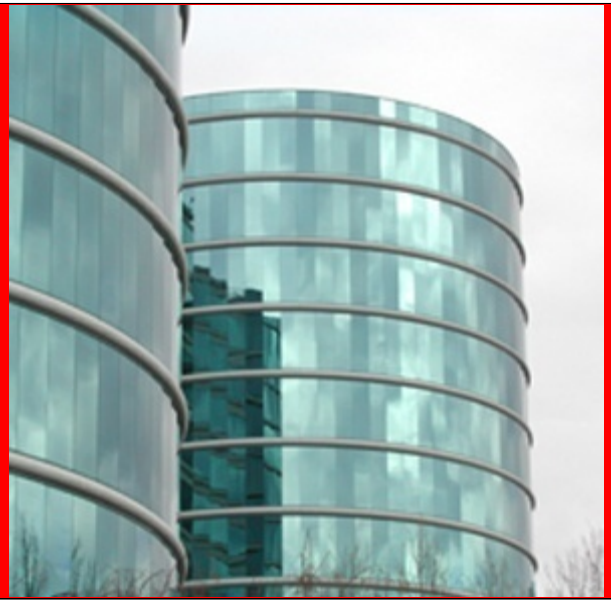

ORACLE®

Results of a Customer POC
SAP on Exadata

- Slide Set for Architects -

Martin Roselieb
Gerald Schreyer

Oracle Solution Center for SAP Competency, Walldorf

Agenda

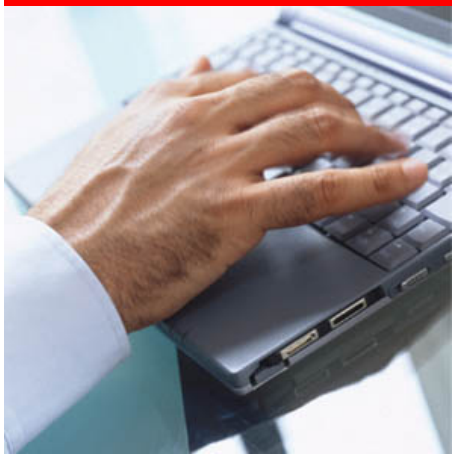

- Account Situation
- Rules For POC
- POC Preparation & Execution
- Results & Achievements
  - Volume tests
  - Demo on High-Availability Features
  - Lessons Learned / Conclusion
- Discussion

Account

Customer's Business

- B-2-B retailer for stationeries and office supplies
- operates in 30+ countries
- 30% growth expected within next 4 years

Oracle footprint

- most of centralized SAP applications are running on Sun hardware
- Oracle as the incumbent database vendor
- some IBM footprint: Websphere, Cognos, few IBM servers

Renewal of entire IT infrastructure is considered

- Strategic IT vendor selection process has started ...

Decision For POC

Respond to a POC which IBM did in Nov 2010

Goals:

- Expose customer's business-critical „real“ SAP workload to an Exadata environment
- Demonstrate significant throughput increase
- Show how Oracle technology can help to minimize the impact of hard-/software failures for the business

Conditions:

- No „Lab“ tuning; all improvements should be as well achievable by customer's staff for their productive systems
- Measure performance „as-it-is“ with a standard Exadata database setup
- Exadata for SAP is not generally available yet
- Exadata only for SAP Unicode systems

Conducting the POC

Preparation (3 weeks):

- Decision taking on volume test cases:
  - Simulation of online users in a web shop
  - Batch job net for invoicing
- Installation of the entire IT infrastructure for SAP
- Loading production database from tape
- 5 TB DB migration from Oracle 10.2 (SPARC) via Oracle 11.2 (x86/Linux) into an Exadata ASM/RAC environment
- Migrating parts of the application from Websphere to Weblogic

POC Execution (1 week):

- on-site at OSC Benchmark Center in Linlithgow, U.K.
- load injection (Loadrunner, Control-M) from customer's site requiring customer's involvement

Hardware Used - *)
Architecture for SAP Workloads

*) without servers for staging, load generation and the application tier for JAVA web shop

5 x Sun Blade 6270 M2
(each 2 x Intel Xeon x5680,
3.3 Ghz) running
Solaris 10 x86
31080 SAPS each
Equivalent to 155.400 SAPS

Exadata Database Machine
Half Rack X2-2
With high Performance
Disks
4 database nodes
(28815 SAPS each : Total
115.000 SAPS
7 storage cell nodes

Industry-standard
hardware running
Enterprise-class OS
for application tier

Database for high
throughput & availability
„out-of-the-box“

Migration Steps – How come from A to B ?

Restore database and upgrade to 11.2

- Restore from tape to M8000 (SPARC)
- Upgrade Oracle software and database to 11.2

Copy database files and convert endianess

- Use RMAN for copying and converting endianess on the fly
- Increase Parallelisation degree until storage saturates
- Use relocatable-tablespace feature
- Structural export (no data)

Create 4-Node RAC DB on exadata

- Graphical tool available – easy to understand. You cannot fail
- Maintain logfiles and REDO tablespace

Copy files into Exadata- Preparation

- NFS-Export the source filesystem and mount into Exadata

Migration Steps – How come from A to B ?

