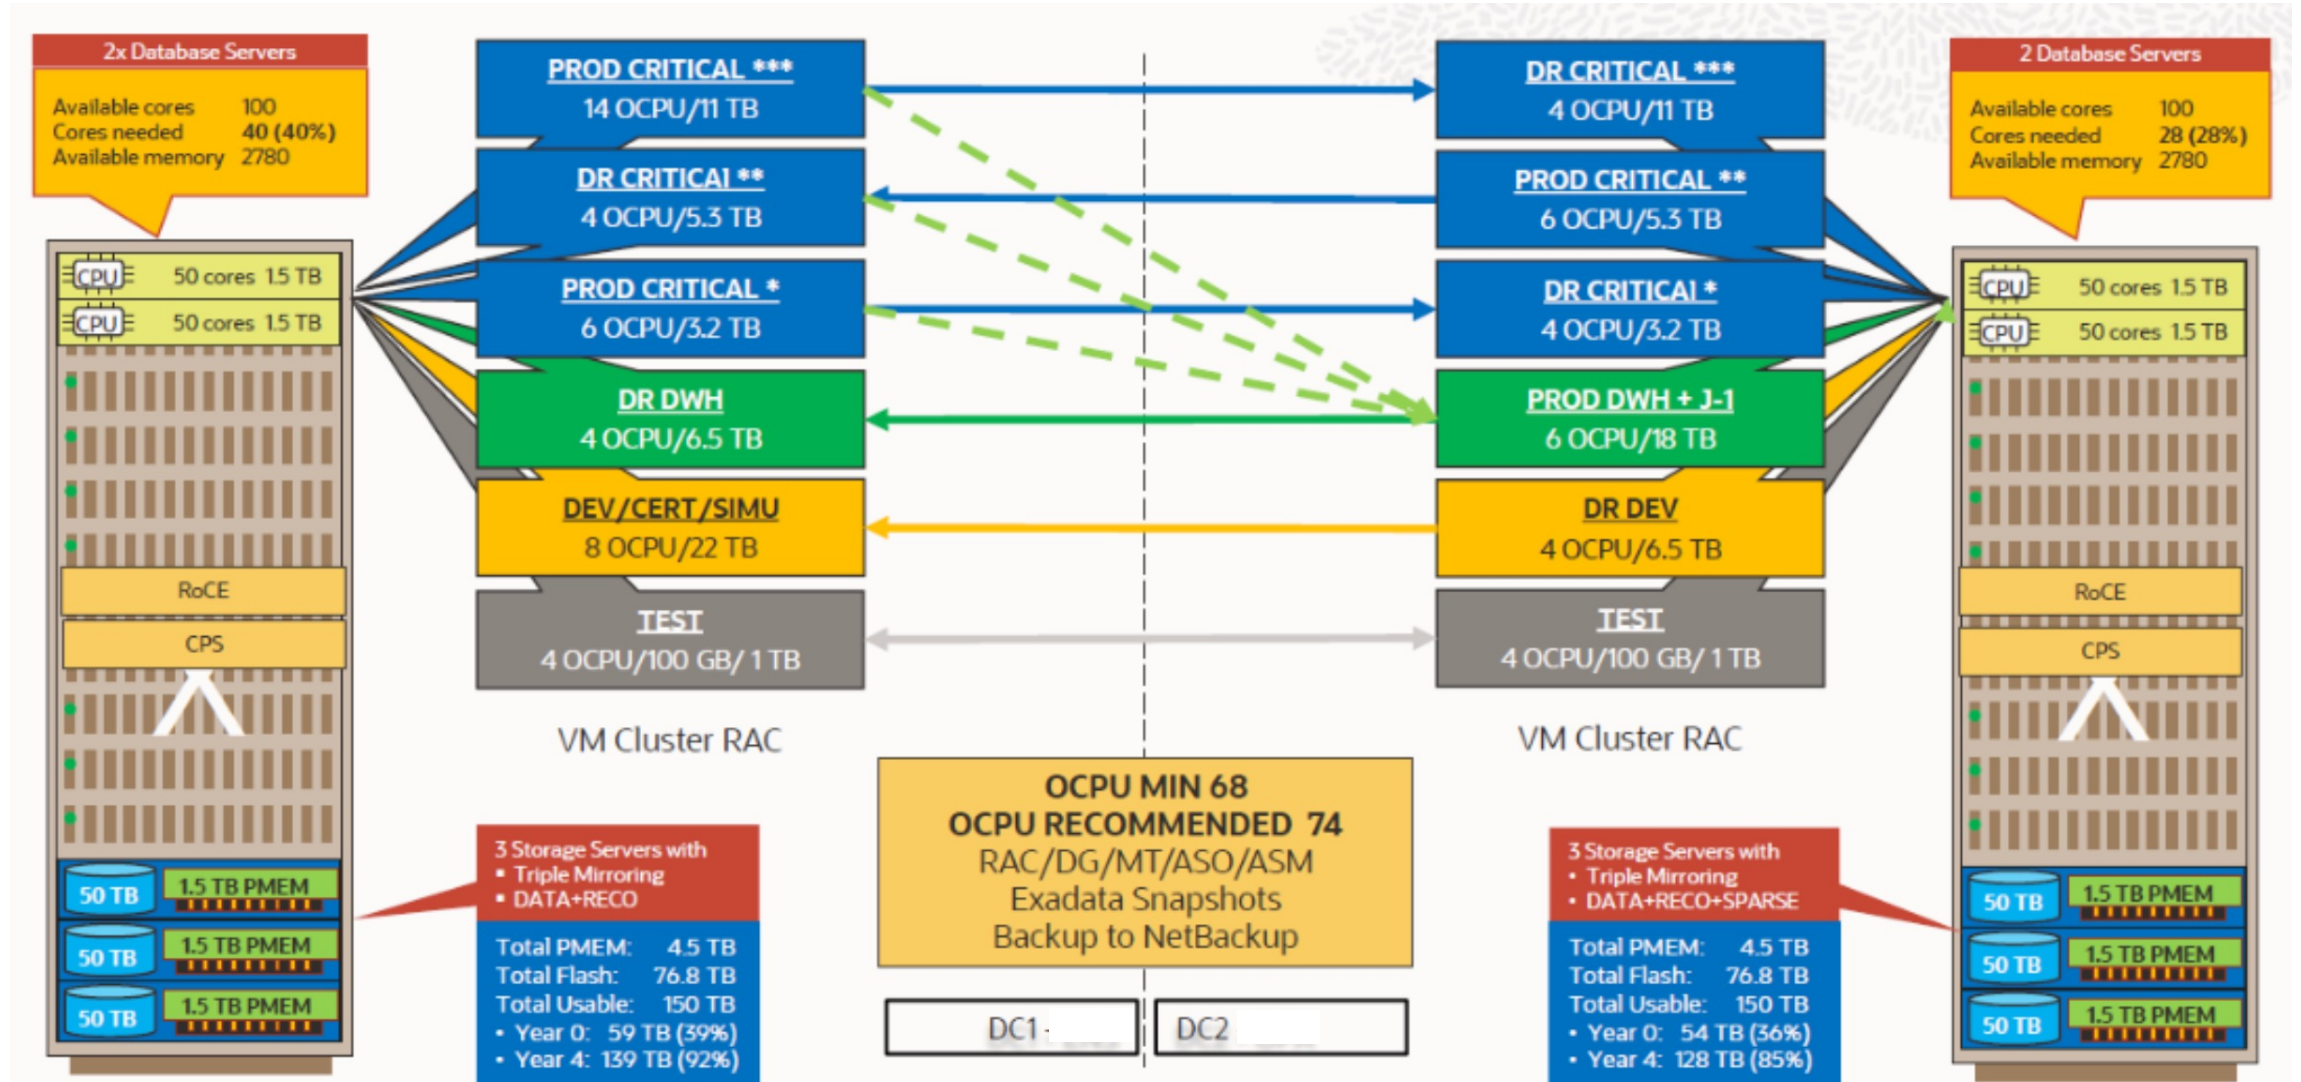
ExaCC **X9M** specs

	Base*	Quarter Rack	Half Rack	Full Rack
DB servers	2	2	4	8
Max OCPU	48	100 124	200 248	400 496
Total Memory	656 GB	2,780 GB	5,560 GB	11,120 GB
Storage servers	3	3	6	12
Usable disk storage	74.8 TB	149.7 TB 192 TB	299.4 TB 384 TB	598.7 TB 769 TB

