
ORACLE®

ORACLE DATABASE 11g

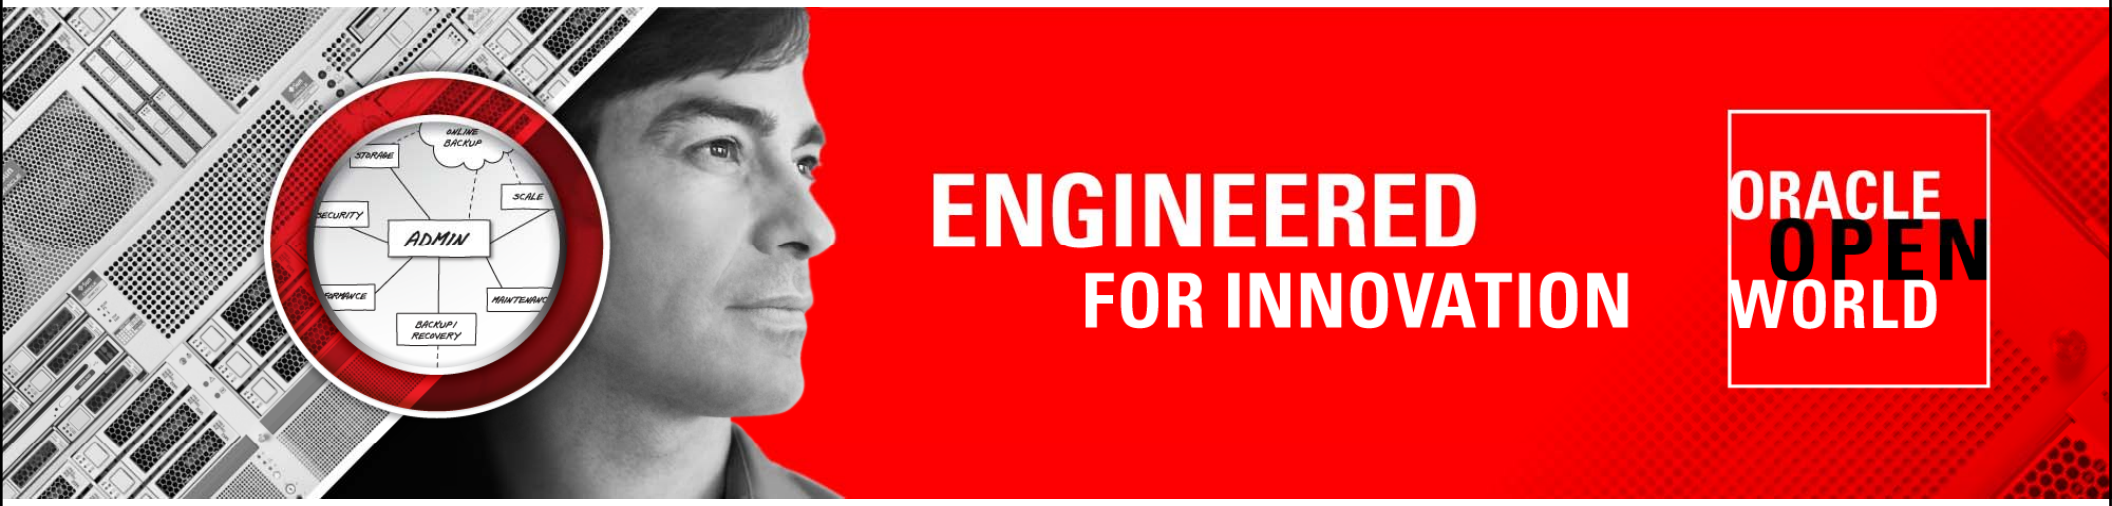

ORACLE

Recovery Manager (RMAN) and Oracle Data Guard: Seven Cool Tips from Oracle

Mahesh Girkar, Oracle
Ashish Ray, Oracle

Presenting with
Lisa Reinheimer
ENTERPRISE HOLDINGS

ORACLE OPEN WORLD

ENGINEERED FOR INNOVATION

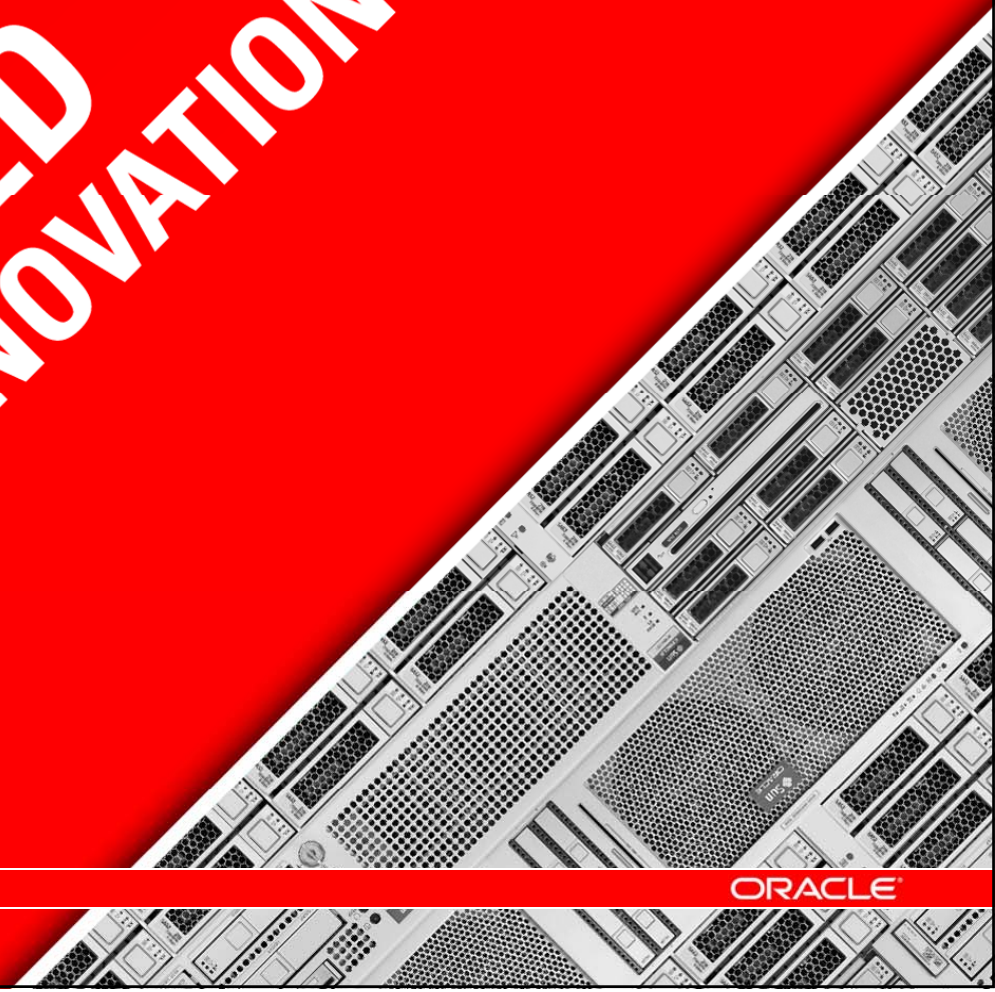

ORACLE

Latin America 2011

December 6–8, 2011

Tokyo 2012

April 4–6, 2012

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Oracle OpenWorld Bookstore

- Visit the Oracle OpenWorld Bookstore for a fabulous selection of books on many of the conference topics and more!
- Bookstore located at Moscone West, Level 2
- All Books at 20% Discount

