


Performance Baseline of SPARC T5 Oracle Database Server

Part II: Server Performance


Technical Presentation

February 2014




- 1 Introduction to Server Performance Tests**
- 2 CPU and Server Configuration
- 3 Server Benchmark Results – In-Memory SQL Operations
- 4 Reviewing CPU and 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
- Server support large memory capacities
 - Intel x86 2 TByte RAM
 - Intel Itanium 8 TByte RAM
 - IBM Power 16 TByte RAM
 - Oracle SPARC 32 TByte RAM
- Memory has become cheap, list price for 1 TByte RAM
 - x86 server: ~ 25'000 USD using 16 GByte DIMM
 - x86 server: ~ 60'000 USD using 32 GByte DIMM
 - RISC server: ~ 55'000 USD using 16 GByte DIMM




Why measure Server Performance?

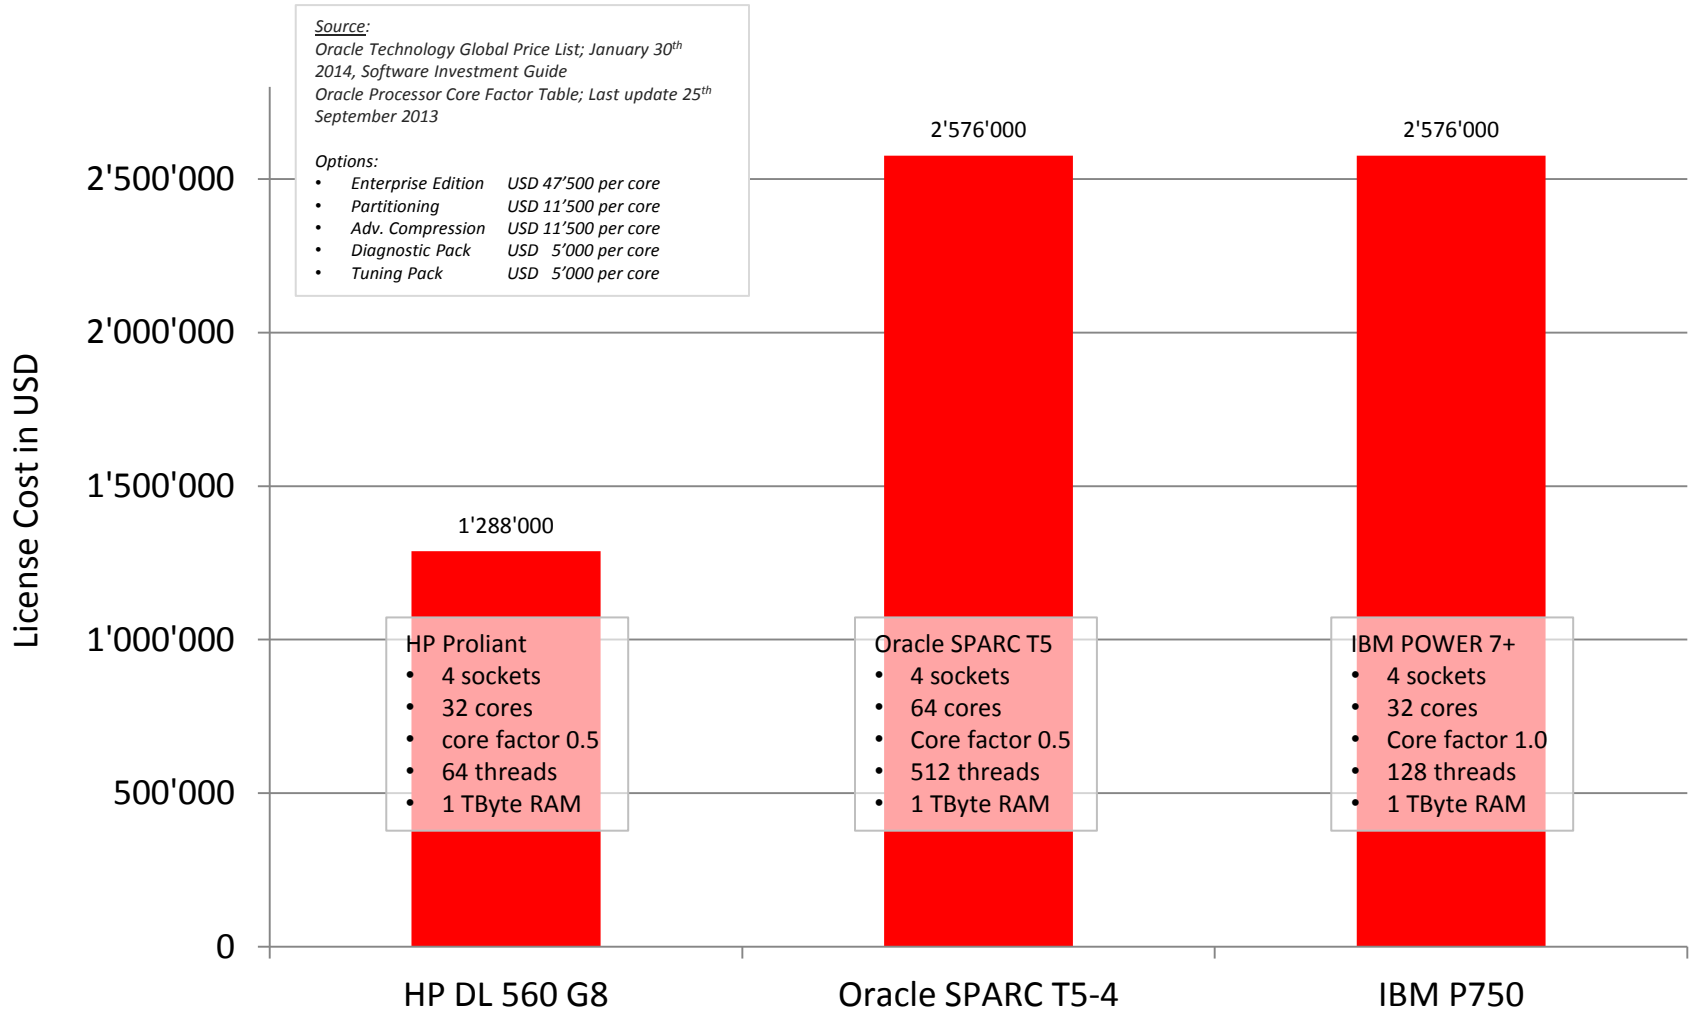
- Oracle recognized this trend and provides specific features for in-memory processing
 - Different Cache types for object pinning
 - Parallel SQL for large in-memory objects
 - New 12c Release 2 In-Memory Option
- These tests are useful to determine performance capabilities of 2 socket server (Oracle SE versus Oracle EE) ¹⁾
 - EE is based on core licensing
 - SE is based on socket licensing, but limited to 4 sockets
 - SE1 is based on socket licensing, but limited to 2 sockets

