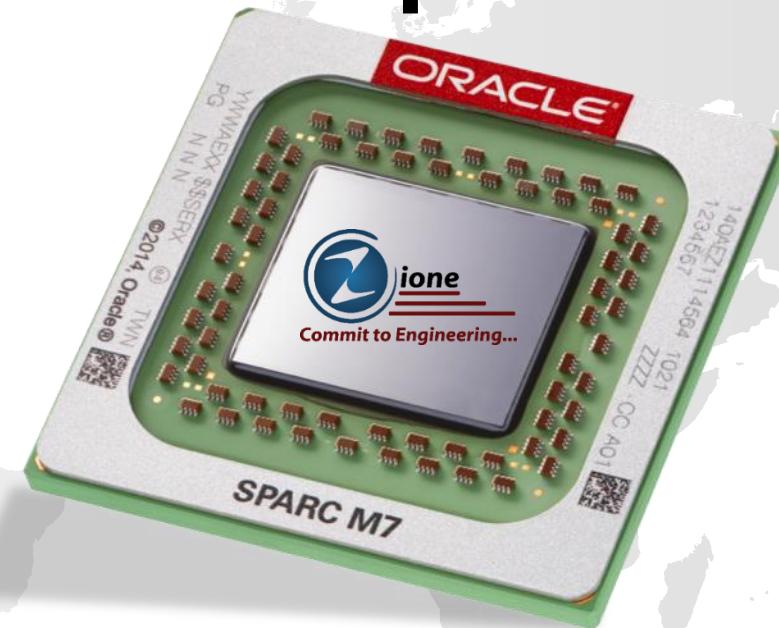


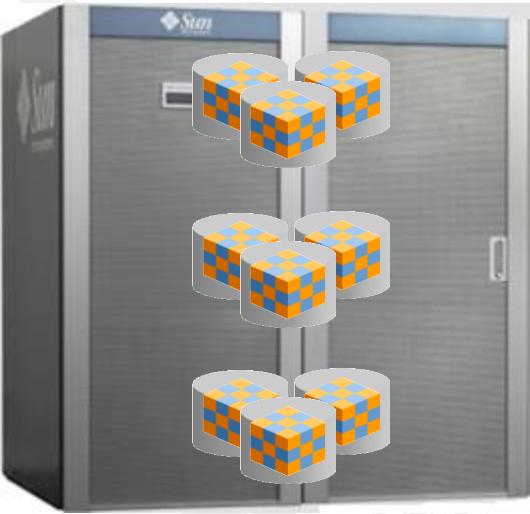
The Future of Sparc Roadmap to Linux



Presented By **Daniel Morgan, Oracle ACE Director Alum**
Phone **+1 612-240-3538**
Email **dmorgan@zionesolutions.com**

- Principal Advisor: Zione Solutions
  Oracle ACE Director Alumnus
- Educator
  Adjunct Professor, University of Washington, Oracle Program, 1998-2009
  Oracle Consultant: Harvard University
 - Guest lecturer at universities in the Americas, EMEA, and APAC
 - Technical Conference Lecturer ... 139 country visits, 43 unique, since 2008
- IT Professional
 - Celebrated 51 years of IT in 2020
 - The Morgan behind www.morganslibrary.org and www.dbsecworx.com
 - Past Chair: Washington Software Association Database SIG
 - Migration Architect for Oracle Consulting Services (OCS)





- From anywhere to anywhere
- Cross-Platform
- Cross-Version
- Risk Averse



- My very first "real" computer was an IBM 370-145
My first language, Fortran IV
- My first UNIX computer was from Sun Microsystems
Sun is still my choice as the best Oracle platform
- Many years ago your organization may have made
the best possible decision ... to choose Sun for your
servers and Solaris for your operating system
- But, technology is ever evolving and now we must
choose a new direction forward

SQL> ORA-03134: Les connexions à cette version de serveur ne sont plus prises en charge.

SQL> ORA-03134: Verbindungen mit dieser Server-Version werden nicht mehr unterstutzt.

SQL> ORA-03134: Forbindelser til denne serverversion understøttes ikke længere.

SQL> ORA-03134: Las conexiones a esta versión del servidor ya no son compatibles.

SQL> ORA-03134: As conexões com esta versão do servidor não são mais suportadas.

SQL> ORA-03134: ไม่รองรับการเชื่อมต่อกับเซิร์ฟเวอร์รุ่นนี้

SQL> ORA-03134: Verbindingen met deze serverversie worden niet langer ondersteund.

SQL> ORA-03134: Kaore i te tautokohia nga hononga ki tenei putanga kaitautoko.

SQL> ORA-03134: Huic servo poema iam non hospites sunt praesto est.

SQL> ORA-03134: Tämän palvelinversion yhteyksiä ei enää tueta.

SQL> ORA-03134: Chan eil ceanglaichean ris an tionndadh frithealaiche seo a 'faighinn taic tuilleadh.

SQL> ORA-03134; このサーバーバージョンへの接続はサポートされなくなりました。

SQL> ORA-03134: Tengingar við þessa netþjónarútgáfu eru ekki lengur studdar.

SQL> ORA-03134: Connections to this server version are no longer supported.

For more than 20 years the best choice has been Sun

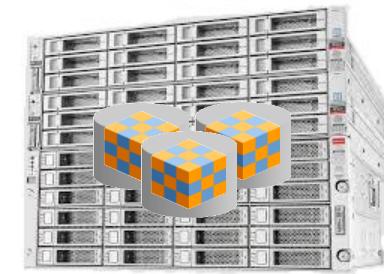
The best choice today is an Oracle Engineered System

Optimized for

- Stability
- Scalability
- Maintainability
- Supportability



The points was driven home to me by a couple of events 8 years ago







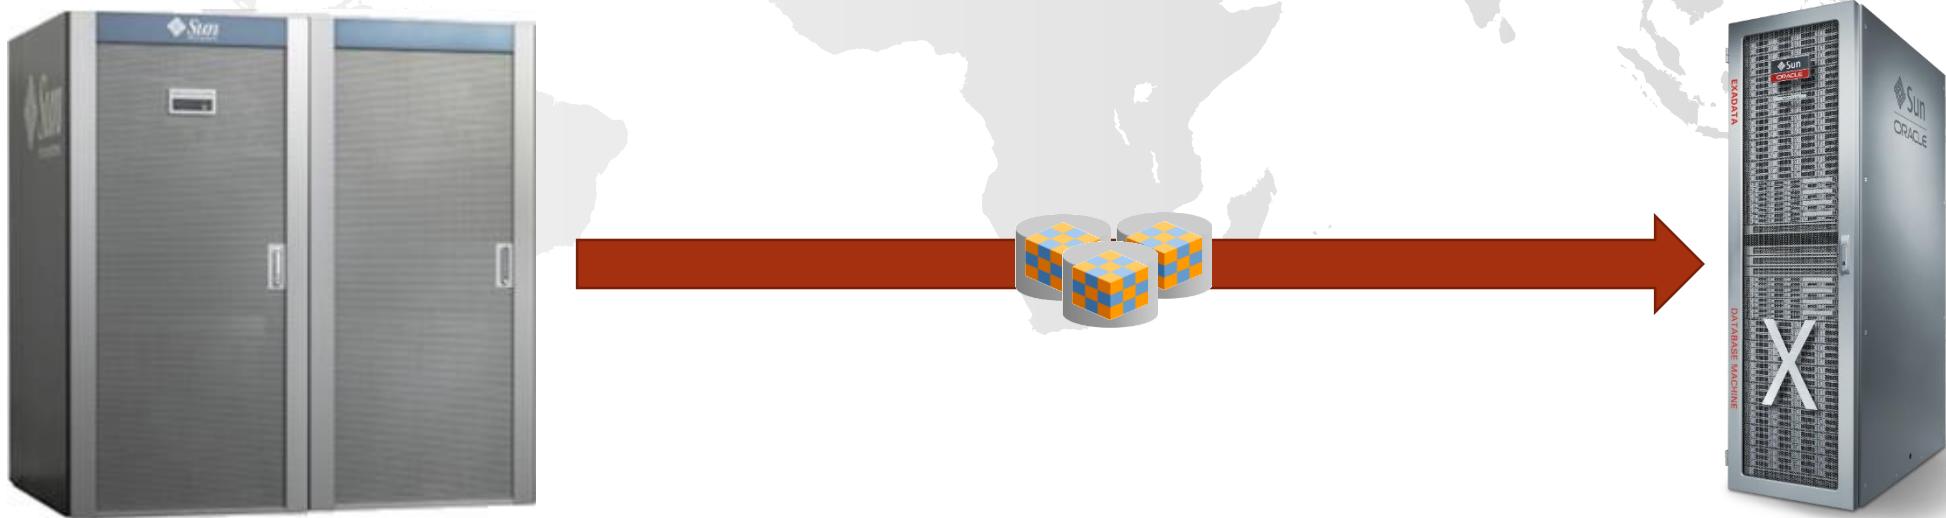
- Even An Untrained Technician Can Perform Maintenance



- Your management only cares about 2 things
 - QoS
 - TCO
- Don't take your eye off the ball



- For every step you take during a migration window
- You must know how to reverse it
- So that the source environment can be brought back up
- No matter where you are in the migration process