Copy database files into Exadata ASM

- Copy into ASM using RMAN
- Choose high parallelisation degree
- Create users
- Import structural information
- Create statistics

Time requirements

- Total runtime : 28 hours – can be further reduced

Result

- 5,7 TB database migrated incl. Endianess conversion
- Identical copy. No optimisation of database files
- Identical DB quality
- Transparent for SAP

Migration Steps – How come from A to B ?

Required knowledge for migration

- „black-box“ approach. You do not need to kwow the details
- No SAP migration certificate required
- No real RAC experience required, just basic know how
- DB creation via graphical tool. Simple and easy

Cookbook

- Whitepaper available for a step-by-step explanation
- Everybody can have it / Everybody can do it

Alternative

- You can do it by standard SAP migration …
- Online migration with Triple-O, downtime < 15 min. #1508271
- Offline migration with O2O (up to 1 TB/h)

Results Volume Testing - Batch

Invoicing of sales subsidiaries:

- job net starts 900+ batch jobs
- varying amount of jobs running in parallel (1 to 130)

Runtimes:

- 11h 30min on customer's current environment
- 7h 20min using our „staging“ database on Linux
  (fast x86 server hardware, single-instance DB)
- 3h 45min on Exadata.
  Two active nodes of a Half-Rack Exadata had only been utilized to 25%.