DigitalGuru
Technical Bookshop

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Oracle Products Available Online

Store

SHOP NOW

Buy Oracle license and support online today at
oracle.com/store

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Program Agenda

- Introduction – Oracle Data Protection
- 3 Cool Tips – RMAN
- 3 Cool Tips – Data Guard
- Putting it Together & 7th Tip: Enterprise Holdings

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Cool Tips: Focusing on Oracle Data Protection

RMAN + Secure Backup + Flashback + Data Guard

Integrated backup & recovery, continuous data protection, disaster recovery

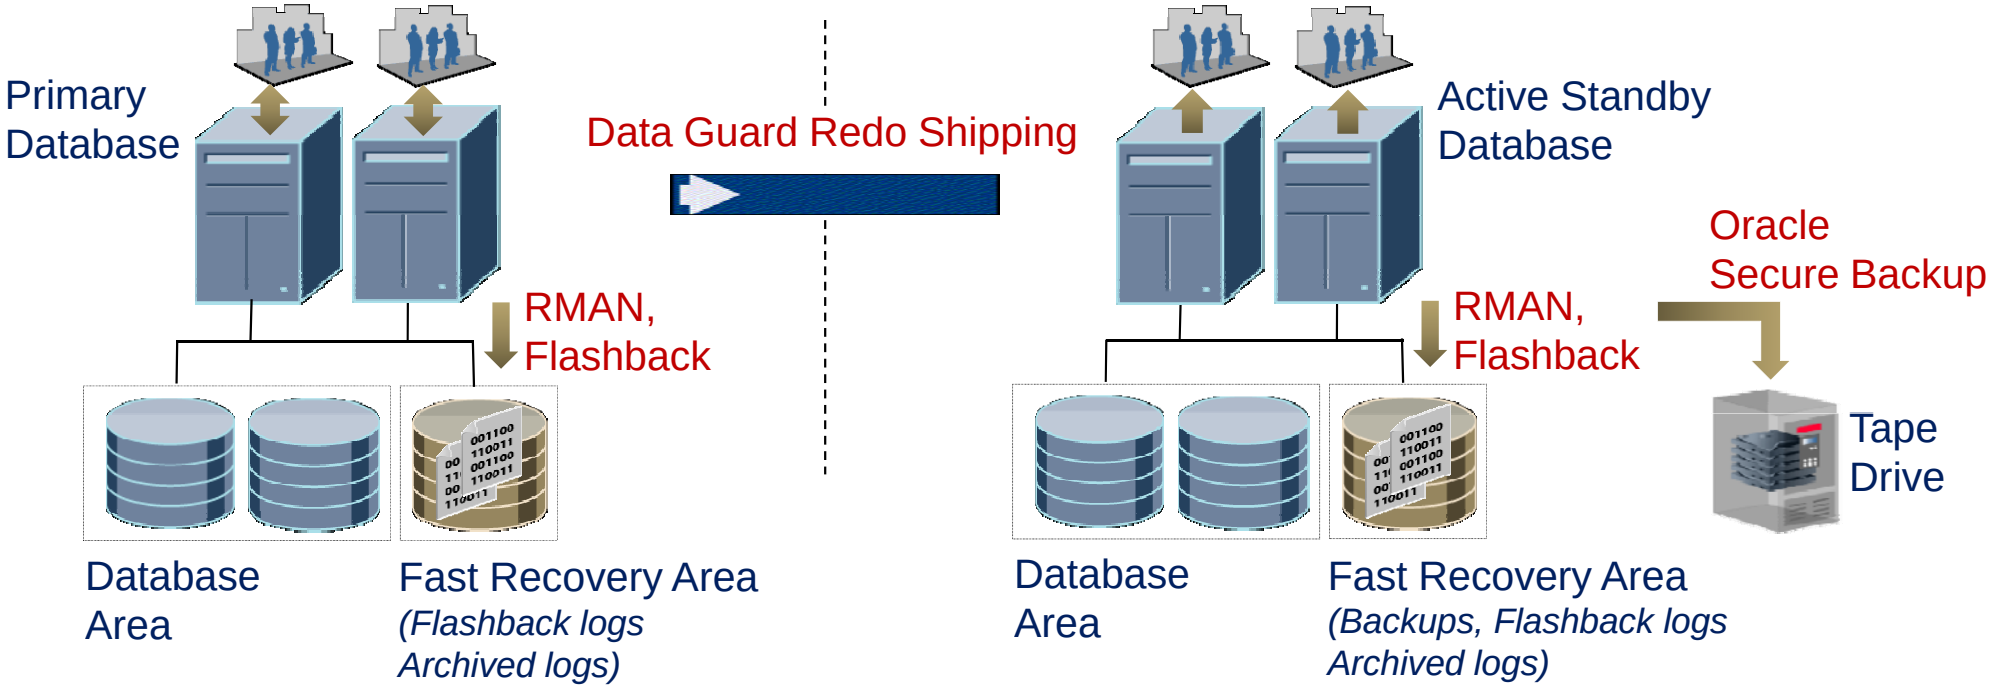

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Program Agenda

- Introduction – Oracle Data Protection
- 3 Cool Tips – RMAN
- 3 Cool Tips – Data Guard
- Putting it Together & 7th Tip: Enterprise Holdings

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Really: Flashback + RMAN Cool Tips

Theme: Combo of Best Practices Plus Tips

1. Guaranteed Restore Point Without Flashback Impact
2. RMAN Switch to Copy – Super Quick Restore
3. Archived Log Deletion Policies with Standby

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Guaranteed Restore Point with no I/O Impact

Why Important?

- Enables creation of discrete point-in-time markers, that are very easy to revert to using Flashback Database

- However, with this tip …
  - *Reduced flashback logging overhead*
  - *No significant additional space consumption*

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Guaranteed Restore Point with no I/O Impact

Flashback Database – Quick Introduction

- Flashback Database: Oracle-integrated efficient way for disk-based continuous data protection (CDP) to quickly undo data changes

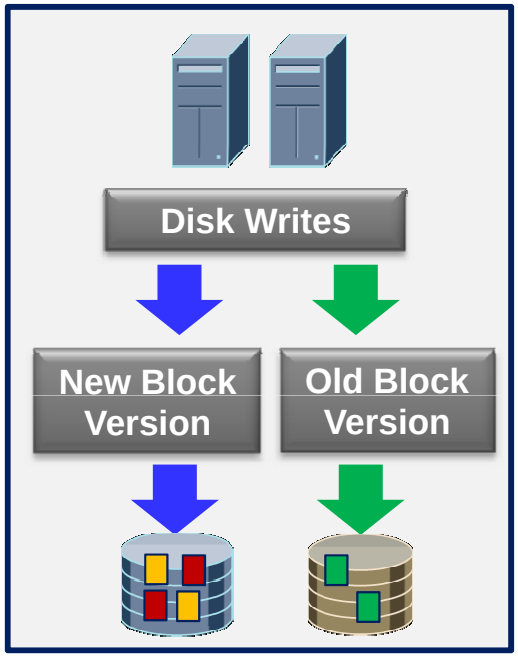

- Tracks data changes efficiently in flashback logs
- Simply rewind to SCN / time / restore point

```
SQL> flashback database to <timestamp>;
```