¹⁾ *Feinberg, Disbrow: Consider Oracle Standard Edition to Reduce Database Management Systems Costs, Gartner RAC Core Research Note G00174599, 3 March 2010*

Server Performance



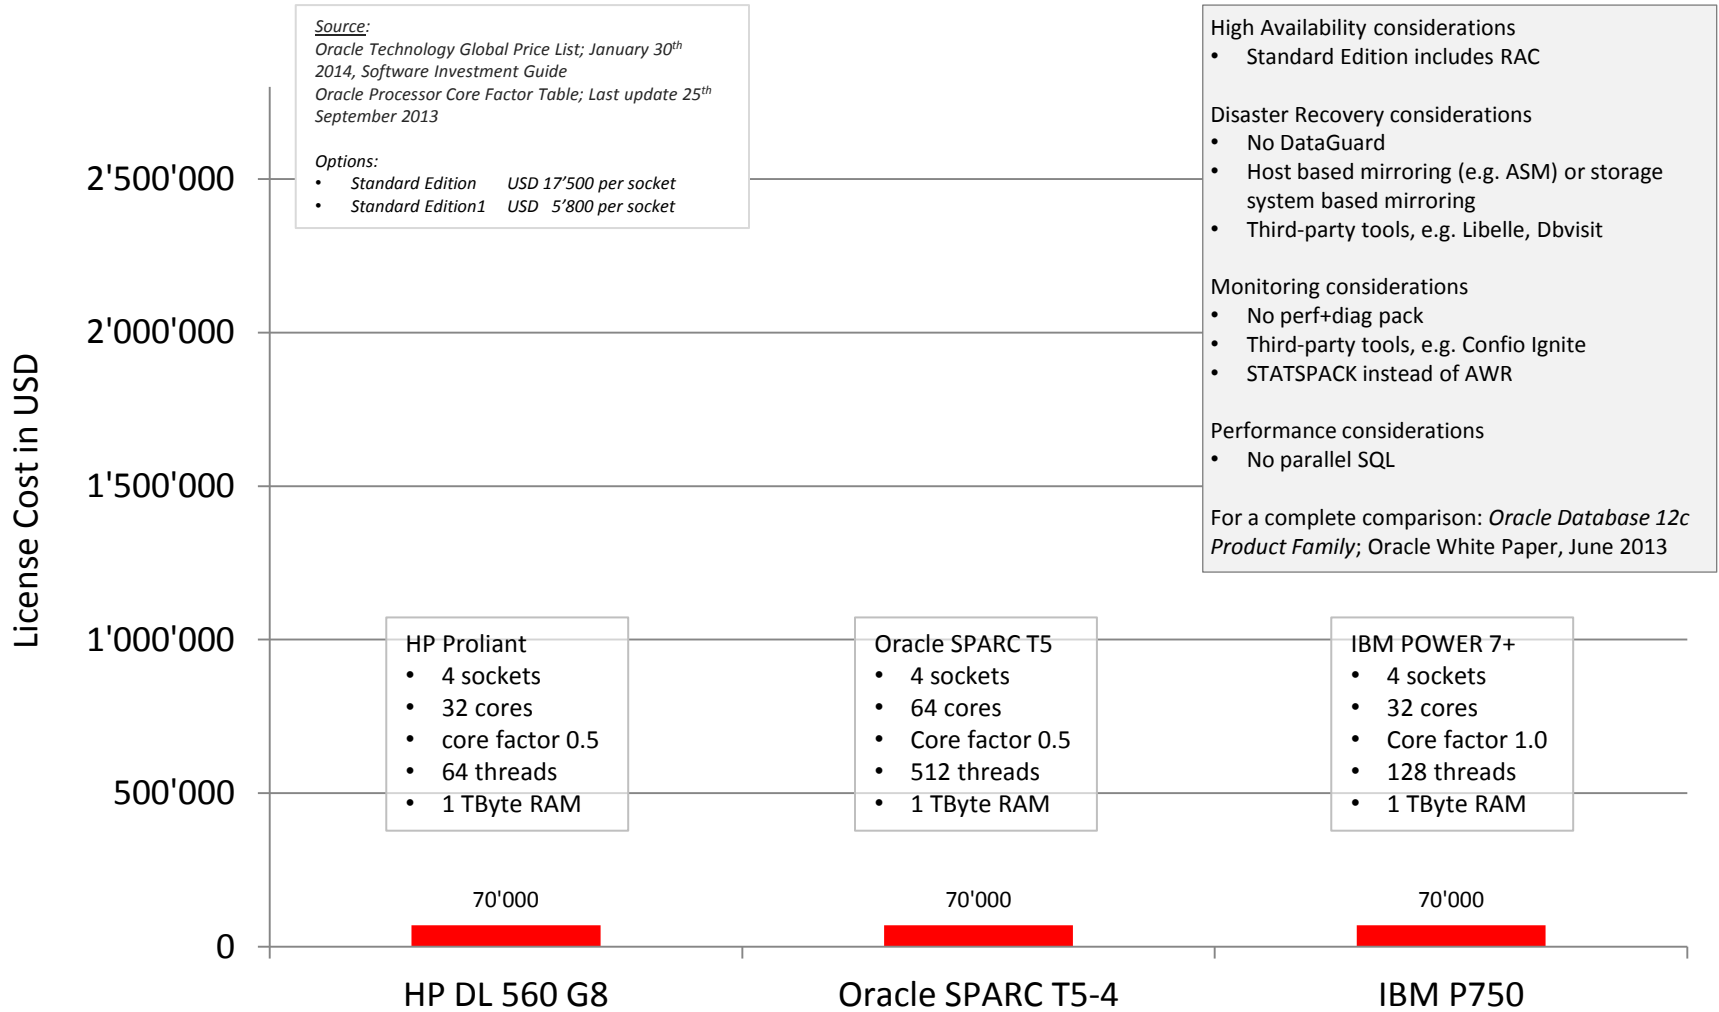
Oracle EE Core Licensing – Price Performance Ratio?




