
ORACLE®


ORACLE®

# Deploy New Features Risk Free Using Database Replay

Prabhaker Gongloor, Oracle Corporation
Paula Camporaso, Rajeev Sethi, Solyndra
Tom Robertson, Nationwide Insurance

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. 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 for Oracle's products remains at the sole discretion of Oracle.

# Agenda

- Deploying New Features - Challenges & Solution
- Database Replay Overview
- Database Replay Enhancements: Oracle Database 11*g* Release 2
- Strategy and Best Practices
- Real-world Customer Case Studies
- Conclusion

## *Please visit us at:*

- **OOW Demo Grounds Moscone West – 038/039**
- **S318966: Database and Application Testing HOL, Wed:4.45-5.45 pm**

# Deploying New Features – Challenges and Solution

## Customers want to take advantage of new features, but

- No easy way to mitigate change risk
- Impossible to test new features with real-world workloads
- No end-to-end testing solution
- Significant risk to production – instabilities, SLAs violated, fire-fighting, etc.

## Real Application Testing makes possible

- Higher quality testing
- Rapid technology adoption

## As a result, businesses can adopt new features at

- Lower cost
- Lower risk

# Real Application Testing Features

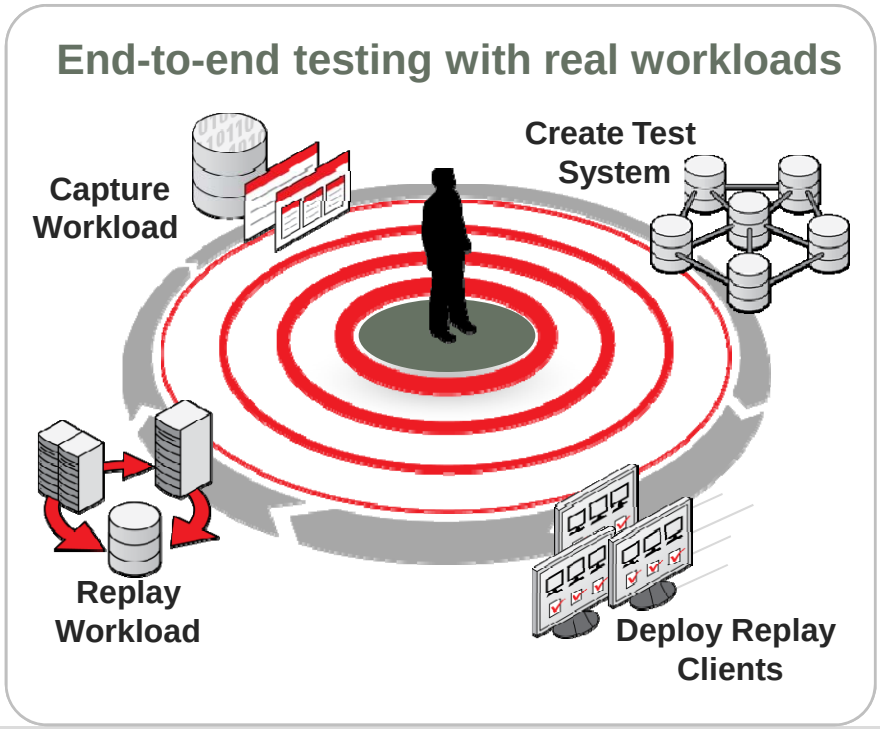


- **SQL Performance Analyzer**
  - SQL unit testing for response time
  - Identify and tune regressed SQL
- **Database Replay**
  - Load, performance testing for throughput
  - Remediate application concurrency problems

# Agenda

- Deploying New Features - Challenges & Solution
- Database Replay Overview
- Database Replay Enhancements: Oracle Database 11*g* Release 2
- Strategy and Best Practices
- Real-world Customer Case Studies
- Conclusion

# Oracle Real Application Testing

## Database Replay

- Database load and performance testing with real production workloads
  - Production workload characteristics such as timing, transaction dependency, think time, etc., fully maintained
- Test and measure transaction throughput improvements
- Identify application scalability and concurrency problems with new features
- Remediate issues pre-production for risk-free change
- Supports migrations from Oracle 9*i*R2 and 10*g*R2

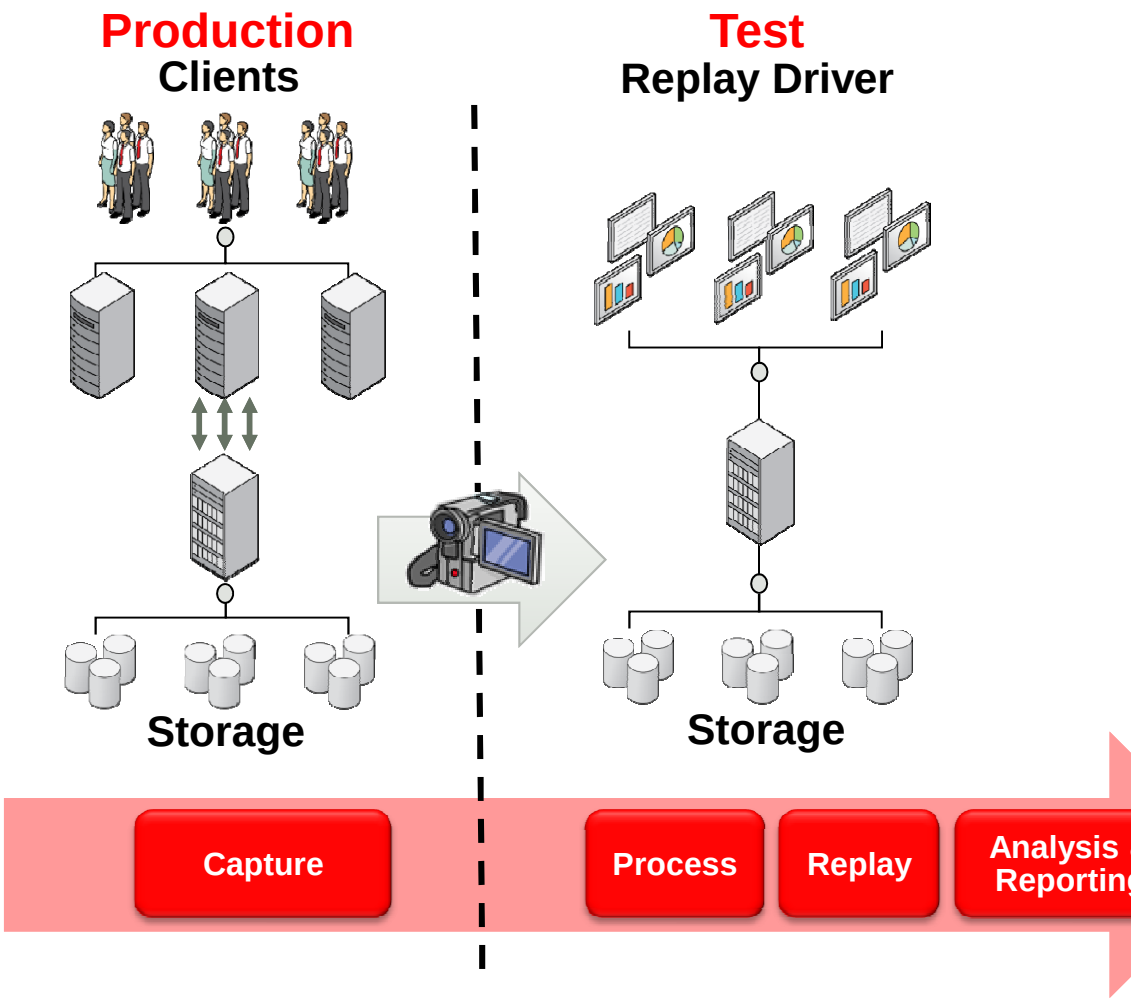


- ***MOS 560977.1: Real Application Testing for Earlier Releases**

# Workload Replay Architecture

- Replay captured workload
  - Replayed operations see the same data and perform the same work
  - Preserve timing and concurrency characteristics
  - Same number of user connections
- Replay Client
  - Multithreaded OCI Client
  - Drives multiple captured processes
  - Scalable Architecture
  - Interprets capture into sequence of OCI calls
  - Functional replay

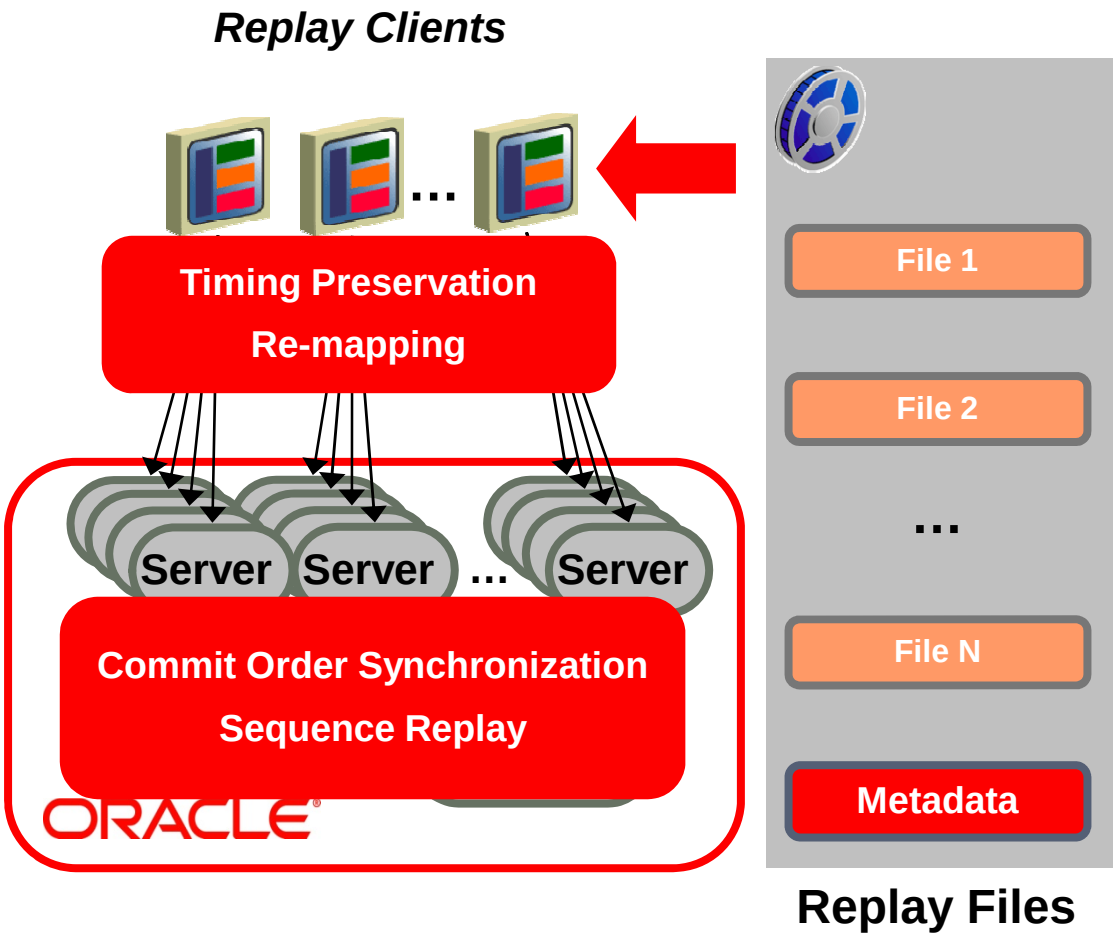

# Database Replay - Supported Changes

**Changes Unsupported**

**Changes Supported**

- Database Upgrades, Patches
- Schema, Parameters
- RAC nodes, Interconnect
- OS Platforms, OS Upgrades
- CPU, Memory
- Storage
- Etc.

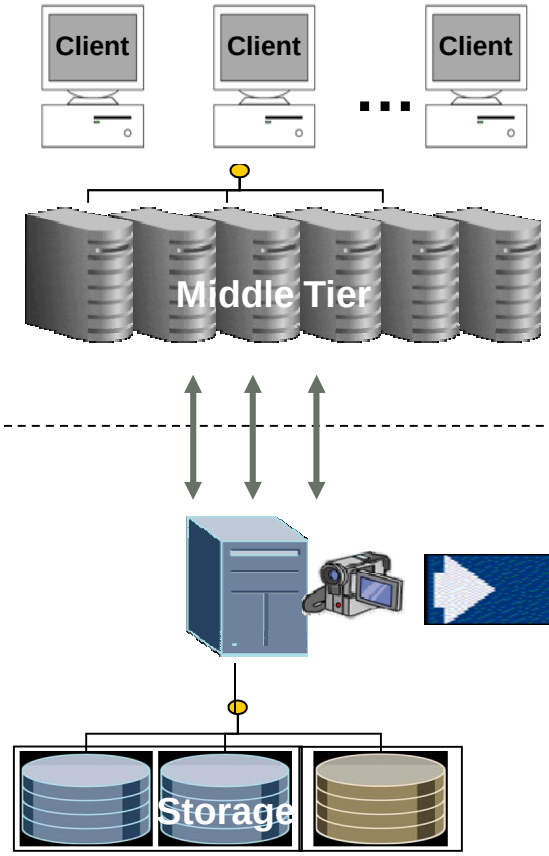


**Recording of External Client Requests**

# Agenda

- Deploying New Features - Challenges & Solution
- Database Replay Overview
- Database Replay Enhancements: Oracle Database 11*g* Release 2
- Strategy and Best Practices
- Real-world Customer Case Studies
- Conclusion

# Database Replay Enhancements
## New in Oracle Database 11*g* Release 2

### Earlier restrictions removed with support for

- Shared Server configuration
- Streams “apply” workload

### Replay filter support to target sub-set workload

- Similar to existing capture filters – Include/Exclude
- Service, Program, Action, Module, etc.
- API support only in current release

### Integration of SPA and Database Replay

- Allows SQL-centric analysis when using Database Replay
- Simultaneously captures SQL workload into two different STS during workload capture and replay
- SPA Report built from the two STS captured helps understand workload drift
- Uses STS Compare functionality to highlight new, missing, top SQL, changes in execution plans, #SQL executions etc.

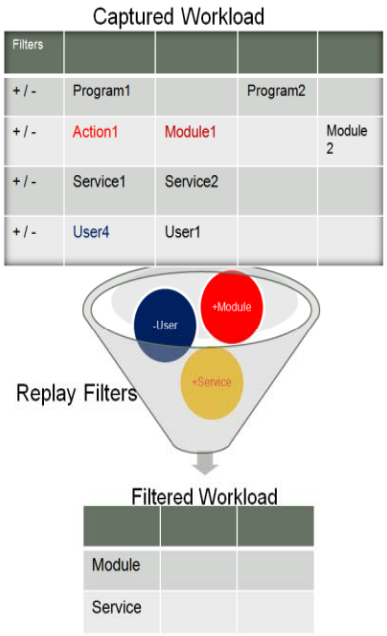

# Database Replay Enhancements

## New in Oracle Database 11*g* Release 2

### Workload Analyzer

- New tool for assessing quality of workload capture and its replayability
- Identifies potential problems and recommends appropriate remediation
- Provides insight into workload capture – quantifies percentage of captured DB Time that is unreplayable
- Rule-based analysis executed as part of pre-processing
- EM and API support (DB release 11.2.0.2 and above)
- OTN download for earlier releases and new rules
- Workload Analyzer recommendations example follows…

# Workload Analyzer: Recommendations

- Run Workload Analyzer on captured workload and follow recommendations to improve replay quality, for e.g.,

| Finding | Maximum Workload Impact |
|---|---|
| In-flight sessions | 36 % |
| SYSDATE and other time-dependent functions | 15 % |
| Missing AWR export | Unknown |

## Findings and Recommendations

### In-flight sessions

**Maximum Workload Impact: 36 % of DB Time**

**Rationale**

A significant part of your captured files have been captured in-flight.
This means that the captured session already existed before capture started, and could have already modified its state (e.g. through 'alter session') or be inside a transaction at this time.
This can cause some divergence (since only a part of some transactions was captured) or even bigger problems if some key session parameters are not set during replay.

