stonebridge

Near Zero Downtime Database Migration

Caiping Yue

Principal Technical Consultant

DBA

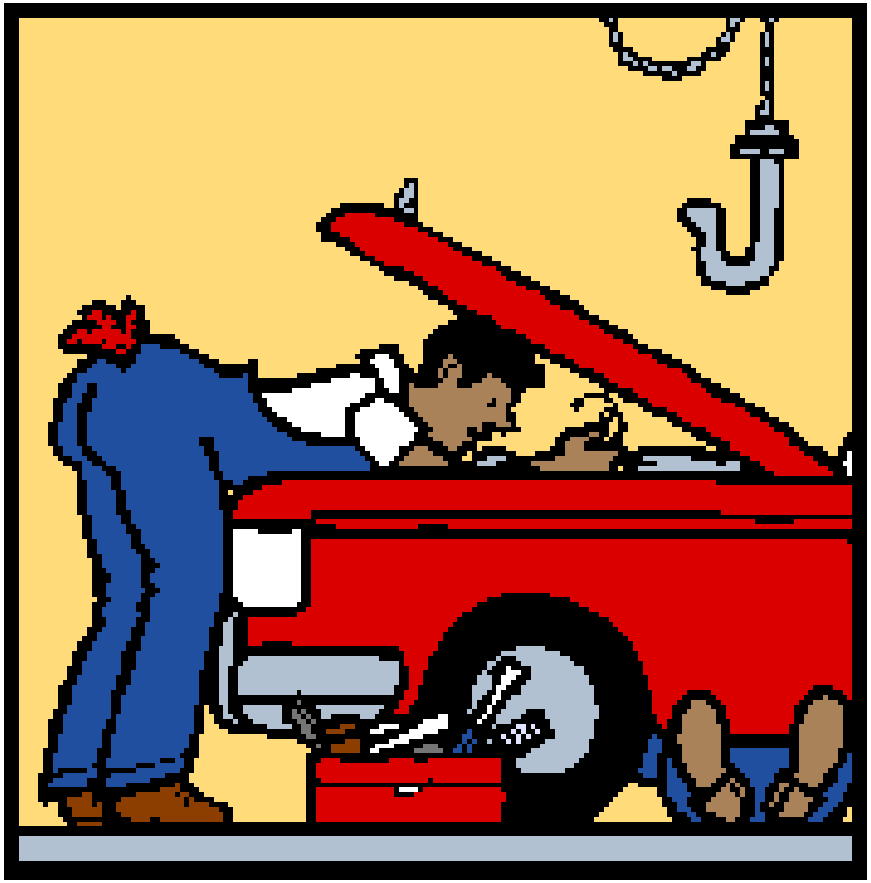

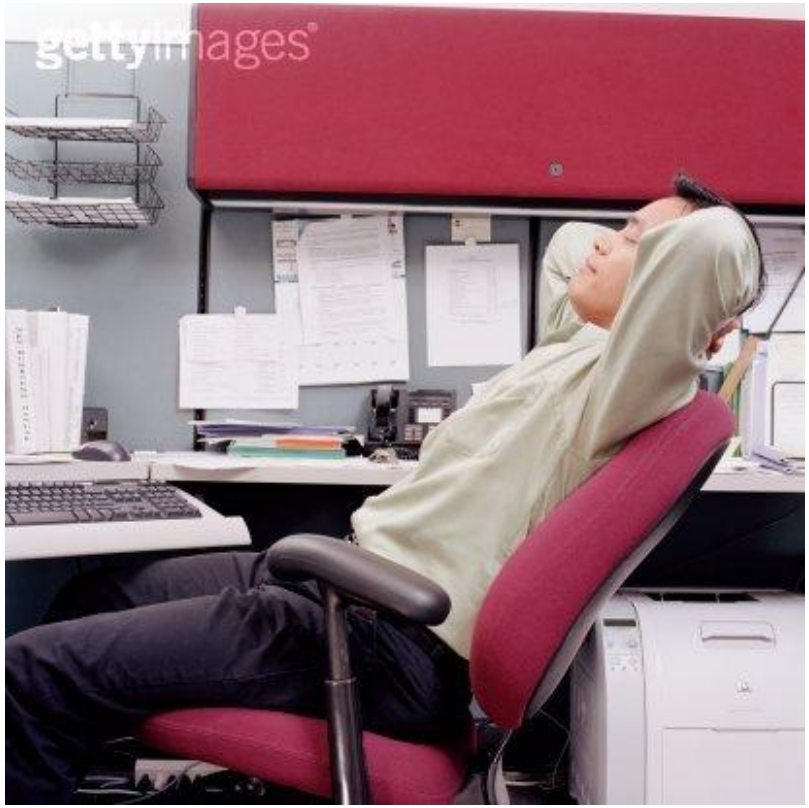

Our Company

Stonebridge is an IT consulting and systems integration services company. We solve business problems caused by data management and data integration issues that negatively affect performance management, corporate governance, workflow effectiveness, and business continuity.

Corporate Background:

- Founded in 1995
- Privately Held / Employee Owned
- Primarily Serving the TOLA Region

Core Competencies:

- Business Intelligence
- Enterprise Content Management
- Custom Development
- IT / Database Services

ORACLE Silver Partner

Microsoft GOLD CERTIFIED Partner

stonebridge
outperform

Database Migration Defined

- Physically move Database from one location to another location
- Location can be different datacenter, or server (rehost)
- OS can be different (re-platform)
- Does not necessarily involve Oracle version change (Upgrade)
- Does not necessarily involve change of database software (i.e., from MS SQL to Oracle)

Driving Forces of Database Migration

- Replacement of old or out of lease equipment
- Databases continue to expand exponentially
- Companies switching x64 architecture
- Initiatives that seek to reduce complexity and costs by moving to commoditized platform
- …

SOURCE: The hidden cost of data migration (http://www.computingsa.co.za/PDFs/AREX3.pdf)

The Client: Calpine

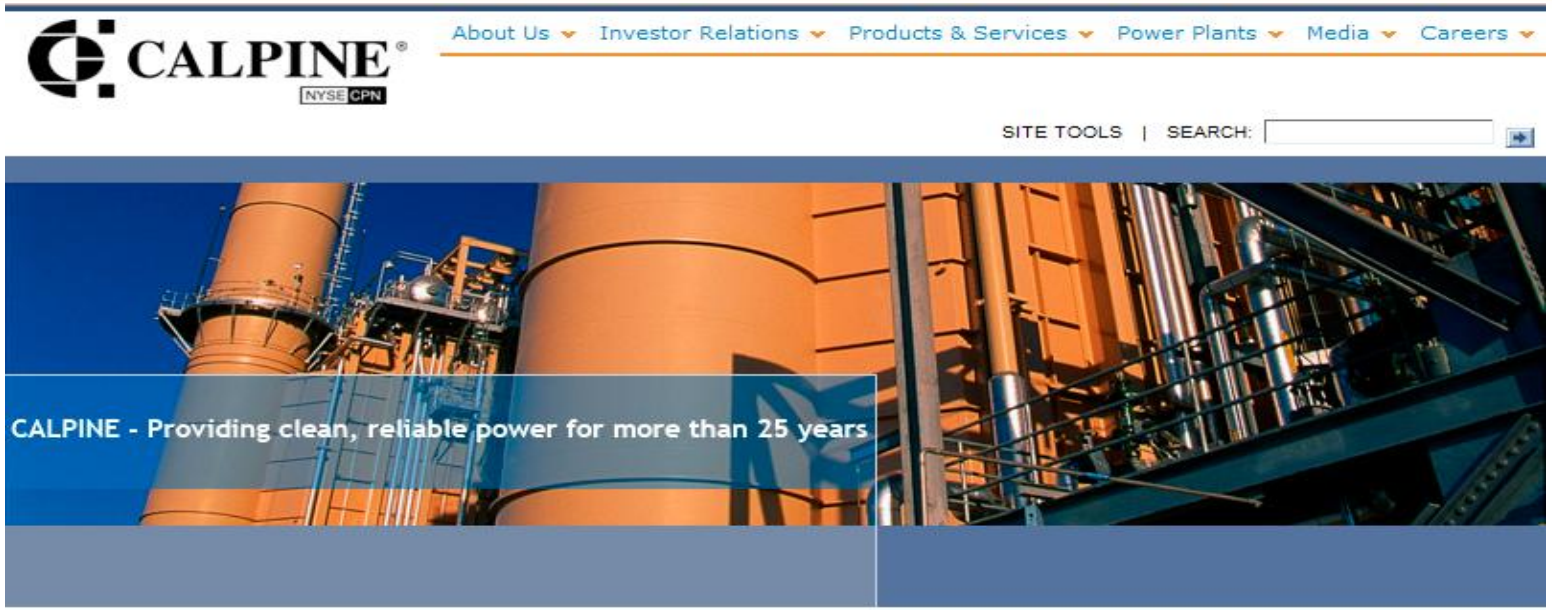

stonebridge
outperform

Calpine's Challenge

- **Technology Refresh** - replace old Solaris hardware with IBM p570 hardware running AIX OS's
- **Datacenter Consolidation** - outsource its datacenter operations to external provider
- **Server Consolidation** - reduce the number of database servers