Server Performance



Oracle SE Socket Licensing – Price Performance Ratio?




Server Performance



What is measured?


- Server performance from the Oracle point of view
 - No I/O operations
 - Oracle SQL transactions on objects pinned in KEEP cache
- 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 $n = \{1, 2, 4, 8, \dots, n\}$
- 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
- 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 (Ø 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 Server Performance	Test Code for	Test Code for	Test Code for	Test Code for
Server-bound SQL database transactions on in-memory data objects - no I/O operations	SELECT	INSERT	UPDATE	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-12), 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-42), 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



Monitoring

■ CPU utilization, speed and throughput

Speed:

- Only 1 process
- No conflicts
- No contention

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 [dbps]	Buffer pread [dbps]	Elap time [s]
6	28	SRV-31	1	1	1	1	1	0	99	2.413E+05	9.652E+03	1.034E-04	2.608E+05	1.000E+00	302
	29	SRV-31	1	2	1	1	1	0	99	4.690E+05	1.876E+04	1.057E-04	5.066E+05	1.000E+00	302
	30	SRV-31	1	4	1	2	1	0	98	8.941E+05	3.576E+04	1.112E-04	9.654E+05	2.000E+00	303
	31	SRV-31	1	8	1	3	2	1	97	1.706E+06	6.825E+04	1.158E-04	1.842E+06	4.000E+00	303
	32	SRV-31	1	16	1	5	4	1	95	3.480E+06	1.392E+05	1.138E-04	3.751E+06	8.000E+00	303
	33	SRV-31	1	32	1	9	8	1	91	5.537E+06	2.214E+05	1.423E-04	5.957E+06	1.500E+01	304
	34	SRV-31	1	64	1	13	12	1	87	6.424E+06	2.568E+05	2.431E-04	6.901E+06	3.000E+01	306
	35	SRV-31	2	64	1	17	16	2	83	1.127E+07	4.507E+05	1.397E-04	1.213E+07	3.000E+01	304
	36	SRV-31	2	128	1	25	23	2	75	1.311E+07	5.243E+05	2.381E-04	1.409E+07	5.900E+01	306
	37	SRV-31	4	128	1	34	31	3	66	2.274E+07	9.094E+05	1.383E-04	2.448E+07	6.000E+01	305
	38	SRV-31	4	256	1	50	45	4	50	2.659E+07	1.064E+06	2.346E-04	2.860E+07	8.900E+01	306
	39	SRV-31	8	256	1	67	61	6	33	4.600E+07	1.840E+06	1.371E-04	4.952E+07	6.000E+01	304
	40	SRV-31	8	512	1	98	90	8	2	5.346E+07	2.138E+06	2.330E-04	5.749E+07	5.900E+01	307


Legend:

- Run benchmark run id
- Tst benchmark test id
- Code benchmark test code
- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [rps] rows per second
- [tps] transactions per second
- [dbps] database blocks per second
- [s] elapsed time in seconds

Max throughput:


- All cpu resources are utilized

- Column 'buffer lread' describes the number of logical I/O in the buffer cache KEEP
- Colum 'buffer pread' describes the number of physical I/O to the buffer cache KEEP – should be zero for best performance
- Column 'SQL service time' describes response time at application level



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

Server Performance




CPU architecture

CPU		SPARC T5-2	SPARC T5-4
Launch date		2013	2013
Clock rate [GHz]		3.6	3.6
Max number of sockets		2	4
#cores per socket		16	16
Multithreading		8-fold	8-fold
Performance Numbers from other Benchmarks			
SPECint_base2006 (speed)		No SPEC results available	
Oracle CPU speed in sys.aux_stats\$			

Remark:

- Oracle has an internal estimation about CPU speed in sys.aux_stats\$, but none estimation about CPU throughput.
- This value does not correlate with SPECint_base2006


Server Performance



Server configuration

Server		SPARC T5-2	SPARC T5-4 Cluster
#sockets		2	8
#cores		32	128
#threads		256	1'024
Comment		-	Equivalent to T5-8 Equivalent to SuperCluster HR
Performance Numbers from other Benchmarks			
SPECint_base_rate_2006 (throughput)		16 cores: 441	3'490
Software			
Operating System		Solaris 11	Solaris 11
Oracle Database System		11.2.0.3	11.2.0.3
Benchware Performance Suite			


