BENCHWARE

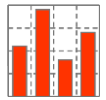

Performance Baseline of Hitachi Data Systems UCP for Oracle

Part II: Server Performance

Benchware Performance Suite Release 8.5 (Build 131015)

October 2013

Contents

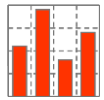

Server Performance

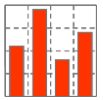

Why measure Server Performance?

- 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
- List price for 16 GByte DIMM's:
  - x86 server: ~ 25'000 USD for 1 TByte
  - Risc server: ~ 55'000 USD for 1 TByte

Remarks:

Currently (September 2013) commercial systems may have following RAM capacities:

- *based on Intel x86 2 TByte RAM*
- *based on Intel Itanium 8 TByte RAM*
- *based on IBM Power 16 TByte RAM*
- *based on Sun SPARC 32 TByte RAM*

Server Performance

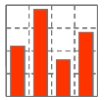

Why measure Server Performance?

- Oracle recognized this trend and provides specific features for in-memory processing
  - Different Cache types for object pinning
  - Parallel SQL even for large in-memory objects
  - New 12c Release 2 In-Memory Option
- These tests are useful to determine performance capabilities of 2 socket server (Oracle SE versus Oracle EE)

Server Performance

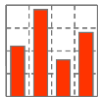

What is measured?

- Server performance from the Oracle point of view
  - No I/O operations
- Speed of single thread
  - Rows per second [rps]
  - Transactions per second [tps]
  - Service time in [s]
- Maximum throughput of system
  - Rows per second [rps]
  - Transactions per second [tps]
  - Service time in [s]
  - Oracle database block gets per second (logical I/O) in [dbps]
- Scalability
  - Throughput per process for n = {1, 2, 4, 8, …, n}
- Efficiency of
  - Huge pages and NUMA architectures when using large RAM capacities
  - Virtualization

Server Performance

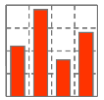

How is Server Performance measured?

- Benchware Loader pins objects in Oracle SGA
- Three typical Oracle transaction profiles
  - Selection of all rows via full table scan (all rows per SQL), e.g. for data analytics
  - Selection of one random row via primary key (1 row per SQL), e.g. searching for bank account, product number, order number
  - Selection of many random rows via secondary key (25 rows per SQL), e.g. part list of order, last 25 transactions of bank account

Server Performance

Overview of Server performance tests with Benchware test codes

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 1 hit per SQL statement	SRV-21	1)	2)	3)
▪ Multi row, secondary key 25 hits per SQL statement	SRV-31	1)	2)	3)

1) Inserting rows generates massive I/O, we use this scenario for the LGWR stress test (test code DBL-11), but not for server tests.
2) 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.
3) SQL delete statements are currently not part of our benchmark suite.

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. Take a look at Gartner Research Note: Consider Oracle Standard Edition to Reduce Database Management System Costs, 3. March 2010*
- *In-memory performance numbers may be important when evaluating Oracle Times Ten versus Oracle RDBMS*

Server Performance

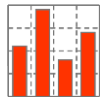

Remarks on other benchmark tools . . .

- TPC, Swingbench, Hammerora, …
  - No specific in-memory performance tests

Contents

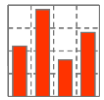

Server Performance

CPU Architecture

CPU	HDS UCP Large SMP	HDS UCP Small SMP
Type	Intel Xeon E7-8870	Intel Xeon E5-2690
Launch Date	2011	2012
Frequency [GHz]	2.4 – 2.8	2.9 – 3.8
#cores per socket	10	8
Multithreading	2-fold	2-fold
Performance Numbers from other Benchmarks	**HDS UCP Large SMP**	**HDS UCP Small SMP**
SPECint_base2006 (speed)	36.4	55.4
Oracle CPU speed in sys.aux_stats$	3'074	2'605

Remark:

Oracle has an internal estimation about CPU speed in sys.aux_stats$, but none estimation about CPU throughput. The Oracle speed estimation does either correlate with SPECint_base2006 numbers nor with Benchware performance results in Oracle 11g Release 2.