Original source: APACOUK


Objectives

Target architecture (X8M Quarter Rack)



Things to consider

- > From Capex to Opex
- > Education
- > Endian change
- > It's owner by Oracle



Things to consider

From Capex to Opex

- > No (less) Capex
 - > No purchase of database server and storage
 - > Still need a bit more capex (networking, rack space) than for public cloud
 - > Easier budgeting and approval process

- > Opex
 - > Subscription
 - > Less Hardware required



Things to consider

Education / Train the DBAs and Developers

Exadata

- > Flash Cache
- > Storage Indexes
- > Smart Scans
- > HCC
- > IORM



RAC / Grid Infrastructure

- > Clusterware
- > ASM

Things to consider

Education / Train the DBAs and Developers

Multitenant

- > Operations
- > Advanced Administration

Snapshots

- > PDB Snapshot Copy
- > Snapshot Hierarchy

Additional features

- > Advanced security
- > Compression
- > In-Memory

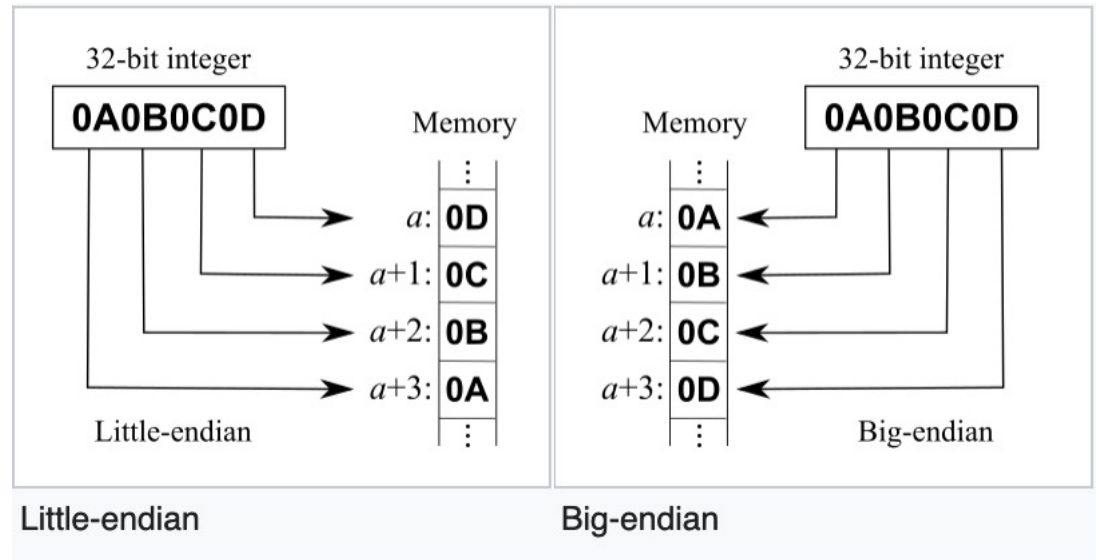


Things to consider

Endian Change

Solaris (Big Endian) -> Linux (Little Endian)

- > Limits the migration possibilities
- > Data Pump (expdp/impdp)
- > Logical replication (e.g. Golden Gate)
- > Transportable Tablespace
- > Full Transportable Export/Import
- > ZDM (Zero Downtime Migration) ?