Server Performance



Oracle Licensing


Oracle Enterprise Edition		SPARC T5-2	SPARC T5-4 Cluster
Oracle core license factor		x 0.5	x 0.5
Oracle license cost <small>(list price 30th of January 2014)</small>			
▪ Enterprise Edition (47'500)		760'000	3'040'000
▪ Partition Option (11'500)		184'000	736'000
▪ Diagnostic Pack (5'000)		80'000	320'000
▪ Tuning Pack (5'000)		80'000	320'000
Total Oracle license cost		1'104'000	4'416'000

Oracle Standard Edition		SPARC T5-2	SPARC T5-4 Cluster
Oracle core license factor		x 0.5	x 0.5
Oracle license cost <small>(list price 30th of January 2014)</small>			
▪ Standard Edition (17'500)		560'000	-
Total Oracle license cost		560'000	-




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

Server Performance



In-memory SQL, full table scan



Server Performance



In-memory SQL, full table scan


SPARC
T5-2

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 [dbps]	Buffer pread [dbps]	Elap time [s]
2	92	SRV-11	1	1	1	0	0	0	100	2.436E+06	1.900E+01	5.113E-02	1.060E+05	0.000E+00	118
	93	SRV-11	1	2	1	1	1	0	99	4.873E+06	3.900E+01	5.118E-02	2.120E+05	0.000E+00	118
	94	SRV-11	1	4	1	2	2	0	98	9.091E+06	7.300E+01	5.407E-02	3.953E+05	0.000E+00	121
	95	SRV-11	1	8	1	3	3	0	97	1.818E+07	1.450E+02	5.427E-02	7.894E+05	0.000E+00	121
	96	SRV-11	1	16	1	6	6	0	94	3.468E+07	2.770E+02	5.621E-02	1.504E+06	0.000E+00	124
	97	SRV-11	1	32	1	12	12	0	88	6.457E+07	5.170E+02	6.019E-02	2.798E+06	0.000E+00	127
	98	SRV-11	1	64	1	24	24	0	76	9.793E+07	7.830E+02	7.895E-02	4.249E+06	0.000E+00	127
	99	SRV-11	1	128	1	49	48	0	51	1.389E+08	1.111E+03	1.119E-01	6.014E+06	0.000E+00	127


SPARC Cluster
T5-4

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 [dbps]	Buffer pread [dbps]	Elap time [s]
7	91	SRV-11	2	1	1	0	0	0	100	1.823E+06	1.500E+01	6.825E-02	8.060E+04	0.000E+00	192
	92	SRV-11	2	2	1	0	0	0	100	3.084E+06	2.500E+01	8.059E-02	1.373E+05	0.000E+00	227
	93	SRV-11	2	4	1	1	1	0	99	6.114E+06	4.900E+01	7.901E-02	2.712E+05	0.000E+00	229
	94	SRV-11	2	8	1	1	1	0	99	1.223E+07	9.800E+01	7.751E-02	5.429E+05	0.000E+00	229
	95	SRV-11	2	16	1	2	2	0	98	2.445E+07	1.960E+02	7.847E-02	1.090E+06	0.000E+00	229
	96	SRV-11	2	32	1	3	3	0	97	4.807E+07	3.850E+02	7.876E-02	2.142E+06	0.000E+00	233
	97	SRV-11	2	64	1	6	6	0	94	9.532E+07	7.630E+02	7.947E-02	4.244E+06	0.000E+00	235
	98	SRV-11	2	128	1	12	12	0	88	1.785E+08	1.428E+03	8.297E-02	7.943E+06	0.000E+00	251
	99	SRV-11	2	256	1	25	24	0	75	2.688E+08	2.150E+03	1.147E-01	1.194E+07	0.000E+00	317
	100	SRV-11	2	512	1	47	46	1	53	4.152E+08	3.322E+03	1.420E-01	1.844E+07	0.000E+00	332