Server Performance

Server Configuration

Server	HDS UCP Large SMP	HDS UCP Small SMP
CPU type	Intel Xeon E7-8870	Intel Xeon E5-2690
#sockets	4	2
#cores	40	16
#threads	80	32
#lcpu	-	-
Server Cost	-	-
Performance numbers from other Benchmarks	**HDS UCP Large SMP**	**HDS UCP Small SMP**
SPECint_rate_base2006 (throughput)	40 cores: 1'000	16 cores: 668
Software	**HDS UCP Large SMP**	**HDS UCP Small SMP**
Operating System	Oracle Linux 6.1	RedHat Linux 6.3
Oracle Database System	11.2.0.4	11.2.0.3
Benchware Performance Suite	8.5 Build 131015	8.4 Build 130731

Server Performance

Oracle Licensing

Oracle Enterprise Edition	HDS UCP Large SMP	HDS UCP Small SMP
Oracle core license factor	x 0.5	x 0.5
Oracle license cost (list price 25th of June 2013)		
▪ Enterprise Edition (47'500)	950'000	380'000
▪ Partition Option (11'500)	230'000	92'000
▪ Diagnostic Pack (5'000)	100'000	40'000
▪ Tuning Pack (5'000)	100'000	40'000
Total Oracle license cost	1'380'000	552'000