The Upshot: All Oracle databases must be converted from previous configuration to the target configuration

stonebridge
outperform

Configuration Changes

	Previous Config	Target Config
OS	Solaris 9	AIX 6.1
hardware	Sun Sparc	IBM p570 (p6)
Storage	EMC Clariion	EMC DMX4 Symmetrix
Database	Oracle 10.2.0.3 EE	Oracle 10.2.0.3 EE
HA	Two node RAC or standalone	standalone
DR	Data Guard (physical)	Data Guard (physical)
Datacenter (PR/DR)	Houston/Cyrus One	Omaha/Tempe

stonebridge
outperform

Data Center Moves

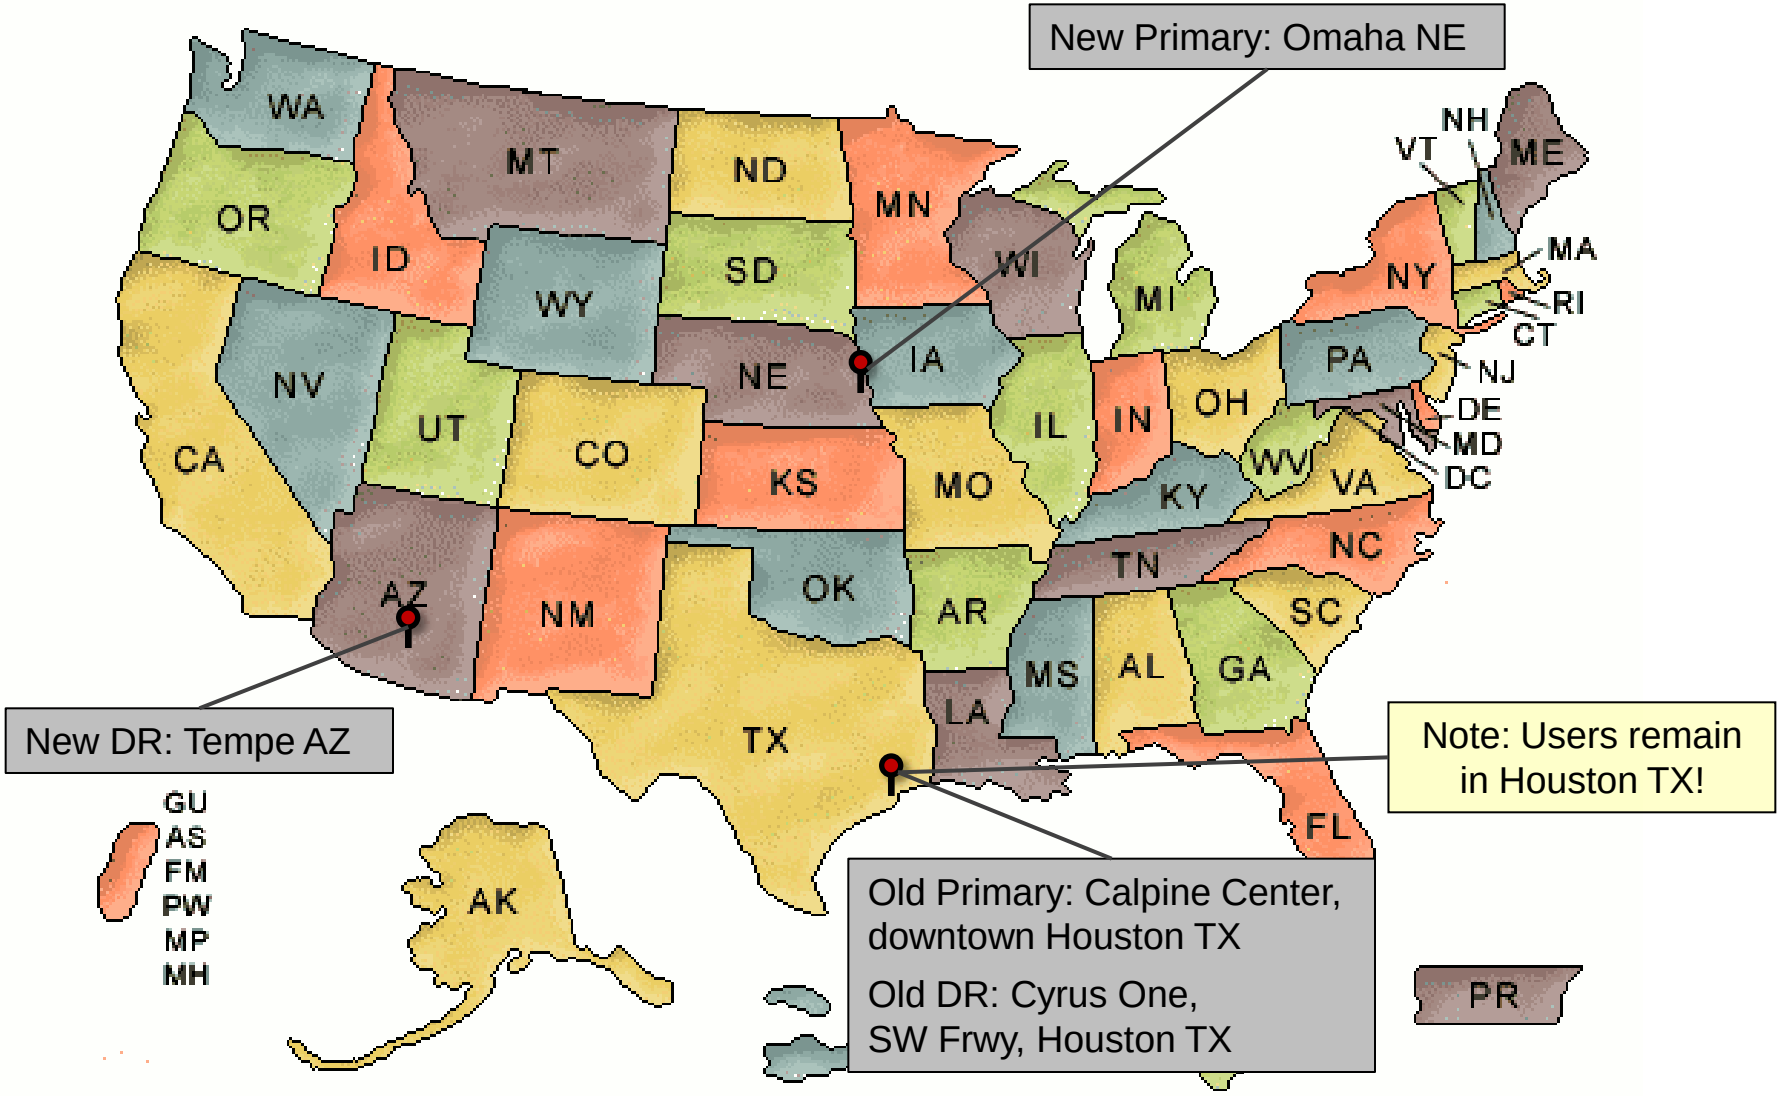

"In-Scope" Mission Critical Databases

Oracle DB	Application Group	Size (GB)	Desired Downtime
DB1	Real Time Trading	23	< 15 minutes
DB2	Real Time Trading	80	< 15 minutes
DB3	Real Time Trading	129	< 15 minutes
DB4	Real Time Trading	157	< 15 minutes
DB5	Real Time Trading	185	< 15 minutes
DB6	Real Time Trading	425	< 15 minutes
DB7	Energy Trading Risk Management	8	< 15 minutes
DB8	Energy Trading Risk Management	27	< 15 minutes
DB9	Energy Trading Risk Management	28	< 15 minutes
DB10	Energy Trading Risk Management	207	< 15 minutes
DB11	Energy Trading Risk Management	313	< 15 minutes
DB12	Energy Trading Risk Management	381	< 15 minutes