Server Performance



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



Server Performance



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


SPARC
T5-2

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 [dbps]	Buffer pread [dbps]	Elap time [s]
2	101	SRV-21	1	1	1	0	0	0	100	1.571E+04	1.571E+04	6.344E-05	4.716E+04	0.000E+00	105
	102	SRV-21	1	2	1	1	1	0	99	3.084E+04	3.084E+04	6.425E-05	9.254E+04	0.000E+00	107
	103	SRV-21	1	4	1	2	1	0	98	6.055E+04	6.055E+04	6.546E-05	1.817E+05	0.000E+00	109
	104	SRV-21	1	8	1	3	3	0	97	1.189E+05	1.189E+05	6.613E-05	3.568E+05	0.000E+00	111
	105	SRV-21	1	16	1	6	6	0	94	2.336E+05	2.336E+05	6.679E-05	7.008E+05	0.000E+00	113
	106	SRV-21	1	32	1	12	11	0	88	4.264E+05	4.264E+05	7.062E-05	1.279E+06	0.000E+00	123
	107	SRV-21	1	64	1	25	24	1	75	7.472E+05	7.472E+05	8.426E-05	2.234E+06	0.000E+00	124
	108	SRV-21	1	128	1	49	47	2	51	1.165E+06	1.165E+06	1.072E-04	3.483E+06	0.000E+00	125