**Action**

The best practice is to avoid in-flight session by restarting the database before capturing.
In case an in-flight session needs some missing session parameters to perform during replay, consider using a login trigger to set those parameters correctly.

# Database Replay Enhancements

## New in Oracle Database 11*g* Release 2

### Replay Compare Period Report

- Provides holistic view of the experiment – covers functional and performance aspects of testing
- "Replay Divergence Summary" categorization indicates if further analysis is necessary: LOW/MED/HIGH
- Two reports are available
  - Capture Vs Replay, Replay Vs Replay
- Identifies interference from other workloads, e.g., maintenance windows or other non-replayed workload
- Automatically runs ADDM
- Reports provide more accurate performance analysis
  - Uses enhanced ASH infrastructure for capture/replay sessions

**Improved in Oracle Database 11g Release 2**

# Agenda

- Deploying New Features - Challenges & Solution
- Database Replay Overview
- Database Replay Enhancements: Oracle Database 11*g* Release 2
- Strategy and Best Practices: Deploying New Features
- Real-world Customer Case Studies
- Conclusion

# Testing New Database Features

***Scenario***

- How can I successfully upgrade to Oracle Database 11*g* and deploy features such as Advanced Compression, RAC, TDE, etc.

***Goal***

- Safely deploy system changes in production without negative impact

# Testing New Database Features

## The Right Approach

### Step 1: Upgrade to Oracle Database 11*g*

- Test the impact of 11*g* Upgrade on the peak workload captured on production system & make sure no negative effects due to the upgrade

### Step 2: Introduce new features

- Then introduce one feature at a time on the workload and test the workload impact - RAC, TDE, Advanced Compression…

*We'll walk through the mentioned scenario using Siebel workload covering:*

- *Recommended testing strategy*
- *Best practices*
- *Replay analysis*

# Siebel Workload Description

- Siebel PSPP workload used for testing DB upgrade scenario
  - Used internally for upgrade certification and new feature uptake
- Siebel 8.0, 1300 users: 700 financial call center, 600 financial partner manager
- Financial call center scenario:
  - Creates new contact
  - Creates new opportunity for the contact
  - Add products to the opportunity
  - Creates quotes
  - Converts quotes to order
- Financial partner manager scenario
  - Creates a new service request
  - Assigns the service request

# Real Application Testing: Recommended Method

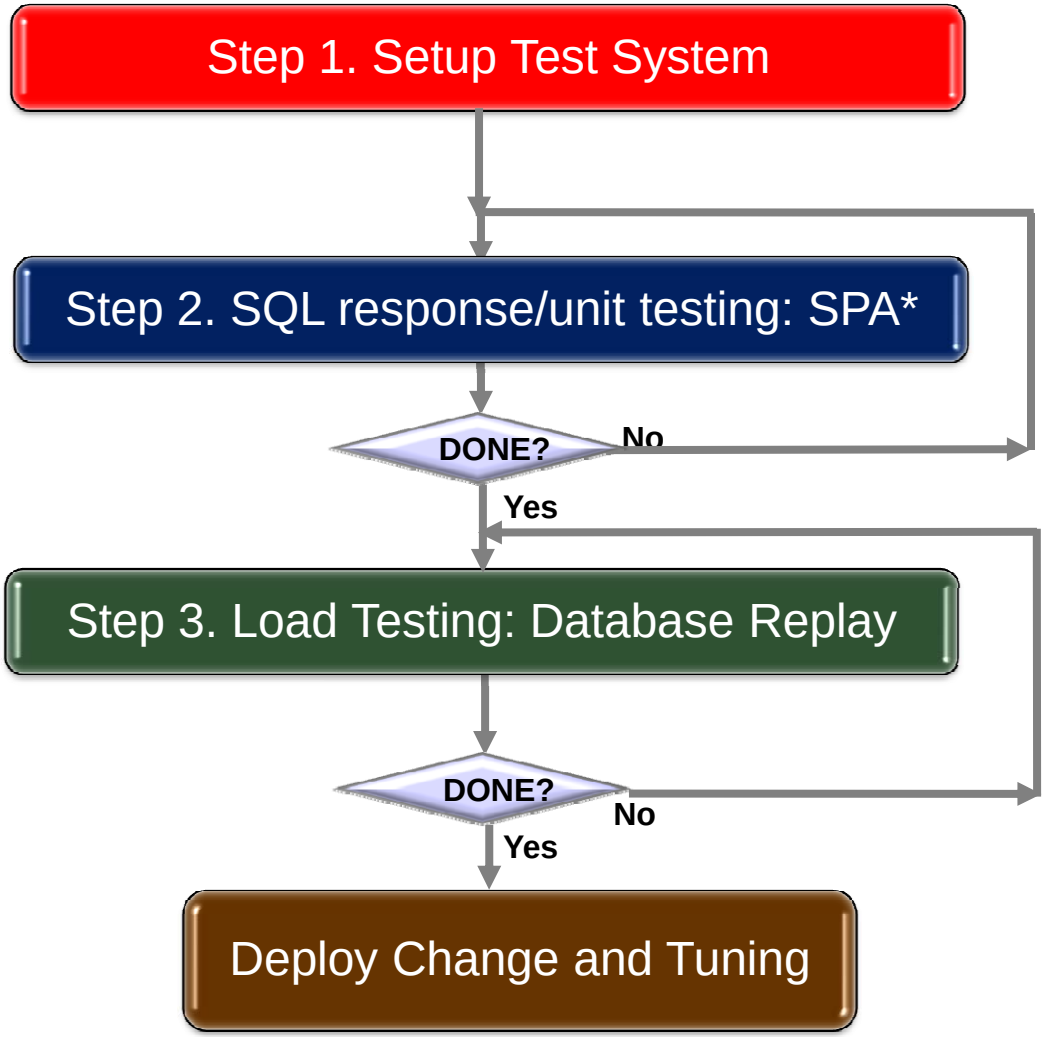

# Real Application Testing: Recommended Method

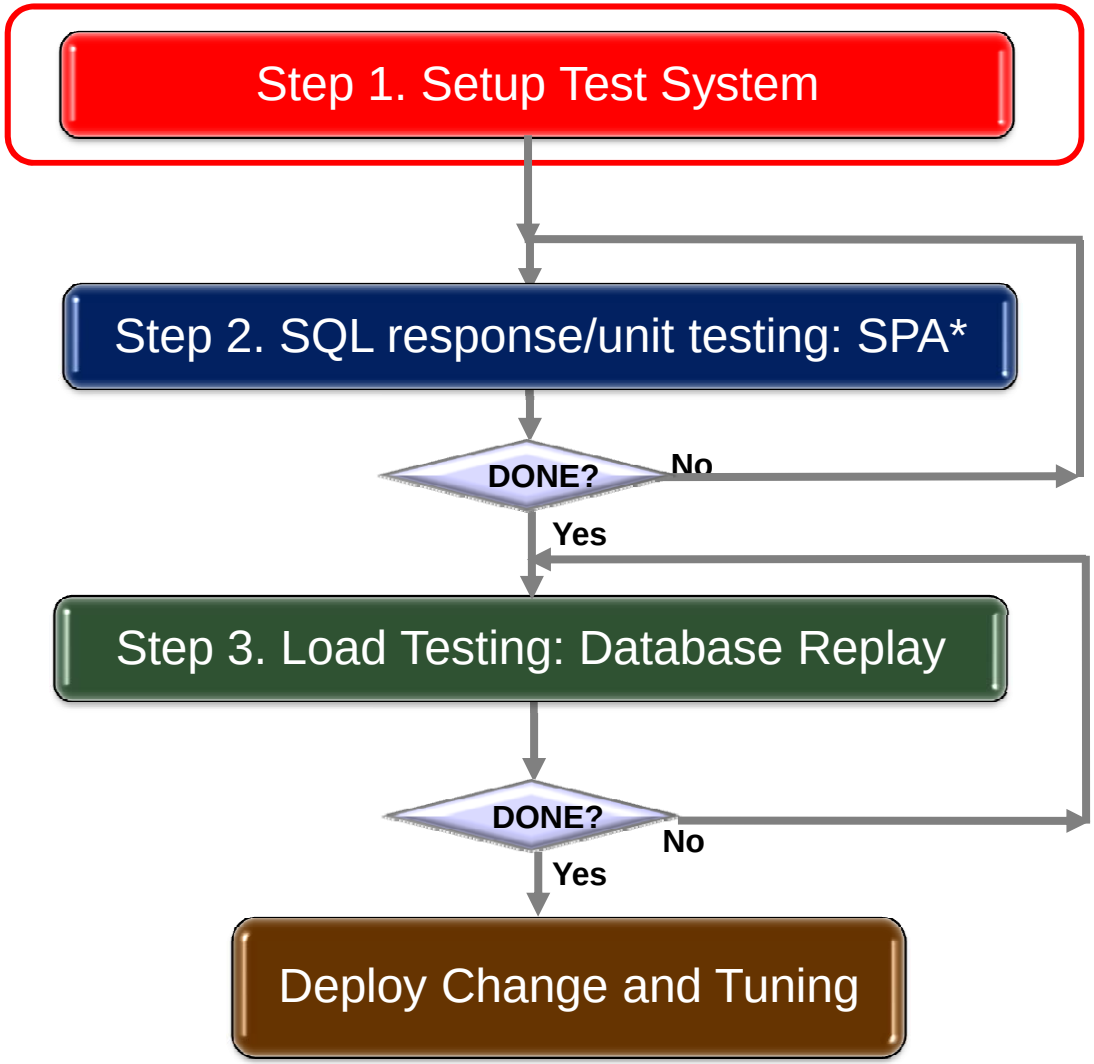

# Step 1: Setup Test System: Best Practices

**Apply recommended patches and use latest software on test and production**

- MOS Note: 560977.1

**Test system should be as close to production as possible**

- Similar HW/OS where possible, unless this is being tested
- Full dataset – should be close or same as production data to avoid divergence

**Validate no missing schema objects (indexes, views, etc.) on test system**

- Use Oracle  Enterprise Manager 11*g* Change and Configuration Management Packs to understand drift between test and production

# Step 1: Setup Test System: Best Practices

## Use Database Flashback, Flash Recovery Area, and Guaranteed Restore Points

- Helps reset database to point of capture

## Use Oracle Enterprise Manager Grid Control Releases 10.2.0.5 or 11.1

- Supports end-to-end workflow including test system creation and cloning
- Best practice based workflows

## Disable maintenance windows and background jobs

- Avoid workload interference, background jobs may already exist on test system

# Real Application Testing: General Recommended Strategy

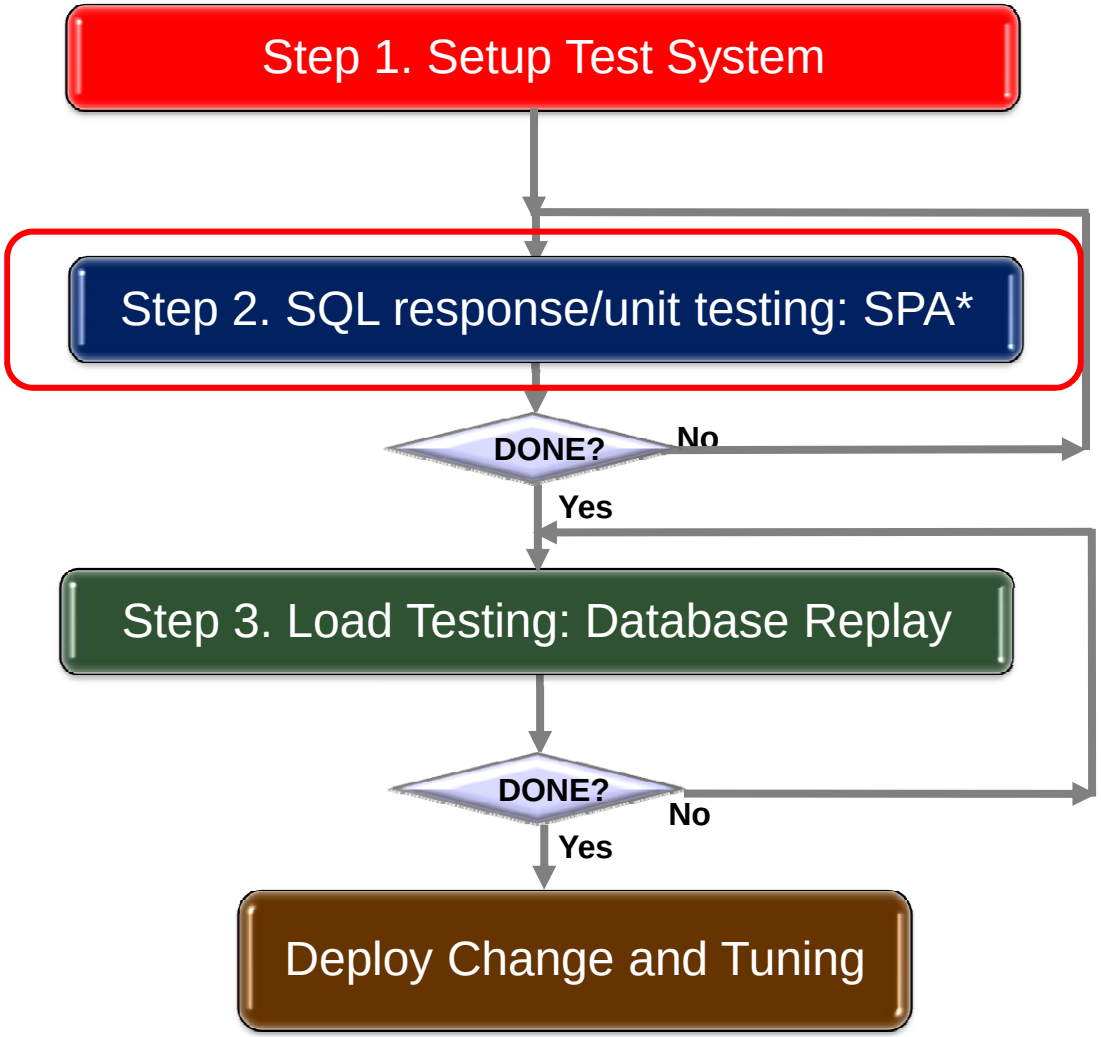


*Session S317300 : Avoiding SQL Performance Regressions – New Techniques for Solving an Old Problem

# Step 1: SQL Performance Analyzer (SPA) Testing

## Always use SPA before Database Replay to help reduce testing cycles

- Most changes such as patch-sets, upgrades may result in plan changes
- SPA is the best tool to perform SQL-centric analysis
- SPA trials complete quickly relative to DB Replay runs
- SPA trials can be repeated without restoring database

## Use SPA to identify regressed SQL and remediate them

- Gives you a mechanism to revert back to old plans if they are better

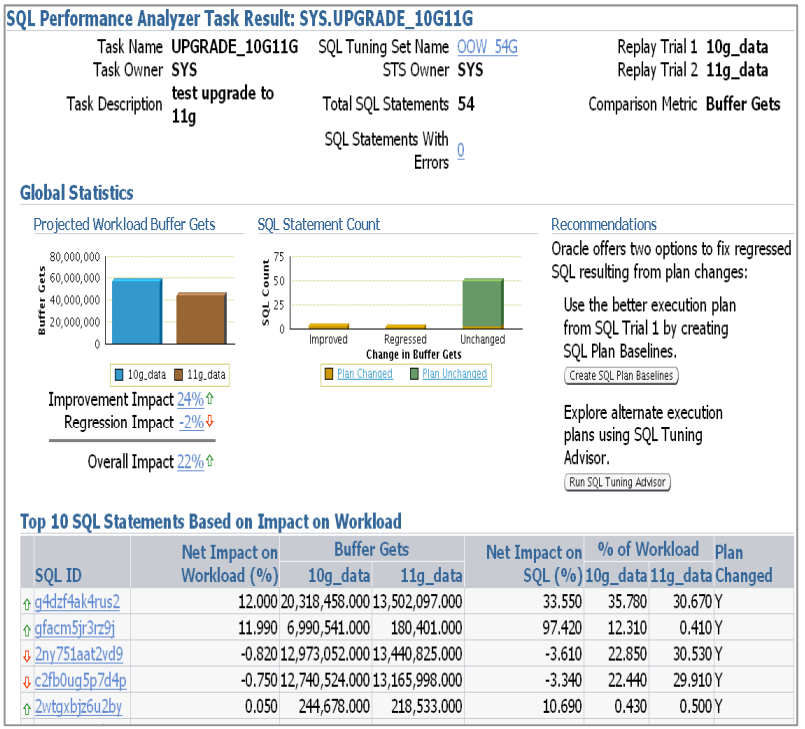

SQL Performance Analyzer Task Result: SYS.UPGRADE_10G11G

Task Name UPGRADE_10G11G
Task Owner SYS
Task Description test upgrade to 11g
SQL Tuning Set Name OOW_54G
STS Owner SYS
Total SQL Statements 54
SQL Statements With Errors 0
Replay Trial 1 10g_data
Replay Trial 2 11g_data
Comparison Metric Buffer Gets

Global Statistics

Projected Workload Buffer Gets



Improvement Impact 24%
Regression Impact -2%
Overall Impact 22%

SQL Statement Count



Recommendations

Oracle offers two options to fix regressed SQL resulting from plan changes:

Use the better execution plan from SQL Trial 1 by creating SQL Plan Baselines.