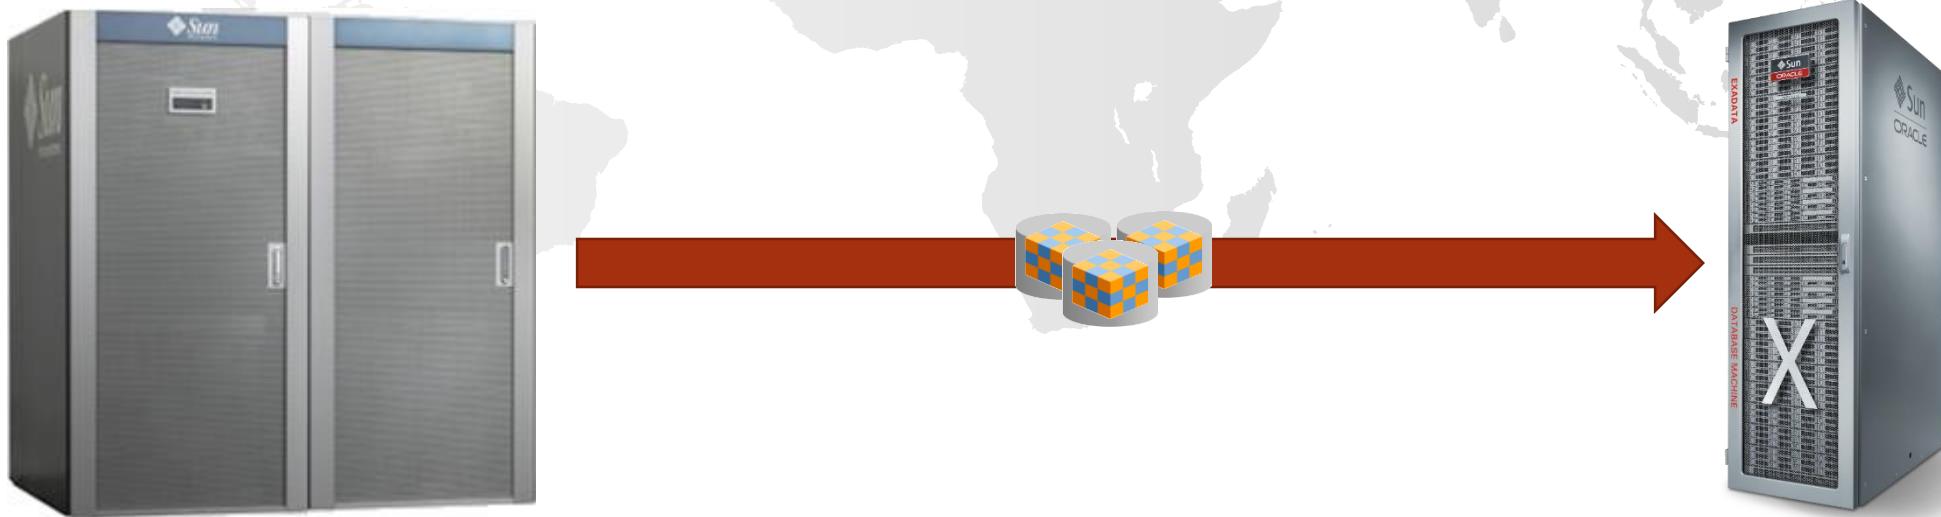


Migration Methodologies

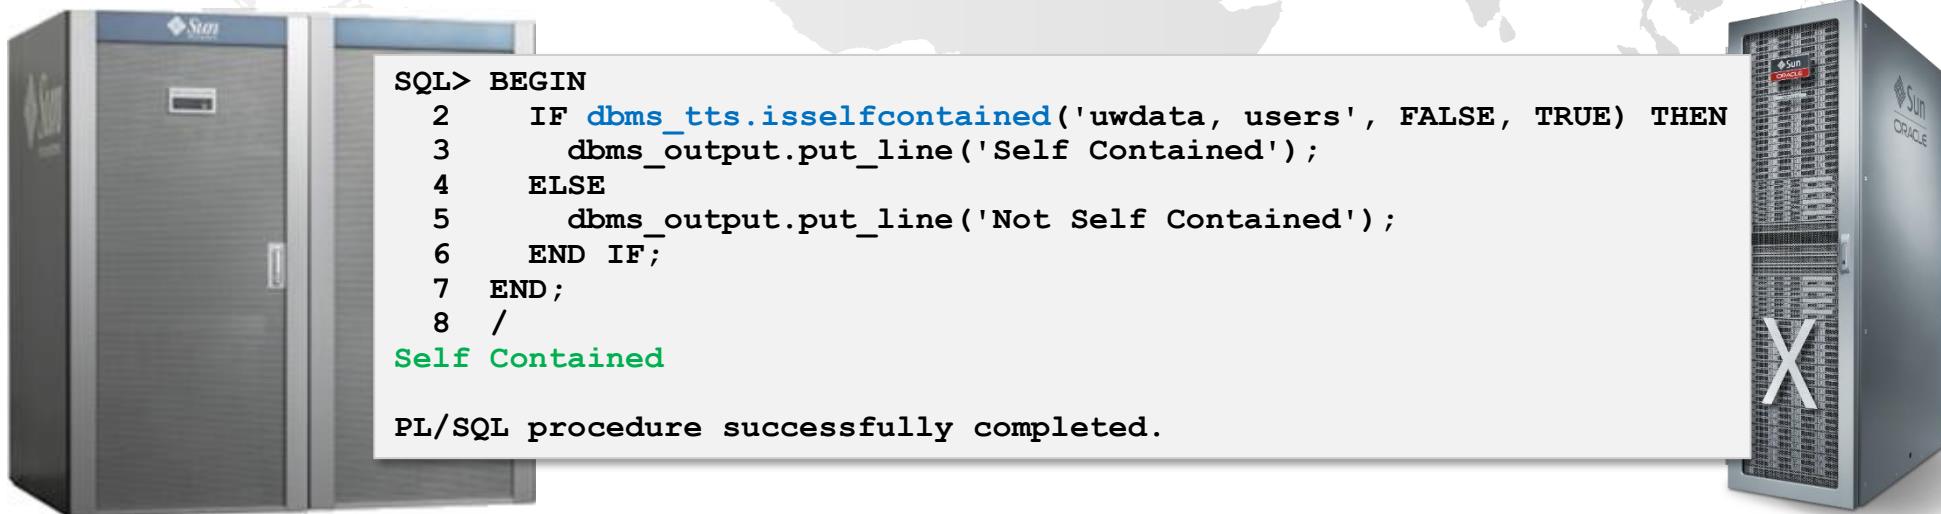
- Database links are the slowest way to migrate data and the method most likely to result in a bad migration
- Cannot be parallelized



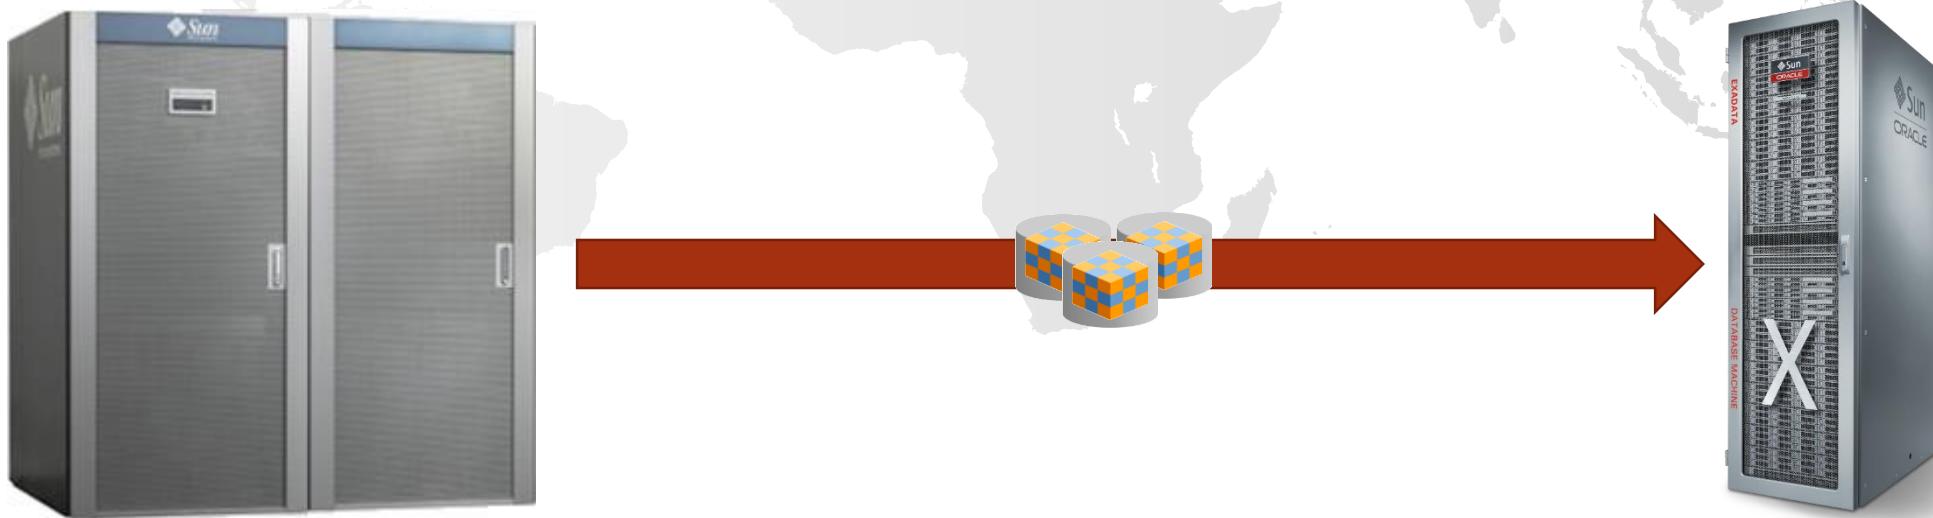
- Can export and import metadata and data
- Primary tool for Transportable Tablespaces
- Supports cross-platform migration
- To the extent permissible by network capacity and cpu ... parallelize
- A database greater than 10TB will likely not fit within the migration window
- Often fails on data loaded from IBM mainframes and NOSEGMENT indexes



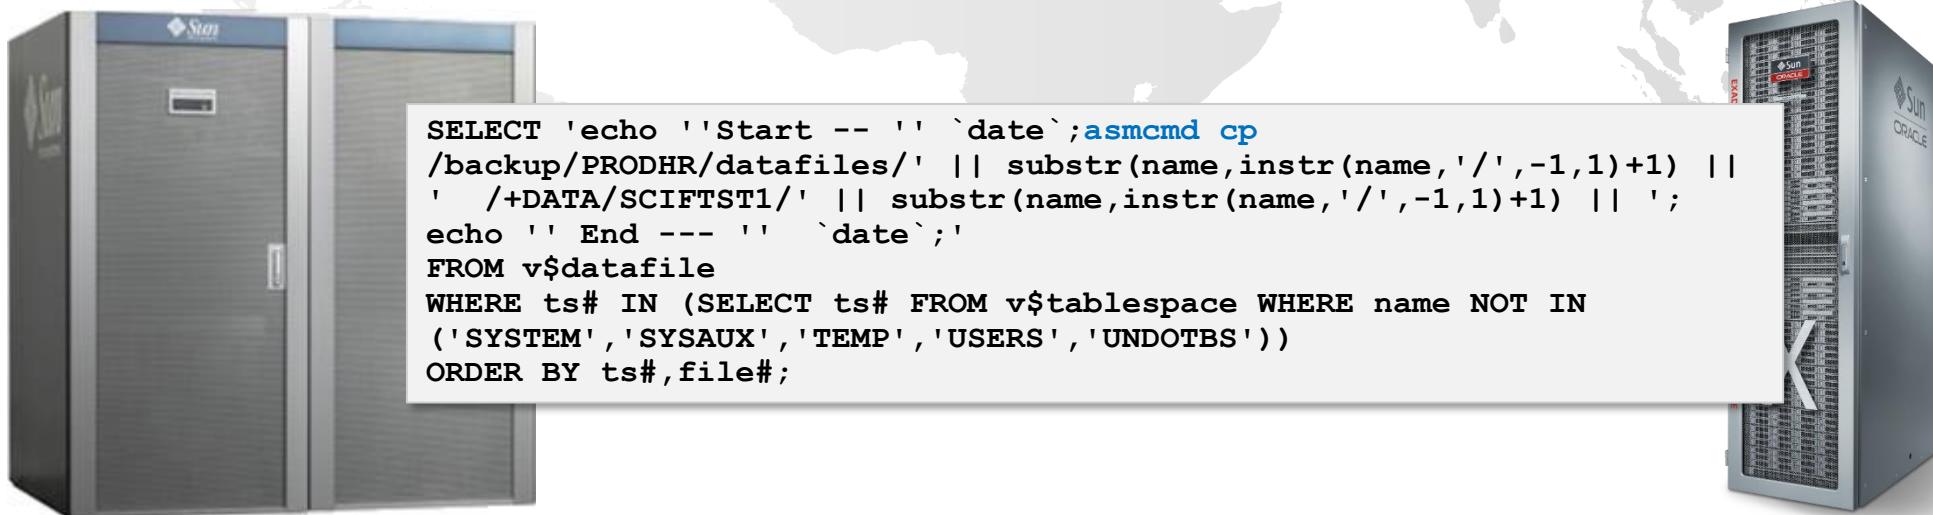
- Good choice for speed but require an explicit endian conversion
 - If ASM is involved ASMCMD and RMAN may be required
- Often has issues with tablespaces not being "self-contained"
- Be sure you are familiar with the supporting built-in packages
 - DBMS_TDB
 - DBMS_TTS



- Physical Data Guard cannot be used for Cross-Platform migrations due to the requirement that source-and-target be byte-for-byte identical
- Logical Data Guard has no advantages over other migration methods and very substantial weaknesses as it cannot guarantee the post-migration database's integrity



- An alternative to DBMS_FILE_TRANSFER and RMAN you can use ASMCMD to copy ASM data files



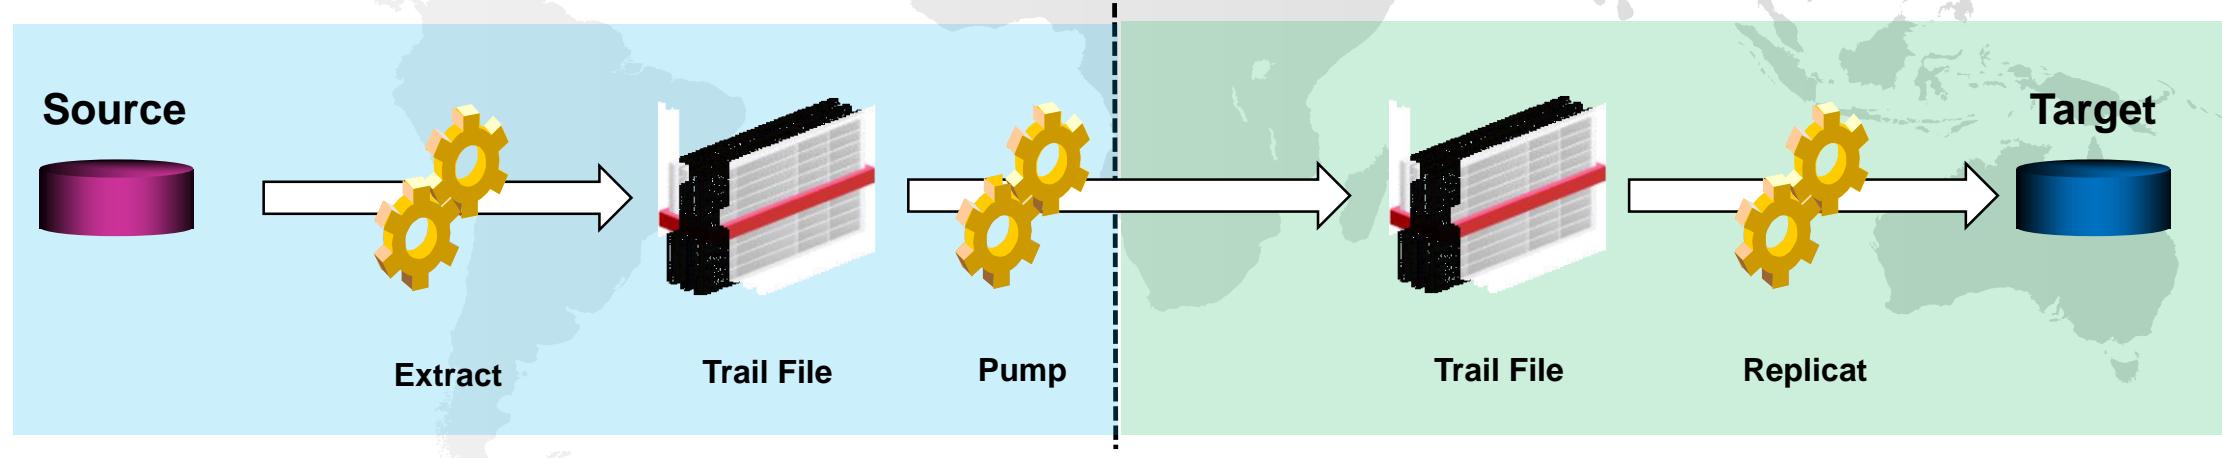
- The most undervalued data migration tool offered by Oracle as it allows for more flexibility than any other
- Copies binary files within a database or between databases
- GET and PUT perform transparent endian conversion on the target
- Cannot copy files in use: Requires an outage or data files put into hot backup mode