- Ref. Support Note: 565535.1 – Flashback Database Best Practices

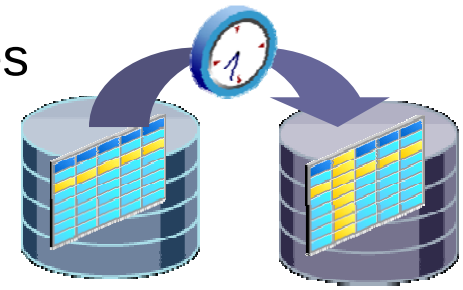

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Guaranteed Restore Point with no I/O Impact

Restore Point Overview & Use Cases

- Normal Restore Point: user-friendly name to associate with a particular point-in-time
  - Can be used to mark point-in-time before starting a workload, application upgrade, etc.: `SQL> CREATE RESTORE POINT before_upgrade;`
- Guaranteed Restore Point (GRP): ensures that relevant flashback logs never age out
  - `SQL> CREATE RESTORE POINT before_upgrade GUARANTEE FLASHBACK DATABASE;`
- Intuitive use with Flashback
  - `SQL> FLASHBACK DATABASE TO RESTORE POINT before_upgrade;`

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Guaranteed Restore Point with no I/O Impact

Use Without Enabling Flashback Database

- If only GRP needs to be retained (i.e. no other point-in-time), no need to explicitly enable Flashback Database

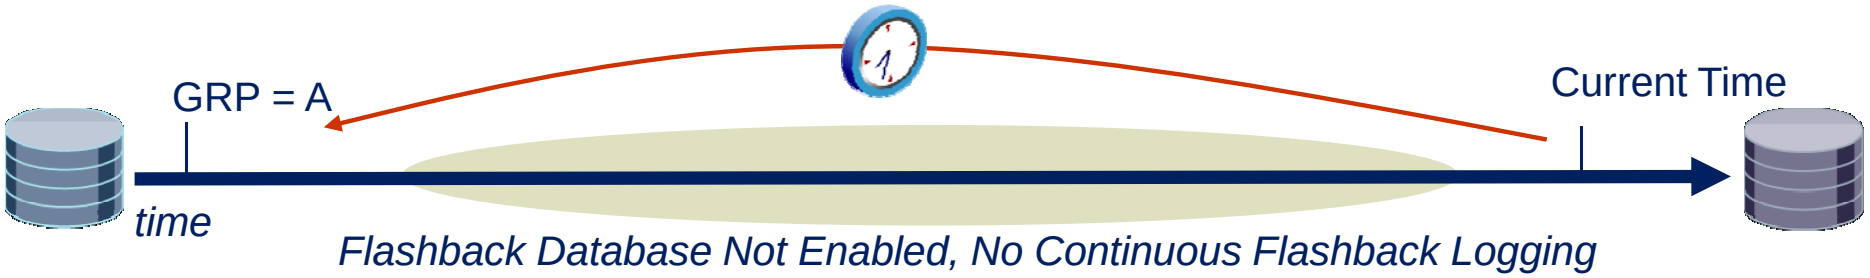

- Ideal use cases: when you do not care about data changes made during a time period, but want ability to go back to start of changes – e.g. reversing an upgrade, snapshot standby, temporary databases (QA / Dev / Training)
- Benefits of Flashback Database without any significant I/O overhead

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Guaranteed Restore Point with no I/O Impact

How Does It Work?

Click Here to Access Super Quick Demo

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Flashback + RMAN Cool Tips

Theme: Combo of Best Practices Plus Tips

1. Guaranteed Restore Point Without Flashback Impact
2. RMAN Switch to Copy – Super Quick Restore
3. Archived Log Deletion Policies with Standby

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

RMAN Switch Datafile to Copy

Why Important?

- If data files get damaged or deleted or are inaccessible, this gives a very quick way to use the backup copy in Fast Recovery Area (FRA), and continue operations
- No impact to the rest of the database
- Doesn’t involve any time consuming restore operation

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

RMAN Switch Datafile to Copy

Fast Recovery Area (FRA)

- Disk location to automate management of Oracle database backup and recovery related files

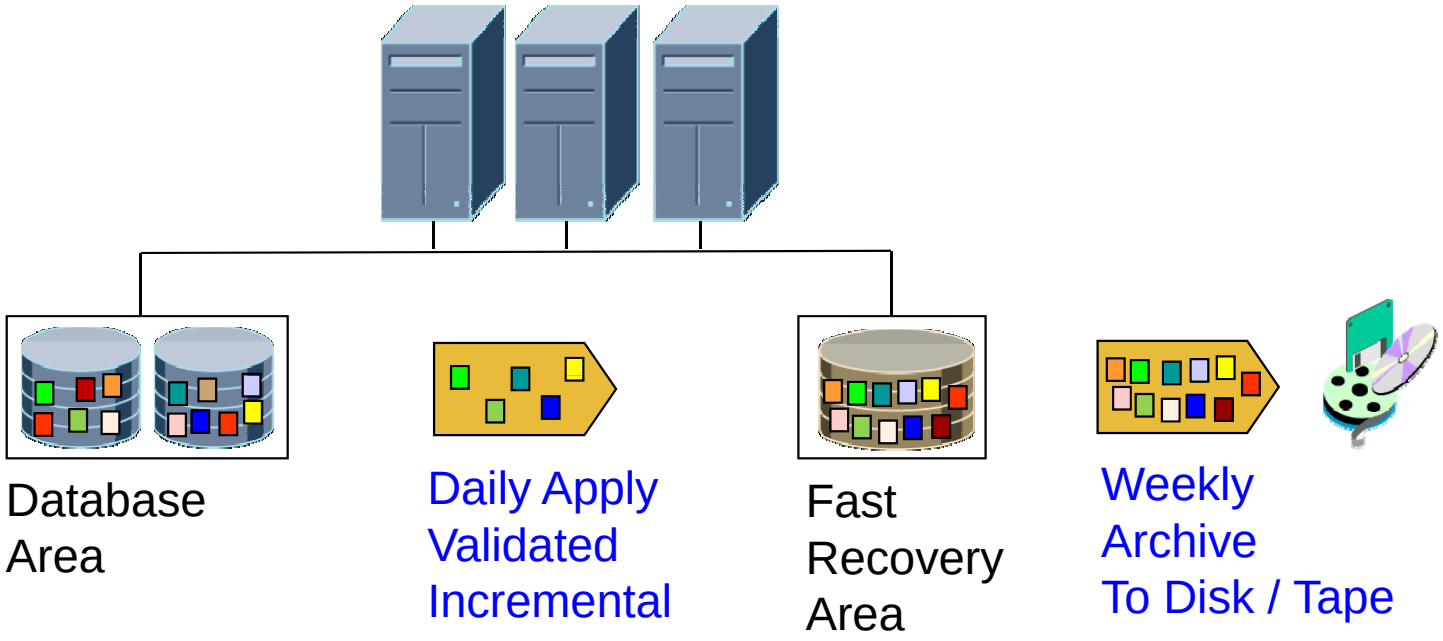

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

RMAN Switch Datafile to Copy

Super Quick Restore Without A Restore

- On data file access problems, eliminate restore time by simply redirecting database to use recent image copy as new data file