Create SQL Plan Baselines

Explore alternate execution plans using SQL Tuning Advisor.

Run SQL Tuning Advisor

Top 10 SQL Statements Based on Impact on Workload

| SQL ID | Net Impact on Workload (%) | Buffer Gets 10g_data | Buffer Gets 11g_data | Net Impact on SQL (%) | % of Workload 10g_data | % of Workload 11g_data | Plan Changed |
|---|---|---|---|---|---|---|---|
| g4dzf4ak4rus2 | 12.000 | 20,318,458.000 | 13,502,097.000 | 33.550 | 35.780 | 30.670 | Y |
| gfacm5jr3rz9j | 11.990 | 6,990,541.000 | 180,401.000 | 97.420 | 12.310 | 0.410 | Y |
| 2ny751aat2vd9 | -0.820 | 12,973,052.000 | 13,440,825.000 | -3.610 | 22.850 | 30.530 | Y |
| c2fb0ug5p7d4p | -0.750 | 12,740,524.000 | 13,165,998.000 | -3.340 | 22.440 | 29.910 | Y |
| 2wtgxbjz6u2by | 0.050 | 244,678.000 | 218,533.000 | 10.690 | 0.430 | 0.500 | Y |

# SPA: Regressed SQL – The Culprit!!

```
SELECT T9.CONFLICT_ID, T9.LAST_UPD, T9.CREATED, T9.LAST_UPD_BY, T9.CREATED_BY, T9.MODIFICATION_NUM,
    T9.ROW_ID, T4.KEY_VALUE, T25.PR_DEPT_OU_ID, T1.INTEGRATION_ID, T1.PRTNR_FLG, T25.PR_REGION_ID,
    T25.CITIZENSHIP_CD, T6.ATTRIB_07, T25.EMAIL_ADDR, T25.EMP_FLG, T25.FAX_PH_NUM, T25.FST_NAME,
    T25.CITIZENSHIP_CD, T25.HOME_PH_NUM, T25.AGENT_FLG, T25.JOB_TITLE, T25.LAST_NAME, T25.SEX_MF,
    T25.MEMBER_FLG, T25.MID_NAME, T25.OWNER_PER_ID, T9.NAME, T25.PERSON_UID, T25.PRIV_FLG, T20.STATUS,
    T8.PR_EMP_ID, T3.NAME, T25.CURR_PRI_LST_ID, T25.PR_OU_ADDR_ID, T1.NAME, T20.PR_ADDR_ID,
    T25.PR_EMAIL_ADDR_ID, T25.PR_ALT_PH_NUM_ID, T25.PR_REP_DNRM_FLG, T25.PR_REP_MANL_FLG,
    T25.PR_REP_SYS_FLG, T25.PR_MKT_SEG_ID, T22.PR_EMP_ID, T8.PR_EMP_ID, T13.LOGIN, T18.LOGIN, T17.PR_FAX_NUM_ID,
    T25.PR_GRP_OU_ID, T25.PR_INDUST_ID, T25.PR_NOTE_ID, T25.PR_OPTY_ID, T25.BU_ID, T25.PR_SYNC_USER_ID,
    T25.PR_PER_ADDR_ID, T25.PR_PER_PAY_PRFL_ID, T25.PR_POSTN_ID, T25.PR_PROD_LN_ID, T25.PR_RESP_ID,
    T17.PR_SMS_NUM_ID, T25.PR_SECURITY_ID, T5.NAME, T25.MED_SPEC_ID, T25.PR_STATE_LIC_ID, T25.PR_TERR_ID,
    T25.PROVIDER_FLG, T12.OWN_INST_ID, T12.INTEGRATION_ID, T11.SHARE_HOME_PH_FLG, T25.CUST_SINCE_DT,
```

**Next Steps:**

- **Tune Regressed SQL Statements using SQL Profiles or Plan Baselines**
- **Replay workload and analyze results**

```
T18, ORAPERF.S_POSTN_CON T19, ORAPERF.S_POSTN_CON T20, ORAPERF.S_ORG_EXT T21, ORAPERF.S_POSTN T22,
ORAPERF.S_ADDR_PER T23, ORAPERF.S_ORG_EXT_FNX T24, ORAPERF.S_CONTACT T25 WHERE T25.PR_DEPT_OU_ID =
T1.PAR_ROW_ID (+) AND T1.PR_POSTN_ID = T22.PAR_ROW_ID (+) AND T25.CURR_PRI_LST_ID = T3.ROW_ID (+) AND
T25.PR_POSTN_ID = T8.PAR_ROW_ID (+) AND T9.ROW_ID = T20.CON_ID (+) AND T20.POSTN_ID (+) = :1 AND T22.PR_EMP_ID
= T13.PAR_ROW_ID (+) AND T8.PR_EMP_ID = T18.PAR_ROW_ID (+) AND T25.PR_PER_ADDR_ID = T23.ROW_ID (+) AND
T25.MED_SPEC_ID = T5.ROW_ID (+) AND T9.ROW_ID = T4.CONTACT_ID (+) AND T9.ROW_ID = T25.PAR_ROW_ID AND
T9.ROW_ID = T17.PAR_ROW_ID (+) AND T9.ROW_ID = T11.PAR_ROW_ID (+) AND T9.ROW_ID = T6.PAR_ROW_ID (+) AND
T9.ROW_ID = T12.PAR_ROW_ID (+) AND T19.POSTN_ID = :2 AND T25.ROW_ID = T19.CON_ID AND T16.ROW_ID =
T19.POSTN_ID AND T19.POSTN_ID = T7.PAR_ROW_ID (+) AND T7.PR_EMP_ID = T15.PAR_ROW_ID (+) AND
T25.PR_DEPT_OU_ID = T2.ROW_ID (+) AND T25.PR_DEPT_OU_ID = T21.PAR_ROW_ID (+) AND T25.PR_DEPT_OU_ID =
T24.PAR_ROW_ID (+) AND T25.PR_SYNC_USER_ID = T10.ROW_ID (+) AND T25.PR_SYNC_USER_ID = T14.PAR_ROW_ID (+)
AND ((T25.PRIV_FLG = 'N') AND (T19.CON_LAST_NAME >= :3)) AND (T9.ROW_ID IN ( SELECT SQ1_T1.PERSON_ID FROM
ORAPERF.S_PARTY_PER SQ1_T1, ORAPERF.S_PARTY SQ1_T2, ORAPERF.S_ORG_EXT SQ1_T3 WHERE ( SQ1_T2.ROW_ID =
SQ1_T3.PAR_ROW_ID AND SQ1_T1.PARTY_ID = SQ1_T2.ROW_ID) AND ((SQ1_T3.INT_ORG_FLG != 'Y' OR SQ1_T3.PRTNR_FLG
!= 'N') AND SQ1_T3.ACCNT_FLG != 'N') AND (SQ1_T3.NAME LIKE :4))) ORDER BY T19.POSTN_ID, T19.CON_LAST_NAME,
T19.CON_FST_NAME
```

# Real Application Testing: General Recommended Strategy

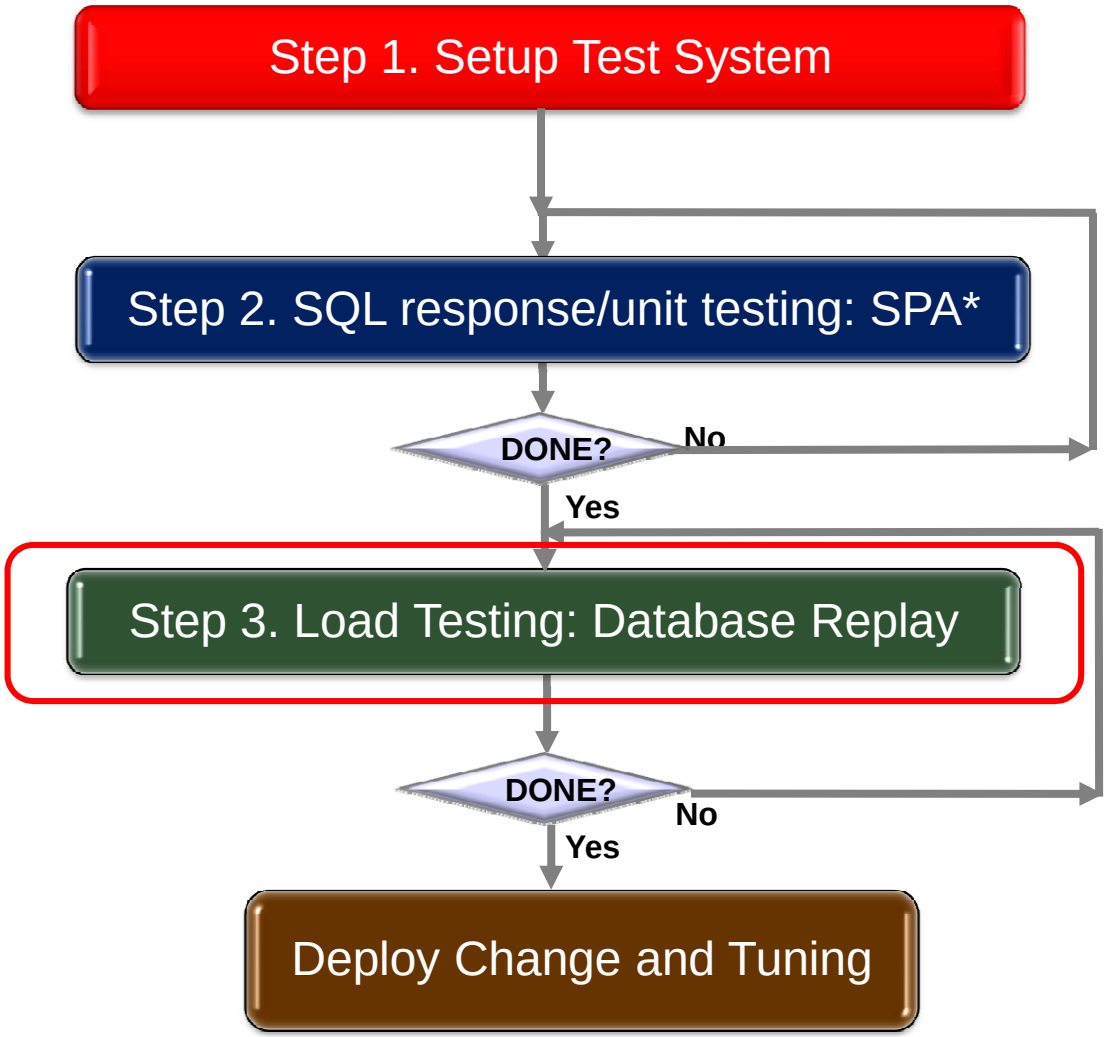

# Recommended Testing Methodology with Database Replay (1)

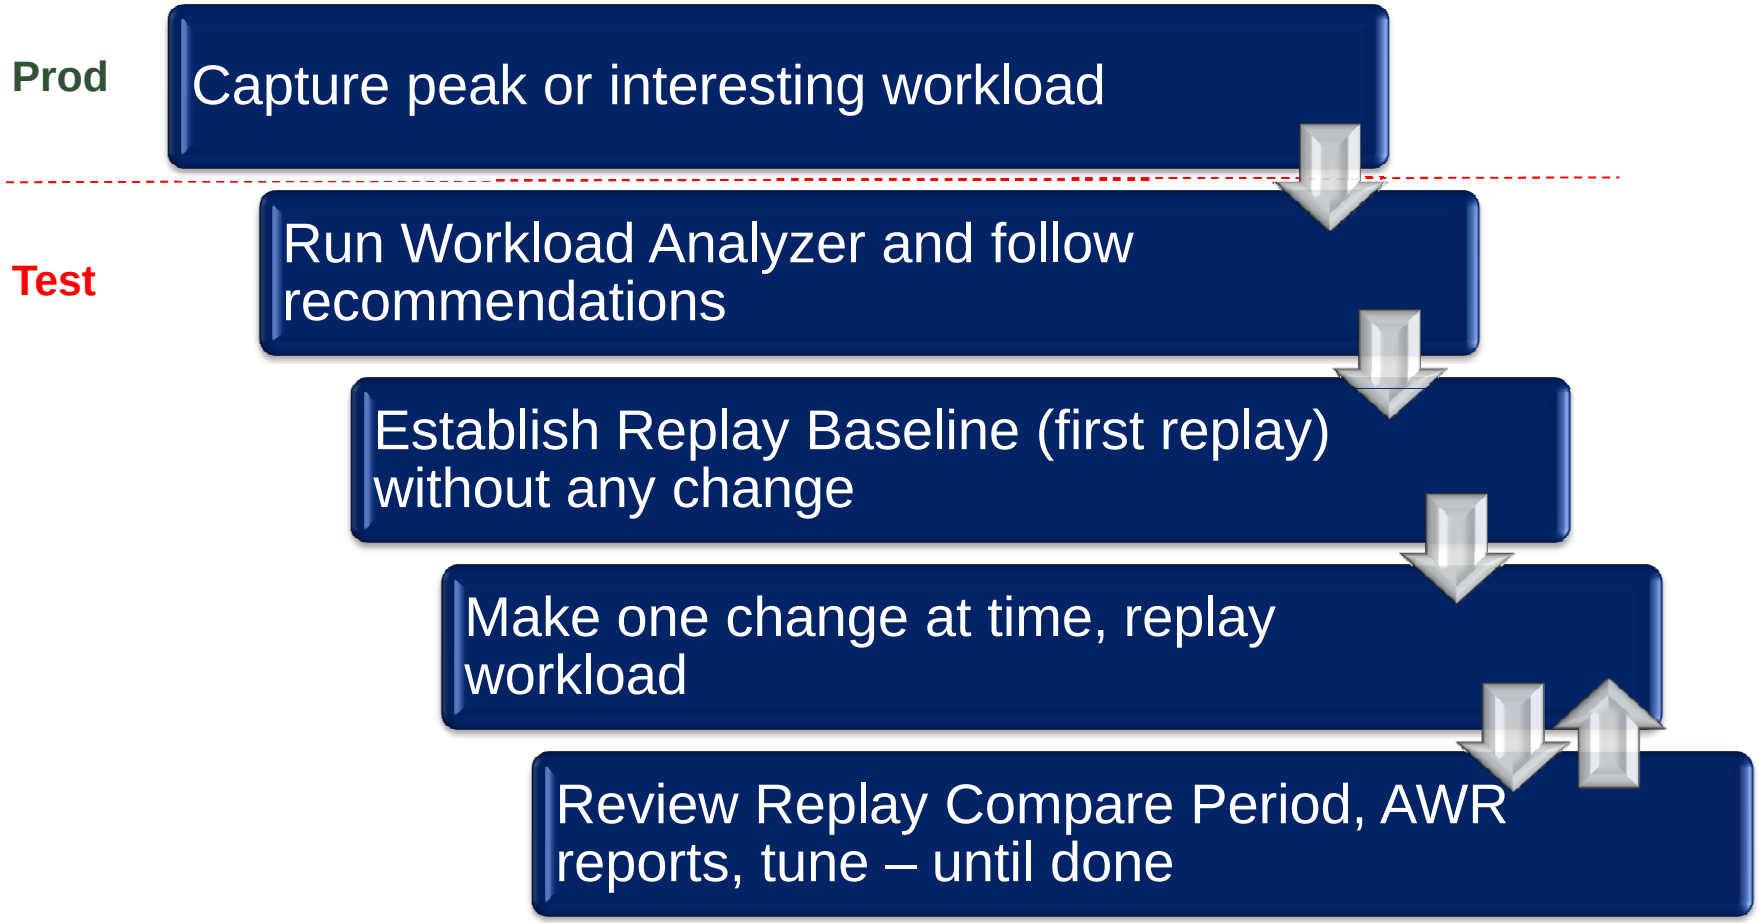

# Recommended Testing Methodology with Database Replay (2)

## Test one change at a time to understand causality

- Exception to this rule is when upgrading to Oracle Database 11g

## Start with small duration capture, e.g., 30-60 min, perform end-to-end testing, then iteratively move on to longer duration testing