NOTE: Per Calpine, DBs in yellow must be moved as a group at the same time

Recap: Calpine's Challenges

- **Change of OS platform**
- **Move ~ 2 TB of production data across WAN (~ 900 Miles between datacenters)**
  - Note: development and stage databases are not included in this value!
- **WAN latency:**
  - Bandwidth between datacenters ~ 155 mbits/second
  - Single stream ftp transfer rate < 5 GB/hour
- **Performance:**
  - User perception
  - Application architecture
  - Datacenter infrastructure

stonebridge
outperform

The Client: Calpine

- Calpine did $8.3 Billion of business in 2008
- ~ $0.94 Million per hour

Calpine Business Requirements

- **Minimal downtime (less than 15 minutes)**
- **Full DR (Data Guard) protection outside of the downtime window**
- **Ability to move a group of 5 databases within the same downtime window**
- **Extensive testing of all mission critical databases at different level (unit, system, integration)**
- **Extensive testing of datacenter infrastructure, Oracle database LPAR performance, Data Guard/DR operation before “Go Live”**

stonebridge
outperform

Reality Check

Can it be done?

How would you do it?

stonebridge
outperform

Common Options

- Export/Import with Data Pump
- Cross-Platform Transportable Tablespace (XTTS)
- Cross-Platform Transportable Database (TDB)
- Application specific methodology to capture changes and replay

Option #1: Export/Import

- Fairly standard process
- Downtime could be days
- Impossible to move several databases at the same time
- Repetitive setup in order to accommodate application testing

Option #2: XTTS

- Cross platform transportable tablespaces (XTTS): tablespaces can be transferred and converted to target platform if they are self contained (10g R1 new feature)