New environment reduced runtime to one third with a lot of spare capacity

CPU Utilization

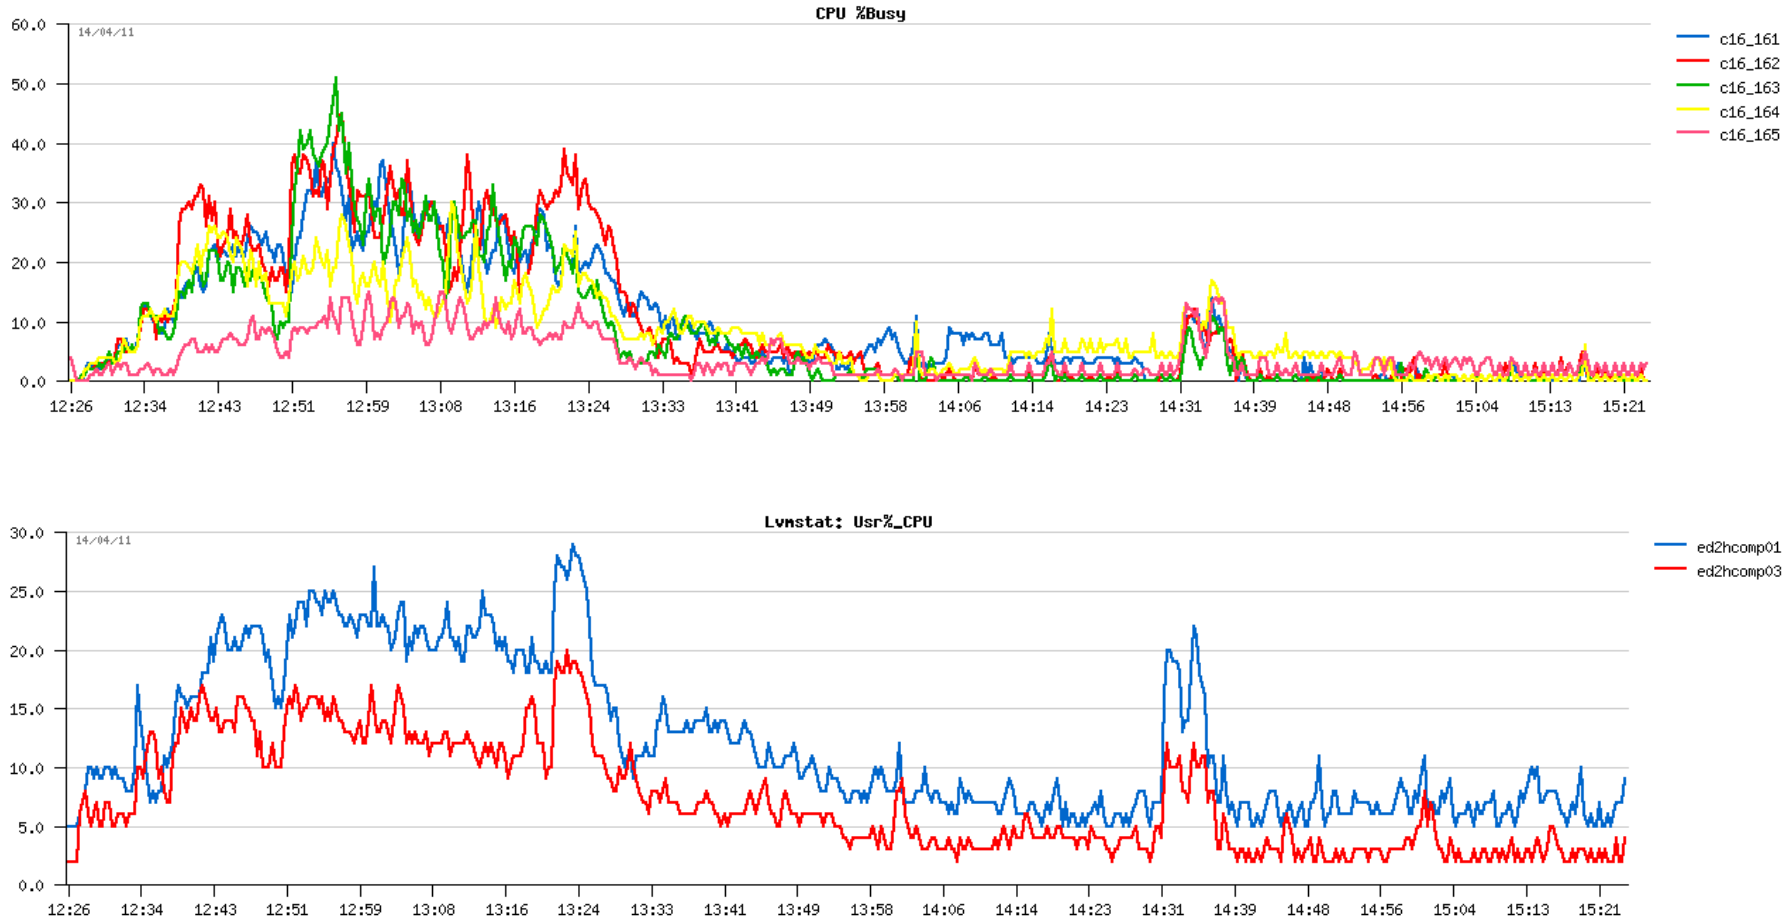

Results Volume Testing - Number of Online Users

Simulation of online orders entered in a web shop:

- Runs with various parameters (#users, thinktime, #RAC nodes)
- After an initial throughput was achieved, scalability bottlenecks within the SAP application had been identified and eliminated

Throughput:

- one node RAC: 7,500 users with 4s thinktime:
  => ~1s response time and 66% CPU utilization (one node)
- four node RAC: 6,000 users with 4s thinktime:
  => ~0.9s response time and 16% CPU utilization (all nodes)

We over-achieved the goal of 2,500 users.
Customer got confidence that the current infrastructure can serve up to 25,000 „real“ web shop users

Results Volume Testing - Number of Online Users Extrapolated

			per Node	Total
	users	util	@100%	@100%
1	7500	66%	11364	11364
2	7000	33%	10606	21212
3				
4	6000	16%	9375	**37500**
5				
6	5500	11%	8333	50000
7				
8	5000	9%	6944	**55556**
		measured		
		extrapolated		

CPU Utilization

Some Workload Characteristics

Keyfigure / per second	Batch	Online
User Calls	39.383	69.285
Physical Writes	6.538	2.952
Physical Reads	14.727	4.404
Logical Reads	348.300	389.741

Demo of High-Availability Features

Shutting down RAC nodes during SAP operations

- Showcasing the automatic reconnect from SAP application servers to RAC database nodes (TAF)
- Significant reduction of unplanned downtimes compared to a traditional server cluster with the need for a subsequent instance recovery

SAP Replicated Enqueue with Sun Cluster

- Demonstrating a solution which ensures continuous operation even in the case when losing the central SAP lock table held in main memory

Simulating Crashes from DB Nodes

Lessons Learned

Exadata provides better performance „out-of-the-box“

- you see instant performance benefits
- no Exadata- or DB specific tweaking/optimization was needed
- however, effect of new technology should be considered as well when making comparisons (e.g. faster CPUs on app server side)

We were not able to saturate Exadata

- no bottlenecks in the hardware/infrastructure were seen
- throughput was limited by contention within the application:
  - assignment of unique document numbers for new postings (SAP)
  - „hot spot“ inserts at the end of the tables (ascending primary key)

  → **in order to achieve even higher throughput numbers and to take full advantage of Exadata you need to resolve these bottlenecks**

Lessons Learned - cont.

User I/O is no more the major wait event

- „cell single block physical read“ (corresponds to „db file sequential read“ in non-Exadata environments) is very fast (< 1msec)

No RAC „Penality“ have been observed

- Almost same throughput numbers in 1-2-4-node configurations

Collecting DB workload & utilization data for sizing

- During test runs AWR reports have been captured along with Exadata utilization data
- Load profile can now be compared with DB workload from other Exadata prospects
- This should allow estimates for Exadata sizing

Lessons Learned - cont.

No Smart Scans in these workloads

- although activated, Smart Scanning did not happen: Well-tuned SQLs with unique index access; little cursor processing with few rows fetches

Discussion