Things to consider

Hardware is owned by Oracle


Exadata owned by Oracle

- > New migration after 4-5 years
 - > New ExaCC
 - > On-Premises
 - > Public Cloud
- > Pro: Easier migration with more possibilities due to all pre-work has been done already
 - > Linux
 - > ASM
 - > Cluster
 - > Multitenant
 - > New tablespace concept



Migration Approach / Planning

- > Migration Approach
- > Planning



Pre-Migration study

- > Use the time before the Exadata is available
- > List all DBs
- > List resource consumption by all DBs
- > Plan Exadata resources required
- > Educate
- > Plan for changes
 - > Characterset
 - > Tablespaces
 - > Used features (e.g. compression)

Migration

- > Doing
- > Plan time and steps for each and every DB-migration



Measure what resources are needed on the Exadata

- > Get resource consumption from AWR-History (avg/peak db-time, db-cpu-time)
- > Compute the hardware requirements on the ExaCC

Put together a list of all DBs

- > Criticality
- > Type
- > Data Guard
- > SGA/PGA-size
- > DB-params (like processes, undo_retention, etc)
- > Characterset
- > Temp-Size Max
- > Migration method
- > Lowest Client version
- > Etc.



Statistics and plans for the future


- > Versions
- > Growth trend
- > VM sizing



Migration Approach

Solaris (Big Endian) -> Linux (Little Endian)

- > Data Pump (expdp/impdp)
- > Logical replication (e.g. Golden Gate)
- > Transportable Tablespace
- > Full Transportable Export/Import
- > ZDM (Zero Downtime Migration) ?

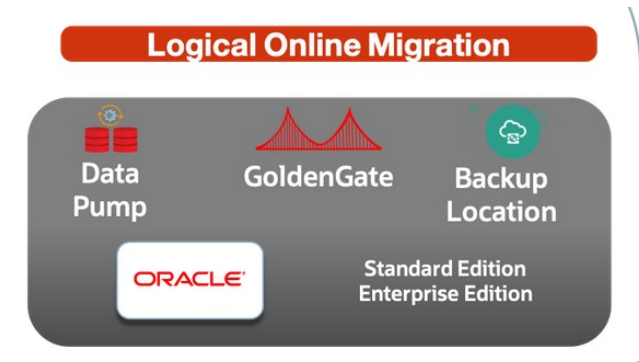
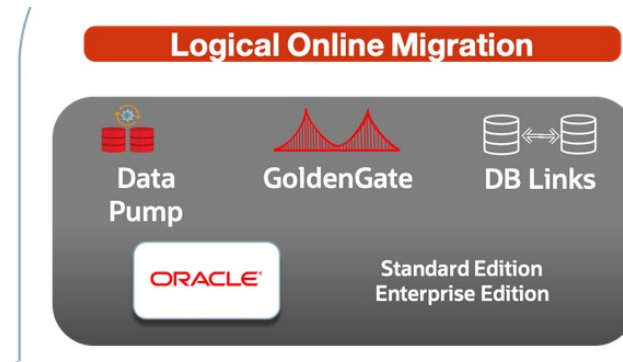


Migration Approach

Zero Downtime Migration






ZDM allows

- > Physical Migration
 - > Offline
 - > Online
- > Logical Migration
 - > Offline (Data Pump)
 - > With Backup location
 - > With Database Links
 - > Online (Golden Gate + Data Pump)
 - > With Database Links
 - > With Backup location



Migration Approach

Zero Downtime Migration

	 Exadata On Premises	 Oracle Database Cloud Services	 Exadata Cloud Service	 Exadata Cloud at Customer	 Autonomous Database Shared	 Autonomous Database Dedicated
Physical Offline Migration	✓	✓	✓	✓		
Physical Online Migration	✓	✓	✓			
Logical Offline Migration		✓	✓	✓	✓	✓
Logical Online Migration		✓	✓	✓	✓	✓

Solaris as
Source was
not supported
when planning
the migration

Migration from Solaris & AIX based Source Databases

Support for Cross-Platform migration is available for Solaris and AIX based Source Databases. Customers can leverage this feature to migrate Oracle Autonomous Database and Co-Managed Cloud Oracle Database targets using the standard offline methodology.



