


---

# Oracle Server Benchmark with In-Memory SQL Processing

Exadata X2-2 half-rack high-capacity

Benchmark Performance Suite Release 8.4 (Build 130630)


June 2013



- 1 Introduction to Server Performance Tests**
- 2 CPU and Server Configuration
- 3 Benchmark Results – In-Memory SQL Operations
- 4 Reviewing Server Benchmark Results

# Server Performance


---



## Why measure Server Performance?


- Applications tend to operate in memory as much as possible to avoid slow I/O operations
  - Some vendors build complete concepts on this idea, e.g. SAP HANA
- Memory capacity of servers has become cheap, less than 10K USD for 128 GByte RAM on x86 systems
- Oracle recognized this trend and provides specific features for in-memory processing
  - Different Cache types for object pinning
  - Parallel SQL for large in-memory objects
- These tests are useful to determine performance capabilities of 2 socket server (Oracle SE vs Oracle EE)

# Server Performance



What is measured?


- Single thread speed
  - Rows per second [rps]
  - Transactions per second [tps]
  - SQL service time [s]
- Maximum throughput of system
  - Rows per second [rps]
  - Transactions per second [tps]
  - Number of Oracle buffer gets (logical I/O) in [bps]
- Scalability
  - Throughput per process for  $n = \{1, 2, 4, 8, \dots, n\}$
- Efficiency of
  - Huge pages and NUMA architectures when using large RAM capacities
  - Virtualization



How is Server Performance measured?

- Benchware Loader pins objects in Oracle SGA and performs different representative database operations on these objects
  - Selection of all rows via full table scan (all rows per SQL)
  - Selection of one random row via primary key (1 row per SQL)
  - Selection of many random rows via secondary key (25 rows per SQL)


# Server Performance



Remarks on other benchmark tools . . .

- SAP, TPC, Swingbench, Hammerora, ...
  - No specific in-memory performance tests

# Server Performance



## Overview of Server performance tests

| Oracle In-Memory Server Performance                     | Test Code for Select | Test Code for Insert | Test Code for Update | Test Code for Delete |
|---|----------------------|----------------------|----------------------|----------------------|
| ▪ All rows, full table scan                             | SRV-11               | 1)                   | 2)                   | 3)                   |
| ▪ Single row, primary key<br>1 hit per SQL statement    | SRV-21               | 1)                   | 2)                   | 3)                   |
| ▪ Multi row, secondary key<br>25 hits per SQL statement | SRV-31               | 1)                   | 2)                   | 3)                   |


<sup>1)</sup> Inserting rows generates massive I/O, we use this scenario for the LGWR stress test (test code DBL-11), but not for server tests.

<sup>2)</sup> Updating rows of in memory tables generates massive I/O, we use this scenario for the DBWR stress test (test code STO-41), but not for server tests.

<sup>3)</sup> SQL delete statements are currently not part of our benchmark methodology.

### Remarks:


- All operations use RAM and cause nearly no I/O operations. Therefore all operations are server bound.
- In some cases cost effective 2 socket servers with Oracle Standard Edition are able to deliver the required performance. These tests are useful to determine the performance border between 2 socket and 2+ socket server.
- Gartner Research Note: Consider Oracle Standard Edition to Reduce Database Management System Costs, 3<sup>rd</sup> March 2010
- In-memory performance numbers may be important when evaluating Oracle TimesTen versus Oracle RDBMS



- 1 Introduction to Server Performance Tests
- 2 CPU and Server Configuration**
- 3 Benchmark Results – In-Memory SQL Operations
- 4 Reviewing Server Benchmark Results



# CPU and Server Configuration




## CPU architecture

| CPU  | Intel Xeon X5675 | System A to compare | System B to compare |
|--|------------------|---------------------|---------------------|
| Launch Date                                      | Q1 2011          |                     |                     |
| Frequency [GHz]                                  | 3.06 - 3.46      |                     |                     |
| #cores per socket                                | 6                |                     |                     |
| #threads per core                                | 2                |                     |                     |
| <b>Performance numbers from other Benchmarks</b> |                  |                     |                     |
| SPECint_base2006 (speed)                         | 40.8             |                     |                     |
| SPECint_base_rate_2006 (throughput)              | 361              |                     |                     |
| Oracle CPU speed in sys.aux_stats\$              | 2795             |                     |                     |

### Remark:


- Oracle has an internal estimation about CPU speed in sys.aux\_stats\$, but none estimation about CPU throughput.