- This strategy will quickly unravel any test system setup issues
- Makes it easier debug potential issues

## Establish Replay Baseline

- Use Replay Compare Period Report to understand Baseline (first replay) deviations from production capture
- Perform replay analysis, understand divergence (covered in Replay Analysis)

# Recommended Testing Methodology with Database Replay (3)

**Once Replay Baseline is established, compare two replays in the same environment**

- Baseline to Replay N
- Replay N-1 to Replay N for incremental changes and tuning

**In addition to replay divergence, use using application metrics to validate replay**

- For e.g., possible application metrics such as calls records processed /hr, orders entered/ min

**Save workload, tuning (SQL Profiles), AWR export, and Replay reports after each run**

- Test systems may need to be refreshed or testing done at a later time
- Better safe than sorry!

# Recommended Testing Methodology with Database Replay (1)

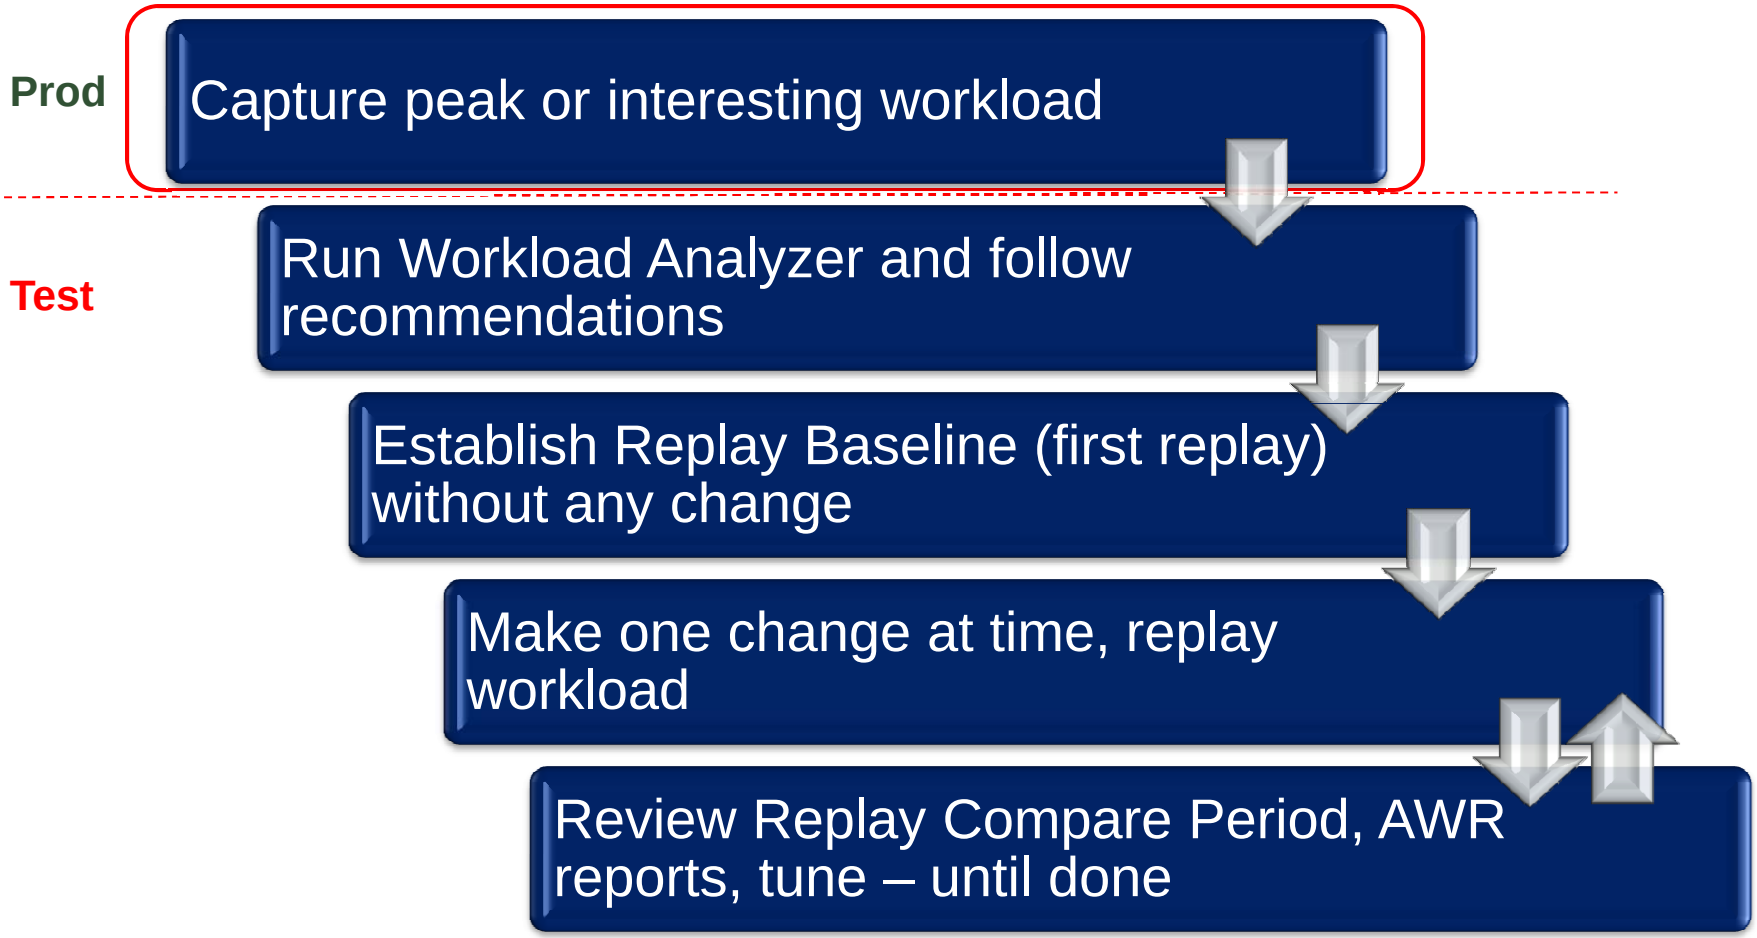

# Database Replay - Workload Capture Tips

## How do I estimate capture disk space?

- Maximum of
  - Extrapolated size from smaller duration capture (10-30min)
  - 2 * Bytes received via SQL*Net from client statistics (from AWR report)

## Should I restart database before capture?

- In general, the answer is NO
- For high number of in-flight transactions/busy system
  - Follow Workload Analyzer recommendations
  - Replay can still be done, but replay analysis should factor possible divergence
  - Application validation for capture duration can help determine if replay quality is good

# Database Replay - Workload Capture: Best Practices

## Filter background activity

- For e.g., monitoring infrastructure - STATSPACK, OMS, EM

## Save AWR performance data

- Create AWR baseline or export AWR after workload capture to avoid purging of AWR data

## Capture SQL workload into STS along with Database Replay workload capture

- Automated workflow in Oracle Database Release 11.2.0.2
- The same can be done manually using API or EM in earlier releases

# Recommended Testing Methodology with Database Replay (1)

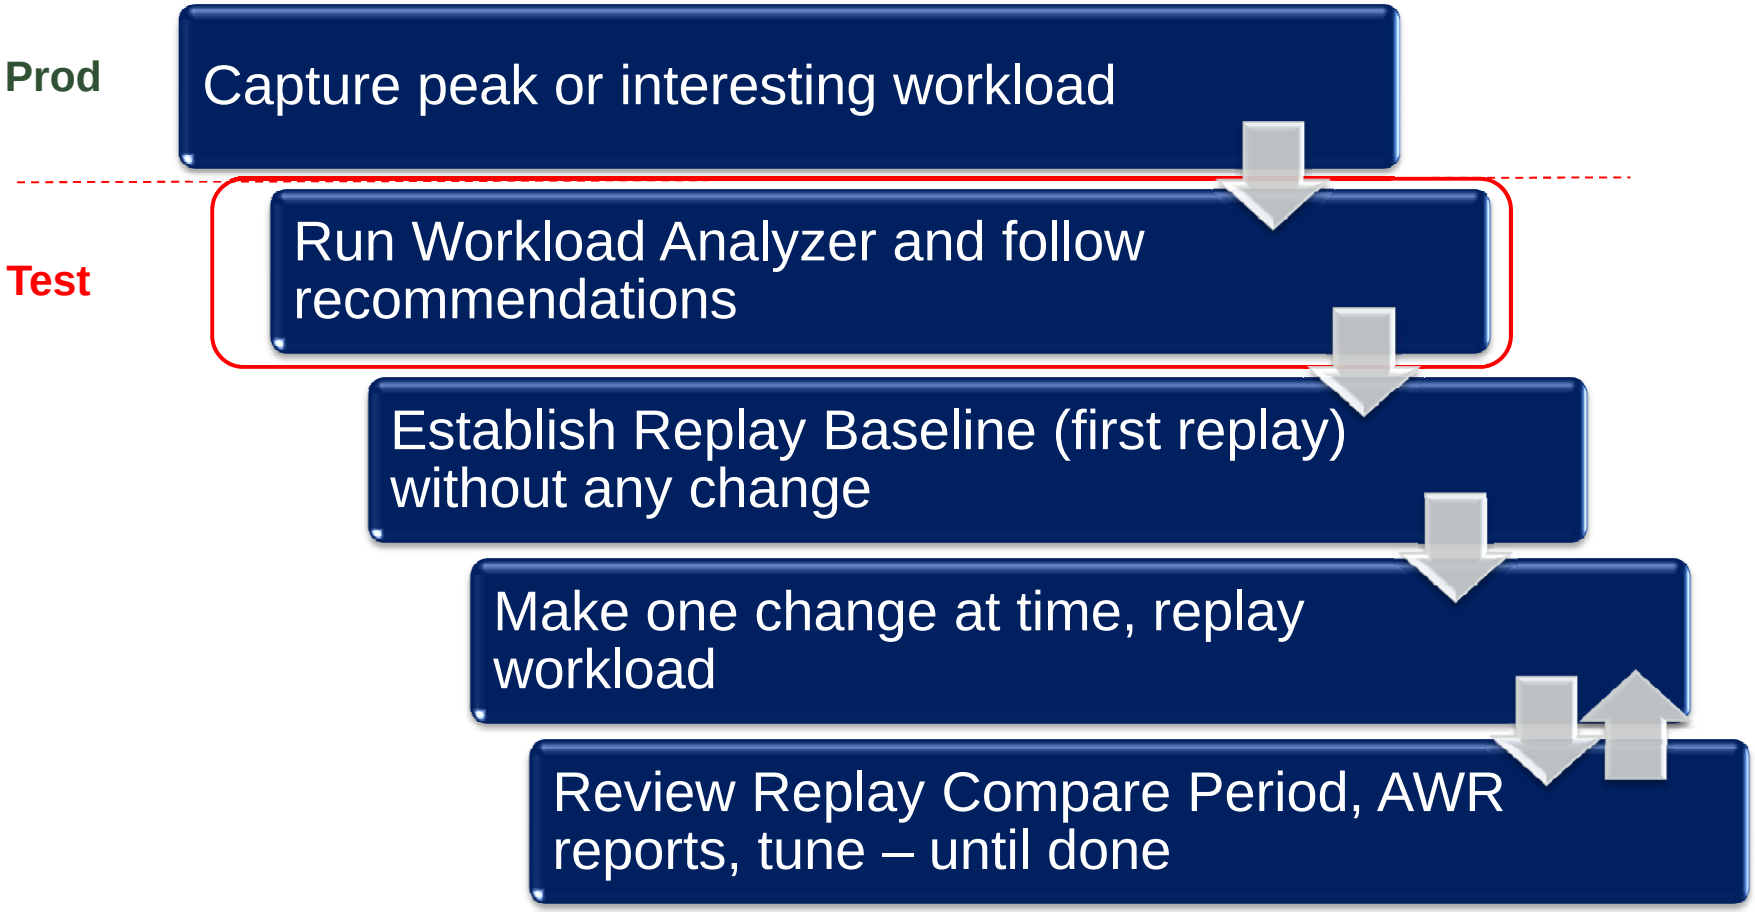

# Run Workload Analyzer and Recommendations

- Run Workload Analyzer on captured workload and follow recommendations to improve replay quality, for e.g.,

| Finding | Maximum Workload Impact |
|---|---|
| In-flight sessions | 36 % |
| SYSDATE and other time-dependent functions | 15 % |
| Missing AWR export | Unknown |

## Findings and Recommendations

### In-flight sessions

**Maximum Workload Impact: 36 % of DB Time**

**Rationale**

A significant part of your captured files have been captured in-flight.
This means that the captured session already existed before capture started, and could have already modified its state (e.g. through 'alter session') or be inside a transaction at this time.
This can cause some divergence (since only a part of some transactions was captured) or even bigger problems if some key session parameters are not set during replay.

**Action**

The best practice is to avoid in-flight session by restarting the database before capturing.
In case an in-flight session needs some missing session parameters to perform during replay, consider using a login trigger to set those parameters correctly.

