

ORACLE®

“This presentation is for informational purposes only and may not be incorporated into a contract or agreement.”

This document is for informational purposes. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle. This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

THE INFORMATION COMPANY

Mike Smith
Joe Meeks
Oracle

Fast-Start Failover

**Best Practices for Automatic Failover
Using Oracle Data Guard 10g Release 2**

Agenda

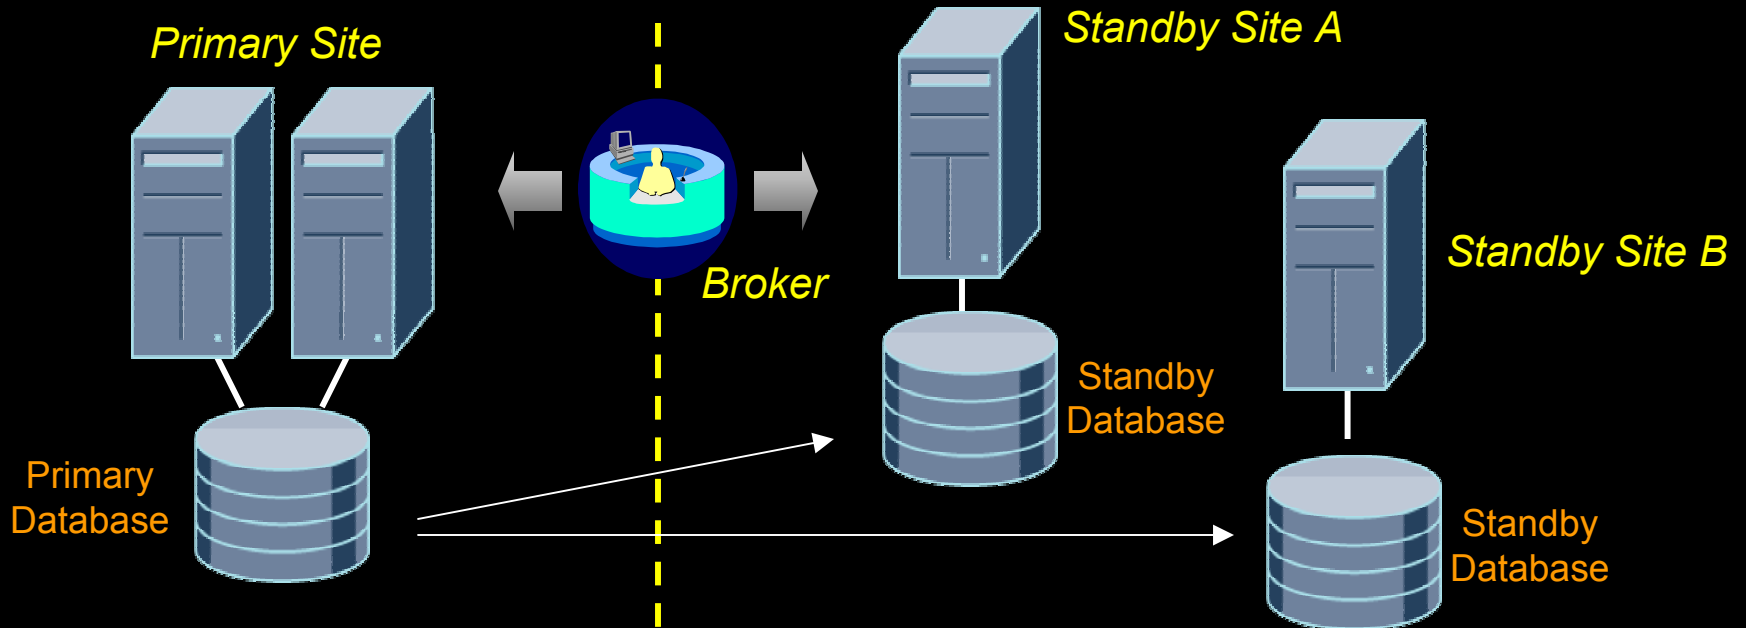
- Fast-Start Failover Overview
- Amazon.com - Addressing User Requirements
- Fast-Start Failover, Feature Details & Best Practices
 - Failover Processing Flow
 - Failover Events and Monitoring
 - Best Practices
 - Test Results
- Automatic Client Failover
- Fast, Simple & Reliable: User Experiences
- Questions?