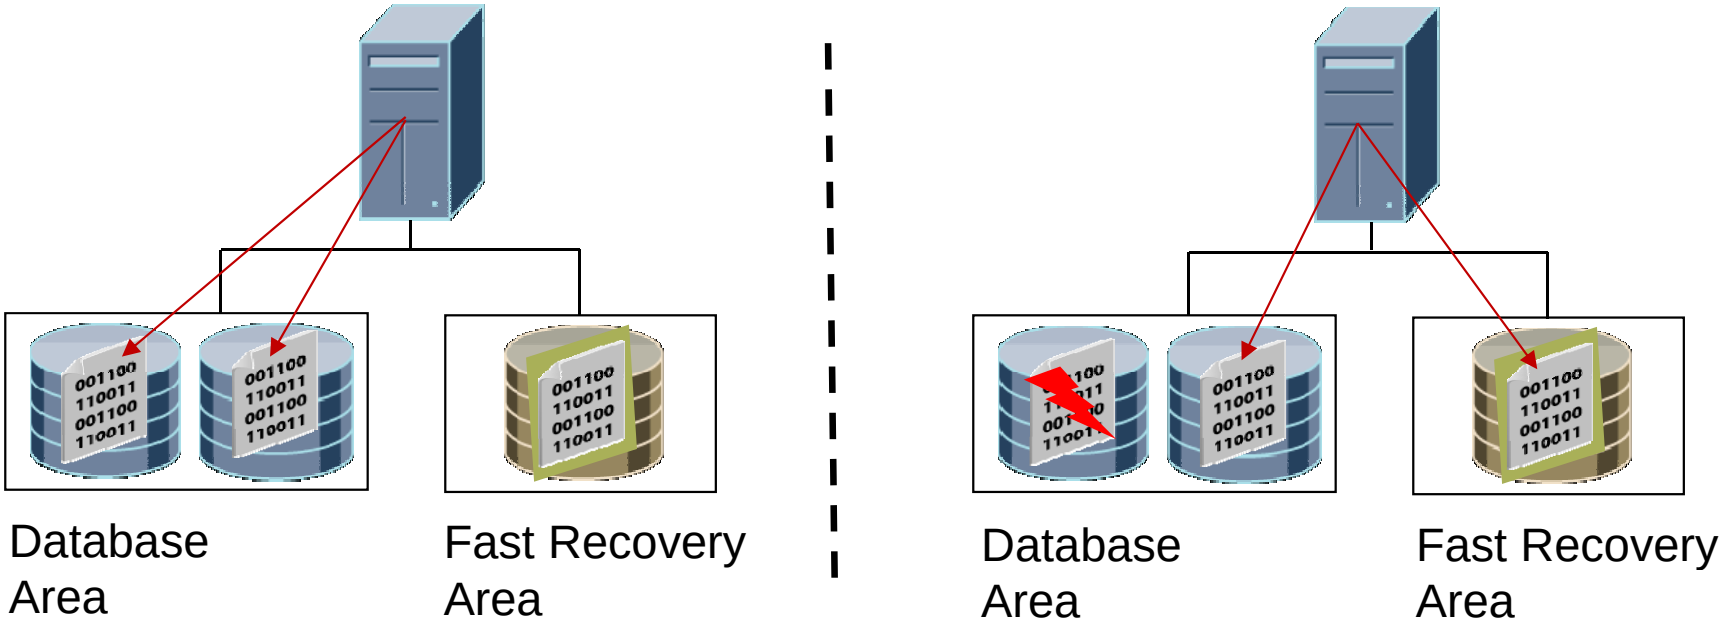

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

RMAN Switch Datafile to Copy

How Does It Work?

Click Here to Access Super Quick Demo

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

RMAN Switch Datafile to Copy

How About Switching Back?

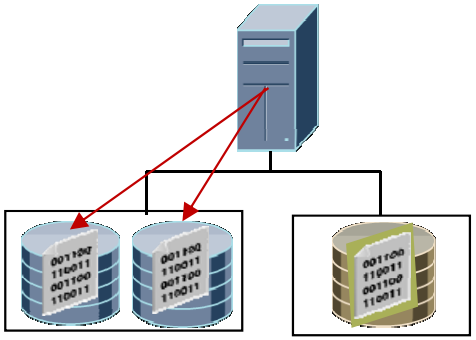

- Steps to switch back:
  - Make an image copy of the datafile to DATA disk group
    - `BACKUP AS COPY DATAFILE 4 FORMAT '+DATA';`
  - Make datafile offline
    - `RMAN> SQL 'ALTER DATABASE DATAFILE 4 OFFLINE';`
  - Switch datafile to copy (however, in this case, the "copy" is on DATA)
    - `RMAN> SWITCH DATAFILE 4 TO COPY;`
  - Recover datafile
    - `RMAN> RECOVER DATAFILE 4;`
  - Make datafile online
    - `RMAN> SQL 'ALTER DATABASE DATAFILE 4 ONLINE';`

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

RMAN Switch Datafile to Copy

Considerations

- Doesn't this mean I am using up my only backup?
  - Yes, but that's what backups are for, to be used in emergencies!
- What if actual restore from disk is super fast – e.g. Exadata?
  - Exadata: image copy backup / restore rate achieved (MAA testing): 17 TB/hr, a restore strategy should be considered if not much data to restore
- If several data files damaged?
  - Consider failing over to Data Guard standby
- How soon should I switch back?
  - This method is for a quick fix – you should switch back soon enough so that you can resume running production from the usual DATA area

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Flashback + RMAN Cool Tips

Theme: Combo of Best Practices Plus Tips

1. Guaranteed Restore Point Without Flashback Impact
2. RMAN Switch to Copy – Super Quick Restore
3. Archived Log Deletion Policies with Standby

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Log Deletion Policies with Physical Standby

Why Important?

- Fast Recovery Area enables efficient backup & space management with auto-deletion of archived log & backup files

- In an integrated RMAN + Data Guard environment, very critical to ensure right automation policies are in place

ORACLE

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

RMAN and Data Guard Physical Standby

Offload Backup Processing

- Offload backups to physical standby database
- Share backups among databases (requires catalog)
- Create standby database efficiently using RMAN

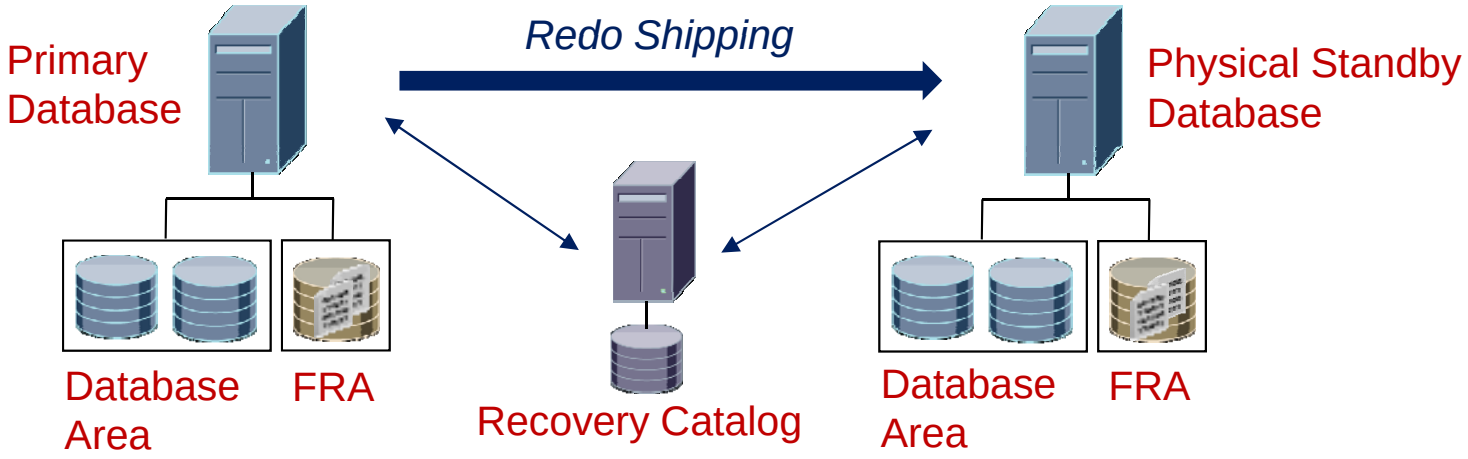

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

RMAN and Data Guard Physical Standby

Basic Configuration

- Primary database
  - **REGISTER DATABASE**
  - **CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF <n> DAYS**
- Database where backups are taken
  - **CONFIGURE CONTROLFILE AUTOBACKUP ON**
  - **CONFIGURE BACKUP OPTIMIZATION ON**
- *What about archived log deletion policies?*