# CPU and Server Configuration




## Server configuration

| Server   | One Exadata X2-2 Database Server | System A to compare | System B to compare |
|--|----------------------------------|---------------------|---------------------|
| CPU type   | Intel Xeon X5675                 |                     |                     |
| #sockets   | 2                                |                     |                     |
| #cores   | 12                               |                     |                     |
| #threads   | 24                               |                     |                     |
| <b>Oracle Licensing</b>  |                                  |                     |                     |
| Oracle core license factor   | x 0.5                            |                     |                     |
| Oracle license cost in USD <small>(list price 25th of June 2013)</small> |                                  |                     |                     |
| ▪ Enterprise Edition (47'500)  | 285'000                          |                     |                     |
| ▪ RAC Option (23'000)  | 138'000                          |                     |                     |
| ▪ Partition Option (11'500)  | 69'000                           |                     |                     |
| ▪ Diagnostic Pack (5'000)  | 30'000                           |                     |                     |
| ▪ Tuning Pack (5'000)  | 30'000                           |                     |                     |




- 1 Introduction to Server Performance Tests
- 2 CPU and Server Configuration
- 3 Benchmark Results – In-Memory SQL Operations**
- 4 Reviewing Server Benchmark Results


# Server Performance



In-memory SQL, full table scan



# Server Performance



## In-memory SQL, full table scan


Exa 2-2 HR

| Run | Tst | Code   | #N | #J  | #T | CPU busy [%] | CPU user [%] | CPU sys [%] | CPU idle [%] | Throughput rows/sec [rps] | Throughput txn/sec [tps] | SQL service time [s] | Buffer lread [bps] | Buffer pread [bps] | Elap time [s] |
|-----|-----|--------|----|-----|----|--------------|--------------|-------------|--------------|---------------------------|--------------------------|----------------------|--------------------|--------------------|---------------|
| 3   | 1   | SRV-11 | 1  | 1   | 1  | 3            | 2            | 0           | 97           | 2.049E+06                 | 1.600E+01                | 6.073E-02            | 1.890E+05          | 2.800E+01          | 122           |
|     | 2   | SRV-11 | 1  | 2   | 1  | 3            | 2            | 0           | 97           | 4.167E+06                 | 3.300E+01                | 5.428E-02            | 1.850E+05          | 0.000E+00          | 120           |
|     | 3   | SRV-11 | 1  | 4   | 1  | 4            | 4            | 0           | 96           | 8.333E+06                 | 6.700E+01                | 5.417E-02            | 3.702E+05          | 0.000E+00          | 120           |
|     | 4   | SRV-11 | 1  | 8   | 1  | 8            | 8            | 0           | 92           | 1.587E+07                 | 1.270E+02                | 5.634E-02            | 7.071E+05          | 0.000E+00          | 126           |
|     | 5   | SRV-11 | 1  | 16  | 1  | 16           | 16           | 0           | 84           | 2.782E+07                 | 2.230E+02                | 6.599E-02            | 1.251E+06          | 0.000E+00          | 133           |
|     | 6   | SRV-11 | 1  | 32  | 1  | 24           | 24           | 0           | 76           | 3.796E+07                 | 3.040E+02                | 9.966E-02            | 1.708E+06          | 0.000E+00          | 135           |
|     | 7   | SRV-11 | 2  | 64  | 1  | 47           | 46           | 0           | 53           | 7.279E+07                 | 5.820E+02                | 1.022E-01            | 3.270E+06          | 0.000E+00          | 136           |
|     | 8   | SRV-11 | 4  | 128 | 1  | 95           | 94           | 1           | 5            | 1.472E+08                 | 1.178E+03                | 1.022E-01            | 6.606E+06          | 0.000E+00          | 136           |