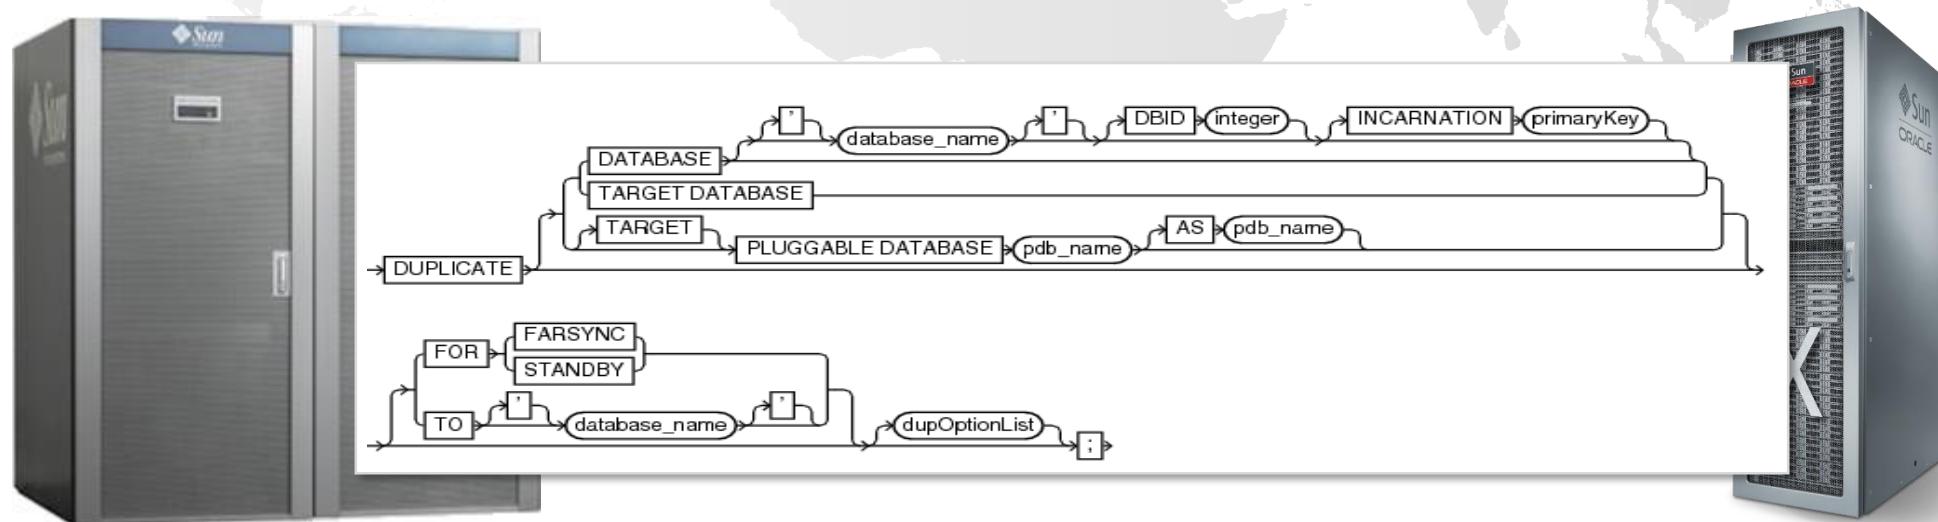
```
dbms_file_transfer.copy_file(  
    source_directory_object    IN VARCHAR2,  
    source_file_name          IN VARCHAR2,  
    destination_directory_object IN VARCHAR2,  
    destination_file_name     IN VARCHAR2);
```



- With Oracle GoldenGate you can often cut database migration times to what we call "near-zero down time" (NZDT)
- Near-zero can range from a literally zero-time no outage to a few minutes
- The determining factor as to whether near-zero can be achieved is whether the source and target databases can both be handling live transactions at the same time: And, the overlap does not cause transactions to fail



- RMAN can be used to migrate from Solaris to Linux
- Pros
 - Guarantees post-migration data integrity
 - Supports endian conversion
- Cons
 - Will bring a Solaris configuration to Exadata or ODA: Not something you want
 - Will bring historical and AWR data that will be irrelevant



- Converts data files to be transported to the destination host format and deposits the results in the specified location

```
CONVERT DATAFILE <datafile_name_list>
DB_FILE_NAME_CONVERT '<source_location> <destination_location>'
FROM PLATFORM <original_platform_name>
```

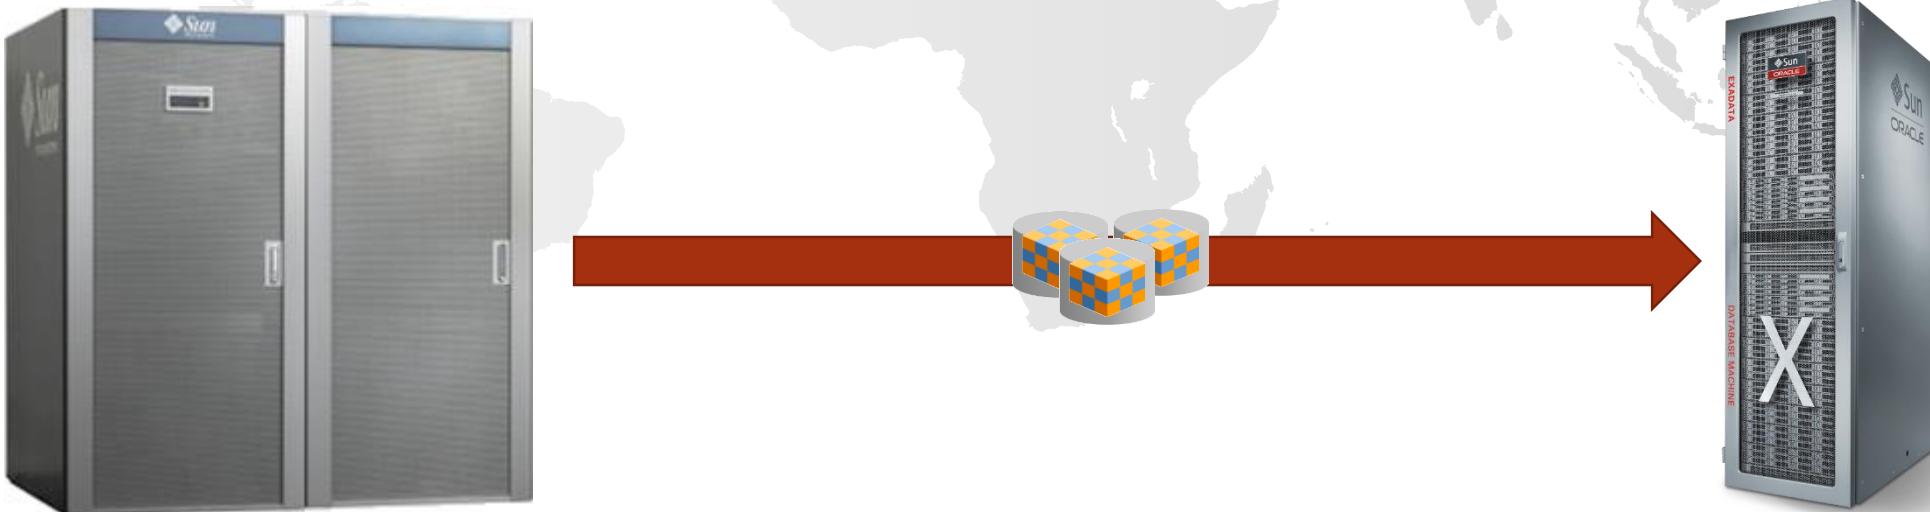
```
CONVERT DATAFILE
'/tmp/transport_from_aix/u01/uwdata01.dbf',
'/tmp/transport_from_aix/u01/uwdata02.dbf',
'/tmp/transport_from_aix/u03/users01.dbf',
'/tmp/transport_from_aix/u03/users02.dbf'
DB_FILE_NAME_CONVERT
'/app/oracle/product/oradata','/stage/oradata',
'/tmp/transport_from_solaris/hr','/stage/oradata'
FROM PLATFORM 'Linux IA (64-bit)'
```



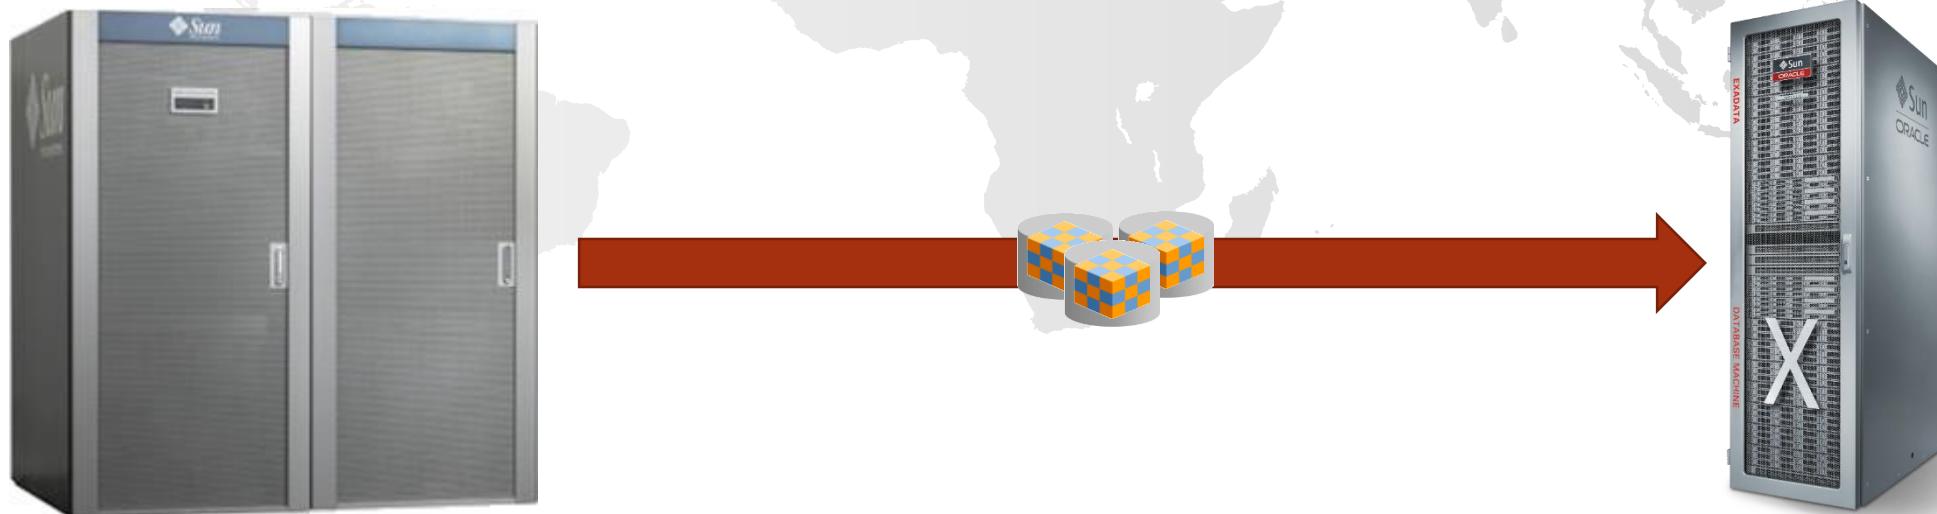


Discovery

- Don't make any decisions until you know the total outage window
 - From that window determine how much time will be required to
 - Move external objects
 - Validate database integrity
 - Migrate applications and tools
 - Perform application testing
 - The database migration window is what's left