Contents

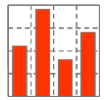

Server Performance

In-memory SQL, full table scan

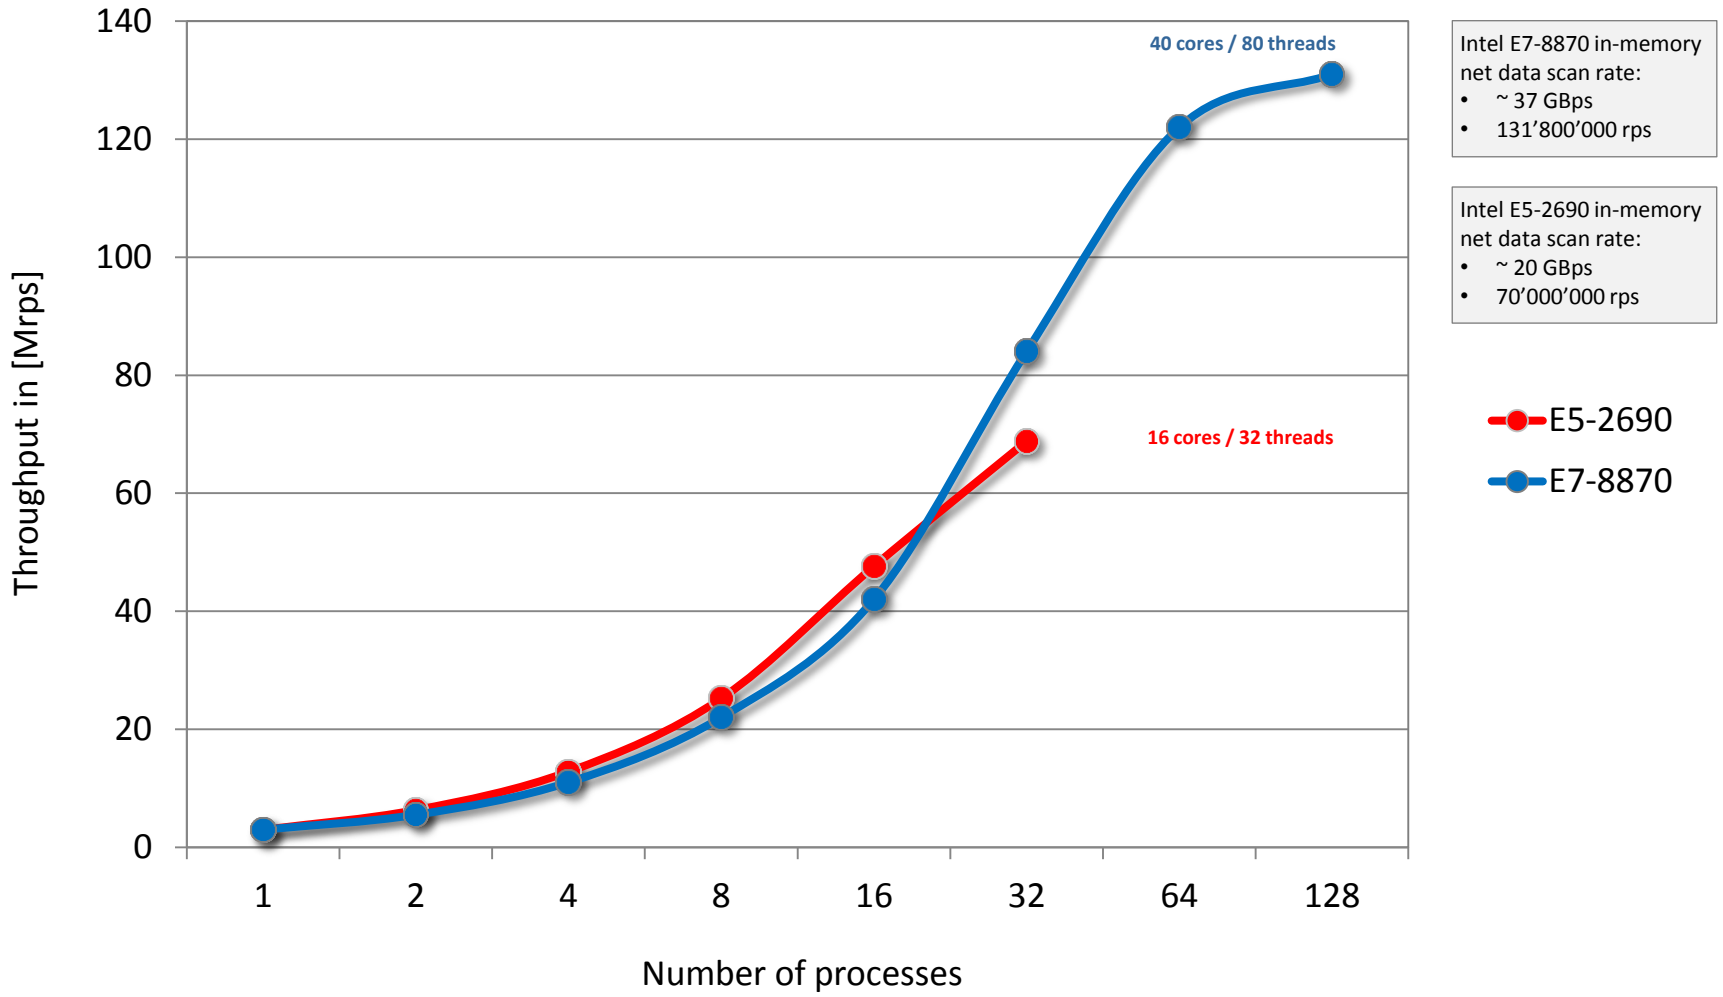

Server Performance

In-memory SQL, full table scan

Intel Xeon E5 2690 2.9 – 3.8 GHz

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]
5	1	SRV-11	1	1	1	4	3	1	96	2.953E+06	2.400E+01	4.240E-02	1.285E+05	0.000E+00	127
	2	SRV-11	1	2	1	7	6	1	93	6.303E+06	5.000E+01	3.848E-02	2.741E+05	0.000E+00	119
	3	SRV-11	1	4	1	13	12	1	87	1.282E+07	1.030E+02	3.815E-02	5.574E+05	0.000E+00	117
	4	SRV-11	1	8	1	25	24	0	75	2.521E+07	2.020E+02	3.878E-02	1.095E+06	0.000E+00	119
	5	SRV-11	1	16	1	49	48	1	51	4.762E+07	3.810E+02	4.071E-02	2.065E+06	0.000E+00	126
	6	SRV-11	1	32	1	94	94	0	6	6.885E+07	5.510E+02	5.615E-02	2.984E+06	0.000E+00	130