- XTTS uses RMAN Convert functionality to perform conversion

Option #2: XTTS Case Study by Oracle

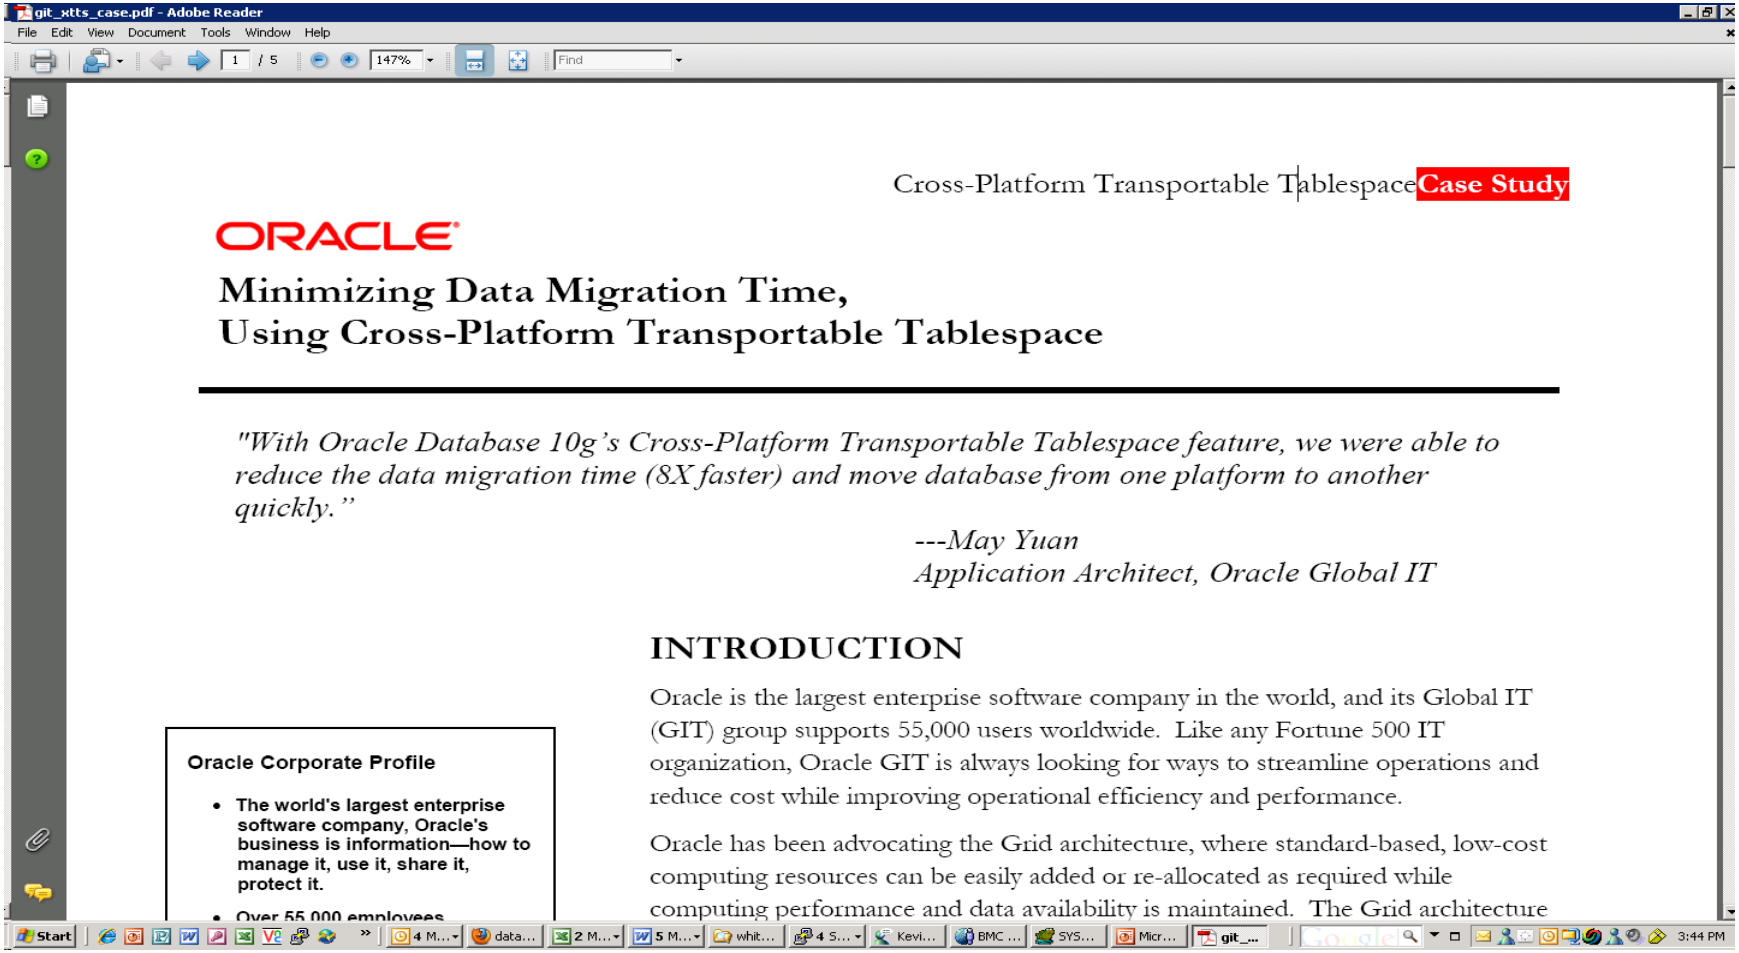

Option #2: XTTS Case Study by Oracle

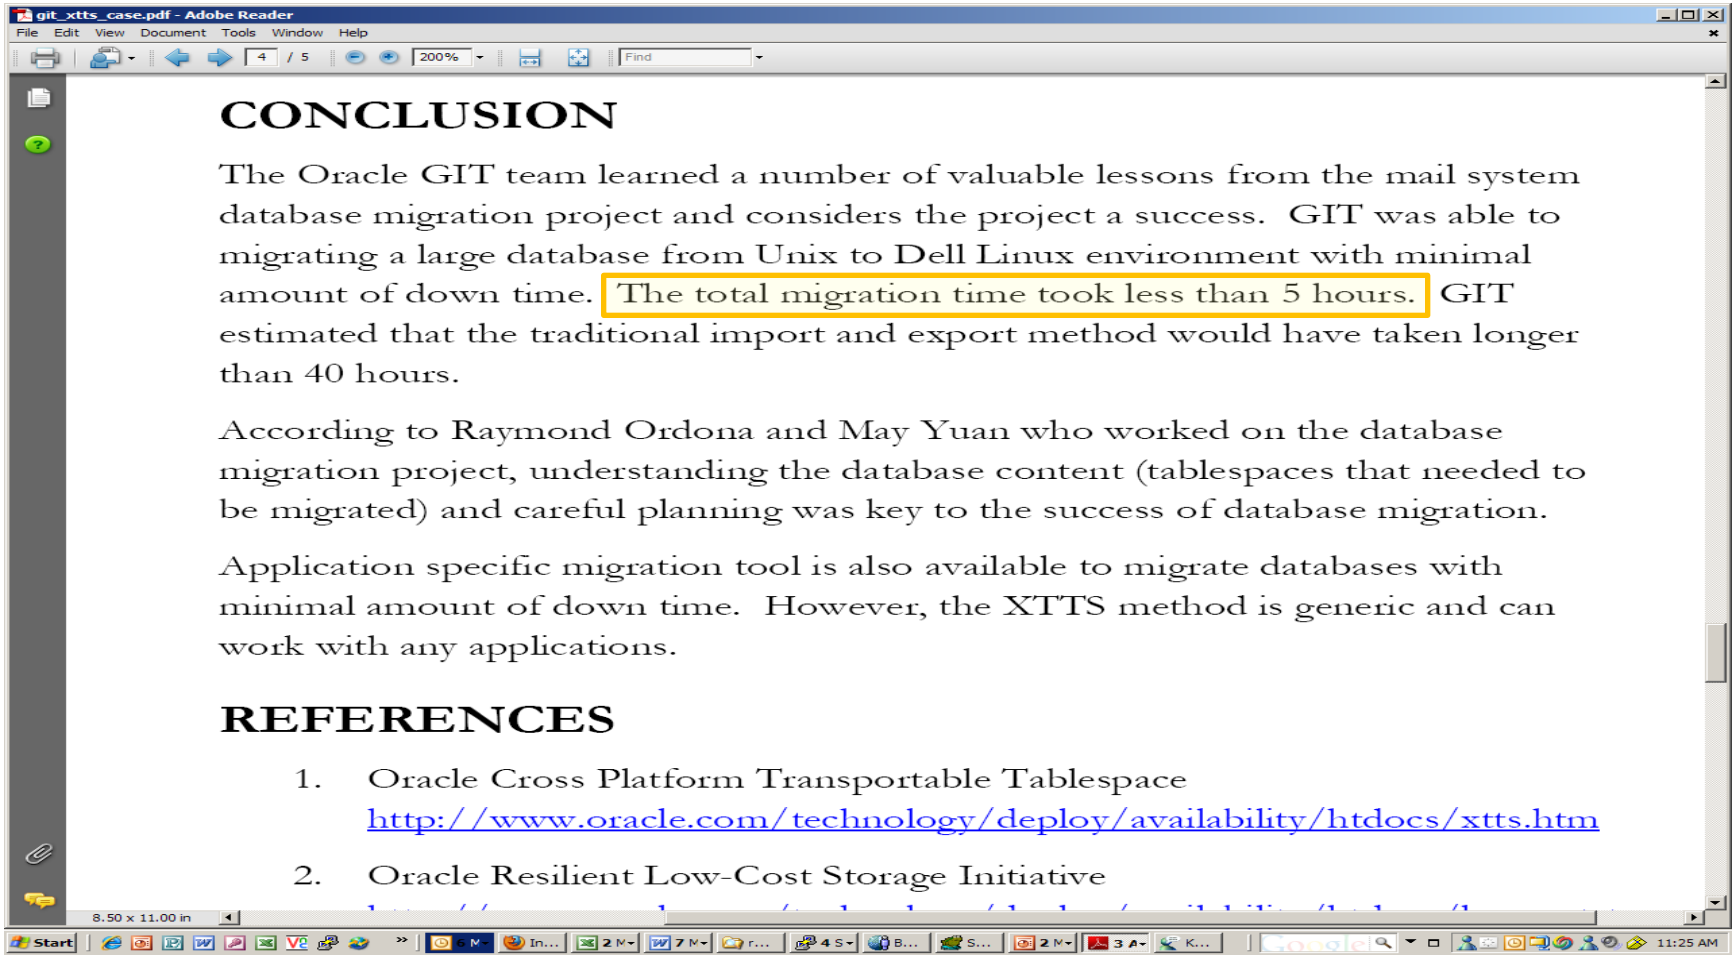

git_xtts_case.pdf - Adobe Reader

CONCLUSION

The Oracle GIT team learned a number of valuable lessons from the mail system database migration project and considers the project a success. GIT was able to migrating a large database from Unix to Dell Linux environment with minimal amount of down time. The total migration time took less than 5 hours. GIT estimated that the traditional import and export method would have taken longer than 40 hours.

According to Raymond Ordona and May Yuan who worked on the database migration project, understanding the database content (tablespaces that needed to be migrated) and careful planning was key to the success of database migration.

Application specific migration tool is also available to migrate databases with minimal amount of down time. However, the XTTS method is generic and can work with any applications.

REFERENCES

1. Oracle Cross Platform Transportable Tablespace
   http://www.oracle.com/technology/deploy/availability/htdocs/xtts.htm
2. Oracle Resilient Low-Cost Storage Initiative

Option #3: TDB

- Transportable database (TDB): whole database can be transferred and converted to target platform if target platform has the same Endian format as the source database (10g R2 new feature)

- Similarly TDB also uses RMAN Convert functionality to perform conversion

Option #3: TDB Case Study in Calpine

Major Steps Contributing to Downtime	Rate	Unit
FTP from Solaris to AIX in same data center (FTP TIME)	120	GB/hr
Converting datafiles (CONVERT TIME)	140	GB/hr
Recompiling (RECOMPILE TIME)	1	hr/db
All other steps (MISC TIME)	2	hr/db
Downtime Estimates Formula	FTP TIME + CONVERT TIME + RECOMPILE TIME + MISC TIME	
Estimated Down Time	4 – 8 hr	
Actual Downtime for 2 Databases (5 and 50 GB)	4 – 6 hr	

Option #4: Apps Specific Capture and Replay

- Certain application may have the ability to capture database changes and replay in the destination database
- Application specific
- Great thinking but very rarely application offers this ability

Reality Check

Calpine's business requirements

CANNOT

be met with typical Options

stonebridge
outperform

A Novel Approach

- Make “Go Live” only a turn of switch
- Convert destination database in advance
- Keep this destination database in sync with source database
- Pre-build DR for this destination
- Allow extensive testing

RMAN Convert

Syntax

convert::=

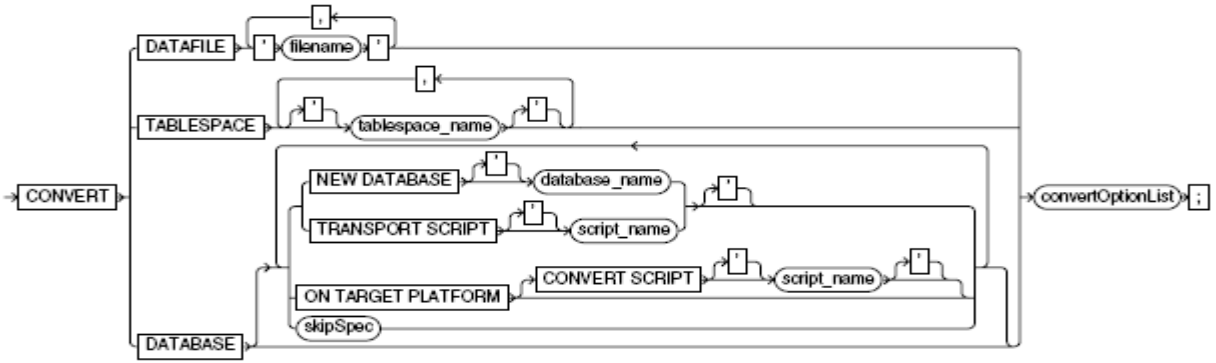

convertOptionList::=

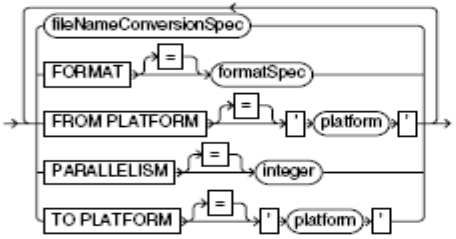

Purpose

To convert a datafile, tablespace or database to the format of a destination platform, in preparation for transport across different platforms.

