
ORACLE®

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.

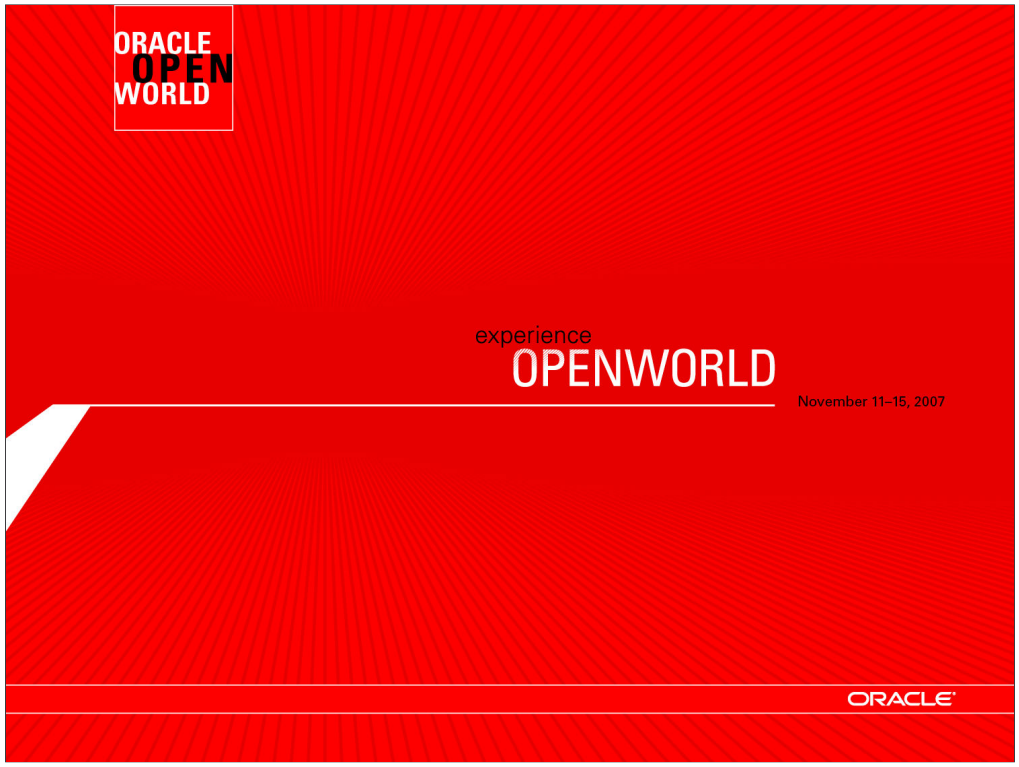
ORACLE
OPEN
WORLD
experience
OPENWORLD
November 11–15, 2007
ORACLE

ORACLE

ORACLE THIRTY YEARS 30

Current Trends in Database Performance

Andrew Holdsworth
Senior Director Real World Performance Server Technologies

Real World Performance Group Sessions

Session ID: S291301
Session Title: Current Trends in Database Performance
Track: TECHNOLOGY; Business Intelligence ; Database; Enterprise Management; Linux and Open Source Solutions
Room: 102
Date: 2007-11-14
Start Time: 09:45

Session ID: S291306
Session Title: Roundtable: The Real-World Performance Group
Track: TECHNOLOGY; Business Intelligence ; Database; Enterprise Management; Linux and Open Source Solutions
Room: 104
Date: 2007-11-14
Start Time: 11:15

Session ID: S294625
Session Title: Performance 2.0
Track: TECHNOLOGY; Database
Room: Yerba Buena Theatre
Date: 2007-11-14
Start Time: 15:00

Session ID: S291304
Session Title: When to Use the Appropriate Database Technology
Track: TECHNOLOGY; Business Intelligence ; Database; Enterprise Management; Linux and Open Source Solutions
Room: 102
Date: 2007-11-15
Start Time: 08:30

ORACLE

Topics for Discussion

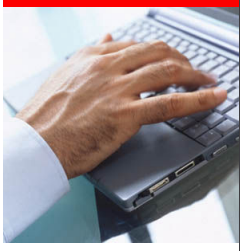

- Background Ideas
- Database Performance Basics
- Today's OLTP Challenges
- Today's DW/BI Challenges
- Benchmarks
- Platform Selection

ORACLE

Background Ideas

- Relational databases became popular as they enabled rapid development of applications.
- This does not imply rapid development of high performance, scalable applications. This takes more skill, effort, care, attention and testing.
- A large number of today's applications are running well despite being sub optimally implemented. This is largely because of the increase in CPU speed and the increase in DB buffer cache sizes over time.
- A large number of customers are only starting to realize how sub optimal their apps are when the active data grows to size exceeding the buffer cache and I/O becomes dominant in the response time of any query.

