Lessons Learned Migrating a Major Application to Exadata v2

Aaron Werman

aaron.werman@gmail.com

Application Use Case Overview

- Complex integration of capital markets trading data
- Hundreds of ETLs, Thousands of tables
- 10K+ ETL executions per day, many highly complex
- Near real –time SLAs
- ODS with data sharing for entire line of business
- Several web applications, each with multiple hundreds of users, doing reporting and analytic queries
- Business, not traditional BI, SLAs for availability and recovery

Disclaimers

- All content is the opinion of the author, not my employer
- No endorsements are intended. All products mentioned are considered favorably by the author!

Some Oracle Context You should Read

- **Master Note (Doc ID 1187674.1)**
  - Database Machine and Storage Server
- **Best Practices Master Note (Doc ID 757552.1)**
  - Performance, Migration, Backup Recovery
  - **Best Practices (Doc ID 1067527.1)**
  - Database Machine Setup/Configuration
- **Supported Versions (Doc ID 888828.1)**
  - Database Machine running Storage Server Software 11g Release 2 (11.2)

Alternatives Considered and Rejected

- Key Issue: mixed workload. Obvious candidates are Teradata, DB2, Netezza, Oracle for ODS and a column database for DW
- General concern: migration from Oracle entails time/cost and risk. Application is highly tuned to Oracle physical design
- Teradata concern: too small a use case for their sweet spot
- DB2 concern: migration, organizational issues, and new support issues
- Netezza: recovery model, fact/fact joins, mixed workload, tuning, LOB support
- New column database for reporting: organizational risk and current ODS I/O bottlenecks
- No other POCs were done! We chose Exadata based on migration risk avoidance

Exadata Justification for the Application

- Migrating from Oracle to another platform migration entails significant schedule risk
- Scaling Oracle ourselves is not justified by cost / risk / technology stack (but may be less painful in terms of corporate architecture)
- Potential loss of business capability and likely miss of critical SLAs if we do not scale adequately
- Current gaps in corporate SAN engineering to support VLDB (and 100TB applications)

Prep Work: Oracle 10g to 11g Conversion/Validation

- We created a small copy of the app in Oracle 11g and tested for functional gaps
- No issues were raised
  - despite some of our stack (Informatica v8.1) not being certified for 11g)

Data Migration

- Key issues are time to migrate and disk space requirements and complexity
- We rejected RMAN Oracle 10g single instance → Exadata
  - Requires migration to 11g, ASM, RAC… too many steps
- We chose to use exports:
  - data pump (network) for almost everything
  - export classic for large LOB tables
- Be careful about considering ASM storage as a file copy target – there are limitations