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

RMAN and Data Guard Physical Standby

Configuration – Archived Log Deletion Policies

- If backups done on Standby
  - Primary database & other by-stander standby databases (if any)
    - `CONFIGURE ARCHIVELOG DELETION POLICY TO APPLIED ON ALL STANDBY`
  - Standby database
    - `CONFIGURE ARCHIVELOG DELETION POLICY TO APPLIED ON ALL STANDBY BACKED UP <n> TIMES TO DEVICE TYPE [SBT | DISK]`
- If backups done on Primary
  - Primary database
    - `CONFIGURE ARCHIVELOG DELETION POLICY TO APPLIED ON ALL STANDBY BACKED UP <n> TIMES TO DEVICE TYPE [SBT | DISK]`
  - All standby databases
    - `CONFIGURE ARCHIVELOG DELETION POLICY TO APPLIED ON ALL STANDBY`
- After switchover / failover, re-execute policies if backup locations switch

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Program Agenda

- Introduction – Oracle Data Protection
- 3 Cool Tips – RMAN
- 3 Cool Tips – Data Guard
- Putting it Together & 7th Tip: Enterprise Holdings

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Data Guard Cool Tips

Theme: Combo of Best Practices Plus Tips

1. Flush Redo: Ensure Zero Data Loss Failover
2. Apply Instance Failover with Data Guard Broker
3. Query Causality with Active Data Guard

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Data Guard Cool Tip – Flush Redo

Switchover: Minimize Planned Downtime

- Coordinated role transition between Primary & Target Standby
- Guaranteed no data loss after role transition
- Requires: Open read-write (healthy) Primary

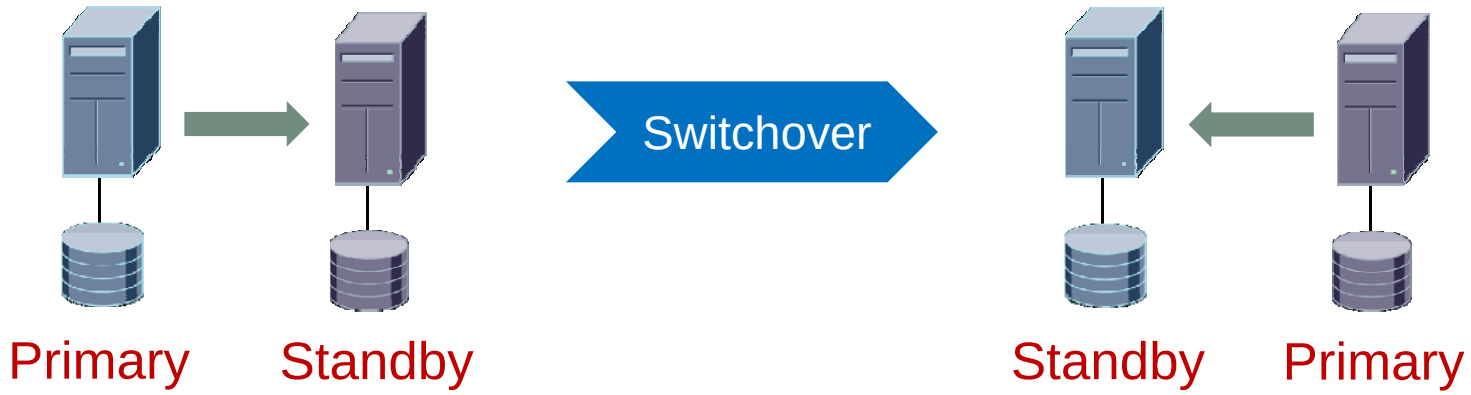

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Data Guard Cool Tip – Flush Redo

Failover: Minimize Unplanned Downtime

- Primary unavailable – because of failure / outage
- Designated standby becomes new primary

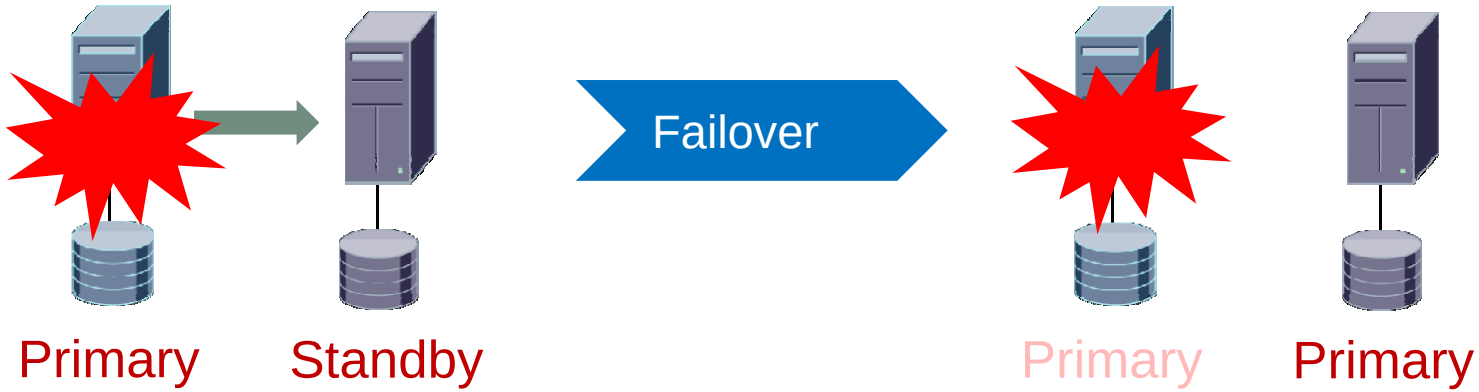

- With Fast-Start Failover, old primary auto-reinstated to new standby

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Data Guard Cool Tip – Flush Redo

Many Faces of “Disaster”

- Example: Primary database cannot open because of system data file corruption

```
SQL> alter database open;  --- primary cannot be opened
alter database open
*
ERROR at line 1:
ORA-01122: database file 1 failed verification check
ORA-01110: data file 1: '/ade/b/707227083/oracle/dbs/t_db1.f'
ORA-01210: data file header is media corrupt
```

- Failover possible, but what about data loss?