Intel Xeon E7 8870 2.4 – 2.8 GHz

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]
12	71	SRV-11	1	1	1	2	1	0	98	2.971E+06	2.400E+01	4.192E-02	1.293E+05	0.000E+00	223
	72	SRV-11	1	2	1	3	3	0	97	5.521E+06	4.400E+01	4.335E-02	2.402E+05	1.000E+00	240
	73	SRV-11	1	4	1	5	5	0	95	1.104E+07	8.800E+01	4.493E-02	4.803E+05	6.000E+00	240
	74	SRV-11	1	8	1	10	10	0	90	2.218E+07	1.770E+02	4.473E-02	9.645E+05	9.000E+00	239
	75	SRV-11	1	16	1	19	19	0	81	4.206E+07	3.370E+02	4.512E-02	1.829E+06	1.700E+01	252
	76	SRV-11	1	32	1	39	39	0	61	8.379E+07	6.700E+02	4.680E-02	3.641E+06	3.300E+01	253
	77	SRV-11	1	64	1	79	79	0	21	1.225E+08	9.800E+02	6.431E-02	5.320E+06	5.400E+01	306
	78	SRV-11	1	128	1	98	97	0	2	1.313E+08	1.050E+03	1.197E-01	5.703E+06	1.128E+03	310

Legend:

#N number of RAC nodes
#J number of jobs
#T number of threads (PX)
[ops] operations per second
[s] elapsed time in seconds

Server Performance

In-memory SQL, primary key access, 1 row hit per transaction (light tx)