- Unless you are doing a pure Lift & Shift where 1 Sparc server = 1 Linux server
 - and not moving from legacy to container architecture
 - and not moving from a file system to ASM/ACFS
 - and not changing edition or version
- Map source servers by name and IP to target servers
- Map source resource requirements (cpu, storage, network, memory) to the target environment



- A CDB runs from a single shared ORACLE_HOME
 - When you shut down the CDB you shut down all PDBs
 - When you patch the ORACLE_HOME you patch all PDBs
 - When you upgrade the ORACLE_HOME you upgrade all PDBs
 - Make a list of init parameters that can be set at the PDB level
 - Management and auditing will need to reflect the new architecture



SQL> desc v\$parameter

Name	Type
NUM	NUMBER
NAME	VARCHAR2 (80)
TYPE	NUMBER
VALUE	VARCHAR2 (4000)
...	
ISSES_MODIFIABLE	VARCHAR2 (5)
ISSYS_MODIFIABLE	VARCHAR2 (9)
ISPDB_MODIFIABLE	VARCHAR2 (5)
ISINSTANCE_MODIFIABLE	VARCHAR2 (5)
...	
CON_ID	NUMBER



- One critical issue with Solaris to Linux migrations is performance issues that may arise due to version and hardware changes
 - ADAPTIVE_QUERY_PLANNING
 - CURSOR_SHARING=similar
 - EVOLVING_BASELINES
- Prepare before the first migration by capturing AWR Reports from all sources and store them in a repository for future reference



```
SQL> show parameter adaptive
```

NAME	TYPE	VALUE
optimizer_adaptive_plans	boolean	TRUE
optimizer_adaptive_reporting_only	boolean	FALSE
optimizer_adaptive_statistics	boolean	FALSE
parallel_adaptive_multi_user	boolean	FALSE



- My last M9000 project, instance 10 of a very large RAC cluster

WORKLOAD REPOSITORY report for

DB Name	DB Id	Instance	Inst num	Startup Time	Release	RAC
OPM01P	782247420	opm01p10	10	05-Apr-11 15:04	11.1.0.7.0	YES

Host Name	Platform	CPUs	Cores	Sockets	Memory (GB)
usp9005b	Solaris[tm] OE (64-bit)	60	32	8	251.58

	Snap Id	Snap Time	Sessions	Cursors/Session
Begin Snap:	20730	05-Apr-11 15:51:15	55	.8
End Snap:	20754	05-Apr-11 22:04:02	54	.9
Elapsed:		372.79 (mins)		
DB Time:		409.79 (mins)		

Report Summary

Cache Sizes

	Begin	End		
Buffer Cache:	35,840M	35,840M	Std Block Size:	32K
Shared Pool Size:	25,600M	25,600M	Log Buffer:	244,776K

- The Oracle Database blocks can be 2K, 4K, 8K, 16K or 32K: 8K is the default
- A single database can simultaneously have multiple block sizes
- Knowing the block size(s) is an essential element of migration planning
- A migration is an excellent time to move everything to 8K blocks
- If migrating from stand-alone to RAC 16K and 32K blocks are will have a substantial negative impact on the interconnect performance



```
SELECT UNIQUE bytes/blocks AS BLOCK_SIZE
FROM dba_data_files
ORDER BY 1;
```



- Verify the source database's character set
- Character set changes are not impossible but they are dangerous
- If you are migrating from, for example WE8MSWIN1252 to AL32UTF8 it should work ... but there are no guarantees
- Do it pre-migration and you might destroy your production system
- Do it post-migration and it must be in our outage window
- Character set conversion needs to part of a test migration (POC)



desc v\$nls_parameters	
Name	Type
PARAMETER	NUMBER
VALUE	VARCHAR2 (4000)
CON_ID	NUMBER



- Compression on a source database is unlikely to impact migration success except in that it affects the ability to predict migration time
- Determine whether there are compressed tablespaces on the source
- If moving to a target with Hybrid Columnar Compression it is recommended that the HCC tablespaces be built prior to the migration
- Be sure that you check for compression at the tablespace, table, and index levels



```
SELECT owner, table_name
  FROM dba_tables
 WHERE compression IS NOT NULL
   AND compression <> 'DISABLED';

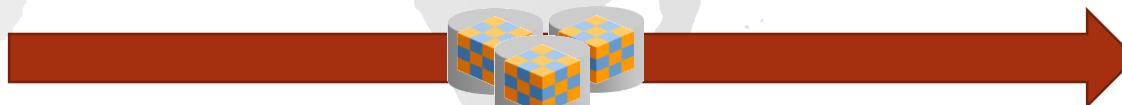
SELECT owner, table_name, index_name
  FROM dba_indexes
 WHERE compression IS NOT NULL
   AND compression <> 'DISABLED';
```



- Compare source server Sparc with target server x86

SPEC® CPU2017 Integer Rate Result																																															
Copyright 2017-2018 Standard Performance Evaluation Corporation																																															
Fujitsu	SPECrate2017 int base = 1180																																														
Fujitsu SPARC M12-2S	SPECrate2017 int peak = 1420																																														
CPU2017 License:	19																																														
Test Sponsor:	Fujitsu																																														
Tested by:	Fujitsu																																														
Benchmark result graphs are available in the PDF report.																																															
<table border="1"> <thead> <tr> <th colspan="2">Hardware</th></tr> </thead> <tbody> <tr> <td>CPU Name:</td><td>SPARC64 XII</td></tr> <tr> <td>Max MHz.:</td><td>4350</td></tr> <tr> <td>Nominal:</td><td>4250</td></tr> <tr> <td>Enabled:</td><td>192 cores, 16 chips, 8 threads/core</td></tr> <tr> <td>Orderable:</td><td>1 to 16 BBs; each BB contains 1 or 2 CPU chips; 2, 3, 4, .. 384 cores</td></tr> <tr> <td>Cache L1:</td><td>64 KB I + 64 KB D on chip per core</td></tr> <tr> <td>L2:</td><td>512 KB I+D on chip per core</td></tr> <tr> <td>L3:</td><td>32 MB I+D on chip per chip</td></tr> <tr> <td>Other:</td><td>None</td></tr> <tr> <td>Memory:</td><td>8 TB (256 x 32 GB 2Rx4 PC4-2400T-R)</td></tr> <tr> <td>Storage:</td><td>1 x 600 GB 10K RPM SAS (for system disk)</td></tr> <tr> <td>Other:</td><td>None</td></tr> </tbody> </table> <table border="1"> <thead> <tr> <th colspan="2">Software</th></tr> </thead> <tbody> <tr> <td>OS:</td><td>Oracle Solaris 11.3 SRU 24.4</td></tr> <tr> <td>Compiler:</td><td>C/C++/Fortran: Version 12.6 of Oracle Developer Studio</td></tr> <tr> <td>Parallel:</td><td>No</td></tr> <tr> <td>Firmware:</td><td>Fujitsu HCP Version 3040 released Oct-2017</td></tr> <tr> <td>File System:</td><td>tmpfs</td></tr> <tr> <td>System State:</td><td>Default</td></tr> <tr> <td>Base Pointers:</td><td>32-bit</td></tr> <tr> <td>Peak Pointers:</td><td>32-bit</td></tr> <tr> <td>Other:</td><td>None</td></tr> </tbody> </table>		Hardware		CPU Name:	SPARC64 XII	Max MHz.:	4350	Nominal:	4250	Enabled:	192 cores, 16 chips, 8 threads/core	Orderable:	1 to 16 BBs; each BB contains 1 or 2 CPU chips; 2, 3, 4, .. 384 cores	Cache L1:	64 KB I + 64 KB D on chip per core	L2:	512 KB I+D on chip per core	L3:	32 MB I+D on chip per chip	Other:	None	Memory:	8 TB (256 x 32 GB 2Rx4 PC4-2400T-R)	Storage:	1 x 600 GB 10K RPM SAS (for system disk)	Other:	None	Software		OS:	Oracle Solaris 11.3 SRU 24.4	Compiler:	C/C++/Fortran: Version 12.6 of Oracle Developer Studio	Parallel:	No	Firmware:	Fujitsu HCP Version 3040 released Oct-2017	File System:	tmpfs	System State:	Default	Base Pointers:	32-bit	Peak Pointers:	32-bit	Other:	None
Hardware																																															
CPU Name:	SPARC64 XII																																														
Max MHz.:	4350																																														
Nominal:	4250																																														
Enabled:	192 cores, 16 chips, 8 threads/core																																														
Orderable:	1 to 16 BBs; each BB contains 1 or 2 CPU chips; 2, 3, 4, .. 384 cores																																														
Cache L1:	64 KB I + 64 KB D on chip per core																																														
L2:	512 KB I+D on chip per core																																														
L3:	32 MB I+D on chip per chip																																														
Other:	None																																														
Memory:	8 TB (256 x 32 GB 2Rx4 PC4-2400T-R)																																														
Storage:	1 x 600 GB 10K RPM SAS (for system disk)																																														
Other:	None																																														
Software																																															
OS:	Oracle Solaris 11.3 SRU 24.4																																														
Compiler:	C/C++/Fortran: Version 12.6 of Oracle Developer Studio																																														
Parallel:	No																																														
Firmware:	Fujitsu HCP Version 3040 released Oct-2017																																														
File System:	tmpfs																																														
System State:	Default																																														
Base Pointers:	32-bit																																														
Peak Pointers:	32-bit																																														
Other:	None																																														
Memory:	768 GB (24 x 32 GB 2Rx8 PC4-2666V-R)																																														
Storage:	960 GB SAS SSD																																														
Other:	None																																														
Base Pointers:	64-bit																																														
Peak Pointers:	32/64-bit																																														
Other:	jemalloc memory allocator library, version 5.0.1																																														

- There is a high likelihood that cron jobs written for Solaris were written in korn shell or utilize commands that are incompatible with bash on Linux
- Rather than rewriting them ... throw them where they belong
- Eliminate the security risks by moving the functionality into a DBM_SCHEDULER internal or external job
- DBMS_SCHEDULER jobs seamlessly handle RAC and Data Guard, and collect performance metadata on every execution

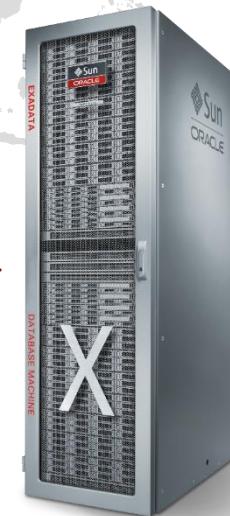


- Oracle started as a relational database but, since 8i, the Oracle database has contained object extensions allow for the creation of types, subtypes, and non-scaler user-defined types such as varrays, nested tables, many of which are defined in the SYSTEM tablespace
- If the data migration technique chose does not migrate the DDL for the types and subtypes they must be manually created prior to starting the migration
- The following SQL will produce a substantial list that should be reviewed



```
SELECT UNIQUE type_owner, type_name
FROM dba_varrays
ORDER BY 1,2;
```

```
SELECT owner, type_name
FROM dba_types
WHERE owner IS NOT NULL
ORDER BY 1,2;
```