Fast-Start Failover – a Feature Of Oracle Data Guard

Data Guard:

- Is Oracle's disaster recovery solution for Oracle data
- Is a feature of Oracle Database Enterprise Edition (EE)
- Automates the creation and maintenance of one or more synchronized copies (standby) of the production (or primary) database
- Provides comprehensive role management services
 - Role transitions – standby to primary and back to standby
 - For planned and unplanned outages

Data Guard Configuration




- Managed as a single configuration
- Primary and standby databases can be Real Application Clusters or single-instance Oracle
- Up to nine standby databases supported in a single configuration

Failover Requirements

- Faster is better
 - Downtime is bad
 - If manual intervention is required, the time it takes to notify administrative staff can be lengthy
- Reliability is a must-have
 - Correct procedure for failover must be followed to minimize data loss (recovery point objective)
- Simplicity is preferred
 - Determining if failure condition warrants a failover adds time & complexity to the failover process
 - Manually rebuilding the old primary following failover consumes time and resources and is error prone

Fast-Start Failover Architecture



- Primary Database
- Target Standby Database
- Observer Process

Characteristics of Data Guard Fast-Start Failover

- Fast
 - Site failover time measured in seconds, not minutes
 - Failover is automatic, no manual intervention
- Reliable
 - Eliminates human error
 - Zero data loss failover
- Simple
 - Automatically determines if failover criteria met
 - Original primary database is automatically reinstated as a new standby database following failover

Fast-Start Failover

Addressing User Requirements

Amazon.com

Amazon.com – Current HA/DR Status

- Major HA & DR requirements:
 - Guaranteed no data loss
 - Failovers should be as quick as possible
 - No impact on primary database's availability due to standby failures
- Data Guard prior to Oracle Database 10g Release 2 had no built-in fully automated support for these requirements
- Amazon built scripts and programs to automatically maintain custom standby databases per requirements
 - Automatically copy, apply and archive redo, ctrl files, etc
 - Actual failover is still human controlled
- Switchover / failover process involves some manual steps

Amazon.com & Data Guard 10g Release 2

- Fast-Start Failover meets Amazon's requirements
 - Now possible to do very fast failovers without any data loss
 - Fast-start Failover will not commence if target standby not synchronized with the primary
 - Synchronization status shown through `fs_failover_status` column in `v$database`
 - Since Fast-Start Failover is based on Maximum Availability, primary database is not impacted by standby failures
 - Old primary automatically reinstated as new target standby

Amazon.com – Benefits of Fast-Start Failover

- Amazon's initial tests with Fast-Start Failover:
 - ✓ Failover time: reduced from minutes to seconds¹
 - ✓ No data loss
 - ✓ No manual intervention

1. Average failover time: 25 seconds, based on one primary/standby system with a network round trip time of 0.5 ms

Amazon.com

The capability of fast, guaranteed zero-data-loss failover with Fast-Start Failover in Oracle Data Guard takes the availability of an Oracle database platform to new levels. Our initial tests running Oracle Database 10g Release 2, show that Fast-Start Failover offers a magnitude of improvement in availability.

Rajesh Sheth

Manager, Database Engineering

Amazon.com

Fast-Start Failover

Feature Details & Best Practices

Mike Smith

Oracle

Fast-Start Failover Requirements

- Primary and Target Standby are managed by the Data Guard Broker
- Primary database must be in Maximum Availability mode – synchronous redo shipping
- Primary and target standby must have Flashback Database enabled
- Observer host must have DGMGRL utility installed and must have Oracle Net connectivity to both the primary and target standby

Fast-Start Failover Configuration Using Enterprise Manager GUI

ORACLE Enterprise Manager 10g
Grid Control

Setup Preferences Help Logout

Home Targets Deployments Alerts Policies Jobs Reports

Databases | Hosts | Web Applications | Services | Systems | Groups | All Targets

Host: north.foo.com > Database Instance: West_Sales > Data Guard > Logged in As SYS

Fast-Start Failover: Configure

TIP The configuration will be automatically upgraded to Maximum Availability mode and the primary database will be restarted.

Target Database Selection

Select a standby database to be the fast-start failover target. The log transport mode for the selected database will be set to SYNC (if not currently set to SYNC).

Select Name	Role	Log Transport Mode
<input checked="" type="radio"/> DR_Sales	Physical Standby	ASync
<input type="radio"/> West_Sales	Logical Standby	ASync

Properties

Fast-start failover requires a Data Guard observer process. For highest availability, Oracle recommends that the observer be on a separate host from the primary and standby databases.

TIP Fast-start failover can be enabled without starting the observer, but cannot actually take place until the observer is running.

Observer Location **Not Set**

There is currently no observer for this configuration, nor has one been specified.

Fast-Start Failover Threshold

Amount of time the primary database must be out of contact with the observer and the standby database before a fast-start failover is initiated.

Fast-Start Failover Configuration Using DGMGRL Command Line

• Configure

- `FastStartFailoverTarget` - DB_UNIQUE_NAME of the database that is intended to be the new primary
 - `DGMGRL> EDIT DATABASE 'North_Sales' SET PROPERTY FastStartFailoverTarget = 'DR_Sales';`
- `FastStartFailoverThreshold` - Number of seconds Observer attempts to reconnect to the primary database before initiating fast-start failover
 - `DGMGRL> EDIT CONFIGURATION SET PROPERTY FastStartFailoverThreshold = 45;`