# Recommended Testing Methodology with Database Replay (1)

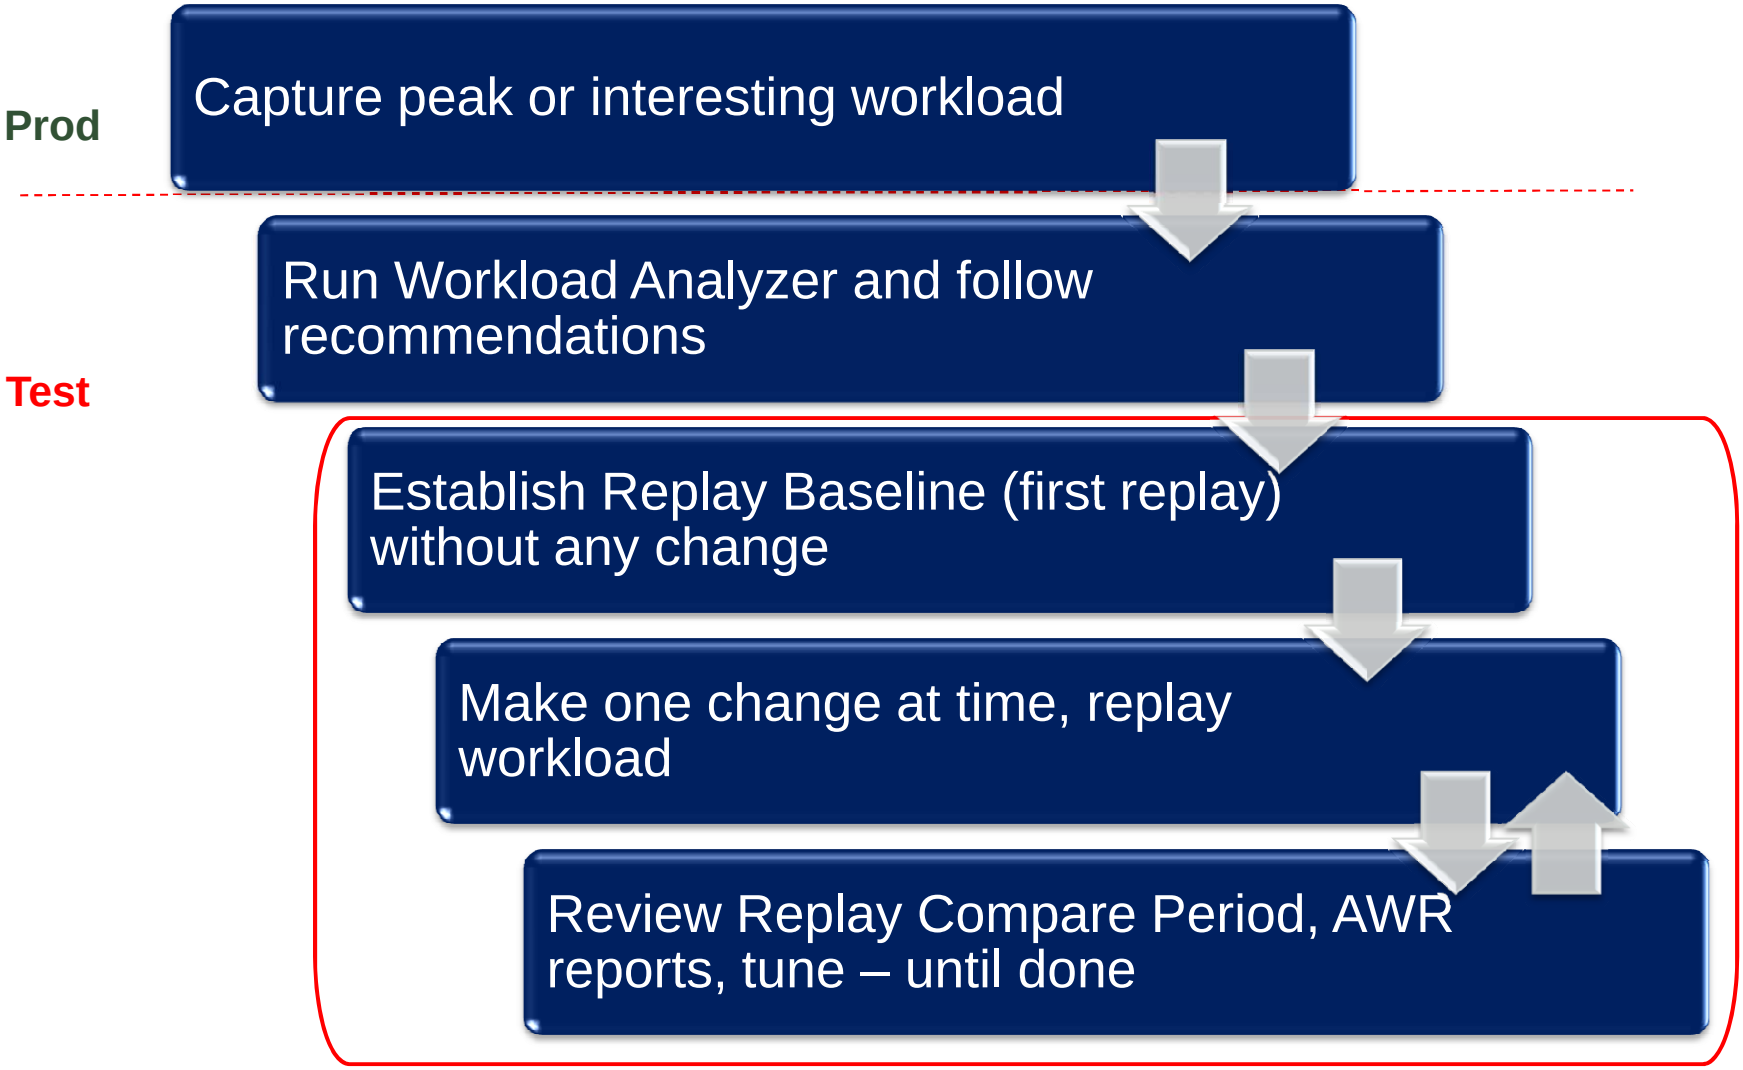

# Database Replay – Workload Replay Best Practices

Resolve and correct external dependencies, e.g., db links, external files

Use at least the number of replay clients as recommended by wrc “calibrate” command

Replay clients should not be co-located with database tier to avoid contention

# Database Replay – Workload Replay Analysis

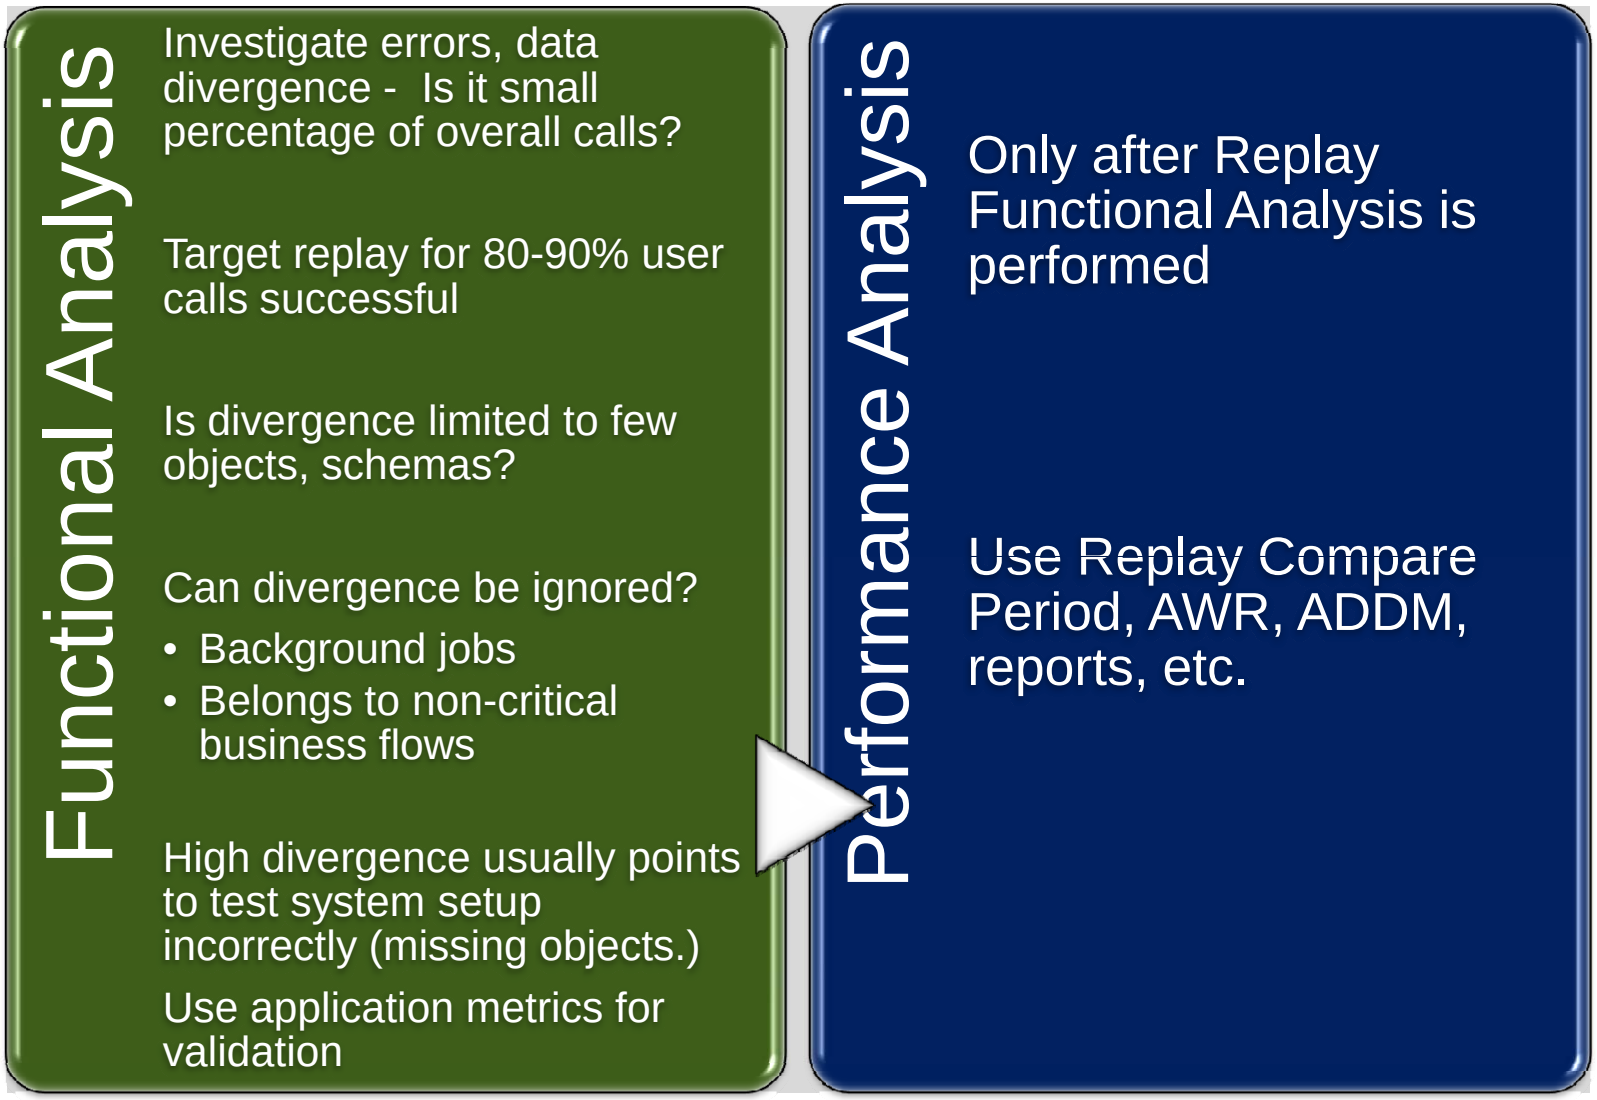

# Step 2: Database Replay Testing and Analysis (new screenshot to be added)

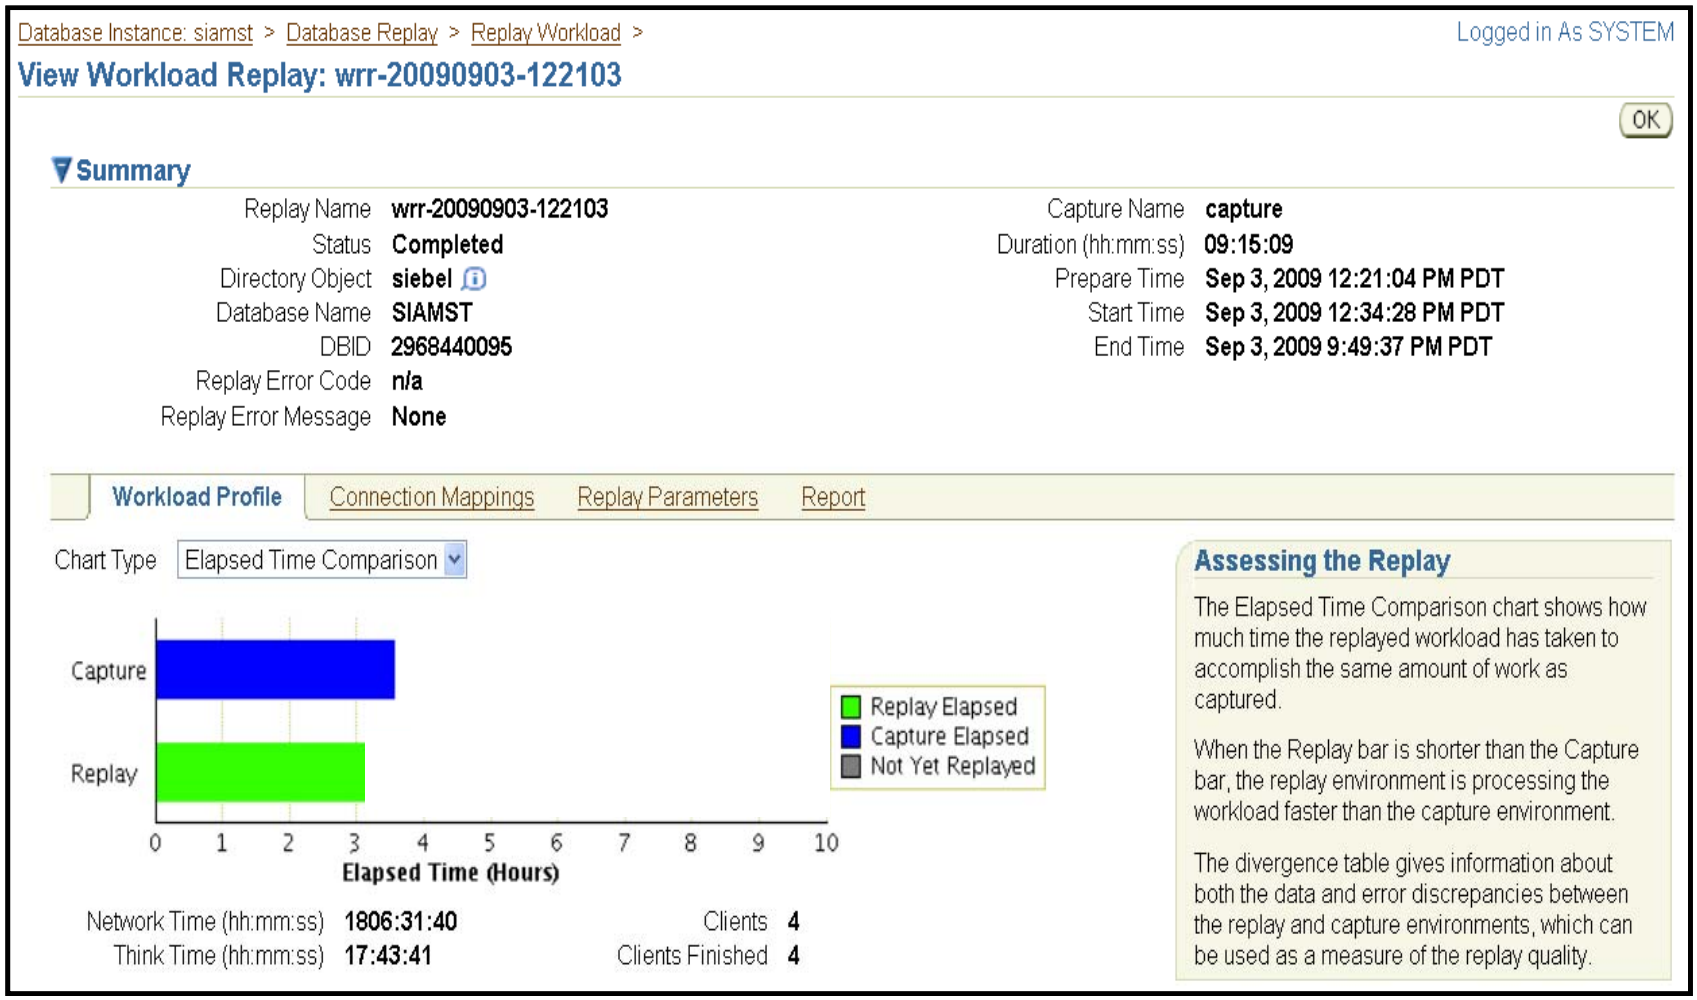

# Step 2: Replay Summary (Contd.): Errors and Data Divergence

## Divergence

| | Number of Calls | Percentage of Total Calls |
|---|---|---|
| **Error Divergence:** | | |
| Session Failures Seen During Replay | 0 | 0.00 |
| Errors No Longer Seen During Replay | 24 | 0.00 |
| Errors Mutated During Replay | 2489 | 0.02 |
| New Errors Seen During Replay | 853 | 0.01 |
| | | |
| **Data Divergence:** | | |
| DMLs with Different Number of Rows Modified | 200 | 0.00 |
| SELECTs with Different Number of Rows Fetched | 425 | 0.00 |