Impact of I/O, Interconnect and SQL Tuning

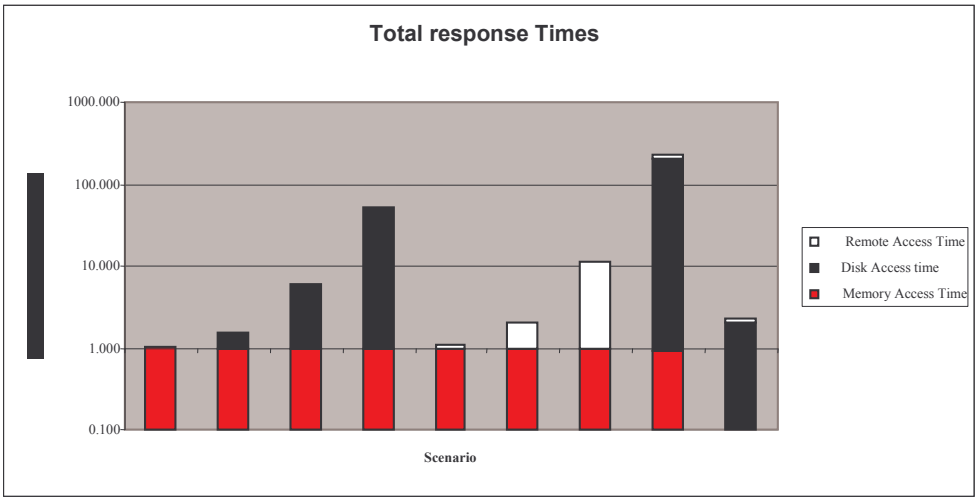

ORACLE

The purpose of this animation sequence is to show what happens when I/O and interconnect block transfers are brought into the response time equation.

Key points

1. **The Y axis is Logarithmic scale**
2. **The reference workload is 1,000,000 buffers. Memory access is set at 1 Microsecond hence response time = 1 Sec**
3. **Disk I/O is introduced by 0.01%, 0.1%, 1% intervals. Note the response time increases**
4. **Disk I/O is eliminated and Interconnect I/O in introduced by 0.01%, 0.1%, 1% Intervals**
   1. Note This is better than disk, Excellent for ERP systems as it reduces I/O
   2. Is also explains why non contentious jobs may take longer on RAC
5. **The Final Bar shows response time with RW I/O(3%) and Interconnect(1%) with optimal HW e.g. 5 ms I/O and 1 ms Interconnect**
6. **Shows the impact of SQL tuning and why it is so important**

Database Performance Basics

Schemas and SQL Statements

- Good schema/data design is the foundation to good database performance.
- Good schema/data design makes downstream activities both simpler and less error prone
  - Selection of Index and Partitioning Columns
  - Gathering and maintenance of schema statistics
  - Writing correct SQL statements
  - Optimization of SQL statements
  - Avoidance of serialization points and transactional bottlenecks

ORACLE

Database Performance Basics

Sessions and Cursors

- The number of sessions within the database should be seen as a design/tuning parameter and NOT a function of the number of active users using the database.
- It is the role of the application server/middleware to manage the connections to the database to ensure consistent db performance and stability.
- High throughput applications should be coded to share SQL statements by use of bind variables and minimize any parsing activity
- Applications should minimize schema object creation/drop/truncate activities

ORACLE

Database Performance Basics

Sessions & Cursors(This slide is over 10 Years old!)