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Data Guard Cool Tip – Flush Redo

Data Loss Issues in Failover

- Failover with Synchronized Standby in Max Availability:
  - No Data Loss
- Failover in Max Performance, or with Unsynchronized Standby in Max Availability:
  - Some Data Loss: unsent data in current online redo logs
  - Missing archived redo logs need manual registration (otherwise, more data loss)

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Data Guard Cool Tip – Flush Redo

Enabling Zero Data Loss Failovers

- Requires:
  - Mounted primary (system / user datafiles do not need to be available)
  - Access to control file, archived & online redo logs
- Flush redo DDL at (mounted) primary database:
  - Ships all redo to target standby
  - Waits until “all redo applied” confirmation (requires recovery at standby)

```
SQL> alter system flush redo to ‘boston';
```

  - Waits until “all redo received” confirmation

```
SQL> alter system flush redo to ‘boston‘ no confirm apply;
```


Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Data Guard Cool Tip – Flush Redo

Target Standby Behavior

- Messages in alert.log:

```
Media Recovery: FLUSH REDO EOR logs encountered.
Media Recovery recovers through FLUSH REDO EOR logs.
Tue Aug 30 21:45:36 2011
……
Standby switchover readiness check: Checking whether recovery
applied all redo..
Physical Standby applied all the redo from the primary.
```

- Failover process can now be completed with no data loss

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Data Guard Cool Tips

Theme: Combo of Best Practices Plus Tips

1. Flush Redo: Ensure Zero Data Loss Failover
2. Apply Instance Failover with Data Guard Broker
3. Query Causality with Active Data Guard

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Apply Instance Failover

Overview

- Multi-node RAC standby: Managed Recovery Process (MRP) runs in only one node – the Apply Instance
- If that node fails, how does MRP failover to another node, and resume Redo Apply?
- Solution: Data Guard Broker

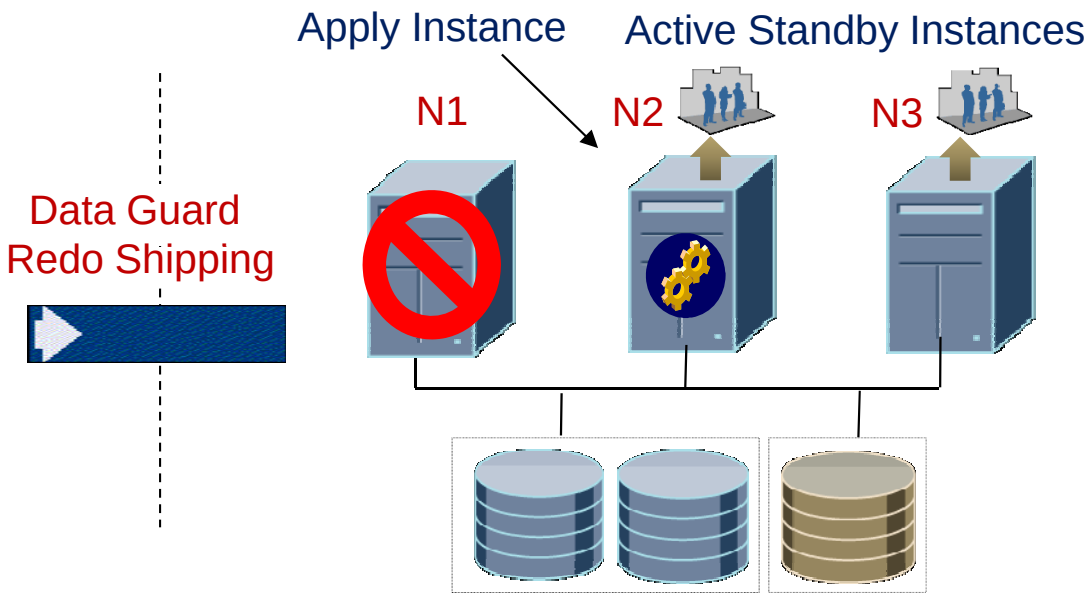

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Apply Instance Failover

How Does The Broker Do It?

- Two simple Broker properties
  - **`PreferredApplyInstance`**: name of the instance you prefer to be the apply instance
  - **`ApplyInstanceTimeout`**: time period that Broker waits after detecting apply instance failure and before initiating the failover; default = 0
- How the Broker chooses the new apply instance
  - If **`PreferredApplyInstance`** indicates an instance that is currently running, Broker selects it as the new apply instance
  - Otherwise Broker picks a random instance that is currently running

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Apply Instance Failover

Nuances: Active Standby Reader Instances

- With Broker, reinstating the pre-failure state is automatic
  - Broker opens instances that had to be closed upon apply instance failure to preserve query consistency (ref. Support Note 1357597.1 - Active Data Guard RAC Standby)
    - Instances are open read-only if they were open prior to failure
    - Otherwise instances left mounted
    - Redo Apply started on one of the remaining instances (based on `PreferredApplyInstance`, if available)

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Data Guard Cool Tips

Theme: Combo of Best Practices Plus Tips

1. Flush Redo: Ensure Zero Data Loss Failover
2. Apply Instance Failover with Data Guard Broker
3. Query Causality with Active Data Guard

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Causality

What is it?

Causality: relationship between an event (the cause) and a second event (the effect), where the second event is understood as a consequence of the first *[Wikipedia]*

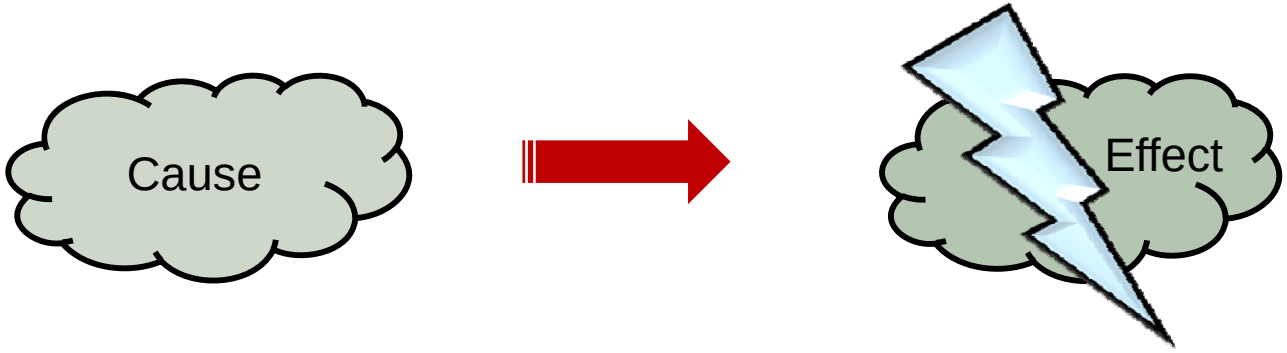

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Query Causality with Active Data Guard

What is it in Data Guard Terms?