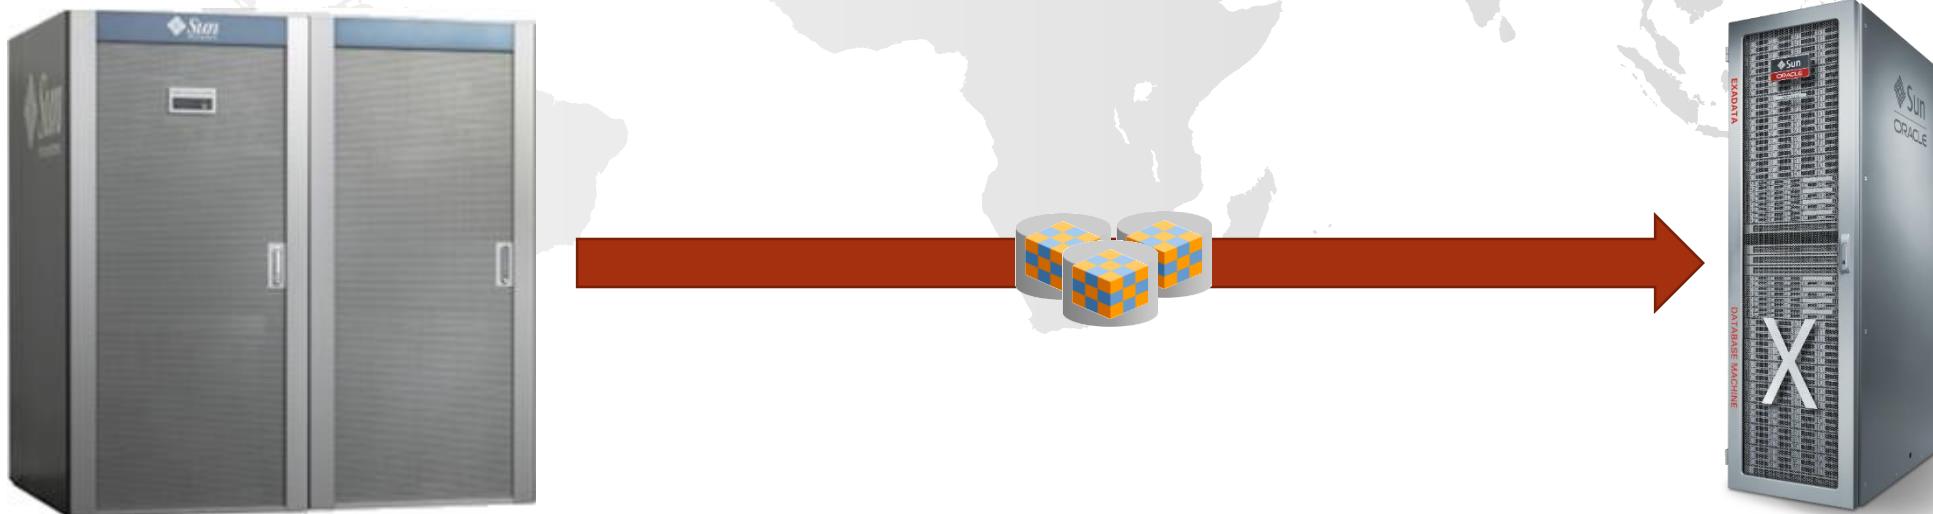
- Determine whether the source database has Database Vault enabled



```
SELECT name, state  
FROM dvsys.dv$policy  
ORDER BY 1;
```



- Oracle Engineered Systems are the optimum target but due to their nature require special care with provisioning and migration methodology
- Your Sparc server did not offer any of the following features
 - Hybrid Columnar Compression
 - Query Offload
 - Storage Indexing
- Be careful you do not overwrite Exa/ODA specific data in the system tablespace



- If migrating from EE to SE
 - Make a list of all enabled EE features not supported in SE such as partitioning
 - Run DBMS_FEATURE_USAGE_REPORT prior to migration



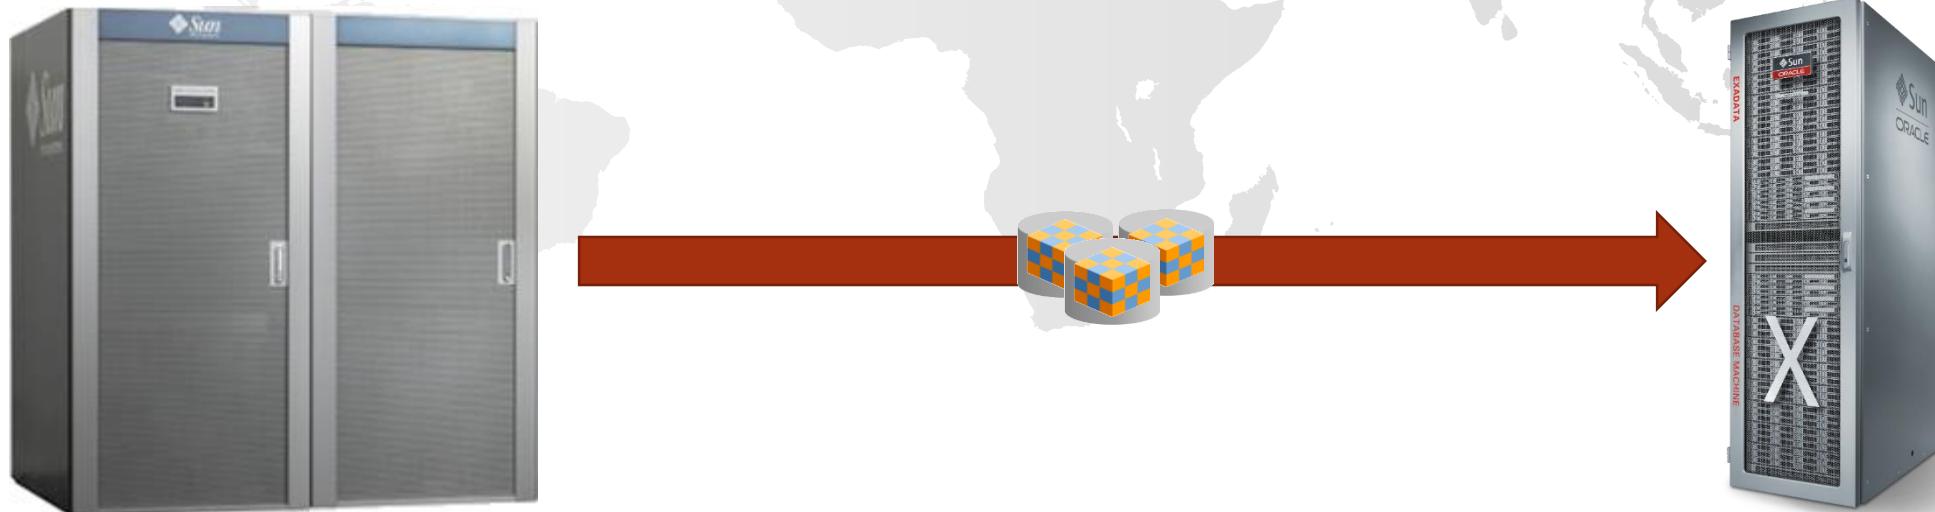
```
spool /home/oracle/usage_report.html

SELECT output
FROM TABLE(dbms_feature_usage_report.display_html);

spool off
```



- Consider the possibility that your Source data files have at-rest encryption implemented by the storage vendor
- If the database is shut down and the files copied ... what is going to decrypt them in the new Target environment?
- Be sure that you perform test migrations that guarantee you ability to access post-migration data
- If encryption at rest is provided by Oracle TDE: Migrate the wallet



- Migrating from a Solaris to Linux some migrations require endian conversion other so do
- Conversion Required
 - RMAN
 - Transportable Tablespaces
- Conversion Not Required
 - Database Link
 - Data Pump (export/import)
 - DBMS_FILE_TRANSFER
 - GoldenGate



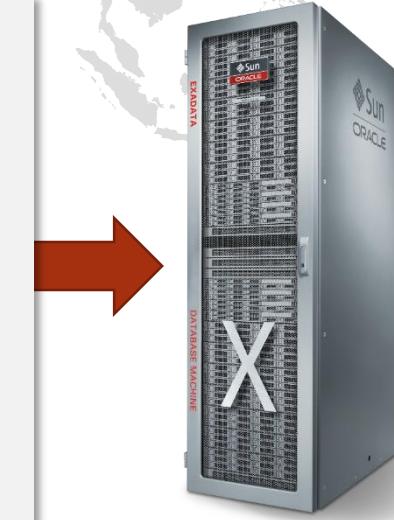
```
SELECT *  
FROM v$db_transportable_platform;  
  
SELECT *  
FROM v$transportable_platform;
```



- External tables are physical files stored on persistent storage that are presented to a database user as read-only tables that can be queried
- The Data Dictionary shows contains their DDL and metadata
- The physical files cannot be migrated by any Oracle tool: They must be migrated by individually copying them between the Source to the Target environments
- If source and target mapping differ the DDL must be rewritten to point to their new location



SQL> desc dba_external_tables		
Name	Null?	Type
OWNER	NOT NULL	VARCHAR2 (128)
TABLE_NAME	NOT NULL	VARCHAR2 (128)
TYPE_OWNER		CHAR (3)
TYPE_NAME		VARCHAR2 (128)
DEFAULT_DIRECTORY_OWNER		CHAR (3)
DEFAULT_DIRECTORY_NAME		VARCHAR2 (128)
REJECT_LIMIT		VARCHAR2 (40)
ACCESS_TYPE		VARCHAR2 (7)
ACCESS_PARAMETERS		CLOB
PROPERTY		VARCHAR2 (10)
INMEMORY		VARCHAR2 (8)
INMEMORY_COMPRESSION		VARCHAR2 (17)



- Some migration technologies such as Transportable Tablespaces do not change index Clustering Factor which is the determinant as to whether the optimizer will use an index
- Some migration technologies such as Data Pump (Export/Import), Database Links and GoldenGate will change the Clustering Factor and likely result in a change in optimizer plans



```
SELECT owner, index_name, clustering_factor
FROM dba_indexes
ORDER BY 1,2;
```



- Pre-migration capture the spfile for all source databases so that they can be referenced if configuration issues arise on a target



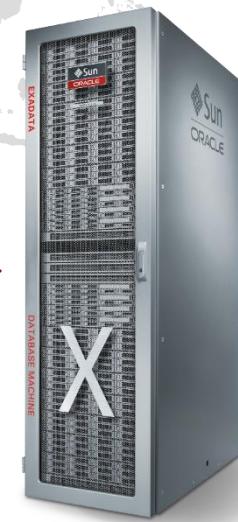
```
CREATE PFILE='/home/oracle/initorabase.ora'  
FROM MEMORY;
```



- Recommended practice is for all databases to be in ARCHIVELOG mode at all times but this may not be the status for pre-prod databases such as DEV, QA, TEST and UAT,
- To improve migration performance it is common practice to disable ARCHIVELOG mode on the target; then to re-enable it post-migration followed by taking an immediate RMAN backup



SELECT log_mode
FROM v\$database



- It is not common to find FLASHBACK DATABASE enabled but when it is that fact is important to properly configuring the target
- It is common practice to disable FLASHBACK DATABASE on the target at the beginning of a migration and re-enable it post-migration to improve migration performance



```
SELECT flashback_on
FROM v$database
```



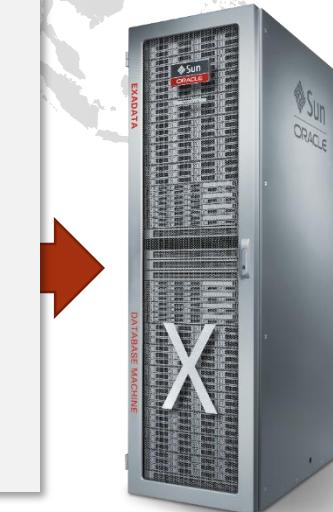
- Oracle In-Memory Database is a two-step process
 1. Allocate part of memory to cache table data
 2. Create or alter a table to place it into the cache



```
SQL> ALTER SESSION SET CONTAINER = pdbdev;  
SQL> ALTER PLUGGABLE DATABASE pdbdev OPEN;  
SQL> ALTER SYSTEM SET inmemory_size=400M;
```

```
SQL> col value format 9999999999  
SQL> SELECT * FROM v$sga;
```

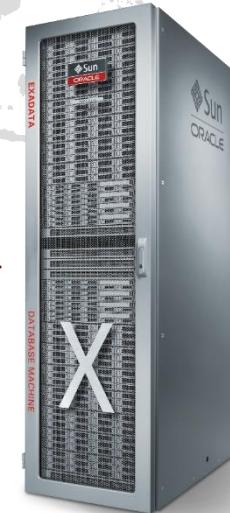
NAME	VALUE	CON_ID
Fixed Size	3047568	0
Variable Size	1086328688	0
Database Buffers	469762048	0
Redo Buffers	13725696	0
In-Memory Area	536870912	0



- Be sure that you obtain an accurate accounting of invalid objects in the source database
- Invalid objects do not belong in any database and prior to starting a migration should either be recompiled successfully or dropped
- To recompile use one or more of the following
 - DBMS.Utility.Compile_Schema
 - UTL_Recomp
 - The UTLRDT.SQL or UTLRP.SQL scripts



```
SELECT owner, object_name, object_type
FROM dba_objects
WHERE status = 'INVALID'
ORDER BY 1,2,3
```