• Start

- `DGMGRL> ENABLE FAST_START FAILOVER;`
- `DGMGRL> START OBSERVER;`

• Reinstate After Failover

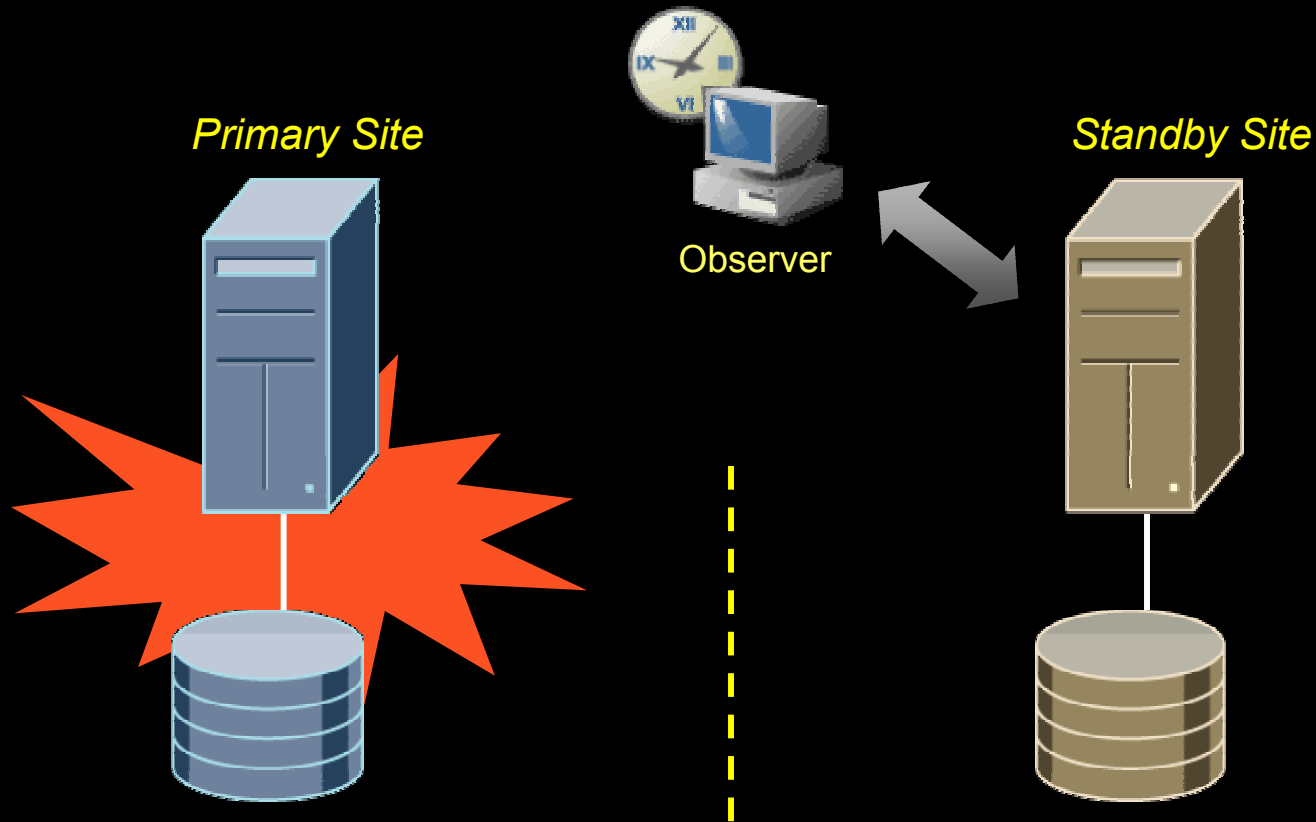
- Auto reinstatement of old primary to be a new standby – or manual:
 - `DGMGRL> REINSTATE DATABASE database_name;`

Fast-Start Failover Processing




- *Data Guard in steady state – transmitting redo*
- *Observer monitoring state of the configuration*

Fast-Start Failover Processing




Disaster strikes the primary – connections lost

Fast-Start Failover Processing




- *Observer <=> primary connection times out (timeout threshold configurable)*
- *Observer asks target standby if it is ready to fail over*
- *Observer begins Fast-Start Failover*

Fast-Start Failover Processing



Target standby automatically becomes new primary

Fast-Start Failover Processing



- *After old primary is repaired, Observer re-establishes connection*
- *Observer automatically reinstates old primary to be a new standby*
- *Redo transmission starts from new primary to new standby*

When is a Fast-Start Failover Triggered?

- Primary Site Failure
- Primary Database Conditions:
 - Instance Failure
 - Last surviving instance if RAC
 - Shutdown abort of the last available instance
 - Datafiles taken offline due to I/O errors
 - Threshold ignored when performing a failover due to offline datafiles

When is a Fast-Start Failover Triggered?

- Network Related Conditions:
 - Failover occurs only if links between primary and observer as well as primary and target standby are down
 - Requires a connection between Observer and standby to enable the Observer to confirm that the configuration is in a synchronized state
- By ensuring that at least two fast-start failover partners are present, conditions such as split-brain scenarios are avoided

Fast-Start Failover Monitoring

ORACLE Enterprise Manager 10g
Grid Control

Setup Preferences Help Logout

Home Targets Deployments Alerts Policies Jobs Reports

Databases | Hosts | Web Applications | Services | Systems | Groups | All Targets

Host: north.foo.com > Database Instance: North Sales > Logged in As SYS

Information
The fast-start failover mode has been successfully changed.