## Detailed Comparison

| | Capture | Replay | Percentage of Capture |
|---|---|---|---|
| Duration (hh:mm:ss) | 03:33:47 | 09:15:09 | 259.68 |
| Database Time (hh:mm:ss) | 09:07:54 | 17:16:07 | 189.11 |
| Average Active Sessions | 2.56 | 1.87 | 72.82 |
| User Calls | 11,212,622 | 11,212,577 | 100.00 |

# Step 2: Replay Errors and Data Divergence Analysis

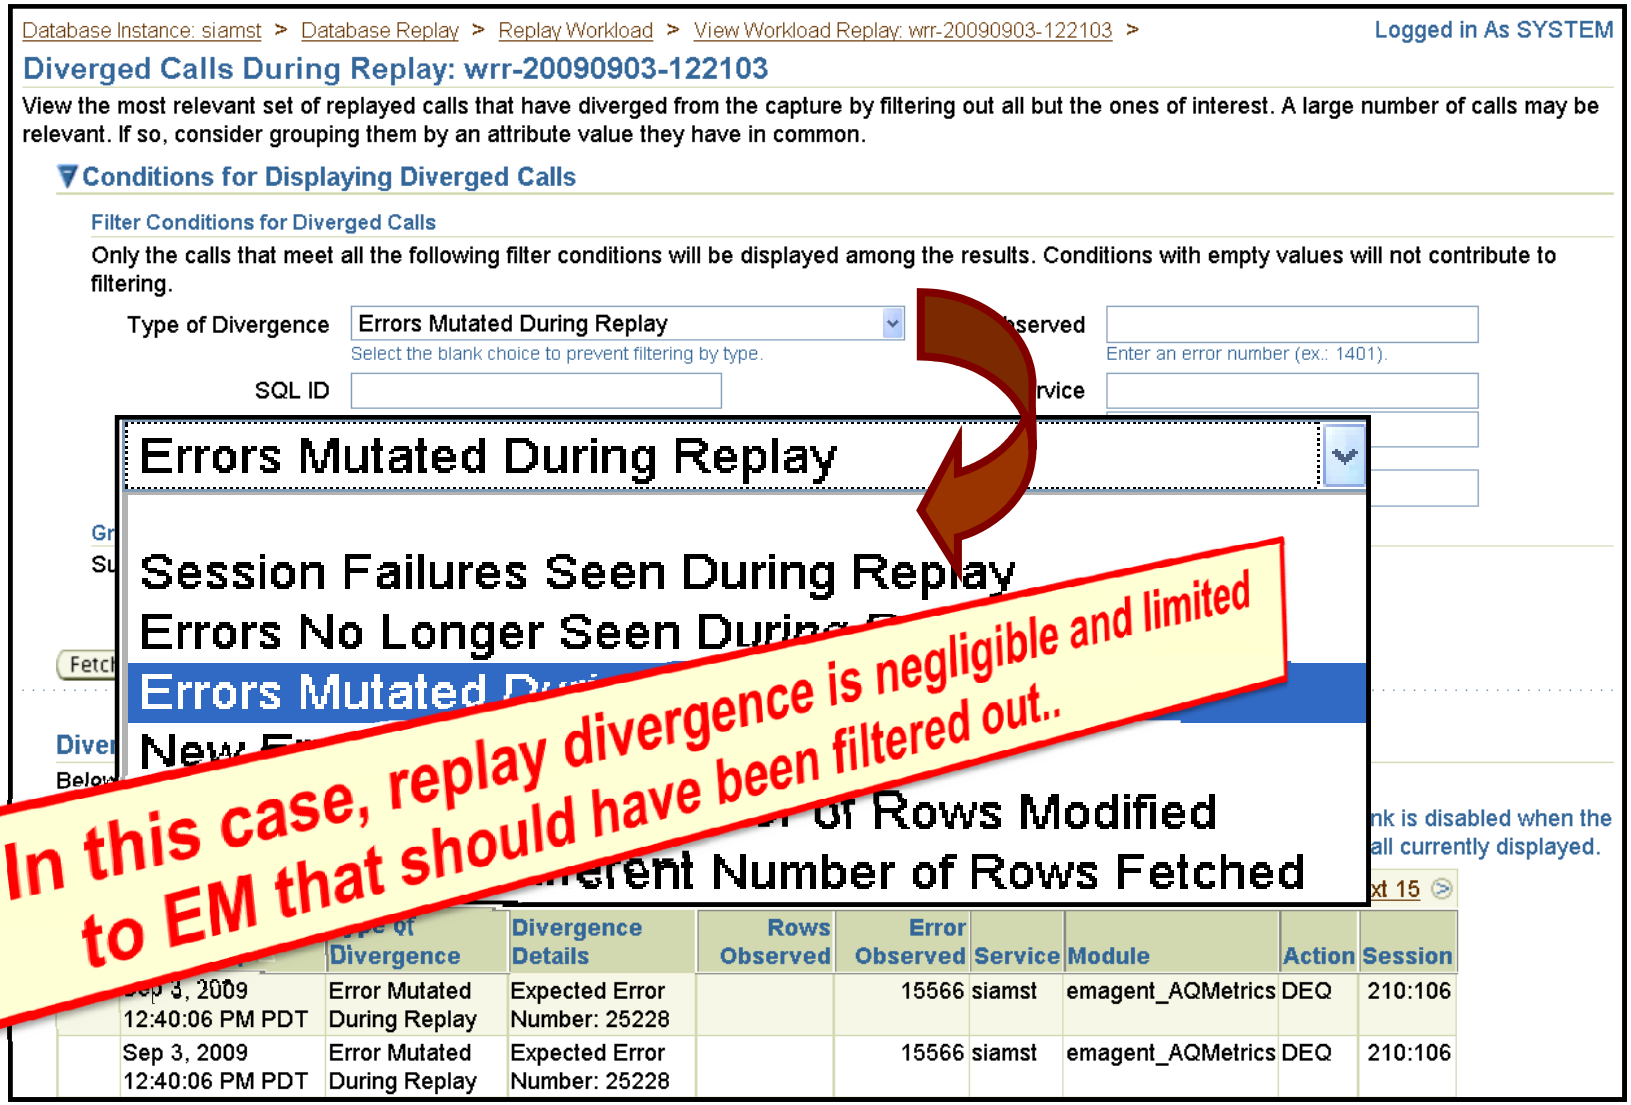

# Step 2: Database Replay Analysis

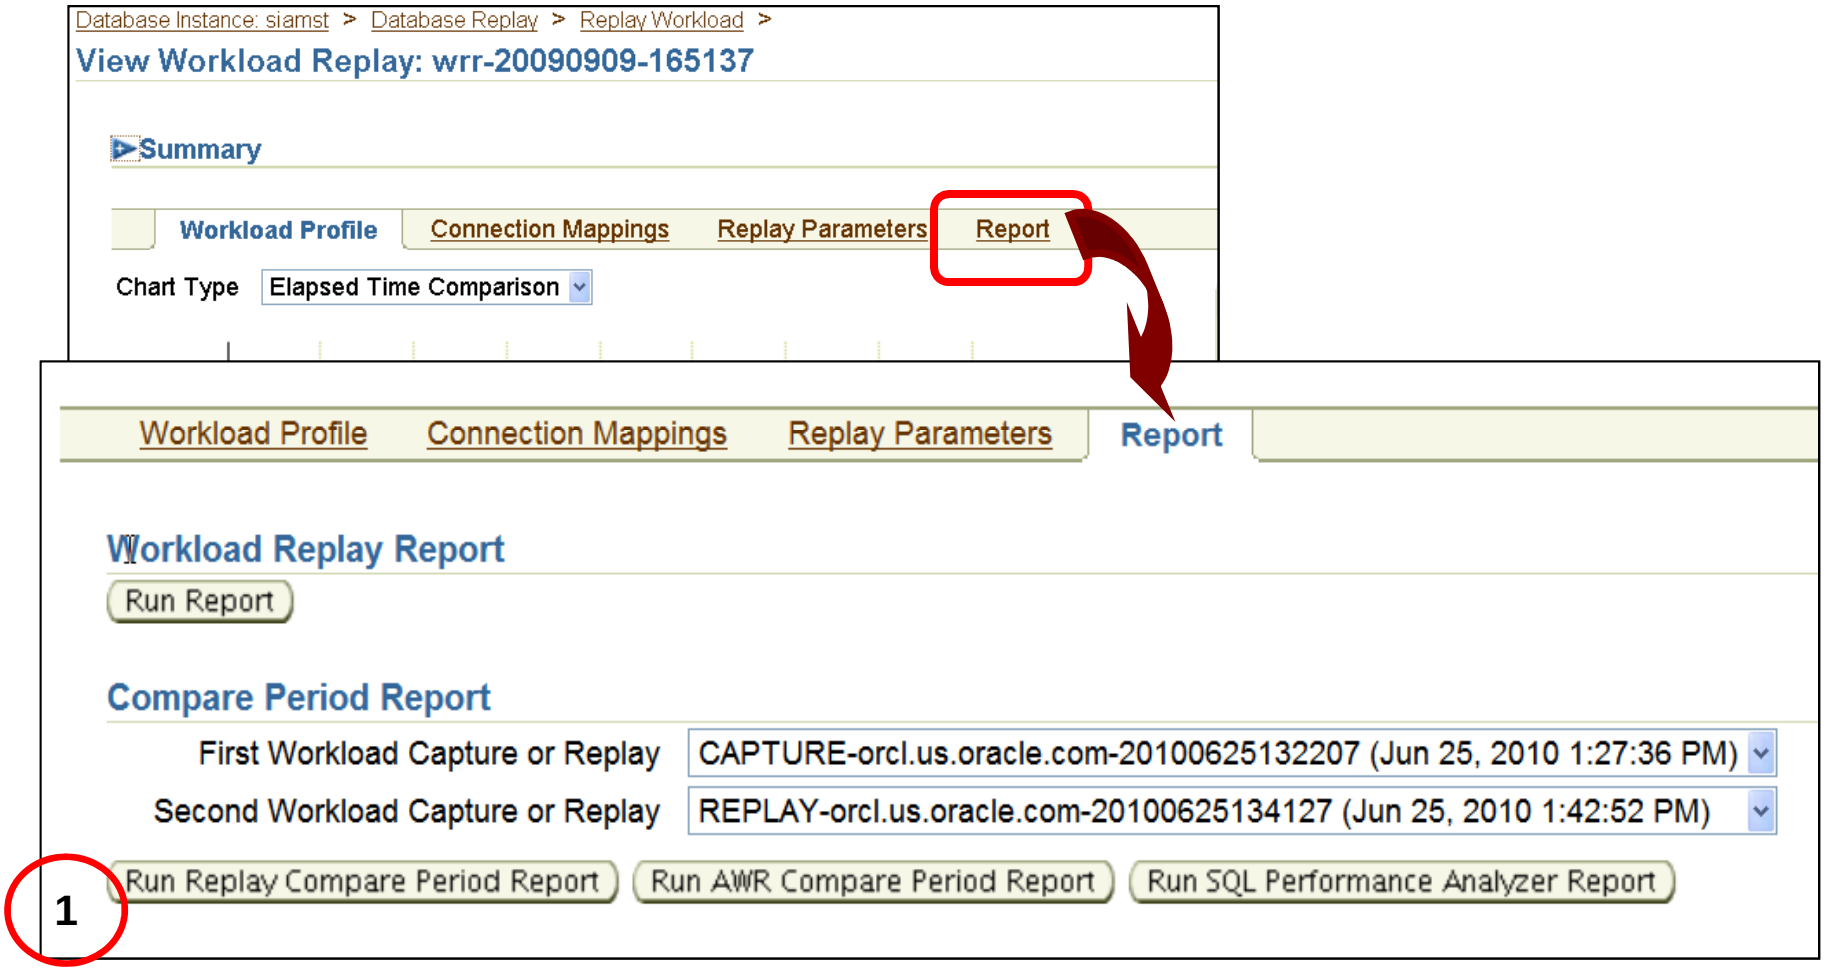

# Replay Compare Period Report

Improved in Oracle Database 11g Release 2

- Your new best friend in analyzing replay information!!
  - Provides holistic view of the experiment – covers functional and performance aspects of testing
  - “Replay Divergence Summary” categorization indicates if further analysis is necessary: LOW/MED/HIGH
- Two reports are available
  - Capture Vs Replay, Replay Vs Replay
- Identifies interference from other workloads, e.g., maintenance windows or other non-replayed workload
- Automatically runs ADDM
- Reports provide more accurate performance analysis
  - Uses enhanced ASH infrastructure for capture/replay sessions

# Replay Compare Period: Capture Vs Replay Baseline

Low Divergence

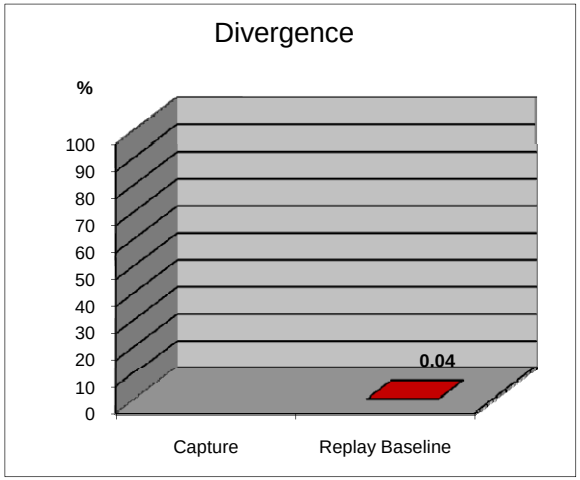


Replay Elapsed Time almost same as Capture

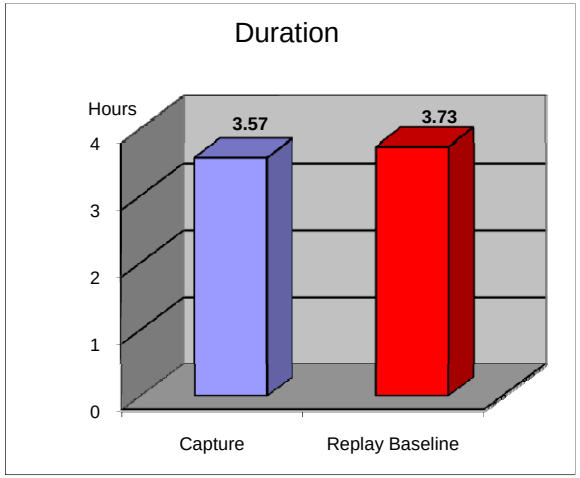


CPU Time is better!