Legend:

- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [rps] rows per second
- [tps] transactions per second
- [bps] blocks per second
- [s] elapsed time in seconds


# Server Performance



In-memory SQL, primary key access, 1 row hit per transaction



# Server Performance



In-memory SQL, primary key access, 1 row hit per transaction


Exa 2-2 HR

| Run | Tst Code  | #N | #J  | #T | CPU busy [%] | CPU user [%] | CPU sys [%] | CPU idle [%] | Throughput rows/sec [rps] | Throughput txn/sec [tps] | SQL service time [s] | Buffer lread [bps] | Buffer pread [bps] | Elap time [s] |
|-----|-----------|----|-----|----|--------------|--------------|-------------|--------------|---------------------------|--------------------------|----------------------|--------------------|--------------------|---------------|
| 1   | 85 SRV-21 | 1  | 1   | 1  | 2            | 1            | 0           | 98           | 1.810E+04                 | 1.810E+04                | 5.520E-05            | 5.436E+04          | 0.000E+00          | 116           |
|     | 86 SRV-21 | 1  | 2   | 1  | 3            | 2            | 1           | 97           | 3.590E+04                 | 3.590E+04                | 5.531E-05            | 1.078E+05          | 0.000E+00          | 117           |
|     | 87 SRV-21 | 1  | 4   | 1  | 5            | 4            | 1           | 95           | 7.059E+04                 | 7.059E+04                | 5.520E-05            | 2.119E+05          | 0.000E+00          | 119           |
|     | 88 SRV-21 | 1  | 8   | 1  | 9            | 7            | 1           | 91           | 1.377E+05                 | 1.377E+05                | 5.533E-05            | 4.132E+05          | 0.000E+00          | 122           |
|     | 89 SRV-21 | 1  | 16  | 1  | 17           | 14           | 2           | 83           | 2.238E+05                 | 2.238E+05                | 6.883E-05            | 6.712E+05          | 0.000E+00          | 126           |
|     | 90 SRV-21 | 1  | 32  | 1  | 25           | 21           | 3           | 75           | 2.630E+05                 | 2.630E+05                | 1.184E-04            | 7.855E+05          | 0.000E+00          | 127           |
|     | 91 SRV-21 | 2  | 64  | 1  | 49           | 42           | 7           | 51           | 5.264E+05                 | 5.264E+05                | 1.180E-04            | 1.570E+06          | 0.000E+00          | 129           |
|     | 92 SRV-21 | 4  | 128 | 1  | 97           | 85           | 13          | 3            | 1.040E+06                 | 1.040E+06                | 1.185E-04            | 3.103E+06          | 0.000E+00          | 129           |


Legend:

- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [rps] rows per second
- [tps] transactions per second
- [bps] blocks per second
- [s] elapsed time in seconds