Data Guard

Page Refreshed June 8, 2005 2:44:45 PM EDT View Data Real Time: Manual Refresh

Overview

Data Guard Status **✓ Normal**
Protection Mode [Maximum Availability](#)
Fast-Start Failover [Enabled to DR_Sales](#)
Observer Location **north**

Primary Database

Name [North_Sales](#)
Host [north_](#)
Data Guard Status **✓ Normal**
Current Log [174](#)
Properties [Edit](#)

Standby Progress Summary

The transport lag is the time difference between the primary last update and the standby last received redo. The apply lag is the time difference between the primary last update and the standby last applied redo.

Database	Transport Lag (seconds)	Apply Lag (seconds)
DR_Sales	0	0
West_Sales	3	35

Standby Databases

[Add Standby Database](#)

Select	Name	Host	Data Guard Status	Role	Last Received Log	Last Applied Log	Estimated Failover Time
<input checked="" type="radio"/>	DR_Sales	north_	✓ Normal	Physical Standby	173	173	0 seconds
<input type="radio"/>	West_Sales	north_	✓ Normal	Logical Standby	173	173	

Performance

[Performance Overview](#)
[Log File Details](#)

Additional Administration

[Verify](#)
[Remove Data Guard Configuration](#)

Fast-Start Failover Monitoring

- Monitor current state of configuration via `FS_FAILOVER_STATUS` column of `V$DATABASE`
 - `SYNCHRONIZED` – Primary and Standby are in sync
 - `UNSYNCHRONIZED` – Standby does not have all of the primary database redo
- Monitor the Observer via the `FS_FAILOVER_OBSERVER_PRESENT` column of the `V$DATABASE` view

Reinstatement after a Fast-Start Failover

- Any attempt to start old primary will stop at the mount state thus preventing split brain
- Once Observer sees the old primary is at the mount state, reinstatement is begun
- The old primary is automatically reinstated as the new standby using flashback database
- New standby is automatically resynchronized with new primary by Data Guard
- Once resynchronized, a switchover can occur if desired – returning all systems to their original roles

Best Practices – Primary Database

- Maximum Availability Protection Mode
 - Redo Transport: LGWR SYNC AFFIRM
- Synchronous Redo Shipping . . . but
 - Primary is not affected by network or standby outages
 - Set `net_timeout` parameter to override TCP timeout
 - Default value is 180 seconds
- Use Flashback Database
 - For Fast-Start Failover, set `DB_FLASHBACK_RETENTION_TARGET = 10 minutes`
 - When Flashback Database is also used to protect against user error & corruption, increase flashback retention to a period of time sufficient to achieve these goals


Best Practices – Network Transport

- Tune OS & network parameters
 - Set SDU=32K
 - Tune network parameters that affect network buffer sizes and queue lengths
 - Ensure sufficient network bandwidth for maximum database redo rate + other activities

Refer to Primary Site and Network Configuration Best Practices

http://www.oracle.com/technology/deploy/availability/pdf/MAA_DG_NetBestPrac.pdf

Impact of Network Tuning



Test Results - Oracle Database 10g Release 1 & 2

Best Practices – Target Standby

- Use Standby Redo Logs
- Use Real-time Apply
- Use Flashback Database
 - Set `DB_FLASHBACK_RETENTION_TARGET` to the same value used on the primary database
- Use MAA best practices to optimize Apply Performance
 - Redo Apply (physical standby):
[Data Guard Redo Apply and Media Recovery](#)
 - SQL Apply (logical standby):
[Oracle Database 10g Data Guard SQL Apply](#)

Best Practices – Observer

- Install in a separate location from Primary & Standby data centers
 - Do not locate the Observer at or near the primary site
 - Proximity to the middle tier is preferred
 - Assumes middle tier is not at or near primary site
 - Oracle Client Administrator install is all that is required for Observer install
 - If using Enterprise Manager, also install the Enterprise Manager Agent on the observer host

Best Practices – Setting FastStartFailoverThreshold

- Failover occurs when observer and target standby lose contact with primary for a specified period of time
 - The observer continuously attempts to reconnect during the threshold time period
- Recommended threshold settings:
 - Single Instance primary with low latency reliable network = 10 – 15 seconds
 - Single Instance primary with high latency network over WAN = 30 – 45 seconds
 - RAC primary = (misscount + reconfiguration time) + 20-40 seconds
 - Misscount value is obtained using `crsctl get CSS misscount`. Oracle Database 10g default is 60 seconds.
 - Reconfiguration time determined by user experience

Best Practices – Multiple Standbys

- Ensure data protection at all times by maintaining a 2nd Data Guard standby at a remote location
 - When regulatory & business requirements mandate that data be protected at all times
- At failover time, the remote standby automatically becomes a standby for the new primary
 - New primary must have begun as a physical standby
- Configure the remote standby for Maximum Performance
 - Eliminates overhead of WAN network latency
 - Recommended redo transport is **LGWR ASYNC**