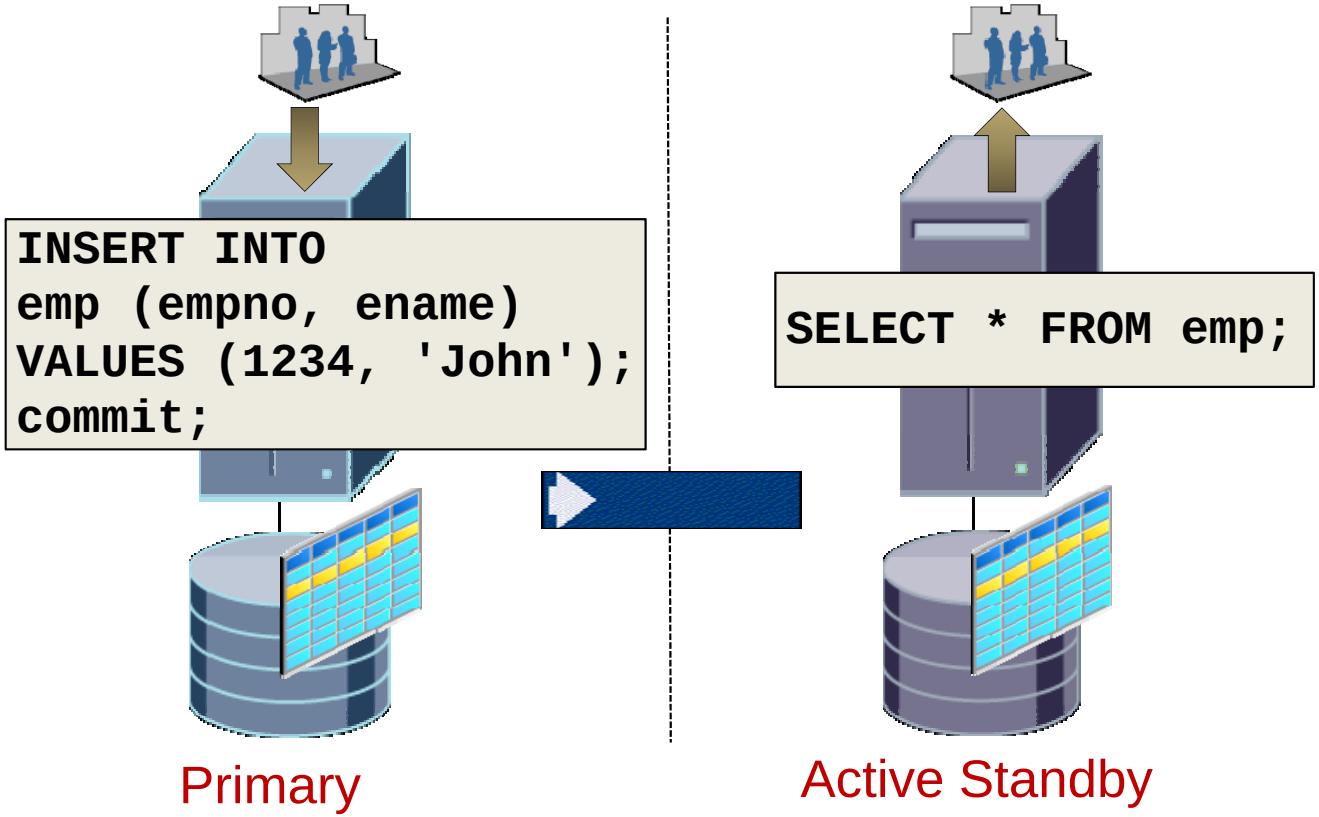

- Ensuring that ..
  - My Active Data Guard Standby has the absolute latest data for my query

    Or
  - My query on Active Standby sees all *committed* transactions that occurred before the start of my query

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Query Causality with Active Data Guard

Requirements: Real Time Apply

- Why Real Time Apply:
  - Reporting on the latest committed data
  - Faster failover times
    - Redo applied directly from standby redo logs instead of archived redo logs => standby synchronized much faster
  - Faster detection of corruptions, if any
    - Don’t need to wait for log switches to discover corruption

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Query Causality with Active Data Guard

Other Requirements

- Maximum Availability with SYNC transport
- Session configurations on Active Standby before issuing query:
  - `alter session set standby_max_data_delay=0;` OR
  - `alter session sync with primary;`

Ensures Active Standby is in lock-step with the Primary!

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Query Causality

Causality with Maximum Performance – Implications

- For an Active Standby in Maximum Performance (ASYNC)
  - Scenario 1:
    - `alter session sync with primary; =>`
    - **`ORA-03173: Standby may not be synced with primary`**
  - Scenario 2:
    - `alter session set standby_max_data_delay=0`
    - `select * from scott.emp; =>`
    - **`ORA-03172: STANDBY_MAX_DATA_DELAY of 0 seconds exceeded`**
  - Both error out immediately: no causality
- How can we simulate Causality with ASYNC?

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Query Causality

Ensuring Causality with Maximum Performance

- Grant **SELECT ANY DICTIONARY** or **SELECT_CATALOG_ROLE** role to session account
- Perform Update / Insert / Delete and Commit at Primary
- Determine values
  - **CURRENT_SCN** from **V$DATABASE** at Primary: = **Commit SCN**
  - **CURRENT_SCN** from **V$DATABASE** at Active Standby: = **Query SCN**
- Reporting application connected to Active Standby waits until:
  - **Query SCN >= Commit SCN**
- Now query data at Standby – has Causality semantics!

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Program Agenda

- Introductions – Oracle Data Protection
- 3 Cool Tips – RMAN
- 3 Cool Tips – Data Guard
- Putting it Together & 7th Tip: Enterprise Holdings

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Enterprise Holdings

Oracle Data Guard - Strategic Overview

Lisa Reinheimer

Systems Architect

Who Is Enterprise Holdings?

- World's largest car rental company in revenue, employees and fleet
- Operator of Alamo Rent A Car, National Car Rental and Enterprise Rent-A-Car brands
- More than twice as many US locations as our nearest competitor

ENTERPRISEHOLDINGS

Two Key Architectures

1. **Fast-Start Failover – Mission Critical Systems**
   - 11gR1, Physical Standby, ASYNC
   - Flashback / Fast Recovery Area
   - GoldenGate – Feeds To Warehouse
   - Transient Logical Standby for Upgrades & PSU's
   - Single Instance per site
2. **Active Data Guard – Business Critical Systems**
   - 11gR1, Primary can be down at times, < 24x7 requirement
   - Read-Only Access 24x7x365
   - Physical Standby, ASYNC, Broker configured
   - Primary can handle all Read-Only load if necessary

ENTERPRISE HOLDINGS

Data Guard at Enterprise Holdings

Fast-Start Failover – Mission Critical Systems

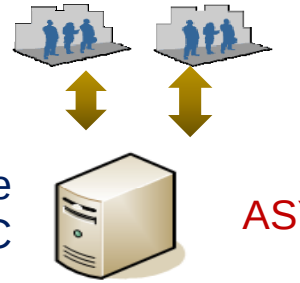

Primary
Database
Non-RAC
< 24x7

>500 Miles < 1000 Miles

< 25 ms latency OC12

ASYNC Data Guard Redo Shipping

Physical Standby
Real Time Apply
FSFO – 180 second threshold
Fast Start Failover Lag Limit – 30

RMAN,
Flashback

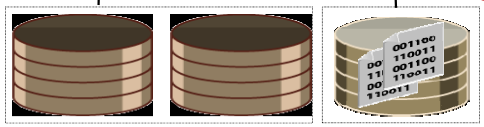

Database
Area

Fast Recovery Area
(Flashback logs
Archived logs)

Flashback

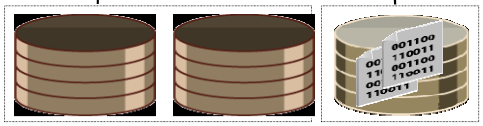

Database
Area

Fast Recovery Area
(Backups, Flashback logs
Archived logs)

GoldenGate
(ALO)