Exadata Target State Architecture

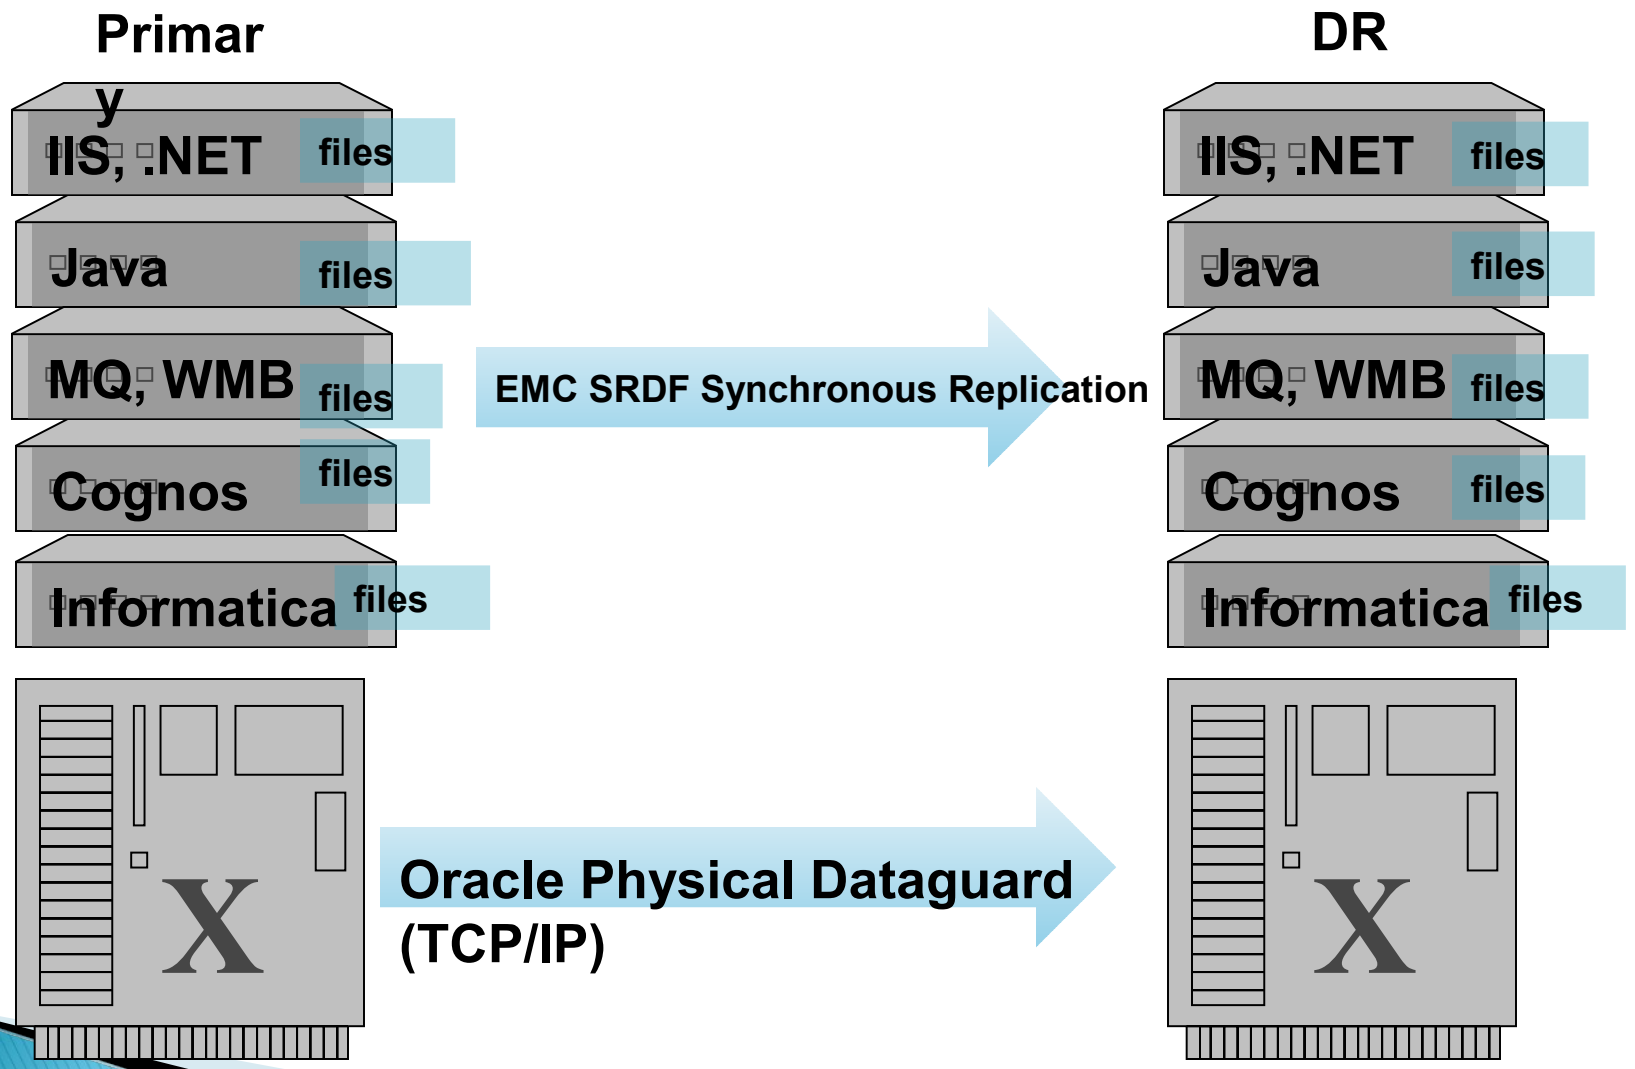

Bugs

Some major bugs encountered:

- 9356344 High CPU utilization of orarootagent.bin process with CRS-2409
- 9338087 ASM AND DATABASE HANG - CONNECT: OSSNET: CONNECTION FAILED TO SERVER, RESULT=5
- 9324531 ORA-00600: internal error code

These are now part of the current Exadata Oracle release

Performance Tuning

- Most due to SQL optimization differences between Oracle 10g and 11g
- ~40,000 SQL statements in app

  68 statements identified as substantially slower

  _____

  37 considered non-SLA relevant and ignored

  31 important SQL statements with significant SLA impact
  - *26 resolved using profiles*
  - *3 resolved using hints*
  - *2 resolved by query rewrite*
- optimizer_use_sql_plan_baselines?
- Note that most statements improved in performance, and improved in proportion to how much work/time they took

RAC Tuning

- Our DBAs, based on prior RAC strategies, initially partitioned the app to segregate load and prevent potential lock/block overhead
- After tuning, we determined there was no gain, and all load was allowed across all nodes
  - Your mileage may vary

Support Model

- Really complex to implement in our enterprise
  - Disruptive technology requires change in strategy for many stakeholders, especially infrastructure support groups
- Include time in your plan to allow for the transition
- Include ALL stakeholders in your planning

Application Design Futures Based on Exadata

- Application changes, such as reducing our real time ETL SLAs by 2/3s
- Index removal
  - We will experiment and remove many “for purpose” indexes
  - Incremental strategy with sufficient testing required
- ILM using Hybrid Columnar Compression
- Reducing duplication of data between operational and reporting requirements
- Likely BI (read-only reporting) against disaster recovery site using Active Data Guard

Questions?