Best Practices – HA & DR

Use RAC & Data Guard Together

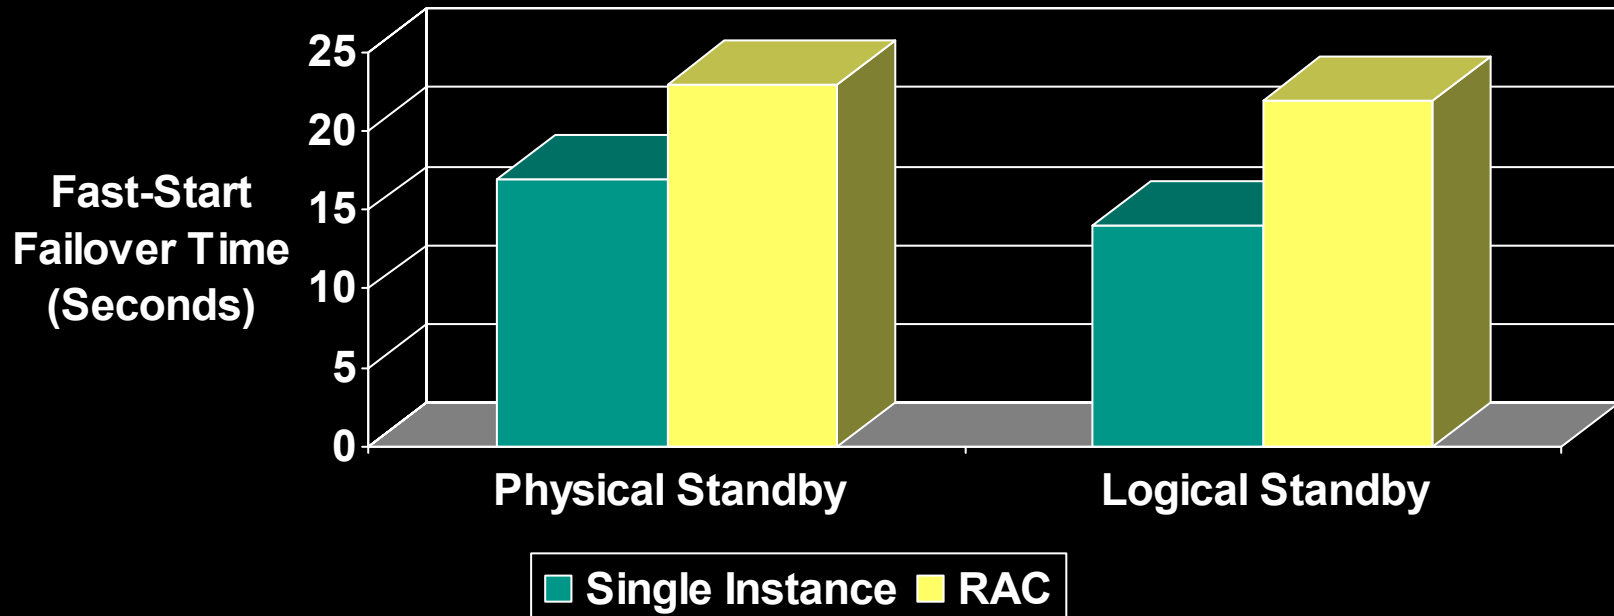
- RAC - automatic failover across nodes in a cluster
- Data Guard - automatic failover between sites
- The best possible combination of HA & DR
 - Scalable, flexible, secure
 - Foundation for Maximum Availability Architecture (MAA)

Fast-Start Failover – How Fast?

Oracle MAA Test Environment

- All three hosts (Observer, Primary & Standby) running Red Hat Linux 3.0
- Primary Database and target standby database approximately 100 GB in size with 1.7 GB of SGA
- Gigabit Network between all three parties
- Primary database producing approximately 3 Megabytes of redo per second
- Network Latency introduced between all three hosts
 - Tests ranged between 0 & 100ms RTT

Fast-Start Failover Test Results



Failover timings were unaffected by increased primary database workload or increased network latency