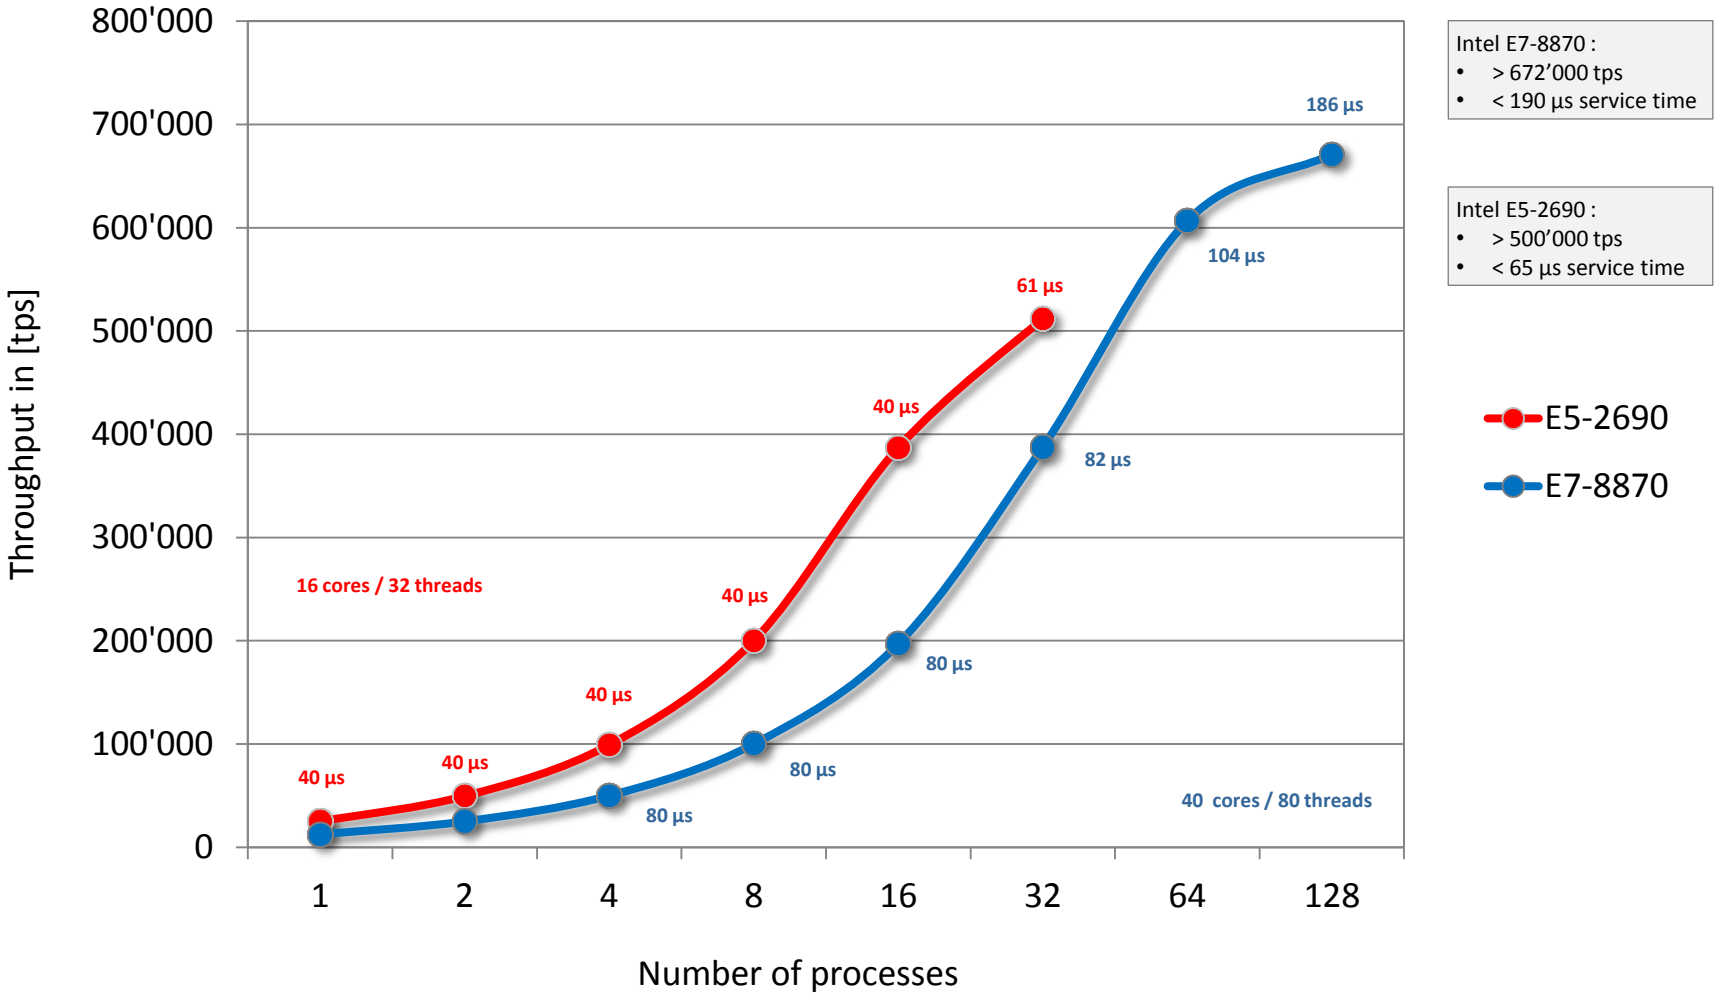

Server Performance

In-memory SQL, primary key access, 1 row hit per transaction (light tx)