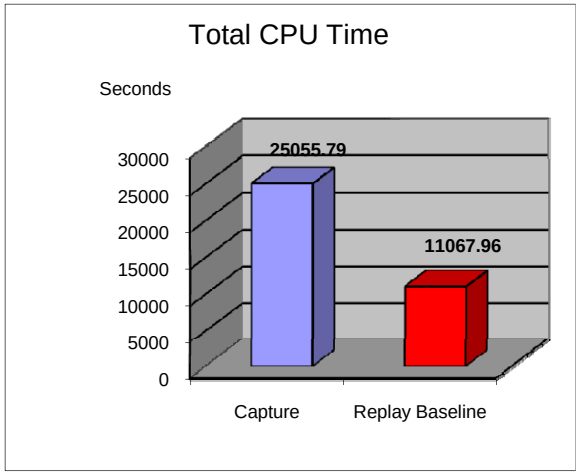


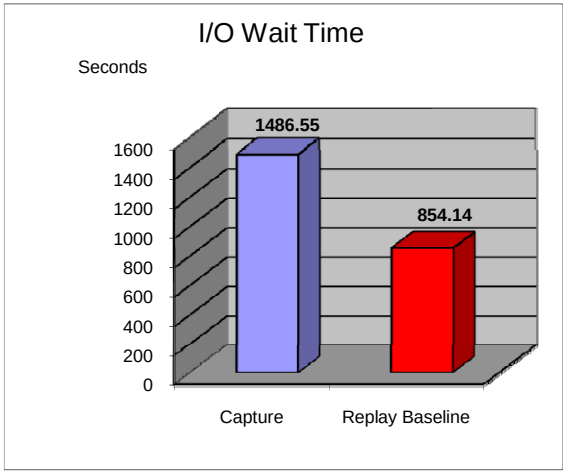


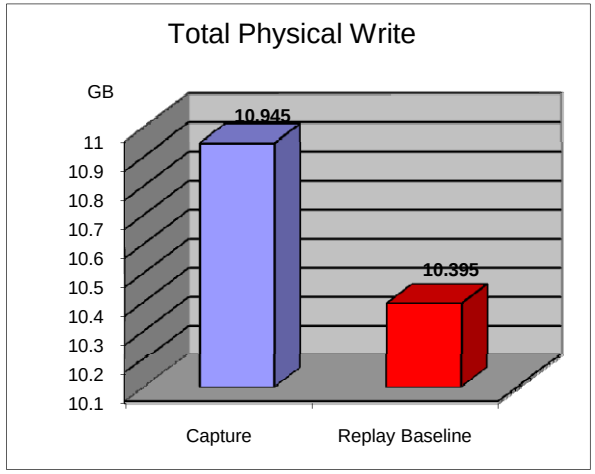


**Compare Period Report Link**

# Excerpts of Replay Compare Period Report

## Important Changes between Capture and Replay Baseline

### (-) Changes to Important Parameters

| | Capture | Replay |
|---|---|---|
| sessions | 566 | 792 |
| compatible | 10.2.0.4 | 11.2.0 |

### (-) Changes to Optimizer-Relevant Parameters

| | Capture | Replay |
|---|---|---|
| optimizer_capture_sql_plan_baselines | NULL | FALSE |
| optimizer_use_invisible_indexes | NULL | FALSE |
| optimizer_use_pending_statistics | NULL | FALSE |
| optimizer_use_sql_plan_baselines | NULL | TRUE |
| optimizer_features_enable | 10.2.0.4 | 11.2.0.1 |
| result_cache_mode | NULL | MANUAL |

### No Changes to Memory Configuration Parameters

### (-) Instances of the Capture Database

| instance number | instance name | host name | number of CPU cores | number of CPU sockets | physical memory | instance type |
|---|---|---|---|---|---|---|
| 1 | siamst | stase07 | 8 | 8 | 31.73 G | RDBMS |

### (-) Instances of the Replay Database

| instance number | instance name | host name | number of CPU cores | number of CPU sockets | physical memory | instance type |
|---|---|---|---|---|---|---|
| 1 | siamst | stase07 | 4 | 4 | 31.73 G | RDBMS |

# Oracle Database 10*g* Release 2 → 11*g* Upgrade: Summary


- After upgrading to 11g without any other change, performance remains almost the same
- Very low divergence rate, limited to background monitoring activity (EM)
- No other issues are noted during replay of the peak workload
- Further tuning can be performed or other new features can be added one at a time….

# Introducing New Features after Upgrade

- After Database Upgrade, introduce new features one at a time
  - TDE
  - Advanced Compression
  - RAC
  - Exadata – real-world example follows

# Agenda

- Deploying New Features - Challenges & Solution
- Database Replay Overview
- Database Replay Enhancements: Oracle Database 11*g* Release 2
- Strategy and Best Practices: Deploying New Features
- Real-world Customer Case Studies
- Conclusion

# Agenda

- Deploying New Features - Challenges & Solution
- Database Replay Overview
- Database Replay Enhancements: Oracle Database 11*g* Release 2
- Strategy and Best Practices: Deploying New Features
- Real-world Customer Case Studies
- Conclusion

# Oracle Real Application Testing
# A Solyndra Success Story

Paula Camporaso

Vice-President, Information Technology

Sept, 2010

SOLYNDRA™

The new shape of solar™

# Solyndra Backgrounder

- 5 year old company – Solar Manufacturer
- 3 years ago..
  - 100 Employees, few contractors
  - No Manufacturing
  - 40K sf. total
- Today..
  - 1200+ Employees, 250+ Mfg Temps
  - 1 Fab at ~Mfg capacity
  - FAB2 in Phase I
  - 1,350,000 sf.




SOLYNDRA

# Oracle Real Applications Testing

## Today's presentation:

**A real Story** about how Solyndra was able to ramp our Factories with minimal risk of change in a 24/7 environment during hyper-growth using Oracle's Real Application Testing Application.

Show of Hands..

- How many have been 'Reproduce Production Workload' in a Test Environment?
  - How many without 'Test Scripts'?
  - How many without Custom Code and/or Many Man hours?





SOLYNDRA

# IT Factory Systems - Automated Factory

"Factory Tools/ Automation are not only interoperable with Factory systems, they are interdependent."

**Key Measures of Success:**

- Factory Systems - 24/7 Production with little to 'No' downtime.
- Ensure 'Production' runs at speed of Tools and Automation
  - Not IT Networks or Information Systems
- Must scale without re-engineering

  ..and in most cases, without 'downtime'..



SOLYNDRA

# Essential 24/7 Factory IT Systems

- MES – Control / Track Product flow (WIP Routes)
- MCS - Automated Material control (AGVs)
- FIS - Quality Tool/ Product data
- SPC - Quality/ Statistical Process Control
- Many System/Tool Interfaces (100+)

  - *Must perform faster than Tools/ Automation (Real-time robotics)*
  - *Be available 24/7 (1 planned downtime per year)*
  - *Must Scale without reengineering and/or downtime*



SOLYNDRA

# Solyndra IT Challenges

## "Hyper-Growth – Change mgmt"

- Fab1 was the first of it's kind. "Everything was New"
- FCS in Q3/08 – Q4 Factory producing product 24/7
- Growth was 2x+ Qtr over Qtr through most of 2009
- IT Factory Systems required to Automate Production
- Transaction Volume and Databases doubling by Qtr
- Growing DBs showing signs quickly of slowing performance
- Applications and Databases were changing almost daily
- No way to truly 'reproduce Production workload' in Test Env



SOLYNDRA

# Phase I
# Fix what we know- “Low hanging Fruit”

- Implemented Load Balancing - Eased change mgmt
- Re-architect MES Application for reduced IO by 125%
- Upgraded all Oracle DBs to 11g – Use DB tools
  - Nehalem Server 35% improvement
  - 10g- 11g ~150% performance gain
  - Oracle DB Compression ~ 40% improvement
  - Oracle Partitioning for Xact tables ~ 50% improvement

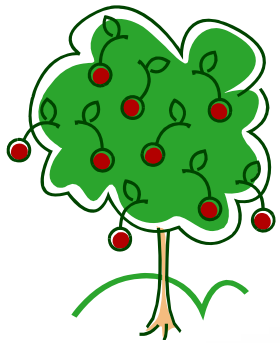

✓ *Yet we still needed significantly more ‘performance runway’ to support even 2009 Production Plan (6x to be exact)*



SOLYNDRA

# Phase II
# Still needed 6x improvement for FAB1

## Bottom-Line

- Needed to perform interative tuning scenarios “In a Safe/Test Env”
- Thus, we needed a tool able to reproduce production workload in Test
- We consulted with Oracle on their Real Application Testing Product

## Oracle Real Application Testing “Eval” deployed

- Captured and Replayed workload under several Scenarios
- Implemented several ‘tuning’ changes with ‘predictable results’
- Able to isolate ‘Storage IO’ as primary remaining bottleneck
- ‘This was huge’ ~ We were seeing the ‘results’ of multiple change Scenarios, previously only possible “In a Production Environment”



SOLYNDRA

# Summary of RAT Analysis/ Findings

- Database Storage was our biggest bottleneck.
- Launched Storage Selection Analysis process
  - NetApp - Baseline
  - HP – 15% improvement
  - Hitachi – 20% improvement
  - Oracle Exadata ’20X’ improvement

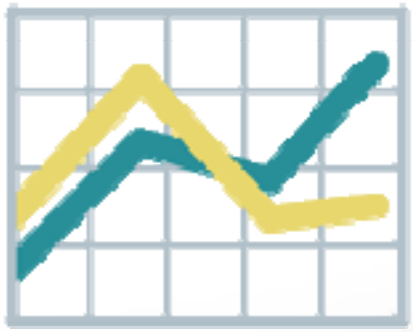



SOLYNDRA

# Oracle Real Application Testing –Results

## "Wow"

| | | Solyndra Single Replay | Exadata - Single Replay | Exadata - 2 Concurrent Replays | Exadata Single Replay + RMAN Bkup + CTAS + SQL Queries |
|---|---|---|---|---|---|
| **DB Time(min)** | | **410** | **114** | **132** | **122 (RAT Replay) 239 (Comprehensive )** |
| **DB CPU (sec)** | | 8,629 | 3,248 | 4,898 | 6,341 |
| **OS Statistics** | | | | | |
| | **IO Wait Time (sec)** | **1,399** | **4** | **3** | **3** |
| | IO Idle_Time (sec) | 150,237 | 281,208 | 275,361 | 260,611 |
| | IO Busy time (sec) | 13,440 | 6,784 | 12,686 | 27,446 |
| **SQL Execute Elapsed Time** | | | | | |
| | **Time (sec)** | **17,618** | **4,870.35** | **5,383** | **12,235** |
| | % DB Time | 72 | 71 | 68 | 85 |
| **PL/SQL execution elapsed time** | | | | | |
| | **Time (sec)** | **43** | **26** | **19** | **40** |

- ✓ *IO Waits 'directly track' to Factory Systems performance on the Mfg Floor*
- ✓ *Eliminating IO Waits, we knew would provide substantial improvements*



SOLYNDRA

# Exadata – Validation of Test Results

## 10X-27X Performance Improvement

| | Netapp | Exadata* | X Factor |
|---|---|---|---|
| • Create Online Copy of 107GB Table(s) | | 120 | |
| • Run monster queries | | | Average 27x |
| Query 1 (s) | 415 | 16 | 26 |
| Query 2 (s) | 849 | 32 | 27 |
| Query 3 (s) | 847 | 49 | 17 |
| Query 4 (s) | 120 | 2 | 60 |
| Query 5 (s) | 215 | 40 | 5 |
| • Online Backup - 1.3TB (m) | | 131 | |

| | NetApp (Production) | | | Exadata* | | | |
|---|---|---|---|---|---|---|---|
| Time Interval | # Execs | % Execs | Average Exec Time (sec) | # Execs | % Execs | Average Exec Time (sec) | X Factor |
| Under 5 (sec) | 15,689,851 | 93% | 0.5 | 12,068,974 | 93% | 0.2 | 2.74 |
| 5 to 10 (sec) | 1,108,001 | 7% | 5.4 | 876,206 | 7% | 1.6 | 3.31 |
| 10 to 20 (sec) | 5,959 | 0% | 18.0 | 4,536 | 0% | 9.4 | 1.92 |
| 20 to 50 (sec) | 43,412 | 0% | 23.7 | 1,779 | 0% | 4.7 | 5.08 |
| 50 to 100 (sec) | 2 | 0% | 59.8 | 16 | 0% | 36.4 | 1.64 |
| 200 to 500 (sec) | 392 | 0% | 388.3 | 43 | 0% | 126.6 | 3.07 |
| 1000 to 2500 (sec) | | 0% | | 20 | 0% | 61.4 | 0.00 |
| More than 2500 (sec) | 4 | 0% | 75,324.1 | 16 | 0% | 7,919.5 | 9.51 |
| Totals | 16,847,621 | | 75,819.8 | 2,951,590 | | 8,159.6 | |

Average 10x



SOLYNDRA

# Real Application Testing results: Exadata

- Overall 4x improvement in statistical performance
- Average 27x improvement in monster query performance
- Almost complete elimination of IO Wait Time
- Super linear scalability

  (2x load takes 1.15x DB time)

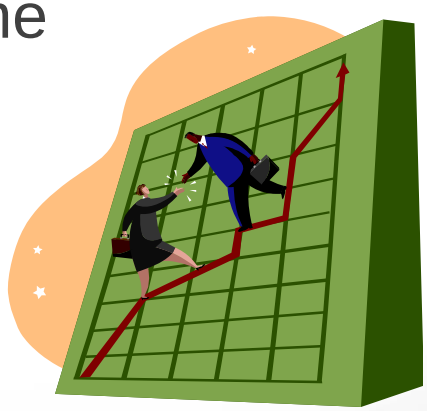

- **Overall ~ 20X improvement on the “Factory Floor”**



SOLYNDRA

# Learnings - incorporated into FAB2 Architecture

(FAB2 = 4x Production volume of FAB1 @ capacity)

- 24/7 Apps must be able to ‘Reproduce Workload production’ in Test
- ORAT is the only tool we know that can truly Capture and Play transactions exactly, providing needed insight to mitigate the ‘risk of change’
- For Solyndra, ‘Flash’ Storage technology is key for our Automated Factory performance needs
- Exadata tested and ‘realized’ 10X-20X performance gains over all other Storage evaluated (Using ORAT)
- Essential to “Proactively” identify issues while in ‘Test’
- *(2) Exadata’s currently in production in FAB2*



SOLYNDRA

# RAT – Business Value/ Impact Summary

- The risk of making changes to our fast growing IT Factory Systems without a tool like Real Application Testing would be prohibitively High.

Final Measure of Success:

"Solyndra has  not had a single Factory System/Database performance issue nor a reboot since April 2009".