Note: RAC result requires Oracle Database 10.2.0.2.0

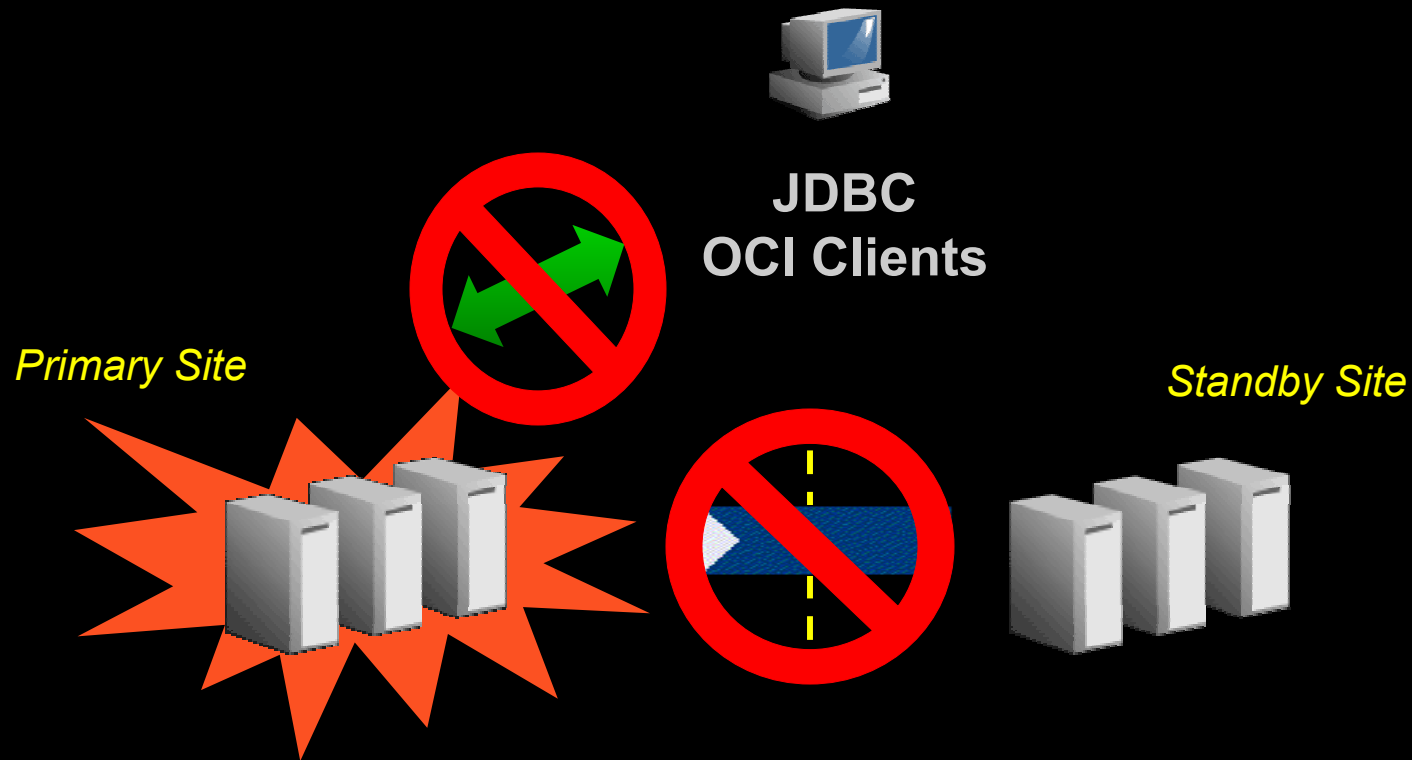
Client Failover

Oracle Data Guard 10g Release 1

vs

Oracle Data Guard 10g Release 2

Oracle Database 10g Release 1 Client Failover




1 Primary site failure - FAN ONS (JDBC) and OCI clients wait for TCP timeout

Oracle Database 10g Release 1 Client Failover