- Unfortunately too late

Solaris (Big Endian) -> Linux (Little Endian)

- > Data Pump (expdp/impdp)
- > Logical replication (e.g. Golden Gate)
- > Transportable Tablespace
- > Full Transportable Export/Import
- ~~ZDM (Zero Downtime Migration) ?~~




Decision was to use the following methods

- > Data Pump (expdp/impdp) for almost all DBs
- > Logical replication (e.g. Golden Gate) for the most critical DBs in terms of downtime


Why not using Transportable Tablespaces / Full Transportable Exp/Imp?

- > Change Tablespace concept
- > Change charsetset during migration
- > Change from consolidated schemas to PDBs



Backup & Recovery, Patching

- > History
- > Alternative methods
- > Other considerations



Backup & Recovery History

- > Backups using shell- and rman-scripts
- > Self developed
- > Limited number of DBAs knowing the details of the scripts used
- > Considerations when RMAN-Catalog or Tape Library is not available



Backup & Recovery

Alternative methods

- > Adjust and use self developed scripts
- > Automatic Backup/Restore provided by Oracle
- > Tool bkup_api
- > DMK DB-Backup




Backup & Recovery

Adjust and use self-developed scripts



Proven, tested and reliable



Has to be maintained



Only limited number of people knowing it




Difficult for operations team to do backup, restore and recovery

Backup & Recovery

Automatic Backup/Restore + bkup_api provided by Oracle



- ✓ Proven by Oracle
- ✓ Backups can easily be setup and configured through GUI and CLI
- ✓ Operations team can do manual backups and restore/recovery through GUI
- ✓ Backups can be adjusted through parameter-files




- ✗ Less flexible
- ✗ Cannot consider unavailability of Tape Library

Backup & Recovery

DMK DB-Backup



- ✓ Proven by dbi services
- ✓ Backups can easily be setup and configured through parameter files
- ✓ Very flexible
- ✓ Easy monitoring and alerting
- ✓ Easy to setup considering availability issues with RMAN catalog and tape library



- ✗ No GUI for operations team
- ✗ DBA Team and operations team need to become familiar with it

Backup & Recovery Decision

- > DMK DB-Backup
- > Use NFS as backup target with 5 days retention
- > Backup Backupsets from NFS to tape library regularly with 31 days retention
- > Backup to tape library to keep them 1 year once a month

> Issue: Backup Job scheduling

- > No oracle crontab on ExaCC
- > DB-Jobs would require additional monitoring
- > ssh only possible to public ip (not SCAN, not VIP)

- > Had to write a script to connect to one of the Cluster nodes and start the backup

<https://blog.dbi-services.com/exadata-cloud-at-customer-considerations-for-backup-recovery/>



Patching Oracle

Patching performed by Oracle

- > Dom0
- > Network switches
- > Power distribution units (PDUs)
- > ILOM
- > Exadata Storage Servers



Patching Customer

Customer responsible for the DomU, GI, DB:


- > VM OS (DomU)
- > Grid Infrastructure
- > DB (ORACLE_HOME)

How to patch?

- > Cloud Service Console or
- > API
- > Manual patching (should be an exception)




Summary



ExaCC is a good solution

- > If you have to keep your data local
 - > Security
 - > Latency
- > For Consolidation
- > To become faster
 - > “Fastest Oracle DB-platform”
- > To use more Oracle features
- > To move from Capex to Opex
- > To Prepare for the future
 - > Data growth
 - > Digitalization




Summary

Summary

Things to consider


- > The migration needs good planning
- > Education necessary
- > Use ZDM if possible
- > Use available features
- > Consolidate
- > Convert to AL32UTF8
- > Use Oracle's Cloud features
 - > Backup (to NFS)/Recovery
- > You have to migrate again



Summary



- ✓ Fast platform
- ✓ Security and low Latency
- ✓ Cloud service and agility
- ✓ Features
- ✓ No (less) Capex
- ✓ Prepared for future needs



- ✗ Migration project
- ✗ Education
- ✗ Key management with TDE and PDBs
- ~~✗ You still patch OS, GI, DB~~
- ~~✗ You have to migrate again~~

Any questions?

Please do ask!



We would love to boost
your IT-Infrastructure
How about you?