Intel Xeon E5 2690 2.9 – 3.8 GHz

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]
4	7	SRV-21	1	1	1	4	3	1	96	2.500E+04	2.500E+04	4.010E-05	7.503E+04	0.000E+00	120
	8	SRV-21	1	2	1	7	6	1	93	4.959E+04	4.959E+04	4.006E-05	1.488E+05	0.000E+00	121
	9	SRV-21	1	4	1	13	12	1	87	9.917E+04	9.917E+04	3.947E-05	2.975E+05	0.000E+00	121
	10	SRV-21	1	8	1	25	23	2	75	2.000E+05	2.000E+05	3.903E-05	5.999E+05	0.000E+00	120
	11	SRV-21	1	16	1	49	45	4	51	3.871E+05	3.871E+05	3.991E-05	1.161E+06	0.000E+00	124
	12	SRV-21	1	32	1	95	88	8	5	5.120E+05	5.120E+05	6.113E-05	1.527E+06	0.000E+00	125

Intel Xeon E7 8870 2.4 – 2.8 GHz

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]
12	79	SRV-21	1	1	1	2	1	0	98	1.254E+04	1.254E+04	7.947E-05	3.765E+04	0.000E+00	303
	80	SRV-21	1	2	1	3	2	0	97	2.508E+04	2.508E+04	7.926E-05	7.527E+04	1.000E+00	303
	81	SRV-21	1	4	1	5	5	1	95	5.000E+04	5.000E+04	7.923E-05	1.500E+05	3.000E+00	303
	82	SRV-21	1	8	1	10	9	1	90	1.002E+05	1.002E+05	7.909E-05	3.005E+05	5.000E+00	303
	83	SRV-21	1	16	1	20	18	2	80	1.970E+05	1.970E+05	8.039E-05	5.912E+05	1.000E+01	305
	84	SRV-21	1	32	1	40	37	3	60	3.875E+05	3.875E+05	8.175E-05	1.163E+06	1.900E+01	305
	85	SRV-21	1	64	1	80	73	7	20	6.067E+05	6.067E+05	1.046E-04	1.810E+06	2.900E+01	305
	86	SRV-21	1	128	1	97	88	8	3	6.719E+05	6.719E+05	1.866E-04	2.009E+06	0.000E+00	310

Legend:

#N number of RAC nodes
#J number of jobs
#T number of threads (PX)
[ops] operations per second
[s] elapsed time in seconds

Server Performance

In-memory SQL, secondary key access, 25 row hits per transaction (heavy tx)