JDBC/OCI Clients

Failed Primary Site

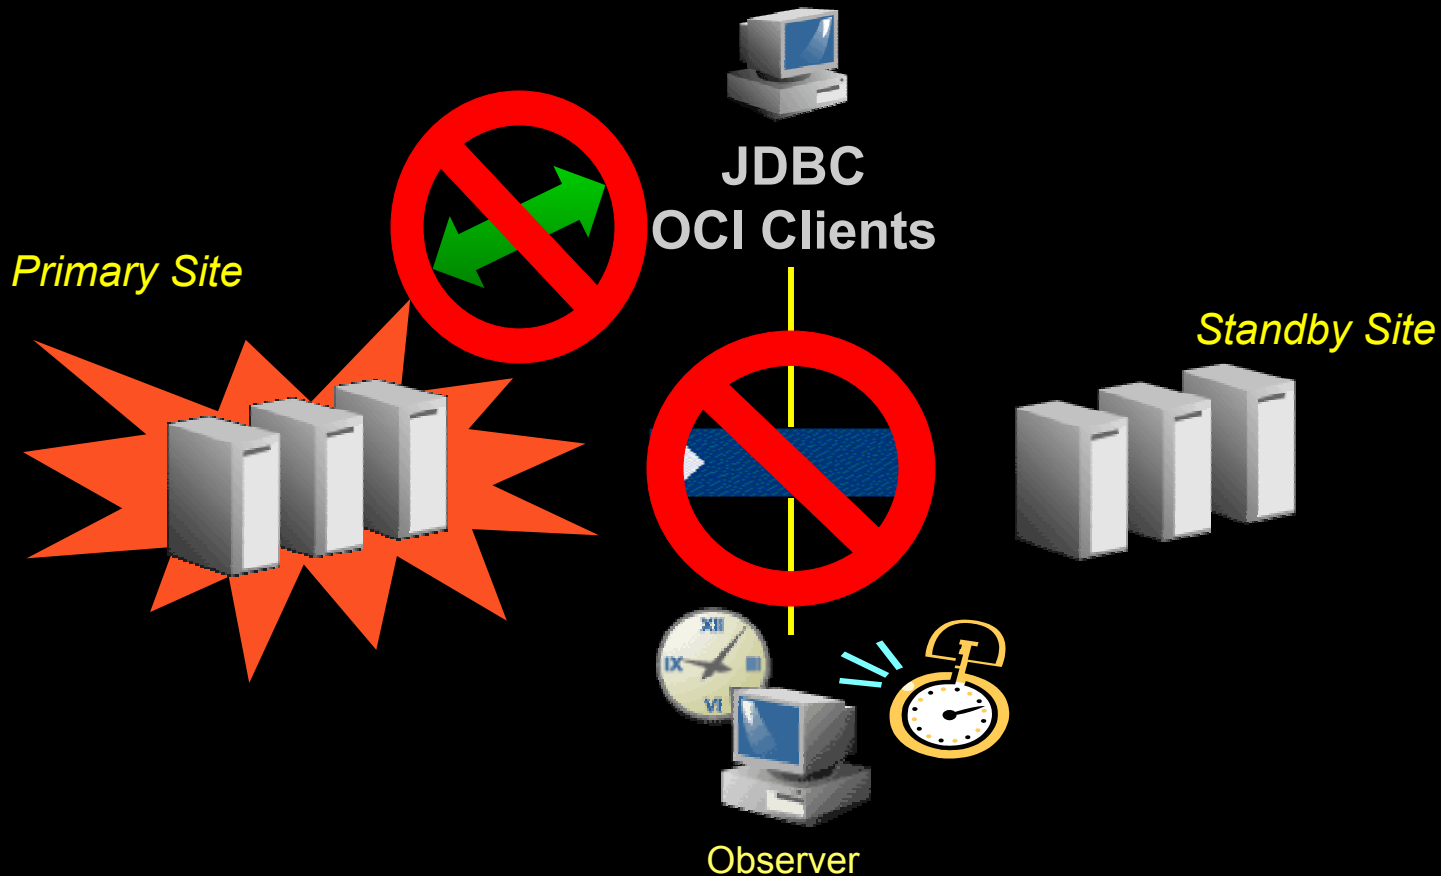


New Primary Site



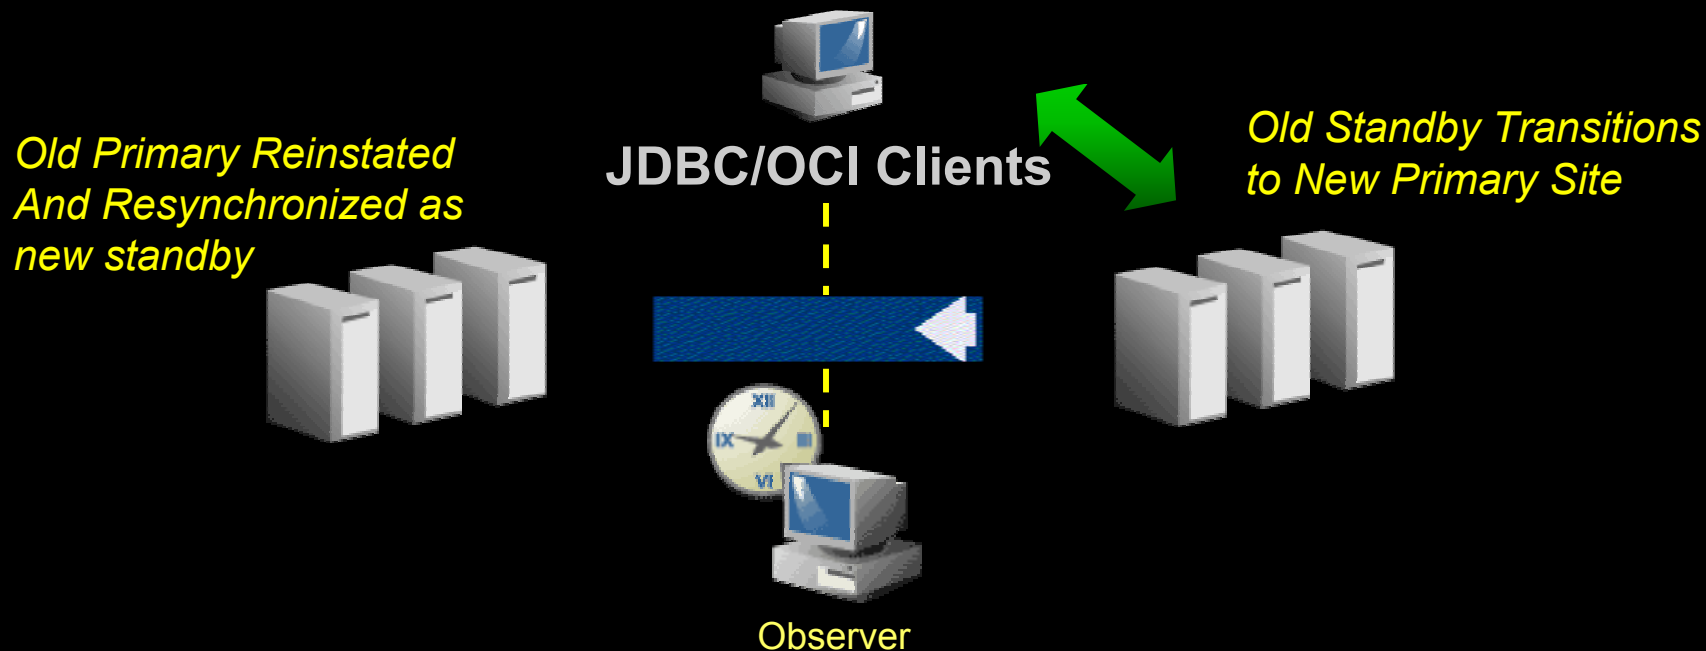
- 2. Data Guard failover is executed manually, standby database transitions to primary role*
- 3. Clients are redirected manually*
- 4. Old primary database must be flashed back manually in order to reinstate as standby*