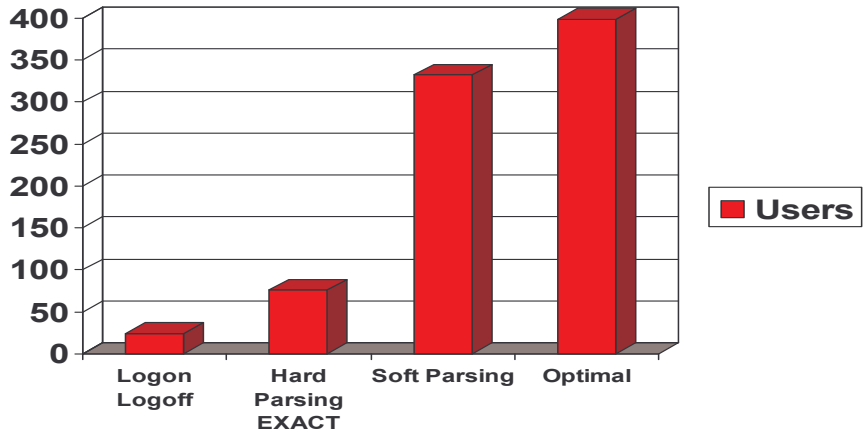

ORACLE

These tests were ran on a 2 CPU machine with 10.2.0.3 The disparity between tests gets worse as you add CPUs

Database Performance Basics

Capacity Planning

- This is a difficult topic as it is often impossible to properly model/test a true production workload.
- It is however possible to increase the probability of success by applying the most basic capacity planning concepts
  - Build a fully balanced system that has no single overloaded HW component.
  - Use intelligent safety margins and factors of ignorance
- Use of intelligent extrapolation from existing systems

ORACLE

Capacity Planning Tips and Ideas Cheat Sheet

CPU and Memory

- For OLTP workloads the CPU utilization should not exceed 65-70%
- Scale Nodes/CPUs/Cores as a power of 2.
- Start with 4Gig Memory per Core and expand according to workload type

I/O

- Target 5-10 millisec response time for disks performing response time critical I/O.
- Start by assuming 30 IOPS per disk for OLTP and 20MBytes per Disk in DSS. This is way below theoretical values but allows for media repair etc.
- You can run HBAs and SAN switches at about 70-90% of theoretical values. No moving parts !
- Allocate and Administer Disk by bandwidth in addition to storage requirements.

Interconnect

- For Data passing intensive DSS Clusters Balance equalize Interconnect bandwidth with I/O bandwidth. Complex queries running across a cluster will perform sorts and joins across the interconnect.
- In the future we may require higher interconnect bandwidth for DW/BI workloads because the data on disk is compressed
- Do not confuse bits for Bytes when capacity planning Interconnect or SAN.

ORACLE

Please not Interconnect requirements for OLTP and DSS are very different. OLTP the focus is on latency for small mesg sizes. For DSS it is all about Bandwidth for larger messages.

Database Performance Basics

Correct Use of The Database Software

- This means using the database software in a manner for which it was designed, developed and tested.
- Use default settings whenever possible for all operations
  - Table creation
  - Statistics gathering
  - Init.ora parameters
- In summary the more you deviate from the defaults the higher the chance of hitting edge conditions.
  - If not on the current release, certainly on the next upgrade.
- Use the appropriate database features correctly and ensure that both application and HW platform is running to full potential.

ORACLE

Today's OLTP Challenges

We Have 2 Types of OLTP Systems

- 1%
  - Single DB Instance per host machine
  - Rigorous system throughput, response time and availability requirements
  - Each Customer is highly visible to Oracle DB development organization
- 99%
  - Multiple Instances per host machine
  - All databases are critical to the business but not running with such rigorous requirements as above
  - The challenge is preventing the anti social instance
  - Minimal visibility to Oracle DB development organization other than via support and bug database.

Today's OLTP Challenges

Predictable Performance

- In many cases getting the initial required performance characteristics of throughput and response time is the easy bit.
- Predictable performance over time is often the hard part. This is a function of:
  - The ability to test changes in configuration such as upgrades and new application logic.
  - SQL plan stability over time
  - Graceful, predictable, debugable behavior on any single component failure
  - SQL statements maintaining identical response times with increasing/changing datasets.

