BENCHWARE

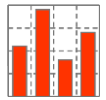

Oracle Performance
on
SPARC T4 versus SPARC VII

Benchmark Report

August 2012

Contents

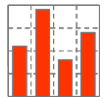

Benchware Ltd

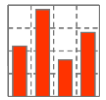

Services and Products

Strong foundation in core technologies like
Oracle database system, server and storage systems

- System Architecture, Component Evaluation, Reviews
- Performance Analysis & Optimization
- Benchmarking
- Database engineering

Benchware Ltd

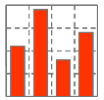

Value proposition

- Vendor-independent company
  - Benchware is completely committed to customers' interests
- Holistic approach in designing, tuning and benchmarking Oracle systems
- Long experience track record
  - Responsible for system architecture of largest DWH and OLTP systems, mainly telecom and finance industry
  - Oracle since 1984 (Oracle Version 3)
  - Performance tuning and benchmarking since 1993 (Oracle Version 7)

Benchware Ltd

Complex architecture of Oracle platforms needs benchmarking

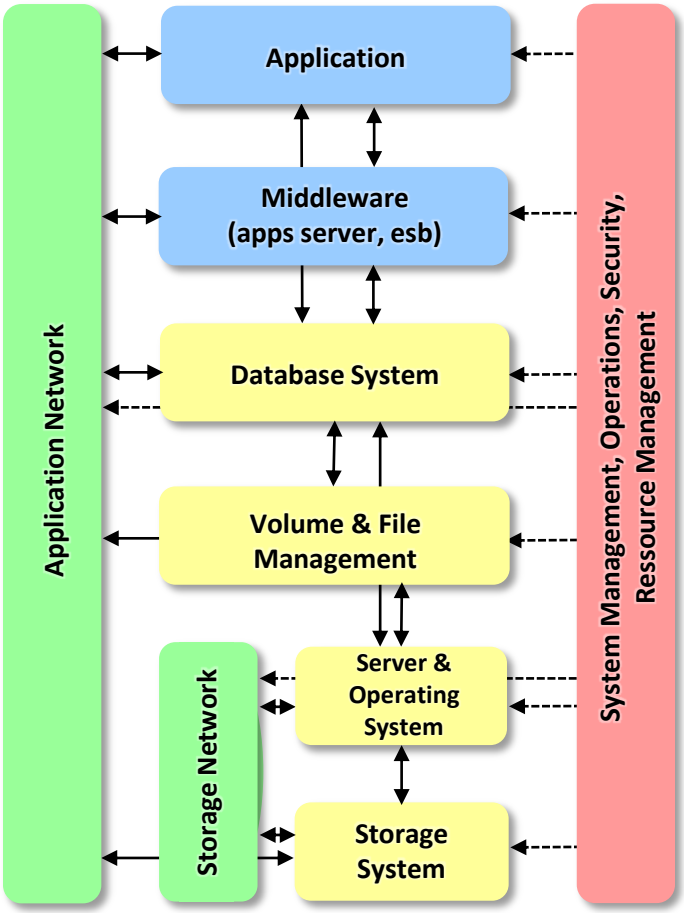

Performance of complex technology stack is NOT predictable – unless running a benchmark

Application Network (IP-based)

Bandwidth, latency during remote database mirroring (sync, async) due to switches and sql*net and tcp/ip stack (frame size, …).

Oracle Database

Different versions, patches and options, about hundred configuration parameters.

Storage Network (FC-, IB- or IP-based)

Bandwidth, latency during remote storage mirroring (sync, async) due to switches, hubs and distance.

Volume & File Management

Different volume managers (VxVM, ASM) and file Systems (UFS, VxFS, ext3, JFS, ZFS, raw devices), different I/O methods (async, direct), a lot of config parameters (#LUNS, queue depth, max i/o unit), software striping and/or mirroring, multipathing.

Server & Operating System

Different server Systems, processors and CPU architectures, (x86, IA-64, UltraSparc, SPARC64, Power), #cores, multithreading, main memory, bus architecture. Different operating Systems and patches, over hundred configuration parameters, virtualization of resources.

Storage System

Different storage Systems, storage tiers and storage technology: spindle count and speed, RAID management, cache management, server interface technology, storage system options like remote copy, hardware striping and/or mirroring, virtualization of resources.