- Most migration methods will not copy jobs deployed using DBMS_JOBS and DBMS_SCHEDULER
- DBMS_JOBS has been deprecated so pre-migration move the jobs to the fully supported DBMS_SCHEDULER package
- Post-migration be sure that your validation process tests the jobs to verify that they will run as designed
- One of the most common migration failures is jobs that fail post-migration



```
SELECT job, next_date, broken, failures
FROM dba_jobs
ORDER BY 2,1;
```

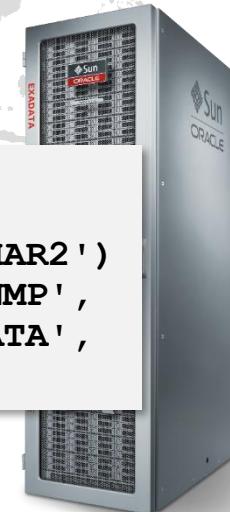
```
SELECT owner, job_name, next_date, enabled, failure_count
FROM dba_scheduler_jobs
ORDER BY 2,1;
```



- As of Oracle 12.1 SecureFile storage is the default for large objects such as BLOBs and CLOBs if migration is performed using DataPump
- Be sure, if your target has SecureFile storage, post-migration you optimize **CHUNK_SIZE**, **DEDUPLICATION** and **RETENTION**
- Configure **COMPRESSION** and **ENCRYPTION** if you have the required licenses



```
SELECT UNIQUE owner, data_type, data_type_owner
FROM dba_tab_cols
WHERE data_type NOT IN ('CHAR', 'DATE', 'NUMBER', 'VARCHAR2')
AND owner NOT IN ('APPQOSSYS', 'AUDSYS', 'CTXSYS', 'DBSNMP',
'DVSYS', 'GSMADMIN_INTERNAL', 'MDSYS', 'OJVMSYS', 'ORDDATA',
'ORDSYS', 'OUTLN', 'SYS', 'SYSTEM', 'WMSYS', 'XDB');
```



- All Solaris environments, by default are NUMA aware
- All Oracle Linux databases, installed by OUI/DBCA have NUMA disabled

```
-- Locality Groups
lgrpinfo [-aceGhIlLmrt] [-u unit] [-C|-P] lgrp ...
lgrpinfo -T [-aceG1Lmr] [-u unit]

-- Kernel Statistics
kstat -m lgrp
```

```
[dmorgan@lxorap1n5 ~]$ numactl --hardware
available: 2 nodes (0-1)
node 0 size: 48457 MB
node 0 free: 269 MB
node 1 size: 48480 MB
node 1 free: 47 MB
node distances:
node 0 1
 0: 10 20
 1: 20 10
```



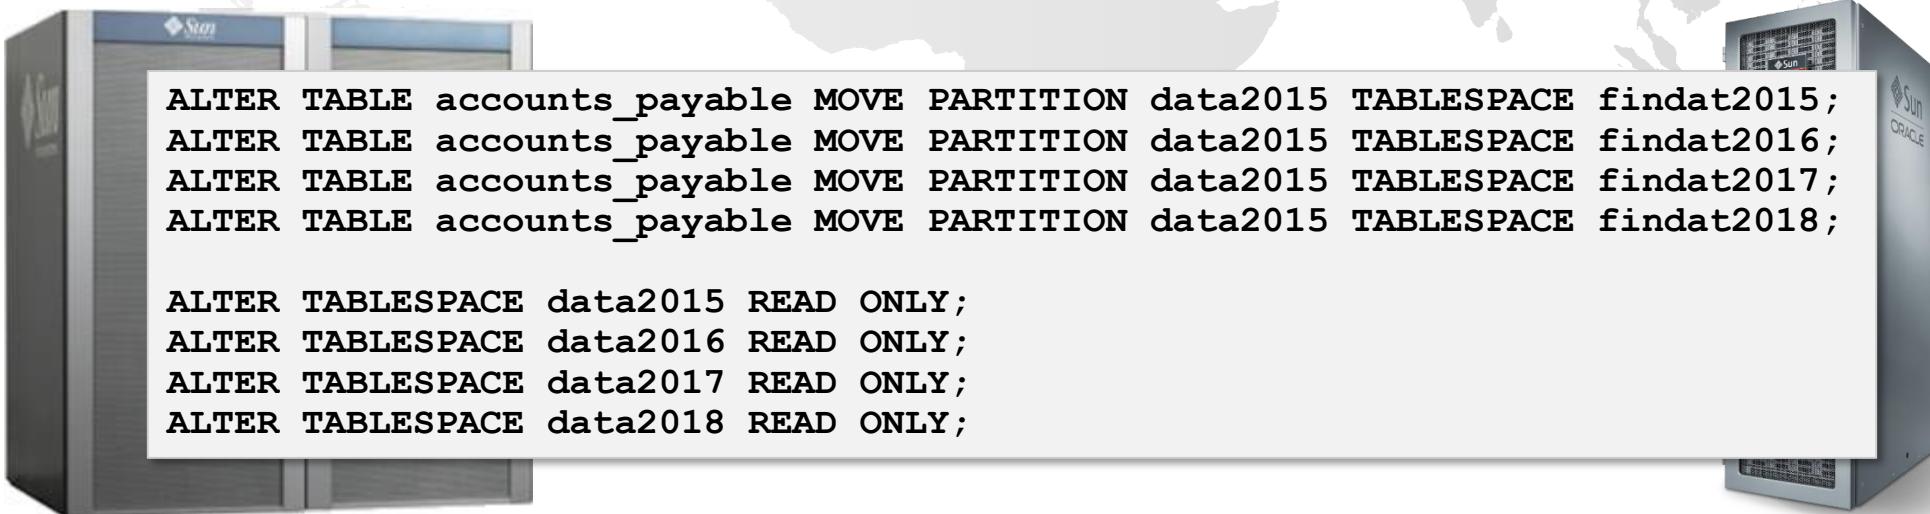
- Oracle 10g and 12c+ use ASMM
- Oracle 11g AMM has numerous issues for the most part centered around automatic resize operations that locked memory and serialized operations resulting in substantial performance issues
- Recommended practice post-migration is to not copy initialization parameters or configurations that enable AMM
- HugePages



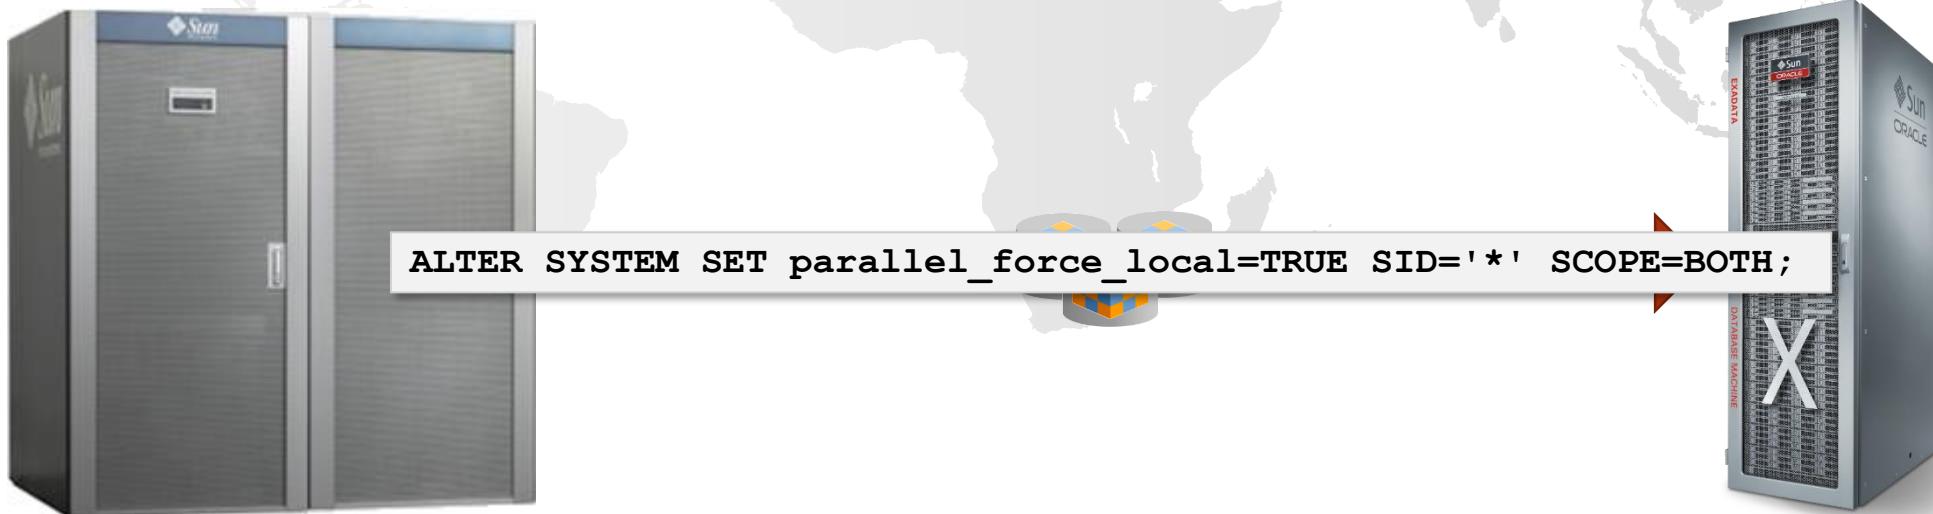
```
ALTER SYSTEM SET memory_max_target=0 SID='*' SCOPE=SPFILE;
ALTER SYSTEM SET memory_target=0 SID='*' SCOPE=SPFILE;
ALTER SYSTEM SET sga_target=<value> SID='*' SCOPE=SPFILE;
ALTER SYSTEM SET pga_aggregate_target=<value> SID='*' SCOPE=SPFILE;
```



- HASH and LIST partitioning will likely not help or hurt a migration
- But RANGE partitioning can be of tremendous value especially if partitions have been built, or can be moved, into their own tablespace
- For very large databases run a POC
 - Convert tablespaces with older partitions to READ ONLY
 - Migrate them prior the outage window



- On a RAC cluster forces parallel query slaves to stay on the local instance and not parallelize across nodes which increases interconnect traffic
- Range of values: {**FALSE** | **TRUE**}
- Migrating from a stand-alone instance significant changes in physical architecture
 - Storage must be UNLOCKED, UNLOCKABLE, SHARED EVERYTHING
 - A separate, private, non-routable, jumbo frames or infiniBand interconnect is mandatory



- The recyclebin is a logical construct that contains objects that have been "dropped" but not "purged"
- Recyclebin objects are limited to
 - Tables
 - Indexes
 - Constraints
 - Triggers



```
SELECT owner, type, COUNT(*)  
FROM dba_recyclebin  
GROUP BY owner, type  
ORDER BY 1,2,3;
```

- Prior to migration check for the following features/options enabled on the source
 - Advanced Security
 - AUDIT_SYS_OPERATIONS
 - Database Vault
 - Label Security
 - Transparent Data Encryption (TDE)



- Metadata you migrate from source to target may lead to poor performance
- Post-migration be sure you collect current
 - Data Dictionary Stats
 - Fixed Object Stats
 - Index Stats
 - System Stats



```
SQL> SELECT pname, pval1
  2  FROM sys.aux_stats$;
```

PNAME	PVAL1
STATUS	
DSTART	
DSTOP	
FLAGS	0
CPUSPEEDNW	1264
IOSEEKTIM	10
IOTFRSPEED	4096
SREADTIM	
MREADTIM	
CPUSPEED	1313
MBRC	
MAXTHR	
SLAVETHR	

Before

```
SQL> SELECT pname, pval1
  2  FROM sys.aux_stats$;
```

PNAME	PVAL1
STATUS	
DSTART	
DSTOP	
FLAGS	0
SBLKRDS	108935
SBLKRDTIM	30962.585
MBLKRDS	13321
MBLKRDTIM	37532.063
CPUCYCLES	11845428
CPUTIM	9216264
JOB	4964
CACHE_JOB	4966
MBRTOTAL	333327
CPUSPEEDNW	1264
IOSEEKTIM	10
IOTFRSPEED	4096
SREADTIM	
MREADTIM	
CPUSPEED	1313
MBRC	
MAXTHR	
SLAVETHR	