Data Warehouse

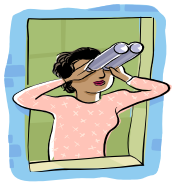

Observer
Process

OEM Grid Control
Manages Observer

ENTERPRISEHOLDINGS

Data Guard at Enterprise Holdings

Active Data Guard – Business Critical Systems

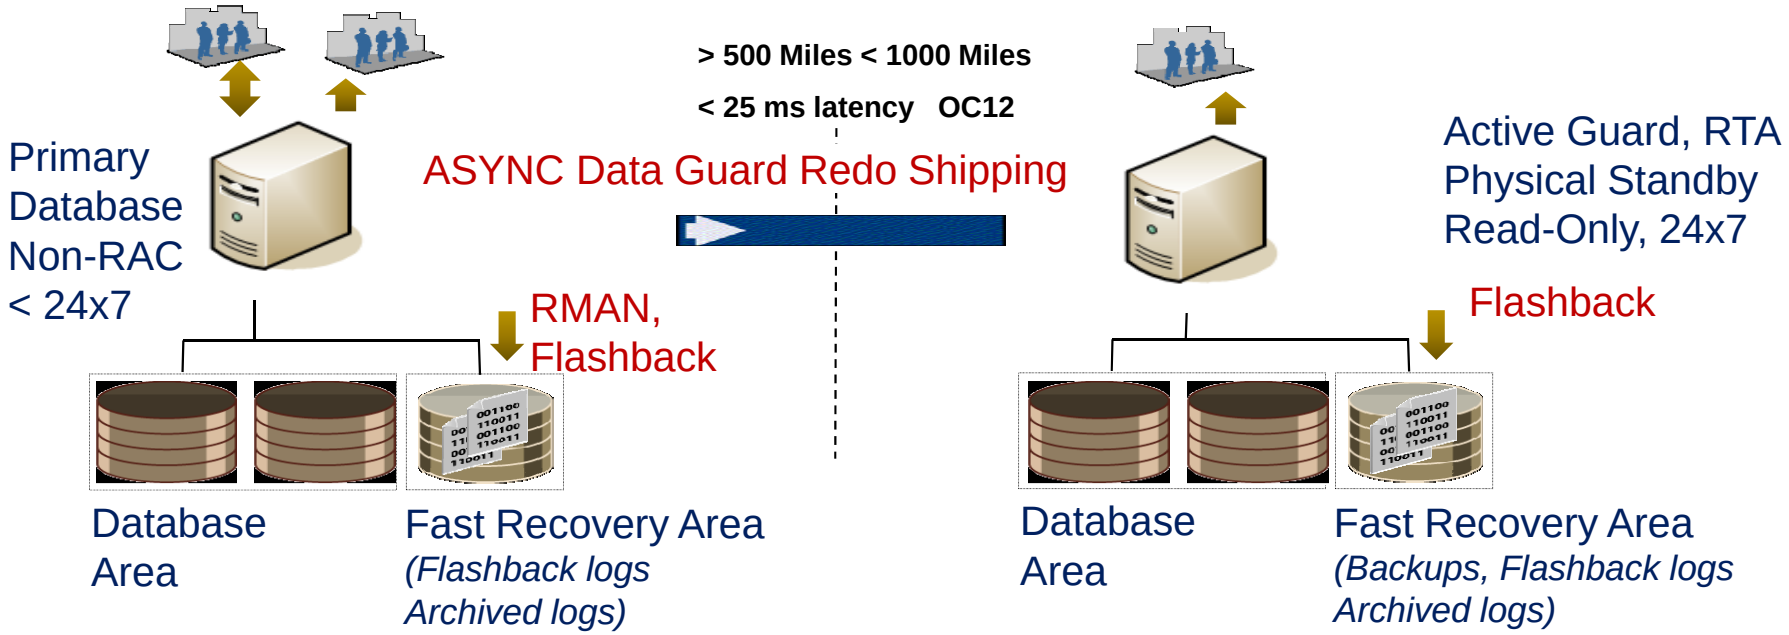

- **used for Customer data, Pricing details**
- **why it works – use of service_names and app separation r/w vs r/o**
- **intelligent network routing, 1st choice always local site**

ENTERPRISEHOLDINGS

Decision Points

- Business Owners are responsible for classifying their systems as Mission Critical or Business Critical. We have an internal document that helps them answer **many** questions leading to the appropriate classification. ASYNC does risk some data loss, we are willing to accept that risk.
- We chose not to implement Real Application Clusters because of the additional complexity. No issues for us at this time, Data Guard has completely met our expectations.

ENTERPRISEHOLDINGS

Data Guard Cool Tips

- Many technologies can provide High Availability. We have used Data Guard reliably for many years now. You should choose HA technologies that are easy to manage and make the most sense for your specific requirements.

And don't forget - leverage OEM Grid Control.

ENTERPRISEHOLDINGS

Steve Jobs Quote:

“When you first start off trying to solve a problem, the first solutions you come up with are very complex, and most people stop there. But if you keep going, and live with the problem and peel more layers of the onion off, you can often times arrive at some very elegant and simple solutions.” Newsweek

ENTERPRISE HOLDINGS

Resources

- **OTN HA Portal:**
  http://www.oracle.com/goto/availability
- **Maximum Availability Architecture (MAA):**
  http://www.oracle.com/goto/maa
- **MAA Blogs:**
  http://blogs.oracle.com/maa
- **Exadata on OTN:**
  http://www.oracle.com/technetwork/database/exadata/index.html
- **Oracle HA Customer Success Stories on OTN:**
  http://www.oracle.com/technetwork/database/features/ha-casestudies-098033.html

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Key HA Sessions, Demos, Labs by Oracle Development

Monday, 3 Oct – Moscone South *

11:00a Auto Detect, Prevent and Repair Data Corruptions, Rm 102
12:30p Future of Oracle Exadata, Rm 104
12:30p RMAN: Not Just for Backups Anymore, Rm 304
2:00p Extreme Data Management, Moscone North Hall D
5:00p Oracle High-Availability System Overview, Rm 104
5:00p GoldenGate Product Update and Strategy, Intercontinental-Sutter

Tuesday, 4 Oct – Moscone South *

10:15a Oracle Secure Backup - Best practices, Rm 304
11:45a Oracle Exadata Technical Deep Dive, Rm 104
3:30p RMAN & Data Guard: Seven Cool Tips from Oracle, Rm 304
3:30p Consolidation on Oracle Exadata, Rm 103

Wednesday, 5 Oct – Moscone South *

10:15a Oracle Active Data Guard - Lessons Learned, Rm 102
1:15p Data Guard for Planned Maintenance, Rm 102
1:15p Understanding Oracle RAC Internals, Rm 103
1:15p Clone Oracle with CloneDB and Direct NFS, Rm 270

Thursday, 6 Oct – Moscone South *

9:00a Exadata Backup and Recovery, Rm 304
10:30a Deduplication and Compression for Backups, Rm 304
12:00p Data Guard Switchover / Failover, Rm 103
3:00p Configure, Size, Monitor Fast Recovery Area, Rm 304
3:00p PeopleSoft with Active Data Guard, Moscone West 2022

Demos Moscone South DEMOGrounds

Mon 10:00a - 5:30p Tue 9:45a - 6:00p Wed 9:45a - 4:00p

Maximum Availability Architecture (MAA)	Exadata
Active Data Guard	Oracle Secure Backup
Recovery Manager & Flashback	GoldenGate
Real Application Clusters	ASM

Hands-on Labs Marriott Marquis, Salon 14 / 15

Monday, Oct 3, 5:00 pm - 6:00 pm Oracle Active Data Guard
Tuesday, Oct 4, 10:15 am - 11:15 am Oracle Active Data Guard

**All session rooms at Moscone South unless otherwise noted*
**After Oracle OpenWorld, ref. http://www.oracle.com/goto/availability*

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Q&A

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Hardware and Software

ORACLE

Engineered to Work Together

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.


ORACLE®