ORACLE

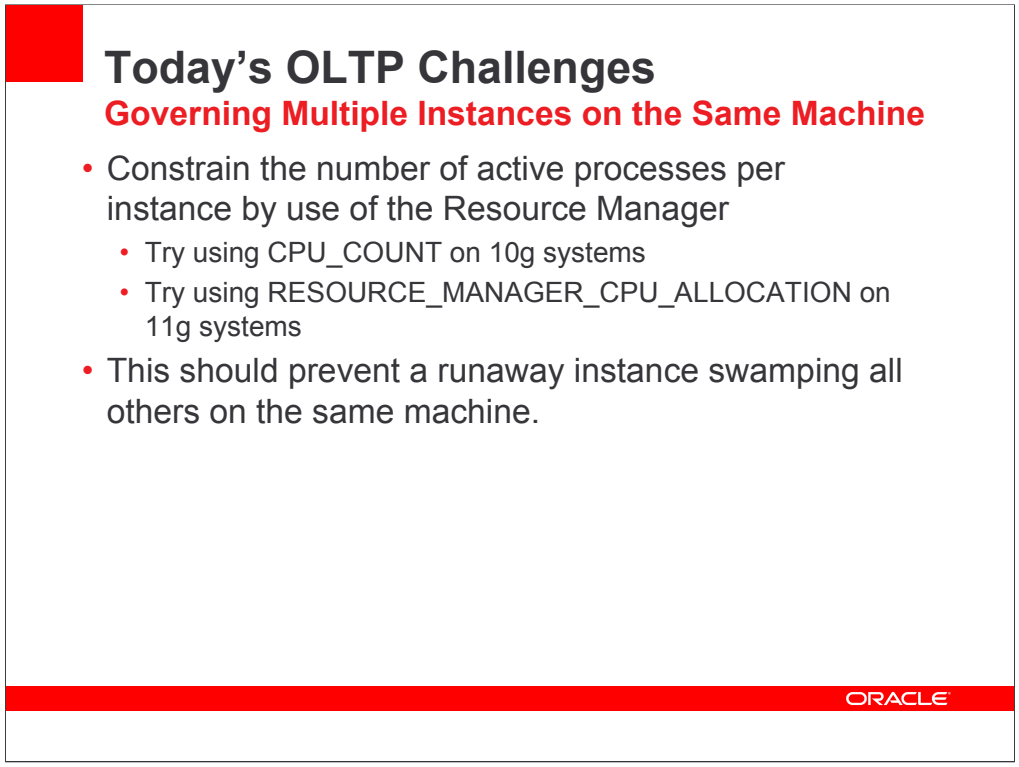

For more information DBRM feel free to check in with DBRM product management

Today's OLTP Challenges

Connection Mangement and Connection Storms

- Connection storms usually occur when an application server is allowed to create new connections to the database when the existing connection pool is exhausted.
- If a system experiences a slowdown for any reason that impacts DB response time this can initiate a feedback loop that can make the number of connections jump from the hundreds to over 10k in a matter of seconds.
- This can render the DB server unstable and may crash it. If however the database server survives the bloat in connections the connection storm interferes so much with the db statistics that it makes root cause analysis of the initial slowdown almost impossible.

Today's DW/BI Challenges

The Capacity Planning/Design Death Spiral

- An initial database is built and deployed in production without a full understanding of correct use of partitioning, parallel query, minimal index design and the correct capacity planning techniques
- By incorrect or poor use of Parallel Query performance is degraded by a factor of 10
- User perception of Query performance is degraded by a factor of 10
- At this point in time the users are not happy so DBAs and consultants try to apply band aids and start building additional indexes, hinting queries etc.

ORACLE

Today's DW/BI Challenges

The Capacity Planning/Design Death Spiral

- After building indexes the query plans are often hinted to use these structures and run sub optimal plans and so degrade performance a further factor of 10
- To compound the agony it is noticed that the CPU utilization on the machine is poor because it is doing continual index range scans. So users launch a large number of jobs to make the CPU look busy.
- The Net of all this is that the query response time is 100-1000 times slower than what should be achieved and the user population is furious and is looking to alternative platforms/solutions

ORACLE

Today's DW/BI Challenges

The Capacity Planning/Design Death Spiral

- Now users are wondering about DB design, platform selection and a competitive benchmark is planned. Corporate politics escalate and vendor relationships degrade.
- At this time the RW Performance team usually gets the call !
- Our approach in most cases is to simplify everything and make sure the HW is balanced, The data is partitioned appropriately, indexes and hints removed and validate the software runs as designed. At this point speed ups of over 100 to 1000 fold is not uncommon.