Oracle Database 10g Release 2 Client Failover



1 Observer detects failure, clients stall until Fast-Start Failover threshold is exceeded

Oracle Database 10g Release 2 Improved Client Failover



2. *Observer automatically executes database failover once threshold has been exceeded*
3. *DB_ROLE_CHANGE trigger fires: enables primary service, updates Oracle Net alias to point to new primary host, restarts JDBC mid-tier clients, calls any other application or pre-failover steps. DB_DOWN event is sent to FAN OCI clients*
4. *Both FAN ONS (JDBC) and OCI clients drop connections and re-attach to the new primary*
5. *Upon restart, the old primary database is automatically reinstated by Fast-Start Failover*

Oracle Database 10g Release 2 Client Redirection

- New DB_ROLE_CHANGE system event fires when the target standby is activated as the new primary
- A Trigger written around DB_ROLE_CHANGE event can automatically:
 - Enable primary service name
 - Modify LDAP or other naming methods
 - Restart JDBC mid-tier clients
 - Start user applications
 - All details w/examples are described in the session paper and a future MAA Best Practices paper to be published on OTN

Oracle Database 10g Release 2

Client Redirection

- New DB_DOWN event is posted after the new primary is open
 - Event notifies FAN OCI clients that the old primary is down
 - Clients reconnect to the new primary/service
 - Done via AQ notifications

Fast-Start Failover

*How Fast, Reliable, and Simple?
User Experiences*

Fannie Mae

Airbus Deutschland GmbH

Thomson Legal & Regulatory

How Fast - Fannie Mae

Fast-start Failover takes the DBA off the critical path. Database failover is automatic. Data Guard can now address recovery time objectives measured in seconds.

Ranjit Singh Veen

Manager, Enterprise Systems Management

Fannie Mae

How Reliable - Airbus

Failover executed automatically in less than a minute. This was much faster than a cold failover using third party cluster technology.

With Data Guard, Airbus can achieve continuous data protection and high levels of availability using a standard feature of the Oracle Database. .

Werner Kawollek

Application Management Operations

Airbus Deutschland GmbH

How Simple - Thomson Legal & Regulatory

Fast-start Failover testing has shown great potential. The original primary database can be reinstated as a new standby in less than 5 minutes once the initial failure has been corrected.

Thomson Legal & Regulatory



QUESTIONS
ANSWERS

High Availability Demos/Sessions From Oracle Development

Demogrounds - Monday, Sep 19 – Thursday, Sep 22

- | | |
|---|---|
| <ul style="list-style-type: none">• <i>Oracle Data Guard</i>• <i>ILM and Storage</i> | <ul style="list-style-type: none">• <i>Oracle Secure Backup</i>• <i>RMAN, Flashback, and Online Redefinition</i> |
|---|---|

Sessions - Monday, Sep 19

- 1:30-2:30 pm, Room 303 - *Optimizing Linux I/O*
- 3:00-4:00 pm, Room 104 - *The Future of Database Information Technology*
- 4:30-5:30 pm, Room 103 - *What They Didn't Print in the DOC - HA Best Practices by Gurus from Oracle's Maximum Availability Architecture Team*

Sessions - Tuesday, Sep 20

- 3:00-4:00 pm, Room 104 - *Logical Standby Unleashed*
- 4:30-5:30 pm, Room 104 - *Best Practices for Oracle Database 10g Backup and Recovery*

High Availability Sessions From Oracle Development

Sessions - Wednesday, Sep 21

- 11:00 am-12:00 pm, Room 104 - *Improve Your Tape Backup Results with Oracle Secure Backup*
- 3:00-4:00 pm, Room 304 - *Implementing Information Lifecycle Management (ILM) using the Oracle Database*

Sessions - Thursday, Sep 22

- 1:00-2:00 pm, Room 104 - *Minimizing Application Development Time Using Flashback: A Customer Case Study*
- 2:30-3:30 pm, Room 104 - *Best Practices To Achieve Business Continuity Using Oracle Applications and Oracle Database Technology*
- 4:00-5:00 pm, Room 104 - *Best Practices for Automatic Failover Using Oracle Data Guard 10g Release 2*

ORACLE®