- `CONVERT TABLESPACE` is used at the source database to produce datafiles for the specified tablespaces in the format of a different destination platform. The converted files can then be transported to the destination platform.
- `CONVERT DATAFILE` is used on the destination database to convert datafiles that are in the format of a different source platform. Once all of the datafiles required for a tablespace have been converted, the datafiles can be transported into the destination database.
- `CONVERT DATABASE` is used to transport an entire database from a source platform to a destination platform, converting the datafiles to the format of the destination platform and ensuring the creation of other required database files. Depending upon the requirements of your situation, `CONVERT DATABASE` on either the source or destination platform.

stonebridge
outperform

Oracle Streams

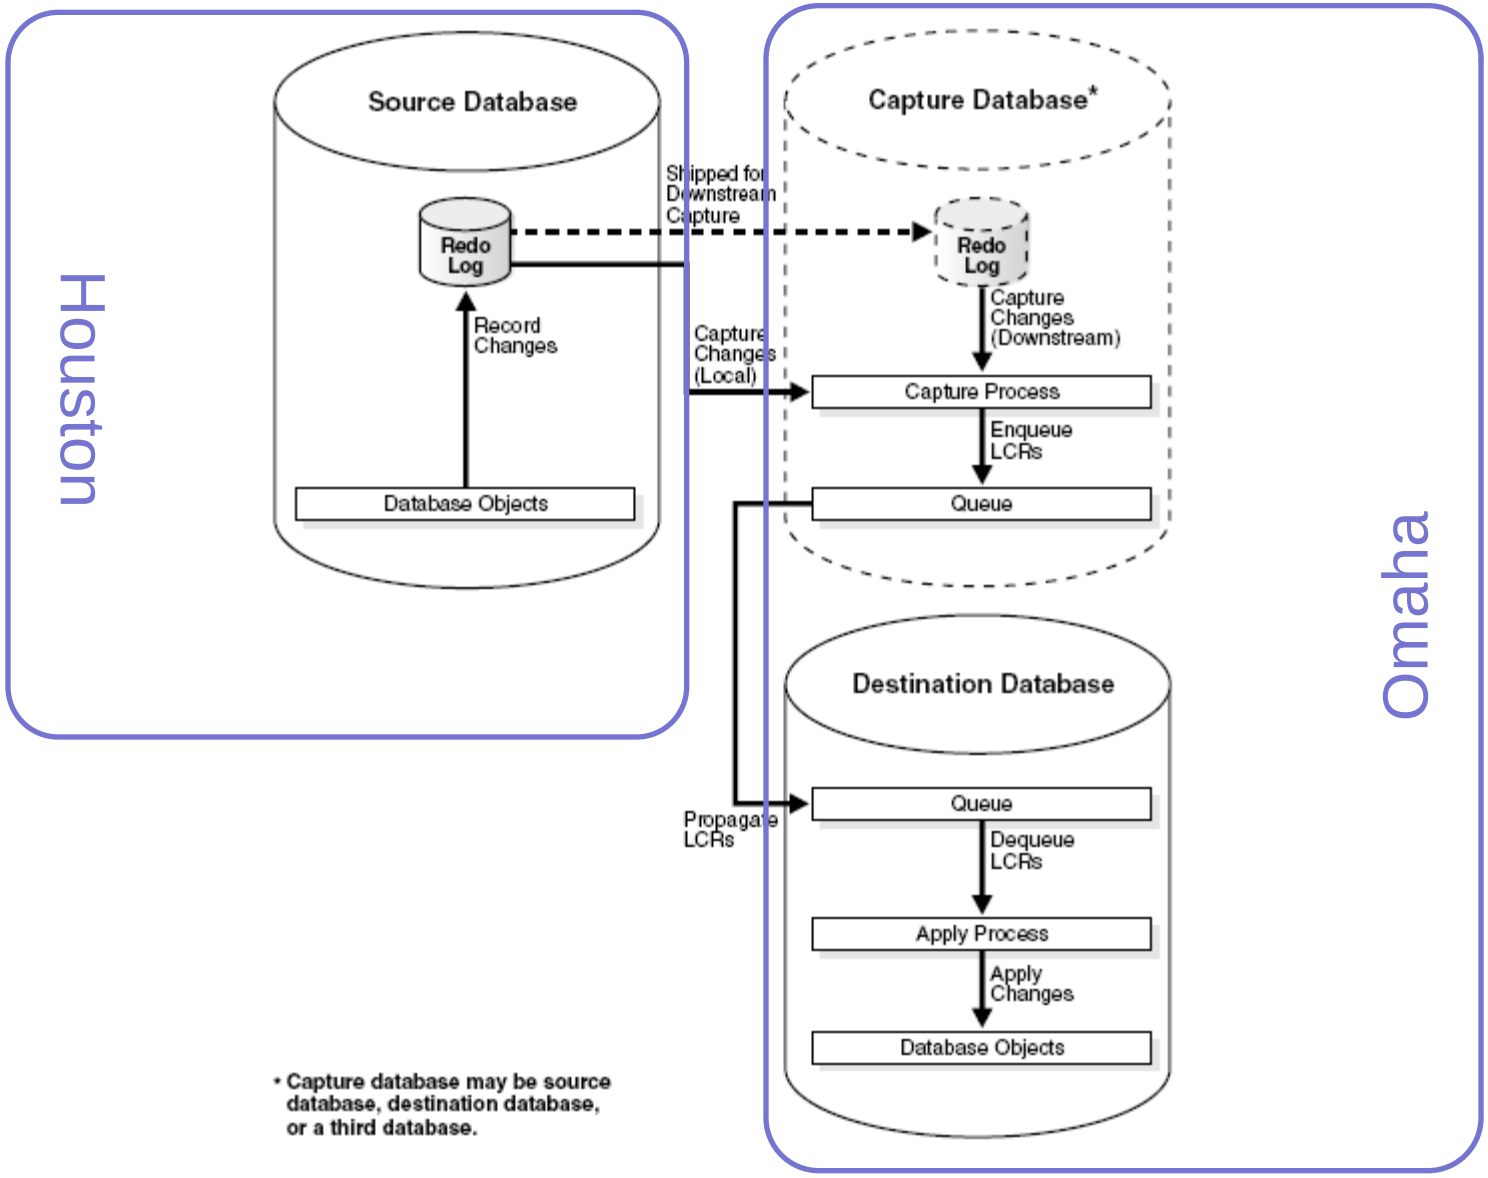

The story of SCDF

when there is only Streams and Convert…

Migrate with Oracle Streams and RMAN Convert

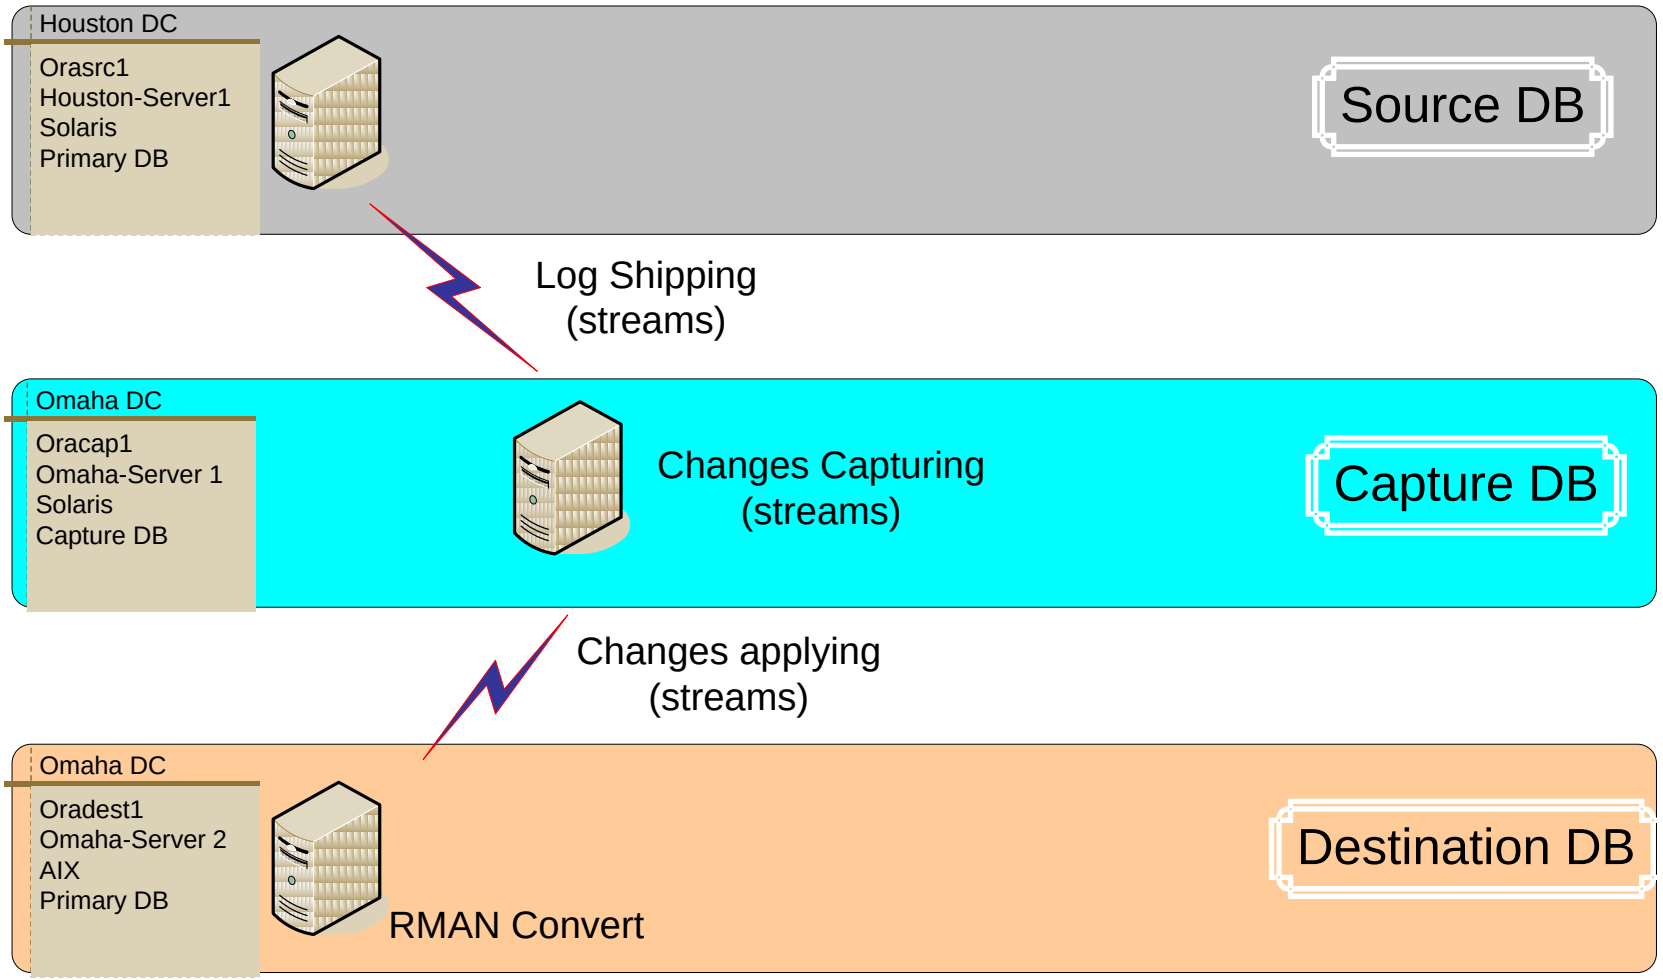

Where S + C Got Us

- Minimal (reduced to a few hrs) downtime (less than 15 minutes)
- Full DR (Data Guard) protection outside of the downtime window
- Ability to move a group of 5 databases within the same downtime window
- Extensive testing of all mission critical databases at different level (unit, system, integration)
- Extensive testing of datacenter infrastructure, Oracle database LPAR performance, Data Guard/DR operation before “go live”

Oracle Data Guard

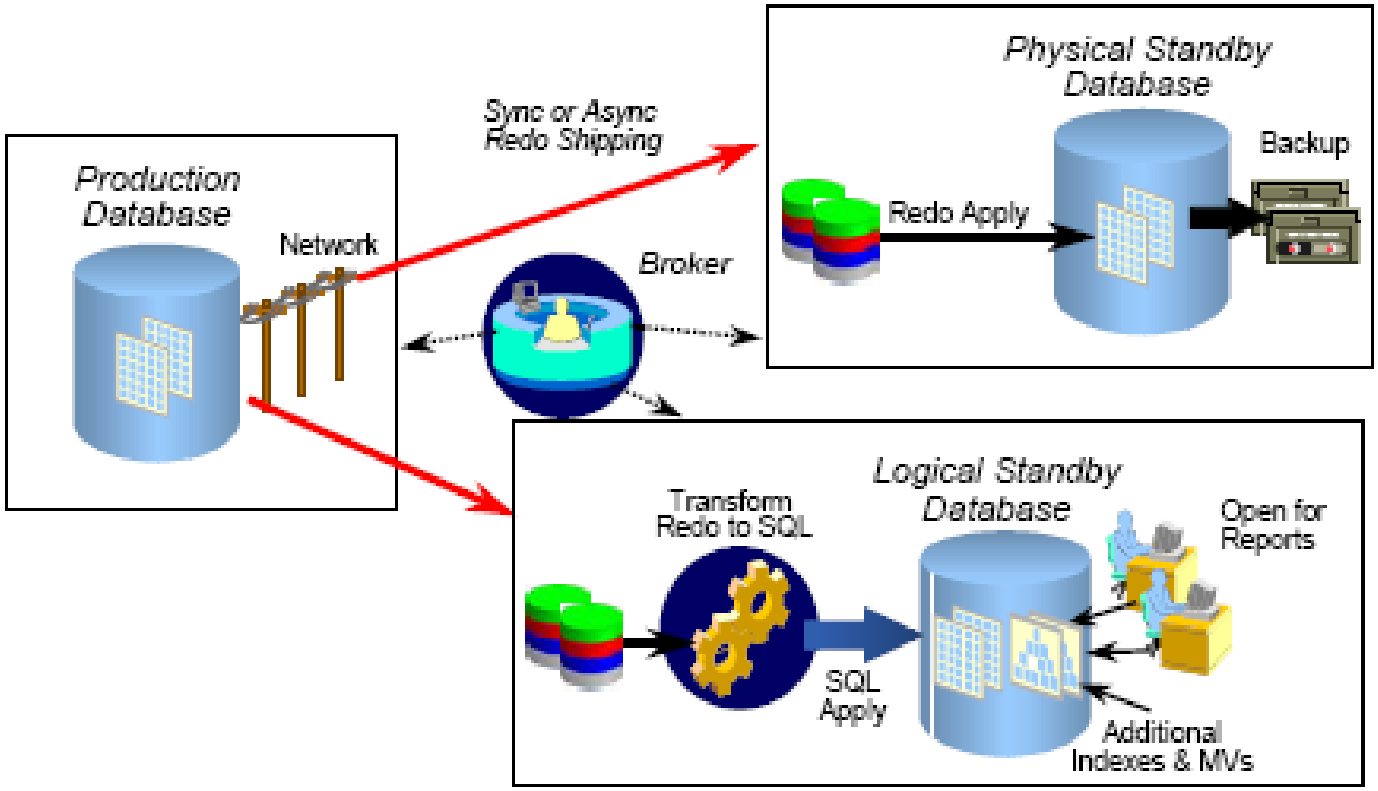

Oracle Data Guard Architectural Components

The story of SCDF

When there is S + C + Data guard...