ORACLE

Benchmarks

TPC Benchmarks

- The public face of Database and HW performance
- The application logic in TPC is implemented in an optimal manner. The applications used in TPC are very simple are in fact over optimized.
- These projects take months to execute
- We see an order of Magnitude more I/O than seen in the Real World (IOPS in TPC-C and I/O Bandwidth in TPC-H)
- The HW used is to optimize the price/performance number. In many cases the HW would be irresponsible to run a highly available database at these transaction rates

ORACLE

I/O rates for TPC-C OLTP benchmark > 100,000 IOPS

I/O rates for TPC-H DSS benchmark > 40 Gigabytes I/O Sec

Benchmarks

Client Benchmarks and Proof of Concept Projects

- Usually executed as part of the sales cycle
- 9 out 10 Benchmark specifications are ill-conceived, are incomplete and bear little relation to the actual required workload.
- In many cases the content and the conduct of the test is designed to support a purchase decision that has already been made. This is often revealed when vendors are restricted to a subset of their database technology.
- This process is lengthy, expensive and extremely stressful for all individuals concerned.
- The Corporate politics, egos involved and the wasted resources and talent required in benchmarks that yield minimal useful engineering data is incredible

ORACLE

Benchmarks

Common Benchmark Mistakes(OLTP)

- Transaction drivers running an invalid transaction mix
- Emulated users running without think time
- Incorrect datasets and input data to transaction drivers
- A focus on db connections as means to achieve concurrency when in reality a connection pool will be used in production
- A benchmark kit that is not repeatable. This common when running concurrent long running batch jobs or use of random numbers to generate workloads.
- Use of Java or C++ threads to emulate users. This does not scale.
- Insufficient application servers and drivers allocated to drive database to its limits.

ORACLE

Benchmarks

Common Benchmark Mistakes(DW/BI)

- Wanting to run 2000 concurrent parallel queries on a 20 CPU machine
- Nonsensical scoring mechanisms that can be manipulated by optimizing a few queries
- The concept of background workloads
- The ready, steady, go approach to query submission
- Badly defined data expansion rules
- Benchmark kits that have not been run in house prior to execution within any benchmark facility.
- Random variables used to seed queries.

Benchmarks

Common Benchmark Mistakes(DW/BI)

- Examples of DW benchmark errors
  - Classification of Queries is very often small, medium, large
    - How is this classification determined ?
      - Length of SQL statement
      - Complexity of statement
      - Production execution frequency
      - Results set returned
      - Actual system resources used
  - We are continually getting benchmarks where the resources required for the small queries is often far exceed the large ones !
    - Really this demonstrates to us that the benchmarking process is flawed and has not been thought through or tested.

ORACLE

Platform Selection

- The topic is actually not that interesting but it seems it gets more focus than fixing bad DB designs and SQL statements !
- However the core elements of building a HW platform are as follows. In some cases the least expensive parts come first !
  - Server Selection
    - SMP
    - Clustering
  - Interconnect
    - Ethernet
    - Infiniband
  - I/O fabric
    - HBAs, Fibre and switches
    - Network cards and switches
  - Storage
    - Disk Arrays

ORACLE

For More Information

search.oracle.com

or

oracle.com

ORACLE

Real World Performance Group Sessions

Session ID: S291301
Session Title: Current Trends in Database Performance
Track: TECHNOLOGY; Business Intelligence ; Database; Enterprise Management; Linux and Open Source Solutions
Room: 102
Date: 2007-11-14
Start Time: 09:45

Session ID: S291306
Session Title: Roundtable: The Real-World Performance Group
Track: TECHNOLOGY; Business Intelligence ; Database; Enterprise Management; Linux and Open Source Solutions
Room: 104
Date: 2007-11-14
Start Time: 11:15

Session ID: S294625
Session Title: Performance 2.0
Track: TECHNOLOGY; Database
Room: Yerba Buena Theatre
Date: 2007-11-14
Start Time: 15:00

Session ID: S291304
Session Title: When to Use the Appropriate Database Technology
Track: TECHNOLOGY; Business Intelligence ; Database; Enterprise Management; Linux and Open Source Solutions
Room: 102
Date: 2007-11-15
Start Time: 08:30

ORACLE


ORACLE®

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.