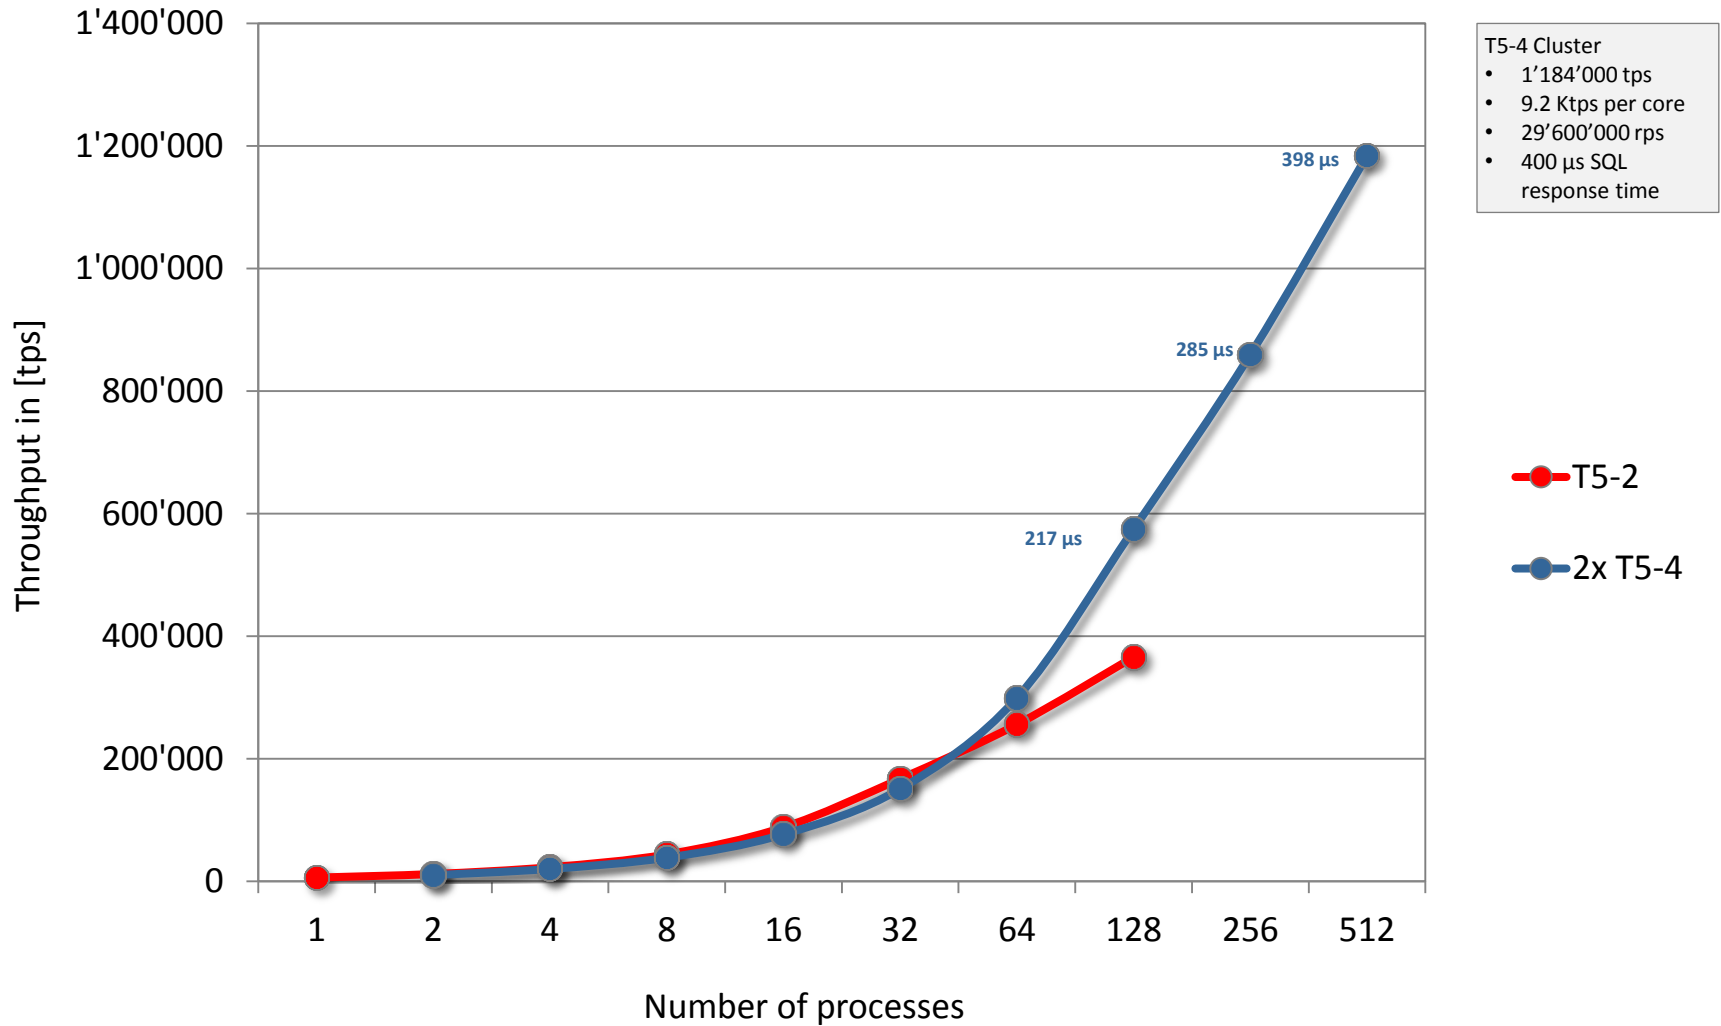
SPARC Cluster
T5-4

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 [dbps]	Buffer pread [dbps]	Elap time [s]
7	101	SRV-21	2	2	1	0	0	0	100	2.655E+04	2.655E+04	7.511E-05	7.970E+04	0.000E+00	275
	102	SRV-21	2	4	1	1	0	0	99	5.233E+04	5.233E+04	7.522E-05	1.571E+05	1.000E+00	279
	103	SRV-21	2	8	1	1	1	0	99	1.028E+05	1.028E+05	7.593E-05	3.085E+05	2.000E+00	284
	104	SRV-21	2	16	1	2	2	0	98	2.049E+05	2.049E+05	7.688E-05	6.148E+05	4.000E+00	285
	105	SRV-21	2	32	1	3	3	0	97	3.973E+05	3.973E+05	7.780E-05	1.192E+06	9.000E+00	294
	106	SRV-21	2	64	1	6	6	0	94	7.892E+05	7.892E+05	7.924E-05	2.367E+06	1.700E+01	296
	107	SRV-21	2	128	1	13	12	1	87	1.527E+06	1.527E+06	8.228E-05	4.576E+06	3.200E+01	305
	108	SRV-21	2	256	1	24	23	1	76	2.516E+06	2.516E+06	9.766E-05	7.533E+06	6.200E+01	315
	109	SRV-21	2	512	1	47	44	3	53	3.805E+06	3.805E+06	1.242E-04	1.139E+07	7.000E+01	328

Server Performance



In-memory SQL, secondary key access, \emptyset 25 rows per transaction (heavy tx)



Server Performance




In-memory SQL, secondary key access, Ø 25 rows per transaction (heavy tx)