stonebridge
outperform

Migrate with Oracle Streams, RMAN convert, and Data Guard

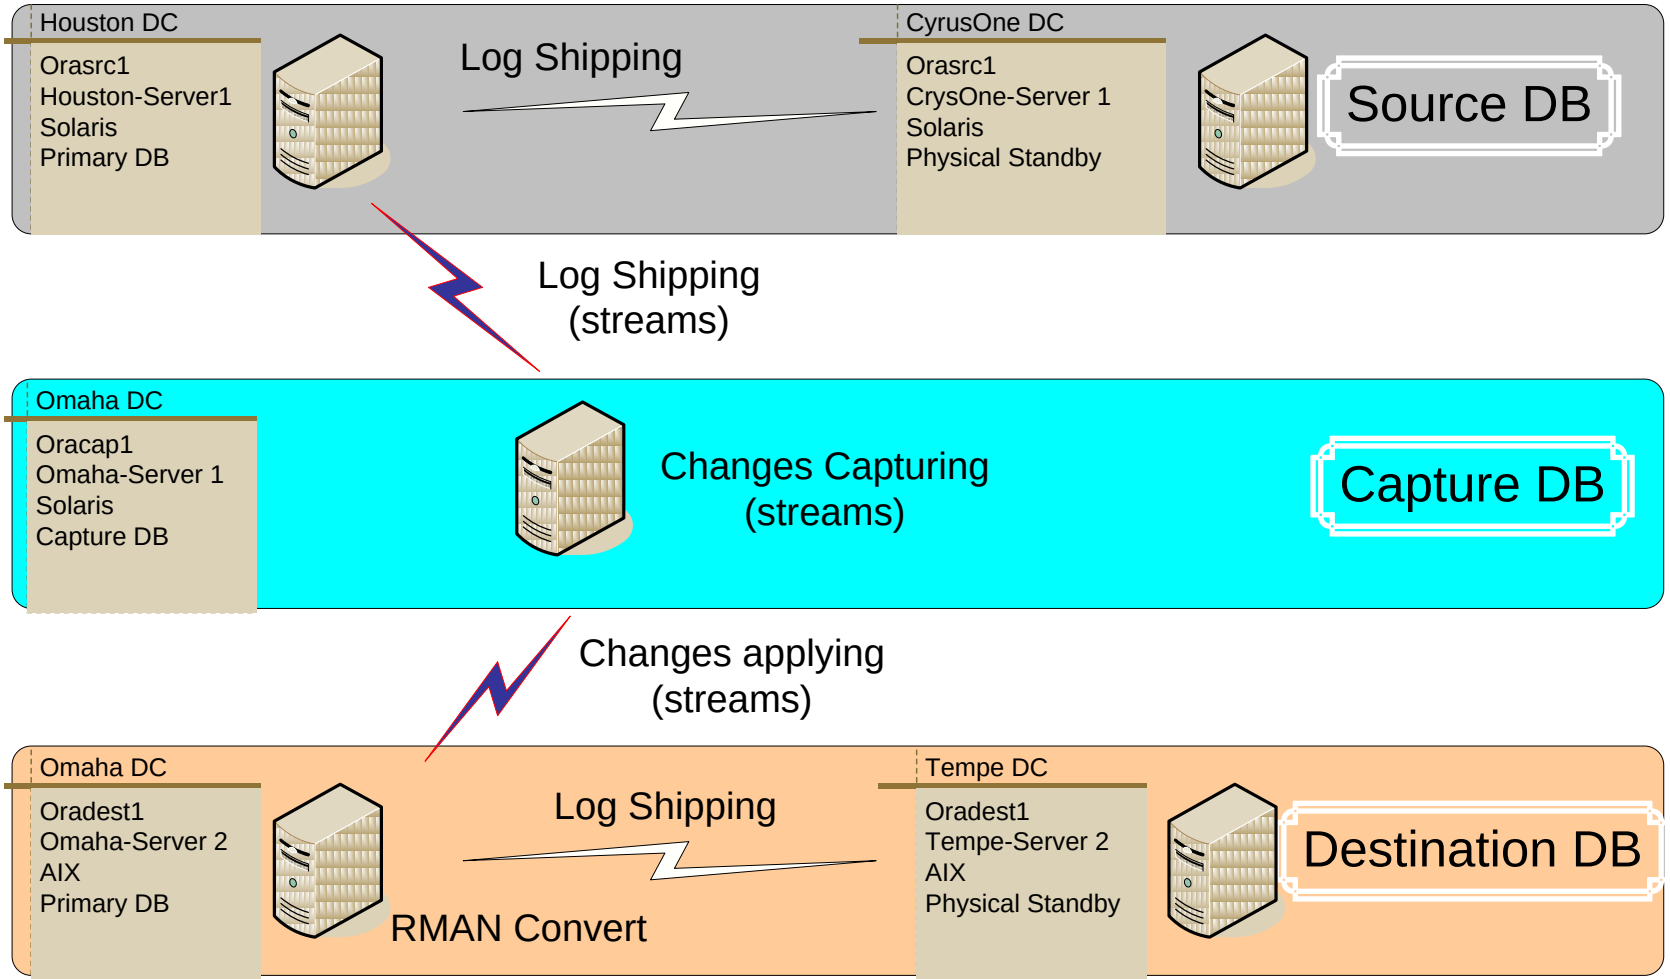

Where S + C + D Got Us

- Minimal downtime (less than 15 minutes)
- Full DR (Data Guard) protection outside of the downtime window
- Ability to move a group of 5 databases within the same downtime window
- Extensive testing of all mission critical databases at different level (unit, system, integration)
- Extensive testing of datacenter infrastructure, Oracle database LPAR performance, Data Guard/DR operation before go-live

Flashback Database

- Rewinds an Oracle database to a previous time, to correct any problems caused by logical data corruptions or user errors. It provides database point in time recovery without requiring a backup of the database to first be restored.

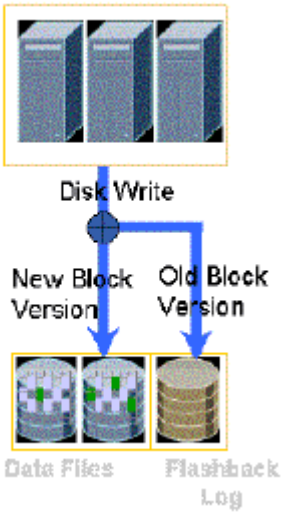

The story of SCDF

When there is S + C + D + Flashback...