# Server Performance



In-memory SQL, secondary key access, 25 row hit per transaction





# Server Performance




In-memory SQL, secondary key access, 25 row hit per transaction

Exa 2-2 HR

| Run | Tst Code   | #N | #J  | #T | CPU busy [%] | CPU user [%] | CPU sys [%] | CPU idle [%] | Throughput rows/sec [rps] | Throughput txn/sec [tps] | SQL service time [s] | Buffer lread [bps] | Buffer pread [bps] | Elap time [s] |
|-----|------------|----|-----|----|--------------|--------------|-------------|--------------|---------------------------|--------------------------|----------------------|--------------------|--------------------|---------------|
| 1   | 93 SRV-31  | 1  | 1   | 1  | 2            | 1            | 0           | 98           | 2.371E+05                 | 9.483E+03                | 1.053E-04            | 2.561E+05          | 0.000E+00          | 116           |
|     | 94 SRV-31  | 1  | 2   | 1  | 3            | 2            | 0           | 97           | 4.662E+05                 | 1.864E+04                | 1.055E-04            | 5.035E+05          | 0.000E+00          | 118           |
|     | 95 SRV-31  | 1  | 4   | 1  | 5            | 4            | 1           | 95           | 9.168E+05                 | 3.667E+04                | 1.065E-04            | 9.895E+05          | 0.000E+00          | 120           |
|     | 96 SRV-31  | 1  | 8   | 1  | 10           | 9            | 1           | 90           | 1.726E+06                 | 6.902E+04                | 1.146E-04            | 2.108E+06          | 1.510E+02          | 122           |
|     | 97 SRV-31  | 1  | 16  | 1  | 18           | 17           | 2           | 82           | 2.637E+06                 | 1.055E+05                | 1.499E-04            | 2.848E+06          | 4.760E+02          | 122           |
|     | 98 SRV-31  | 1  | 32  | 1  | 24           | 22           | 2           | 76           | 2.949E+06                 | 1.180E+05                | 2.672E-04            | 3.176E+06          | 1.120E+02          | 123           |
|     | 99 SRV-31  | 2  | 64  | 1  | 47           | 43           | 3           | 53           | 5.861E+06                 | 2.344E+05                | 2.520E-04            | 6.307E+06          | 0.000E+00          | 122           |
|     | 100 SRV-31 | 4  | 128 | 1  | 85           | 79           | 6           | 15           | 1.089E+07                 | 4.354E+05                | 2.301E-04            | 1.172E+07          | 0.000E+00          | 123           |


Legend:

- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [rps] rows per second
- [tps] transactions per second
- [bps] blocks per second
- [s] elapsed time in seconds



- 1 Introduction to Server Performance Tests
- 2 CPU and Server Configuration
- 3 Benchmark Results – In-Memory SQL Operations
- 4 Reviewing Server Benchmark Results**

# Performance Results



## Reviewing Server Performance


|                          | Metric | One Exadata X2-2 Database Server | System A to compare | System B to compare |
|--------------------------|--------|----------------------------------|---------------------|---------------------|
| #cores                   |        | 12                               |                     |                     |
| #threads                 |        | 24                               |                     |                     |
| In-memory SQL operations | Metric | One Exadata X2-2 Database Server | System A to compare | System B to compare |
| Single thread speed      |        |                                  |                     |                     |
| ▪ Full table scan        | [Mrps] | 2.4                              |                     |                     |
| ▪ Primary key access     | [tps]  | 18'100@56μs                      |                     |                     |
| ▪ Secondary key access   | [Mrps] | 0.237                            |                     |                     |
| Throughput               |        |                                  |                     |                     |
| ▪ Full table scan        | [Mrps] | 37.9                             |                     |                     |
| ▪ Primary key access     | [tps]  | 263'000@120μs                    |                     |                     |
| ▪ Secondary key access   | [Mrps] | 2.949                            |                     |                     |

Legend:

[Mrps] million rows per second

[tps] transactions per second

# Performance Results




## Reviewing Server Performance

---

- In-memory full table scan
  - Speed (1 process): 2.4 million rows per second
  - Speed (1 process): < 500 nsec processing time per row
  - Throughput (saturation): 37.9 million rows per second
  - Throughput (saturation): > 10 GByte per second user data
  - Throughput (saturation): < 900 nsec processing time per row

# Performance Results



## Reviewing Server Performance

- In-memory primary key access (1 row hit per transaction)
  - Speed (1 process): 18'100 transactions per second
  - Speed (1 process): 56  $\mu$ sec service time
  - Throughput (saturation): 263'000 transactions per second
  - Throughput (saturation): 120  $\mu$ sec service time
- Times Ten is factor 28 faster for this kind of test

### Remark:

- *TimesTen provides 2  $\mu$ s service time for single row look up on Intel Xeon 5670 with 2.93 GHz*
- *Oracle published some performance numbers in the Oracle TimesTeen data sheet from 2011*  
*<http://www.oracle.com/technetwork/products/timesten/overview/ds-timesten-imdb-129255.pdf?ssSourceSiteId=ocomen>*

**BENCHWARE**

*swiss precision in performance measurement*

*[www.benchmarkware.ch](http://www.benchmarkware.ch)*

*[info@benchmarkware.ch](mailto:info@benchmarkware.ch)*