After

- On the source database determine whether Smallfile or Bigfile tablespaces are in use: Factor this into your migration planning
- Check PCT_FREE and PCT_USED for tables: Changing PCT_FREE to zero (0) post migration can often lead to substantial improvements in performance and storage usage
- It is very likely you are wasting about 10% of all the storage allocated to your databases



```
SELECT UNIQUE owner, pct_free, pct_used
FROM dba_tables
ORDER BY 1,2,3;
```



- A database can have multiple TEMP tablespaces
- A TEMP tablespace can consist of multiple data files
- Most database migration methods do not create or migrate the TEMP tablespace
- Be sure that you factor into migration planning how to create the TEMP tablespaces you will need
- If moving from Legacy to Container architecture recommended practice is to configure a TEMP tablespace per PDB



```
SELECT tablespace_name, file_name
FROM dba_temp_files
ORDER BY 1,2;
```



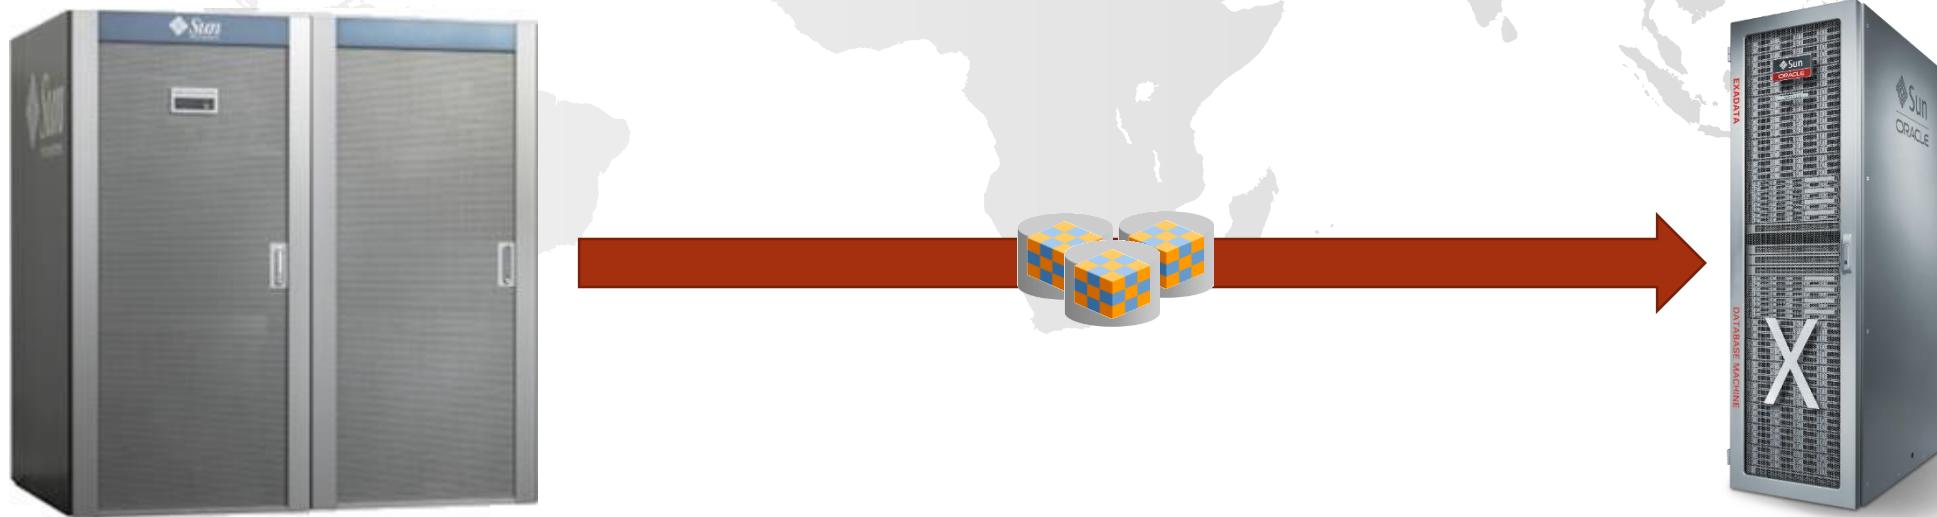
- A Container database can have multiple undo tablespaces
- An UNDO tablespace can consist of multiple data files
- Most database migration methods do not create or migrate the TEMP tablespace
- Be sure that you factor into migration planning how to create the TEMP tablespaces you will need
- The Oracle default retention of 900 seconds is far from optimum: Be sure that you configure proper size and retention on the target



```
SELECT tablespace_name, retention
  FROM dba_tablespaces
 WHERE contents = 'UNDO'
 ORDER BY 1
```



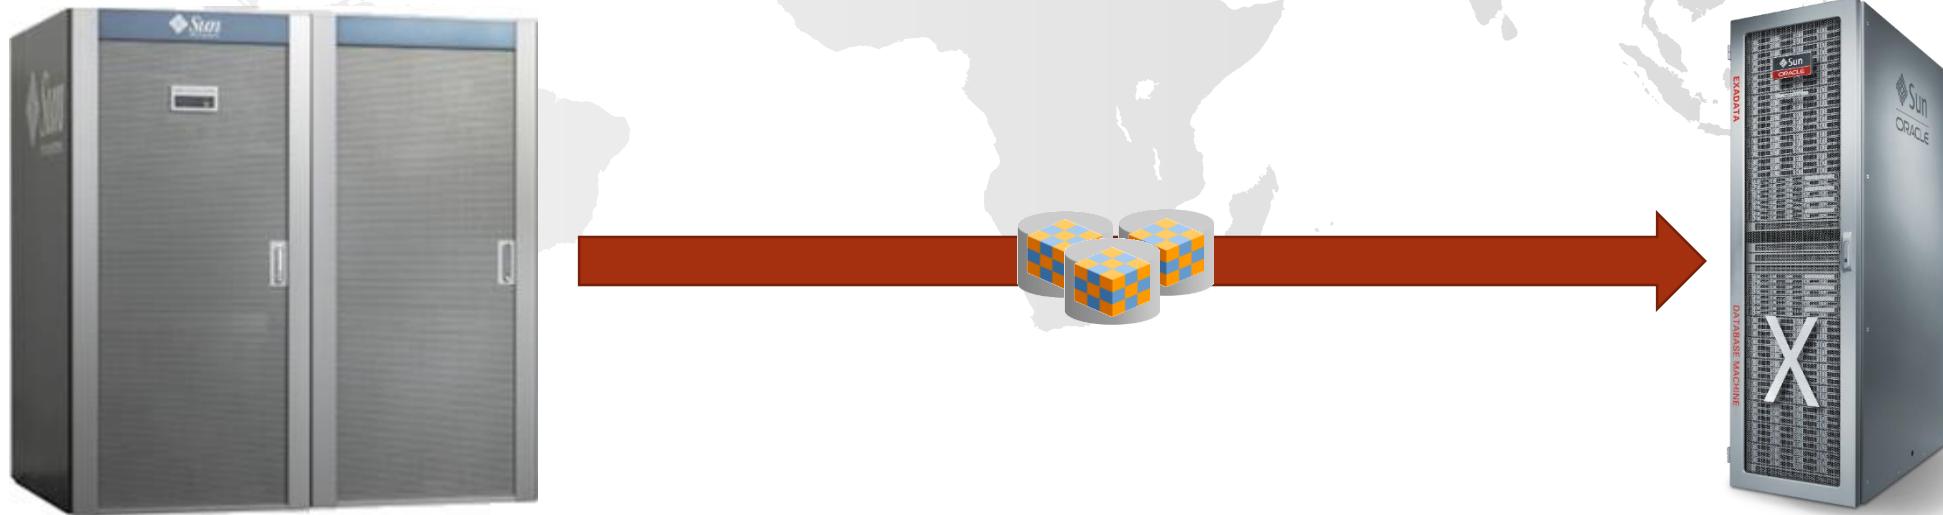
- What is the time zone of the source server and database?
- Are they the same as the time zone in the target environment?
- Does the database contain date-range partitioning?
- Does application code use functions such as LOCALTIMESTAMP?
- If there is a Data Guard Physical Standby are the server and database configured to use the same time zone as the primary production environment?



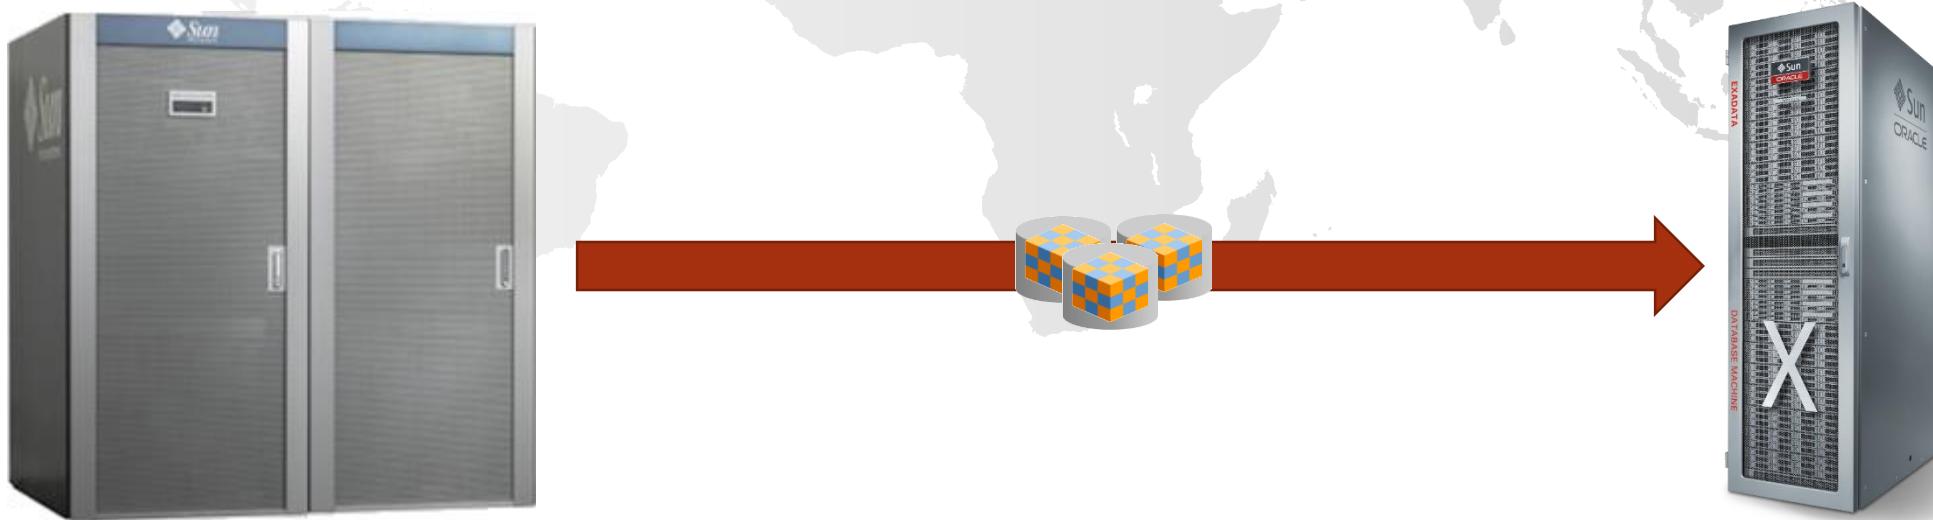
- Examine every line of the following files in \$GRID_HOME and \$ORACLE_HOME as appropriate
 - SQLNET.ORA
 - LISTENER.ORA
 - TNSNAMES.ORA
 - CMAN.ORA
- Be sure that the target is properly configured



- It is best to never copy an init.ora/spfile from source to target
- If recreating source parameters be sure that they are not deprecated
- During pre-migration planning, if changing versions, verify that deployed features such as Advanced Replication and Streams are not in use
- During pre-migration testing (POC) be sure that the migration team has worked with developers and DBAs to identify new features such as Adaptive Query Plans and Evolving Baselines and will not be ambushed by them



- Does the current database use a wallet for Transparent Data or other forms of encryption security?
- If it does there are migration risks if the wallet keys and certificates are not migrated or the data decrypted prior to migration