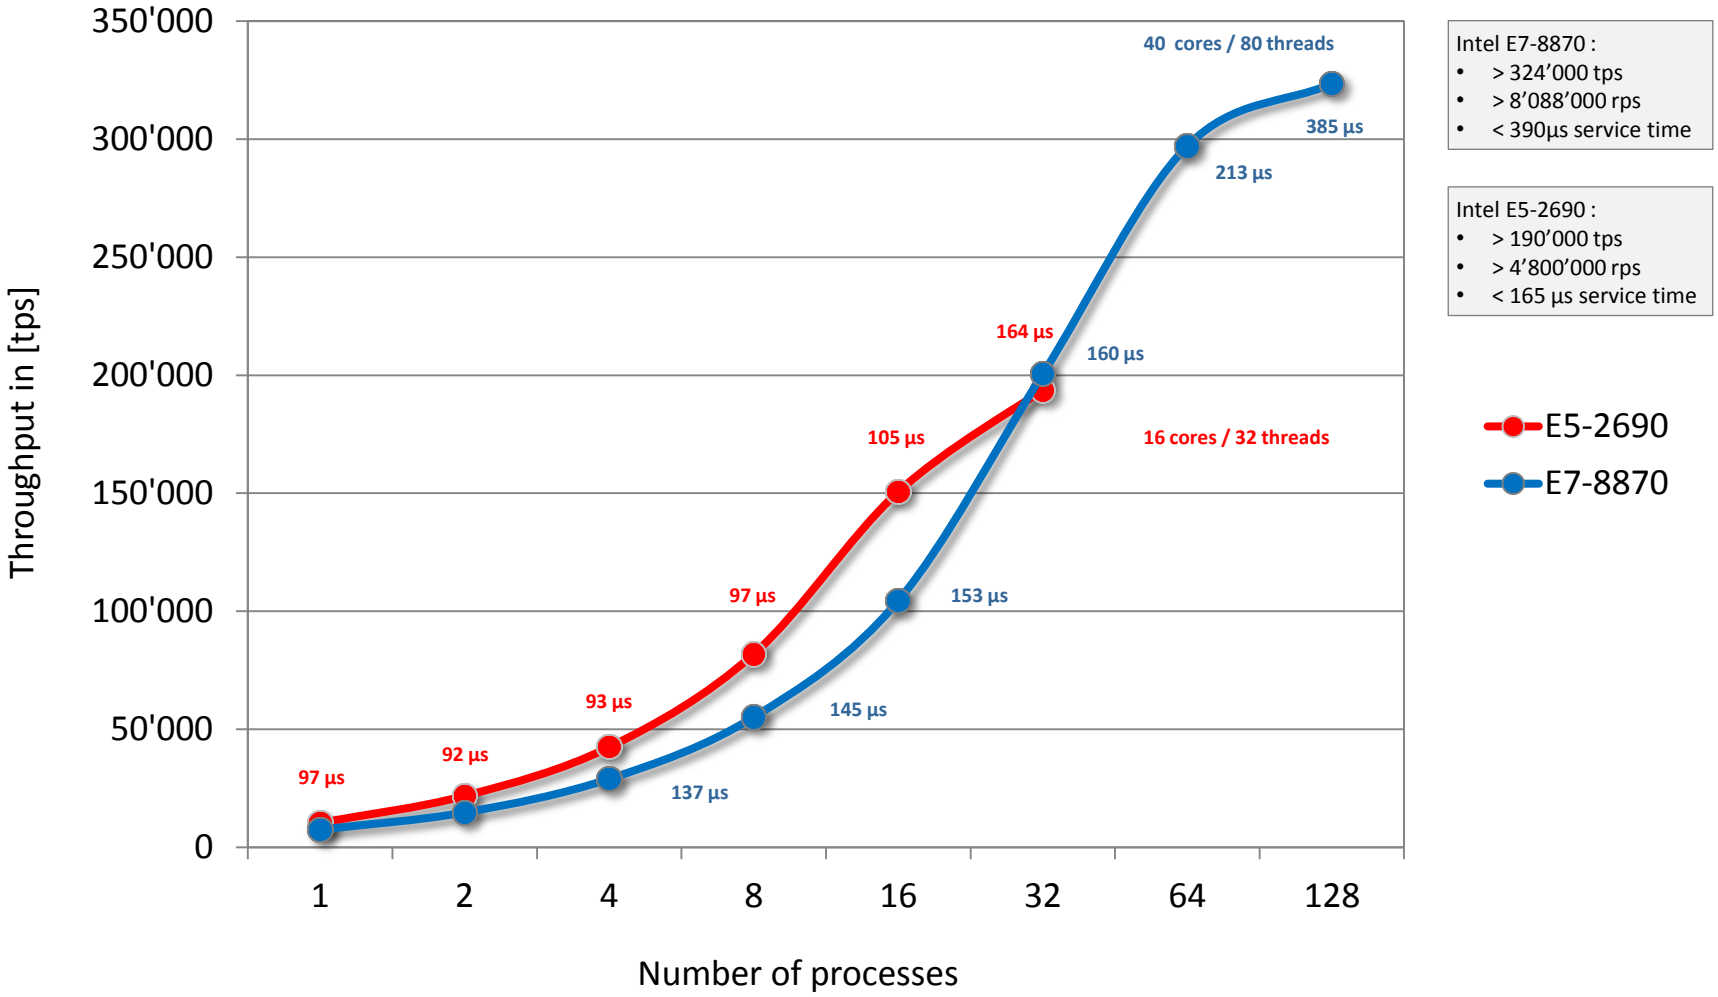

Server Performance

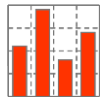

In-memory SQL, secondary key access, 25 row hits per transaction (heavy tx)