SOLYNDRA

# Real Application Testing

**Tom Robertson**
**Database Technology Architect**
**Infrastructure and Operations**
**September, 2010**

Nationwide®
*On Your Side*

# Overview

| | |
|---|---|
| **Challenges** | • Applications and systems close to maximum utilization OR due for hardware refresh<br>• Complexity: reporting instances, increased storage costs and capacity, information not available in timely manner<br>• Can Real Application Testing provide a method to validate new commodity hardware, platform strategy?<br>• Can commodity platforms handle high OLTP and mixed workloads? |
| **Solution Approach** | • Use Database Replay to capture batch and OLTP workloads for peak periods<br>• Use Database Replay to execute workload on new systems, measure system performance, resource utilization, and any SQL regression |
| **Benefit** | • Replay workload on new systems: 2x to 12x production volumes<br>• Process performance improvements noted: 2x to 10x<br>• Validated new features – up to 67x reduction in space through Advanced Compression*<br>• Validated new hardware handles peak workload with excess capacity and dramatically improved performance |

# Use Case and Load Testing Results

## System and Workload

Packaged Application and Custom Reporting Databases

Oracle Database 9.2.0.8 Release: 2 unique databases and Instances

Moving to RAC Platform, x86, database upgrade 11g: single database with multiple instances

Validated that migrated system can handle peak/mixed workload at 20% utilization, obtained range of 2x - 10x DB time improvement for various databases

x86 and RAC platforms here generate an average savings on hardware of 85% and 25%+ for software

***Legacy Server Model*** *- Capture*
RISC CPUs: 8 x 2 @2150 MHz
Memory: 64 Gb

***Replacement Server Model*** *- Replay*
CISC CPUs: Intel Xeon 2.27 Ghz 2 x 4
Memory: 72 Gb RAM

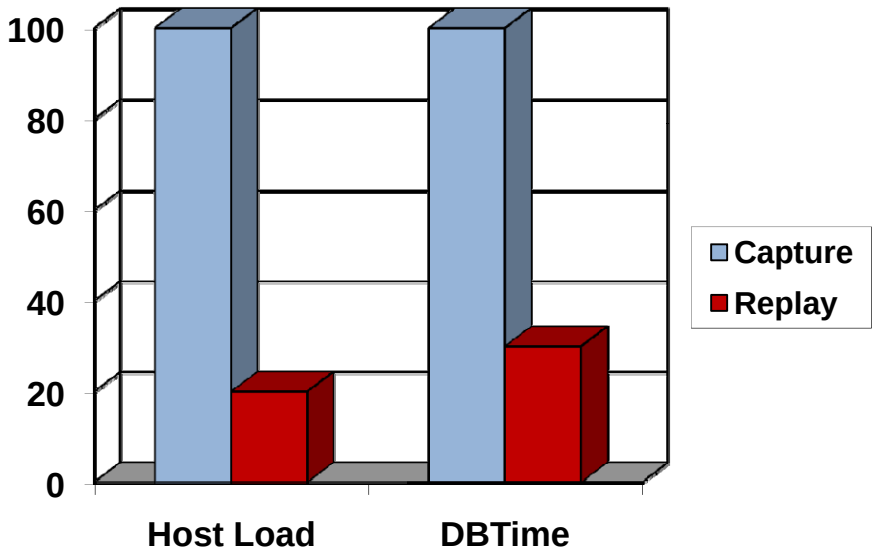

# Compression Results

Compression Advisor used to evaluate various compression techniques – OLTP, Archive High on key tables

Able to reduce disk footprint by up to 67x*

Performance overhead was negligible for OLTP compression

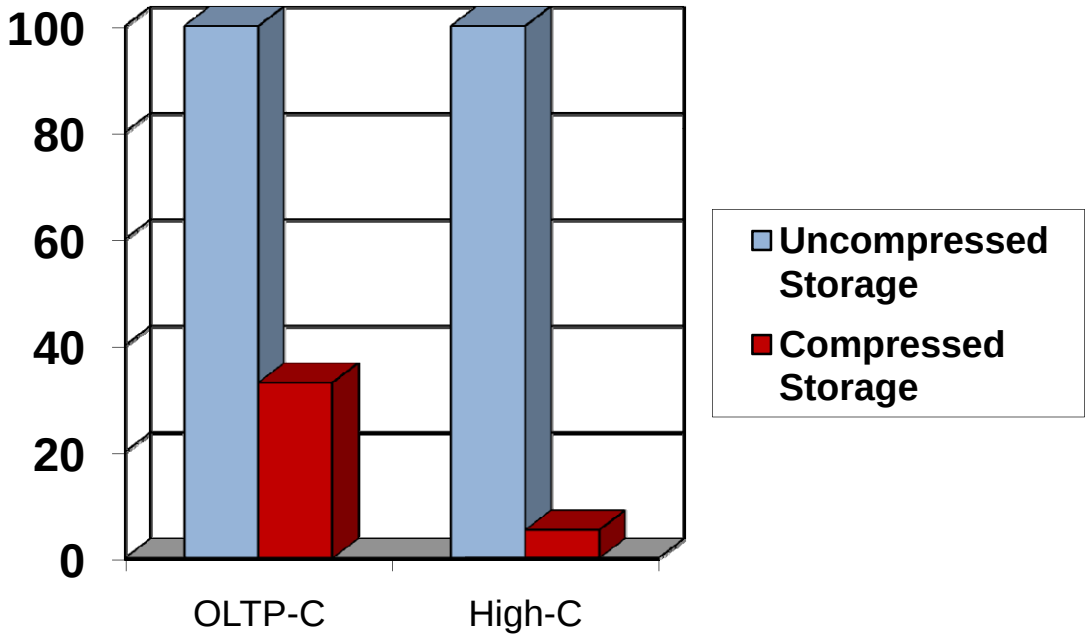


*Maximum compression achieved (67x) based on Exadata testing,

# Conclusion

- Real Application Testing enables businesses to safely deploy new database features using real workloads
- Increases business agility and uptime
- Increases staff productivity – less firefighting
- More focus on strategic planning and execution
- Increased capital expenditure savings

- 232% ROI over 4 years*
- <12 months payback period *

*Based on POC and business case from a Fortune 500 Insurance Company

Stadtwerke Munich Services GmbH

## What Munich Services GMBH Is Saying

"The positive results from using the RAT option for the first time enable us to look to the future and the upcoming changes in a much calmer manner than before"

"Facts instead of surprises!"

Christian Duschle, Munich Services GMBH, SAP customer

ORACLE

Source: Oracle for SAP Technology Update, Vol 19, May 2010

# Backup Slides

# Database Replay Terminology

- “wrc”: Workload Replay Client – These recreate production workload on test system. Multiple wrc clients can be used to drive large workloads many replay hosts
- Workload Capture Files: Workload that is captured on production and moved to test system for replay, files in binary format
- Workload Analyzer: Analyzes captured workload and provides information on how to improve replay quality

# Workload Capture

- Minimal overhead
- Platform and protocol independent
- Workload filters
- Capture interesting workload

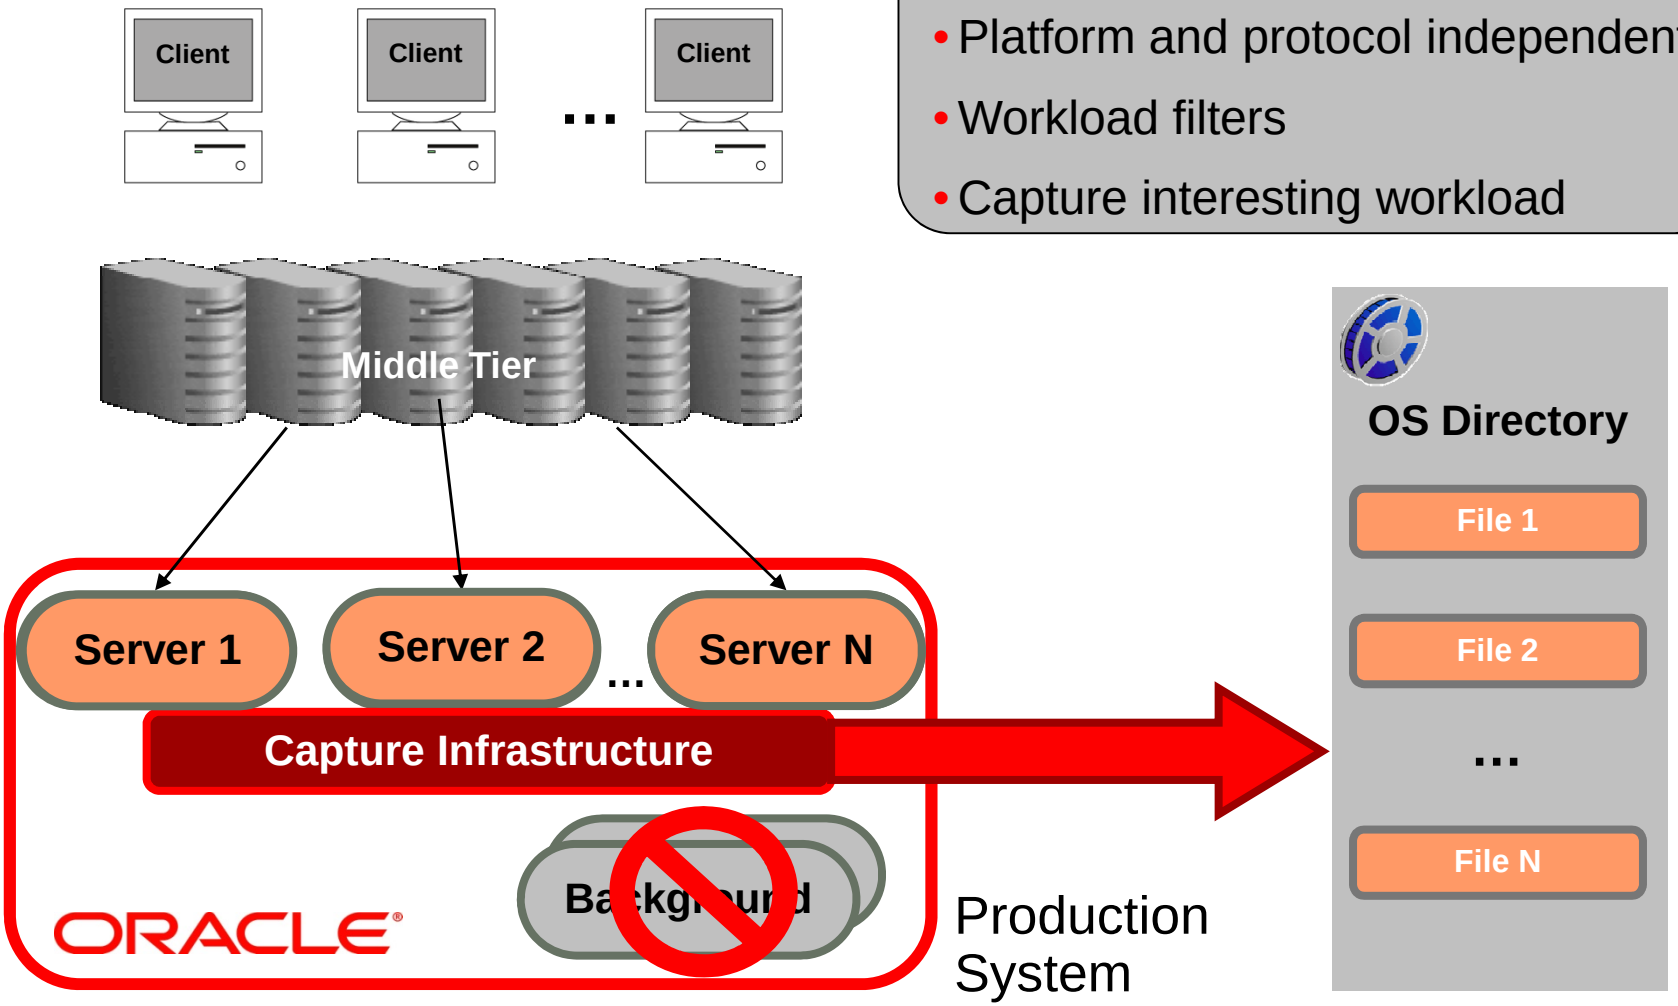

# Workload Capture Overhead

- Performance overhead
  - Workload dependent
  - Proportional to the data sent from the client
  - TPC-C throughput degradation about 4.5%
- Workload capture size
  - TPC-C 20min 100 users 10 warehouses: 1.2G
  - Maximum of (a, b)
    a) Extrapolated size from smaller duration capture (15-30min)
    b) 2 * Bytes received via SQL*Net from client AWR statistic
  - Enable capture for few minutes to assess size

# Step 2: Process Workload Files

- Setup test system
  - Application data should be same as production system as of capture start time
  - Use RMAN, Snapshot Standby, imp/exp, Data Pump, etc. to create test system
  - Make change: upgrade db and/or OS, change storage, migrate platforms, etc.
- Processing  makes captured data into replay ready format
- Once processed, workload can be replayed many times
- For RAC copy all capture files to single location for processing

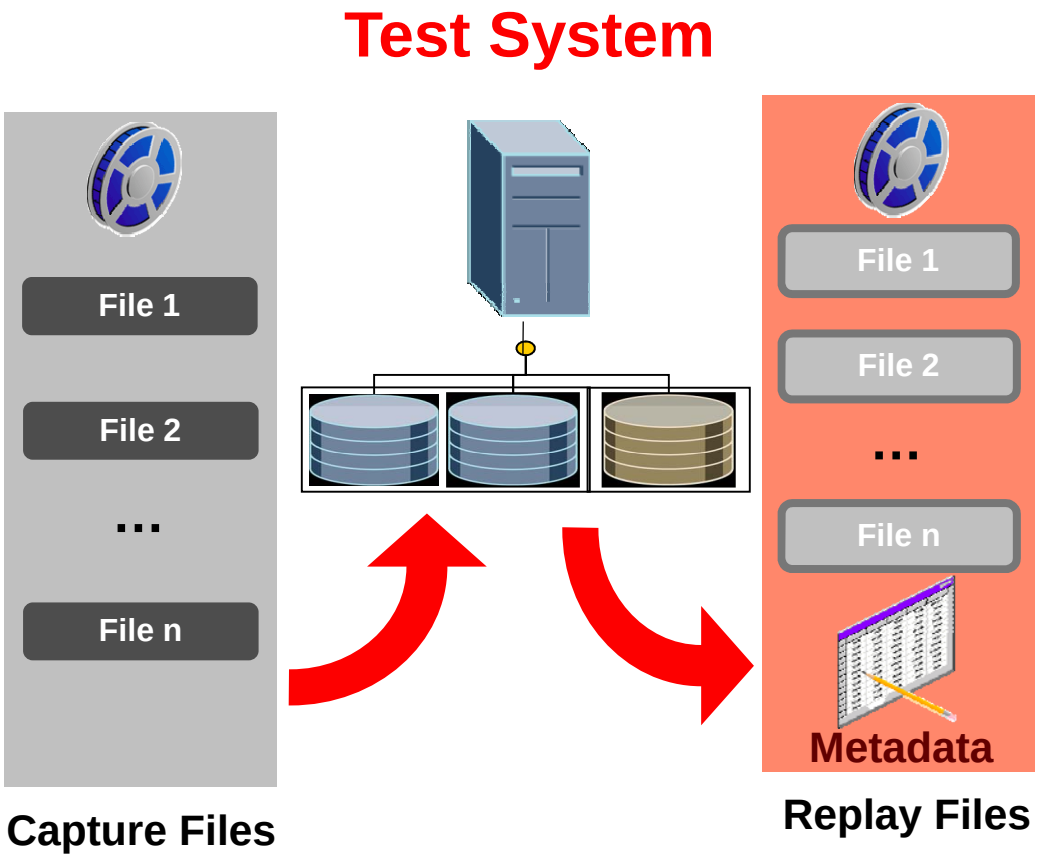

# Workload Replay Architecture

- Replay captured workload
  - Replayed operations see the same data and perform the same work
  - Preserve timing and concurrency characteristics
  - Same number of user connections
- Replay Client
  - Multithreaded OCI Client
  - Drives multiple captured processes
  - Scalable Architecture
  - Interprets capture into sequence of OCI calls
  - Functional replay