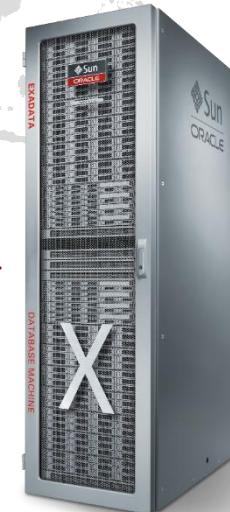
Wrap Up

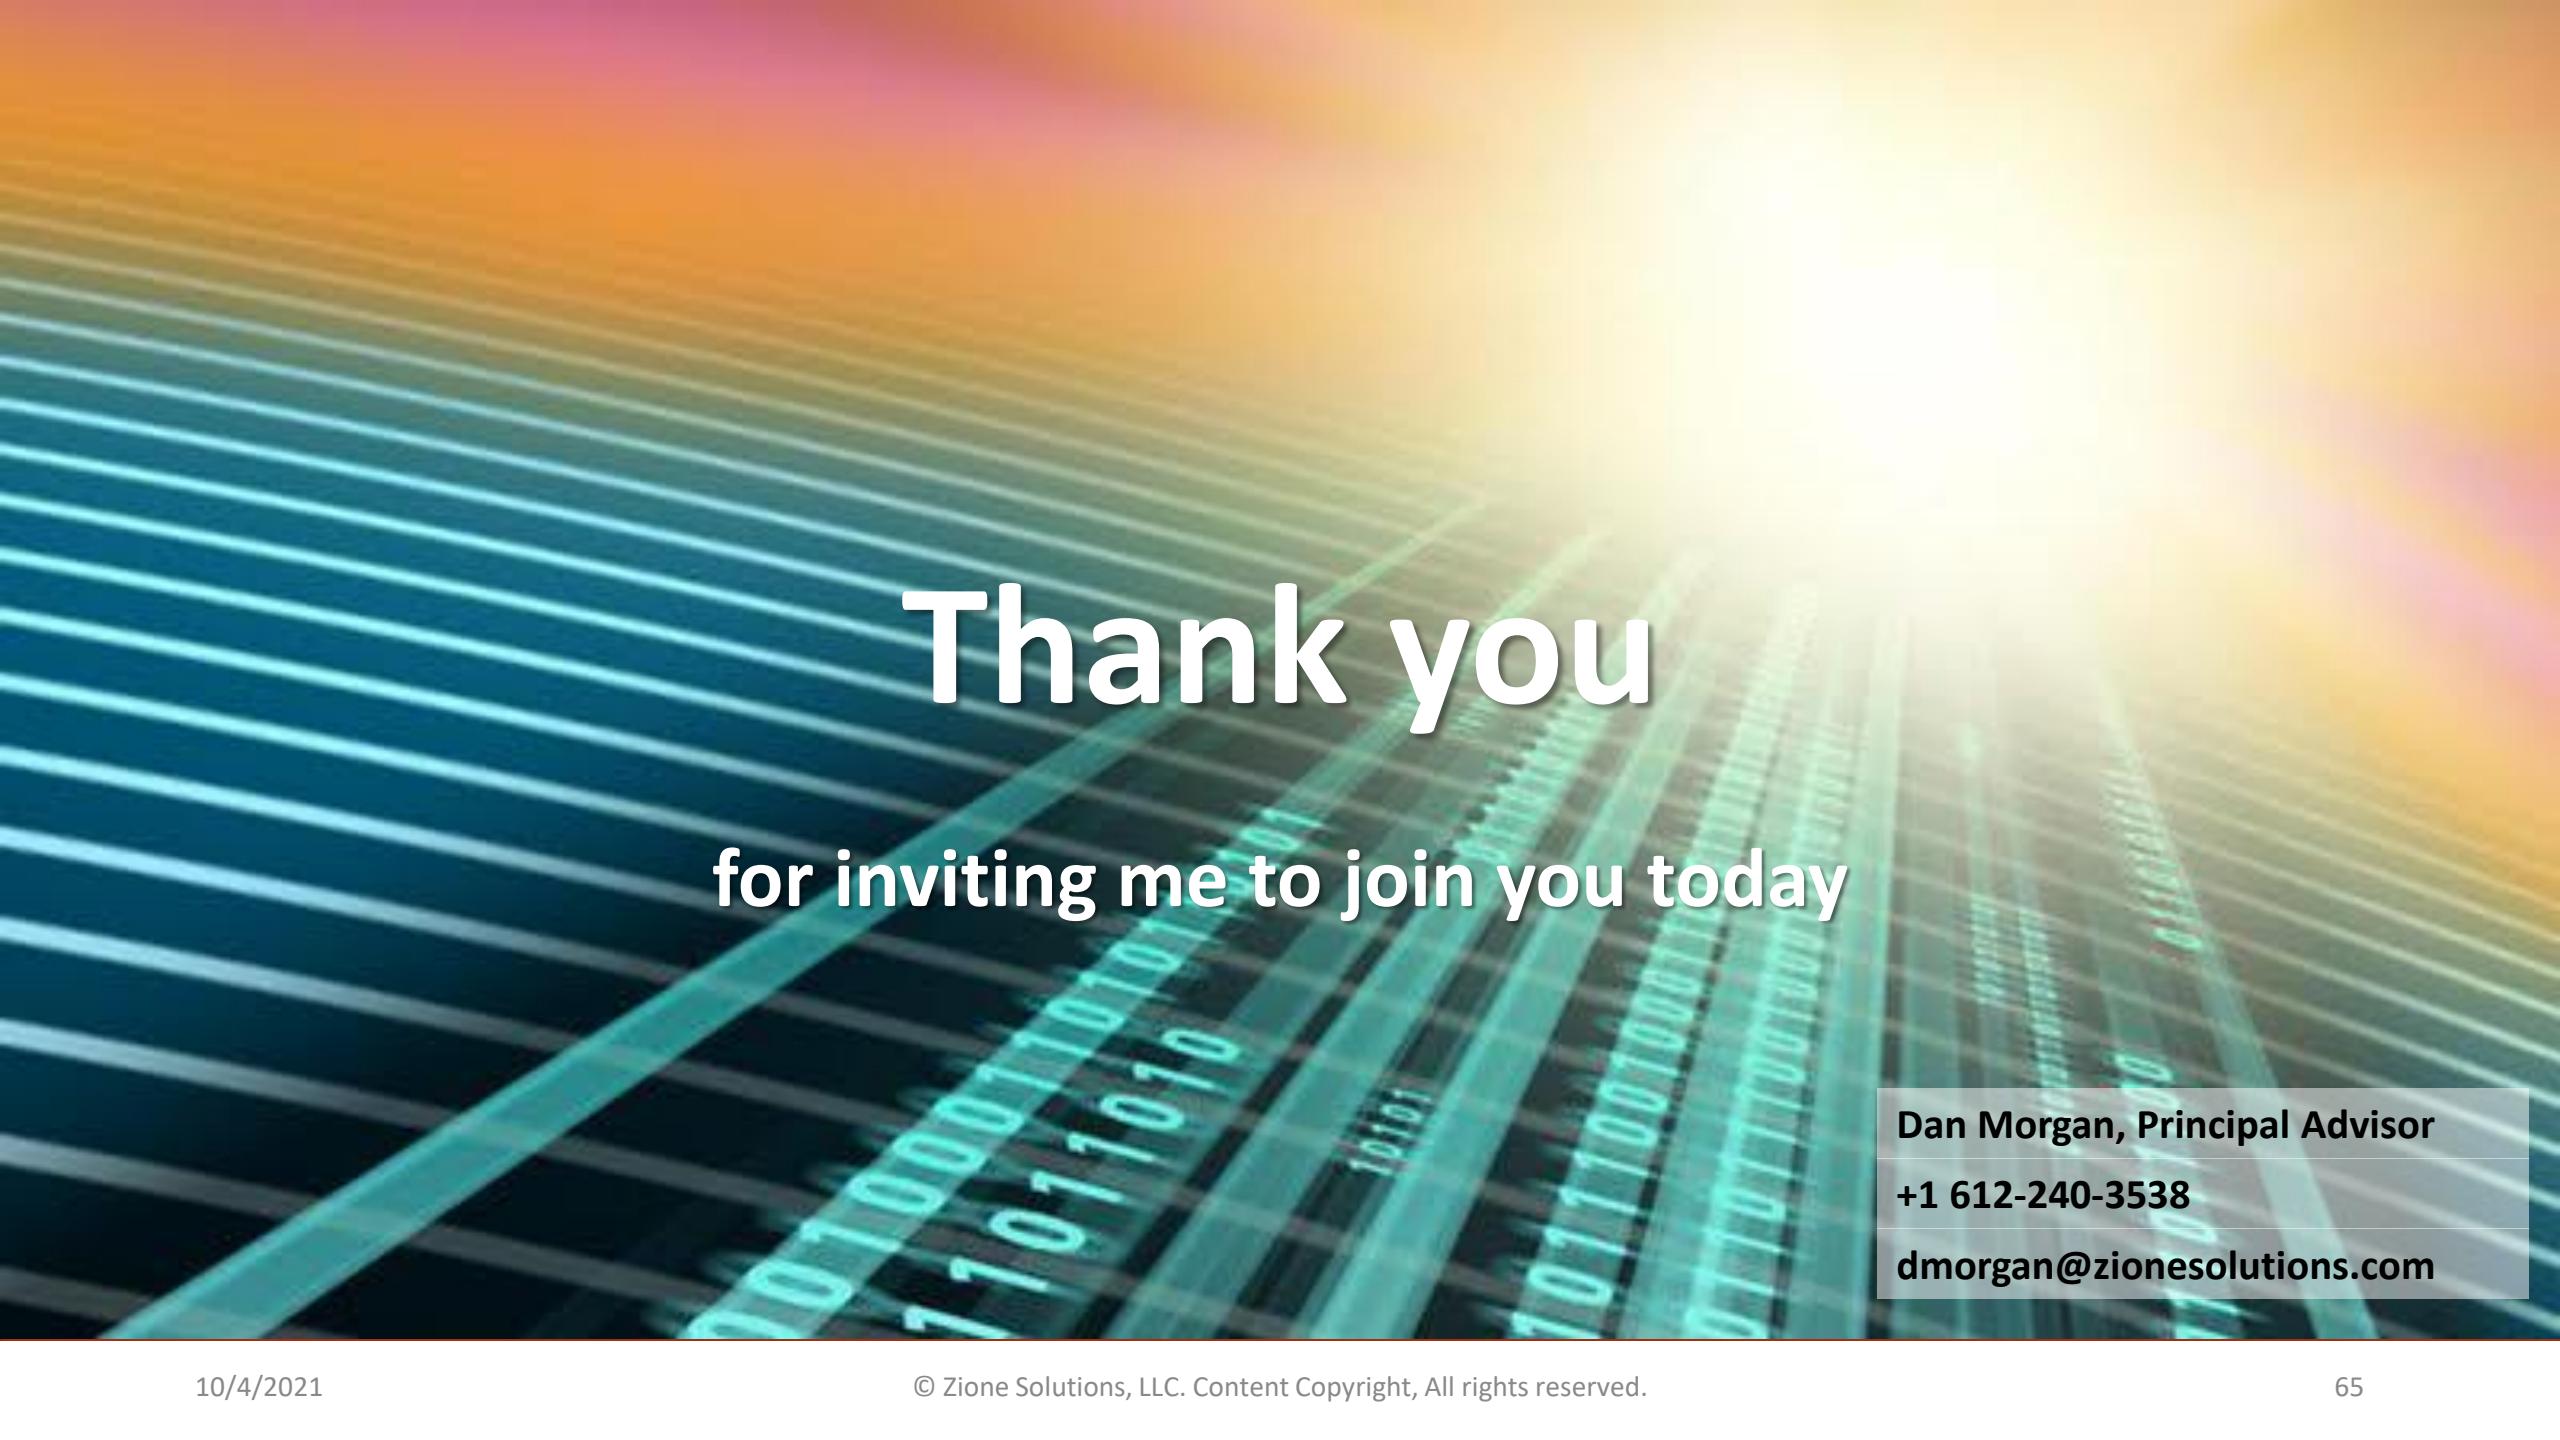
```
SELECT UNIQUE solaris_migration_expertise
FROM zione_solutions
WHERE mig_date > SYSDATE
AND skills_required = 'HIGH'
AND source IN ('Sun', 'Sparc', 'Solaris')
AND target IN ('Data Center', 'Colo', 'Cloud', 'Hybrid')
AND tolerance_for_failure = 0;
```



A copy of this presentation will
be given to your user group for
your website or distribution.

Email any questions to me at:
dmorgan@zionesolutions.com





**Thank you
for inviting me to join you today**

Dan Morgan, Principal Advisor
+1 612-240-3538
dmorgan@zionesolutions.com

```
sqlplus19

SQL*Plus: Release 19.0.0.0.0 - Production on Wed Jul 22 12:39:51 2020
Version 19.3.0.0.0

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

Enter user-name: / as sysdba

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.3.0.0.0

SQL> ERROR at line 1:
SQL> ORA-00028: your session has been killed
```