SPARC
T5-2

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 [dbps]	Buffer pread [dbps]	Elap time [s]
2	110	SRV-31	1	1	1	0	0	0	100	1.571E+05	6.283E+03	1.590E-04	1.696E+05	0.000E+00	113
	111	SRV-31	1	2	1	1	1	0	99	3.035E+05	1.214E+04	1.646E-04	3.277E+05	0.000E+00	117
	112	SRV-31	1	4	1	2	2	0	98	5.805E+05	2.322E+04	1.698E-04	6.266E+05	0.000E+00	121
	113	SRV-31	1	8	1	3	3	0	97	1.103E+06	4.410E+04	1.790E-04	1.190E+06	0.000E+00	122
	114	SRV-31	1	16	1	6	6	0	94	2.207E+06	8.828E+04	1.798E-04	2.380E+06	0.000E+00	122
	115	SRV-31	1	32	1	12	12	0	88	4.185E+06	1.674E+05	1.898E-04	4.506E+06	0.000E+00	122
	116	SRV-31	1	64	1	25	24	0	75	6.403E+06	2.561E+05	2.461E-04	6.889E+06	0.000E+00	123
	117	SRV-31	1	128	1	49	48	1	51	9.157E+06	3.662E+05	3.392E-04	9.827E+06	0.000E+00	125


SPARC Cluster
T5-4

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 [dbps]	Buffer pread [dbps]	Elap time [s]
7	110	SRV-31	2	2	1	0	0	0	100	2.572E+05	1.029E+04	1.911E-04	2.779E+05	0.000E+00	278
	111	SRV-31	2	4	1	1	0	0	99	5.108E+05	2.043E+04	1.936E-04	5.517E+05	1.000E+00	280
	112	SRV-31	2	8	1	1	1	0	99	9.696E+05	3.878E+04	2.005E-04	1.047E+06	2.000E+00	295
	113	SRV-31	2	16	1	2	2	0	98	1.920E+06	7.678E+04	2.012E-04	2.073E+06	4.000E+00	298
	114	SRV-31	2	32	1	3	3	0	97	3.774E+06	1.510E+05	2.064E-04	4.074E+06	7.000E+00	303
	115	SRV-31	2	64	1	6	6	0	94	7.483E+06	2.993E+05	2.094E-04	8.068E+06	1.300E+01	304
	116	SRV-31	2	128	1	13	12	0	87	1.437E+07	5.747E+05	2.176E-04	1.546E+07	2.500E+01	308
	117	SRV-31	2	256	1	24	24	1	76	2.149E+07	8.596E+05	2.857E-04	2.307E+07	4.900E+01	314
	118	SRV-31	2	512	1	47	45	1	53	2.960E+07	1.184E+06	3.977E-04	3.177E+07	2.700E+01	328



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

Benchmark Results




Reviewing Server Performance

	Metric		SPARC T5-2	SPARC T5-4 Cluster
#cores			32	128
#threads			256	1'024
In-memory SQL operations	Metric		SPARC T5-2	SPARC T5-4 Cluster
Single thread speed				
▪ Full table scan	[Mrps]		2.436	1.823
▪ Light transaction (primary key access)	[Ktps]		15.710@63µs	26.550@75µs
▪ Heavy transaction (secondary key access)	[Ktps]		6.283@160µs	10.290@191µs
	[rps]		157'000	257'000
Throughput				
▪ Full table scan	[Mrps]		138.900	415.200
▪ Light transaction (primary key access)	[Ktps]		1'165@107µs	3'805@124µs
▪ Heavy transaction (secondary key access)	[Ktps]		363@339µs	1'184@397µs
	[Mrps]		9.157	29.600

Legend:

[rps] rows per second [Krps] thousand rows per second [Mrps] million rows per second
 [tps] transactions per second [Ktps] thousand transactions per second


Benchmark Results



Reviewing Server Performance

- Chip performance is described by two characteristics
 - Speed
 - Throughput
- SPARC offers
 - Almost competitive speed within the market of scalable servers (4, 8 or more sockets)
 - Very high throughput
- SPARC servers are ideal for in-memory processing because of its large memory capacities

Benchmark Results



Reviewing Server Performance

- The high number of cores per socket may lead to high Oracle license cost
 - The T5-2 can be operated with Standard Edition
 - All other models need Enterprise Edition

BENCHWARE

swiss precision in performance measurement

www.benchmarkware.ch

info@benchmarkware.ch