Migration with Oracle Streams, RMAN Convert, Data Guard, and Flashback Database

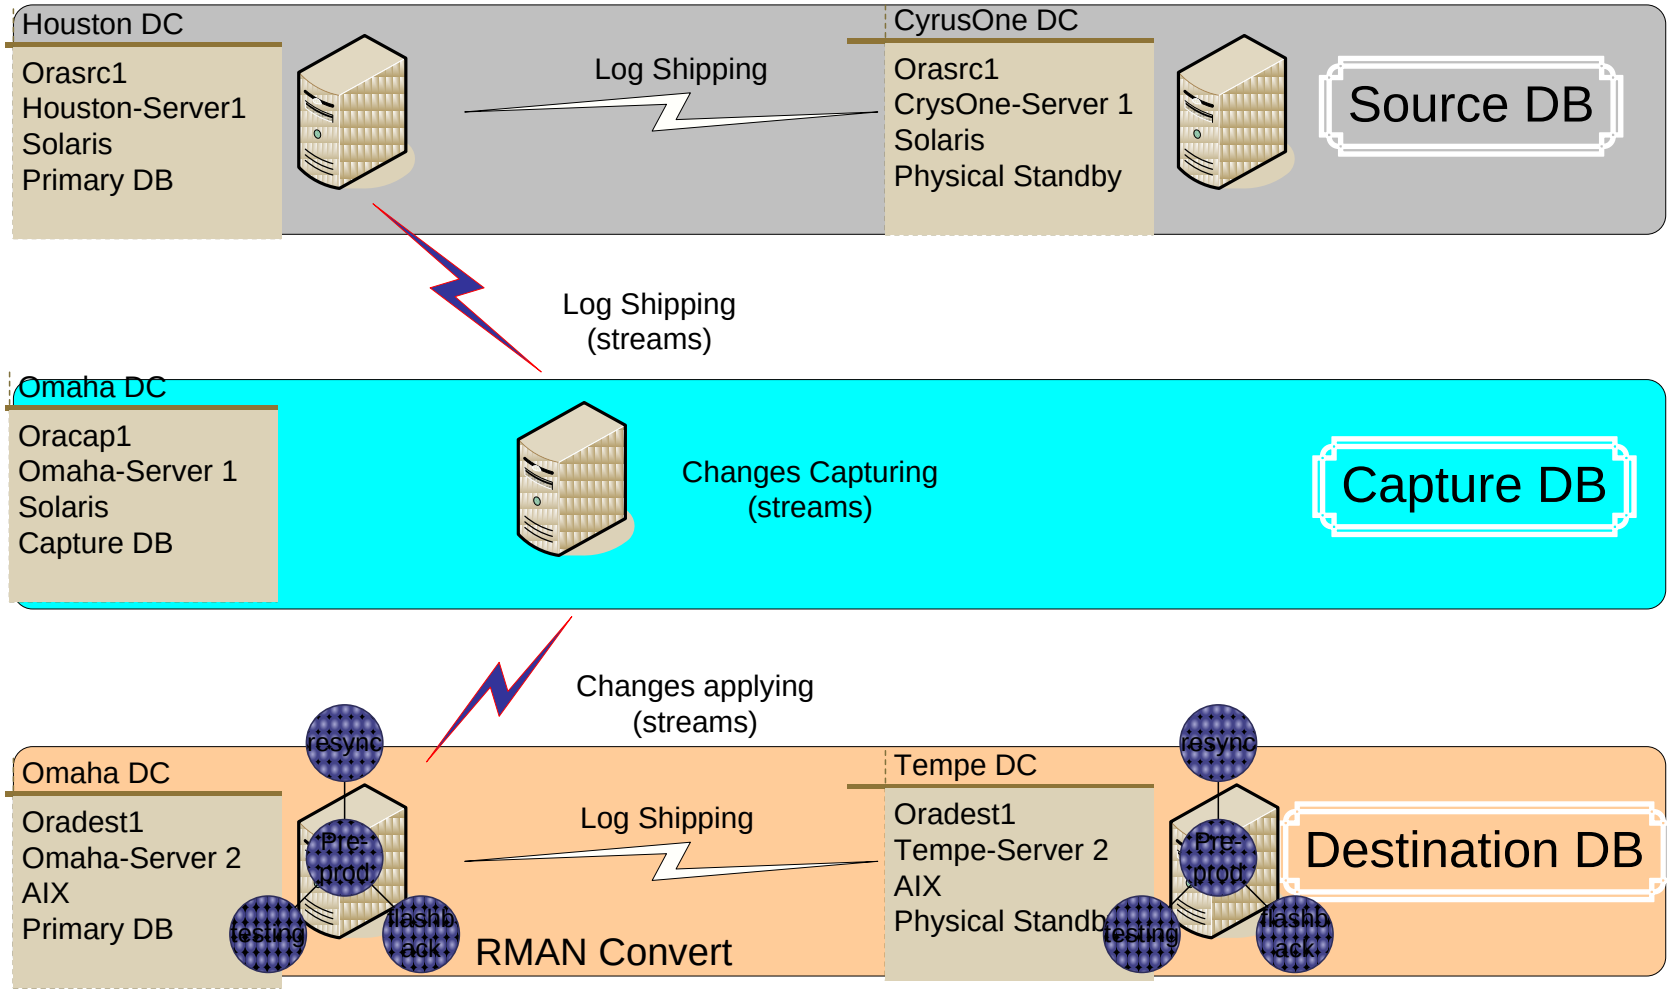

Where S + C +D + F Got Us

- Minimal downtime (less than 15 minutes)
- Full DR (Data Guard) protection outside of the downtime window
- Ability to move a group of 5 databases within the same downtime window
- Extensive testing of all mission critical databases at different level (unit, system, integration)
- Extensive testing of datacenter infrastructure, Oracle database LPAR performance, Data Guard/DR operation before go-live

Our Novel Approach: SCDF

- SCDF: Our novel Oracle Streams-centric approach

  S - Oracle **S**treams

  C - RMAN **C**onvert

  D – **D**ata Guard

  F - **F**lashback Database

SCDF In Action

Phases	Disruptive	Downtime
Setup	No*	0 or < 15 minutes
Application testing	No	0
“Go live”	Yes	< 15 minutes

stonebridge
outperform

SCDF Set-up (Non-disruptive)

- Create capture database in Omaha (4 – 8 hours)
- Instantiate destination database in Omaha using RMAN convert (24 – 36 hours)
- Establish Oracle streams between source/Houston and destination/Omaha (4 – 8 hours)
- Build standby database in Tempe (8 – 12 hours)

Target time to complete: 2 -3 days for each database

stonebridge
outperform

SCDF Application Test (Non-disruptive)

- Code freeze
- Stop Streams' change Capture and Apply processes
- Users conduct application test on destination database
- Flashback destination database once user test is done
- Start Streams' change Capture and Apply processes again
- Allow destination to catch up with Source database
- This testing/resynchronization process can be repeated until "go live".

Target time to complete: is highly application specific, but each test cycle was limited to 2 -3 days.

stonebridge
outperform

SCDF "Go Live" (Disruptive)

- Destination database final resync with source database
- Verify Tempe Standby in Sync with Primary in Omaha
- Shut down apps server/database in Houston
- Switch user to apps server/databases in Omaha
- Clean up Oracle Streams in source and target databases.

Target down-time for "go live": less than 15 minutes

stonebridge
outperform

SCDF Migration Results

Oracle DB	Application Group	Size (GB)	Desired Donwtime	Actual Downtime
DB1	Real Time Trading	23	< 15 minutes	< 15 minutes
DB2	Real Time Trading	80	< 15 minutes	< 15 minutes
DB3	Real Time Trading	129	< 15 minutes	< 15 minutes
DB4	Real Time Trading	157	< 15 minutes	< 15 minutes
DB5	Real Time Trading	185	< 15 minutes	< 15 minutes
DB6	Real Time Trading	425	< 15 minutes	2 hour *
DB7	Energy Trading Risk Management	8	< 15 minutes	< 15 minutes
DB8	Energy Trading Risk Management	27	< 15 minutes	< 15 minutes
DB9	Energy Trading Risk Management	28	< 15 minutes	< 15 minutes
DB10	Energy Trading Risk Management	207	< 15 minutes	< 15 minutes
DB11	Energy Trading Risk Management	313	< 15 minutes	8 hour *
DB12	Energy Trading Risk Management	381	< 15 minutes	< 15 minutes

SCDF Migration Results-Continue

- Due to Oracle streams limitations DB6 and DB11 were migrated using Oracle transportable database (TDB)

- With 2 – 8 hours downtime

Oracle DB	Application Group	Size (GB)	Desired Donwtime	Actual Downtime
DB6	Real Time Trading	425	< 15 minutes	2 hour *
DB11	Energy Trading Risk Management	313	< 15 minutes	8 hour *

stonebridge
outperform

SCDF Limitations/Restrictions

- The database is not part of an existing Oracle Streams environment
- The database is not part of an existing logical standby environment
- The database is not part of an existing Advanced Replication environment
- Materialized views or logs are not supported
- Any user-created queues are read-only during the maintenance operation
- All database tables should have primary or unique key (DB6)
- Certain datatypes may hit Oracle bugs (DB11)
- Endian format needs to be considered

Recap

- Combining Oracle **S**treams, RMAN **C**onvert, **D**ata Guard, and **F**lashback database technologies, we successfully migrated most of Calpine's mission critical databases with less than 15 minutes of downtime.

- The SCDF approach is highly stable, and flexible, but does require deep knowledge to implement.

Retail Experience

Rossi was the manager of an upscale men's wear store in a wealthy section of town and was interviewing Abe for the recently advertised salesman role.

Rossi looks at Abe's resume and notices that Abe has never worked in retail before.

Rossi says to Abe, "What chutzpah, if you don't mind me saying. For someone with no retail experience, you are certainly asking for a high salary."

"Well I suppose I am," Abe replies, "but you must understand that the work is so much harder when you don't know what you're doing."

stonebridge
outperform

Questions

?

References

- The Hidden Costs of Data Migration
- The hidden costs of data migration (IBM)
- Minimizing Data Migration Time, Using Cross-Platform Transportable tablespace
- Cross-Platform Oracle Database Migration Using Transportable Tablespaces and EMC Open Replicator for Symmetrix
- Platform Migration using Transportable Tablespaces: Oracle Database 10g Release 2
- Platform Migration Using Transportable Database Oracle Database 11*g* and 10*g* Release 2
- Oracle® Streams Concepts and Administration 10*g* Release 2 (10.2)

Contacts

- Caiping Yue, Principal Technical Consultant
- Stonebridge
- www.sbti.com
- caiping.yue@sbti.com
- (cell): 713-385-5846

- Dale Young, Director of Database Services
- Stonebridge
- www.sbti.com
- dale.young@sbti.com
- (cell): 972-365-9877