Intel Xeon E5 2690 2.9 – 3.8 GHz

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]
4	13	SRV-31	1	1	1	4	3	1	96	2.583E+05	1.033E+04	9.661E-05	2.790E+05	0.000E+00	120
	14	SRV-31	1	2	1	7	6	1	93	5.440E+05	2.175E+04	9.131E-05	5.874E+05	0.000E+00	114
	15	SRV-31	1	4	1	13	12	1	87	1.060E+06	4.239E+04	9.243E-05	1.144E+06	1.000E+01	117
	16	SRV-31	1	8	1	26	24	1	74	2.046E+06	8.182E+04	9.673E-05	2.205E+06	0.000E+00	121
	17	SRV-31	1	16	1	51	48	2	49	3.765E+06	1.506E+05	1.050E-04	4.046E+06	0.000E+00	122
	18	SRV-31	1	32	1	98	95	3	2	4.842E+06	1.937E+05	1.635E-04	5.200E+06	0.000E+00	122

Intel Xeon E7 8870 2.4 – 2.8 GHz

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]
12	87	SRV-31	1	1	1	2	1	0	98	1.860E+05	7.439E+03	1.341E-04	2.008E+05	0.000E+00	285
	88	SRV-31	1	2	1	3	3	0	97	3.681E+05	1.472E+04	1.355E-04	3.975E+05	0.000E+00	288
	89	SRV-31	1	4	1	5	5	0	95	7.236E+05	2.894E+04	1.368E-04	7.814E+05	1.000E+00	293
	90	SRV-31	1	8	1	10	10	1	90	1.378E+06	5.513E+04	1.443E-04	1.488E+06	1.000E+00	302
	91	SRV-31	1	16	1	20	19	1	80	2.610E+06	1.044E+05	1.526E-04	2.818E+06	3.000E+00	302
	92	SRV-31	1	32	1	40	38	2	60	5.012E+06	2.005E+05	1.590E-04	5.405E+06	6.000E+00	302
	93	SRV-31	1	64	1	79	76	4	21	7.428E+06	2.971E+05	2.126E-04	8.007E+06	8.000E+00	305
	94	SRV-31	1	128	1	97	92	5	3	8.088E+06	3.235E+05	3.847E-04	8.722E+06	0.000E+00	310

Legend:

#N number of RAC nodes
#J number of jobs
#T number of threads (PX)
[ops] operations per second
[s] elapsed time in seconds

Contents

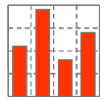

Benchmark Results

Summary Server Performance

	Metric	HDS UCP Large SMP	HDS UCP Small SMP
#cores		40	16
#threads		80	32
In-memory SQL operations	**Metric**	**HDS UCP Large SMP**	**HDS UCP Small SMP**
Single thread speed			
▪ Full table scan	[rps]	2’971’000	2’953’000
▪ Light transaction (primary key access)	[tps]	12’544@80µs	25’000@40µs
▪ Heavy transaction (secondary key access)	[tps]	10’333@96µs	7’439@134µs
	[rps]	186’000	258’300
Throughput			
▪ Full table scan	[rps]	131’300’000	68’850’000
▪ Light transaction (primary key access)	[tps]	671’900@187µs	512’000@61µs
▪ Heavy transaction (secondary key access)	[tps]	193’700@164µs	323’500@385µs
	[rps]	8’088’000	4’842’000

Legend:
[rps] rows per second
[tps] transactions per second

Benchmark Results

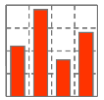

Reviewing Server Performance

- HDS offers its x86 servers with different Intel processors
- This benchmark compares the Oracle in-memory server performance on a smaller 2 socket server versus a larger 8 socket server

Benchmark Results

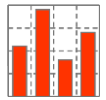

Reviewing Server Performance

- The smaller SMP server with E5-2690 processor
  - Less Oracle license cost (factor 2.5)
  - More speed
  - Less throughput, but not factor 2.5
  - Better transactional service times
- Final conclusion
  - The smaller SMP provides better price/performance ratio

swiss precision in performance measurement

www.benchware.ch

info@benchware.ch