Benchware Ltd

Benchware Performance Suite

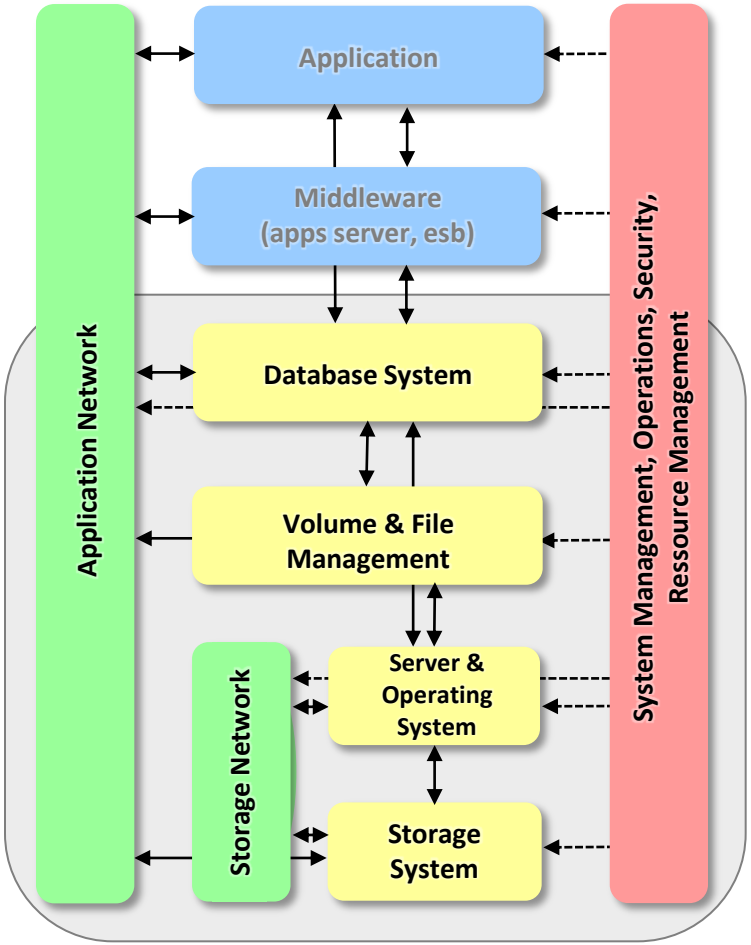

Object of measurement

- Benchware Performance Suite
  - Benchware Monitor
  - Benchware Loader
- Performance measurement at the interface between application and technology stack
- Key Performance Metrics can be used for SLA between IT operation and business
- Benchware uses Oracle Database stack to generate all kind of loads for cpu, server, storage and database

Benchware Ltd

Library of Oracle benchmark tests - implemented in PL/SQL, Java and SQL

CPU Performance CPU-bound Oracle operations All operations in Level 1, 2, 3 CPU cache	OLTP systems	DWH systems	Efficiency	Metrics	Unit
• pl/sql basic operations	★★	★★	multithreading virtualization	speed throughput	[s] [ops]
• pl/sql algorithms fibonacci, prime numbers	★★★	★★			

Server Performance Server-bound Oracle operations All operations in RAM - no I/O operations	OLTP systems	DWH systems	Efficiency	Metrics	Unit
• in-memory SQL	★★★	★★	scalability cc-numa virtualization	speed throughput	[µs] [s] [bps] [tps] [rps]
• pl/sql algorithms quicksort	★★	★			

[s]	seconds	[bps]	buffers per second	[MBps]	mega bytes per second	★	less important
[ms]	milli seconds (10^{-3})	[rps]	rows per second	[GBps]	giga bytes per second	★★	important
[µs]	micro seconds (10^{-6})	[tps]	transactions per second	[iops]	i/o operations per second	★★★	very important
[ns]	nano seconds (10^{-9})	[ops]	operations per second	[qpm]	queries per minute		

Benchware Ltd

Library of Oracle benchmark tests - implemented in PL/SQL, Java and SQL

Storage Performance I/O-bound Oracle operations	OLTP systems	DWH systems	Efficiency	Metrics	Unit
• sequential I/O 1 MByte, read and write	★★	★★★	RAID tiering striping virtualization replication	service time throughput	[ms] [MBps] [GBps] [iops]
• random I/O 8 kByte, read and write	★★★	★			

Database Performance Mixed resource usage: CPU, memory, storage	OLTP systems	DWH systems	Efficiency	Metrics	Unit
• data load uncompressed, compressed	★★	★★★	scalability	speed throughput service time	[ms] [s] [rps] [tps] [qpm]
• data scan	★	★★★			
• data aggregation & reports	★★	★★★			
• OLTP transactions insert, select, update	★★★	★			

[s]	seconds	[bps]	buffers per second	[MBps]	mega bytes per second	★	less important
[ms]	milli seconds (10^{-3})	[rps]	rows per second	[GBps]	giga bytes per second	★★	important
[µs]	micro seconds (10^{-6})	[tps]	transactions per second	[iops]	i/o operations per second	★★★	very important
[ns]	nano seconds (10^{-9})	[ops]	operations per second	[qpm]	queries per minute		

Contents

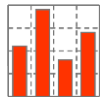

Benchmark Environment

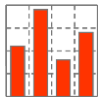

Database System

Installation	M5000	T4-2
Oracle Edition	Enterprise	Enterprise
Oracle Release	10.2.0.4	11.2.0.1
Real Application Cluster	No	No
Diagnostic Pack	Yes	Yes
DataGuard	No	No
Flashback	No	No

Configuration	M5000	T4-2
SGA capacity [GByte]	64	16
PGA capacity [GByte]	16	4
Block size [kByte]	8	8

Benchmark Environment

Benchmark Suite

	M5000	T4-2
Release, Build	OBS 6.9	BPS 8.0, Build 111201
Benchmark Database size	V - 1 TByte	M - 256 GByte
Small table • #rows • Capacity [GByte]	 32’000’000 10	 8’000’000 2.5
PL/SQL code	interpreted	interpreted

- *In this benchmark we used interpreted PL/SQL code for compatibility reasons*
- *Newer Benchware benchmarks use compiled PL/SQL code*

Contents

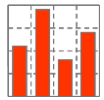

Oracle Database Platform

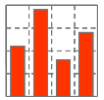

CPU

CPU	SPARC VII	SPARC T4
Frequency [GHz]	2.4	2.85
#cores	4	8
Multithreading per Core	2-fold	8-fold
Server	**SPARC VII**	**SPARC T4**
#sockets	4	2
#cores	16	16
#threads	32	128

CPU has huge impact on performance of many database operations - but also on Oracle license cost!

CPU Performance

PL/SQL string processing (data type VARCHAR2)

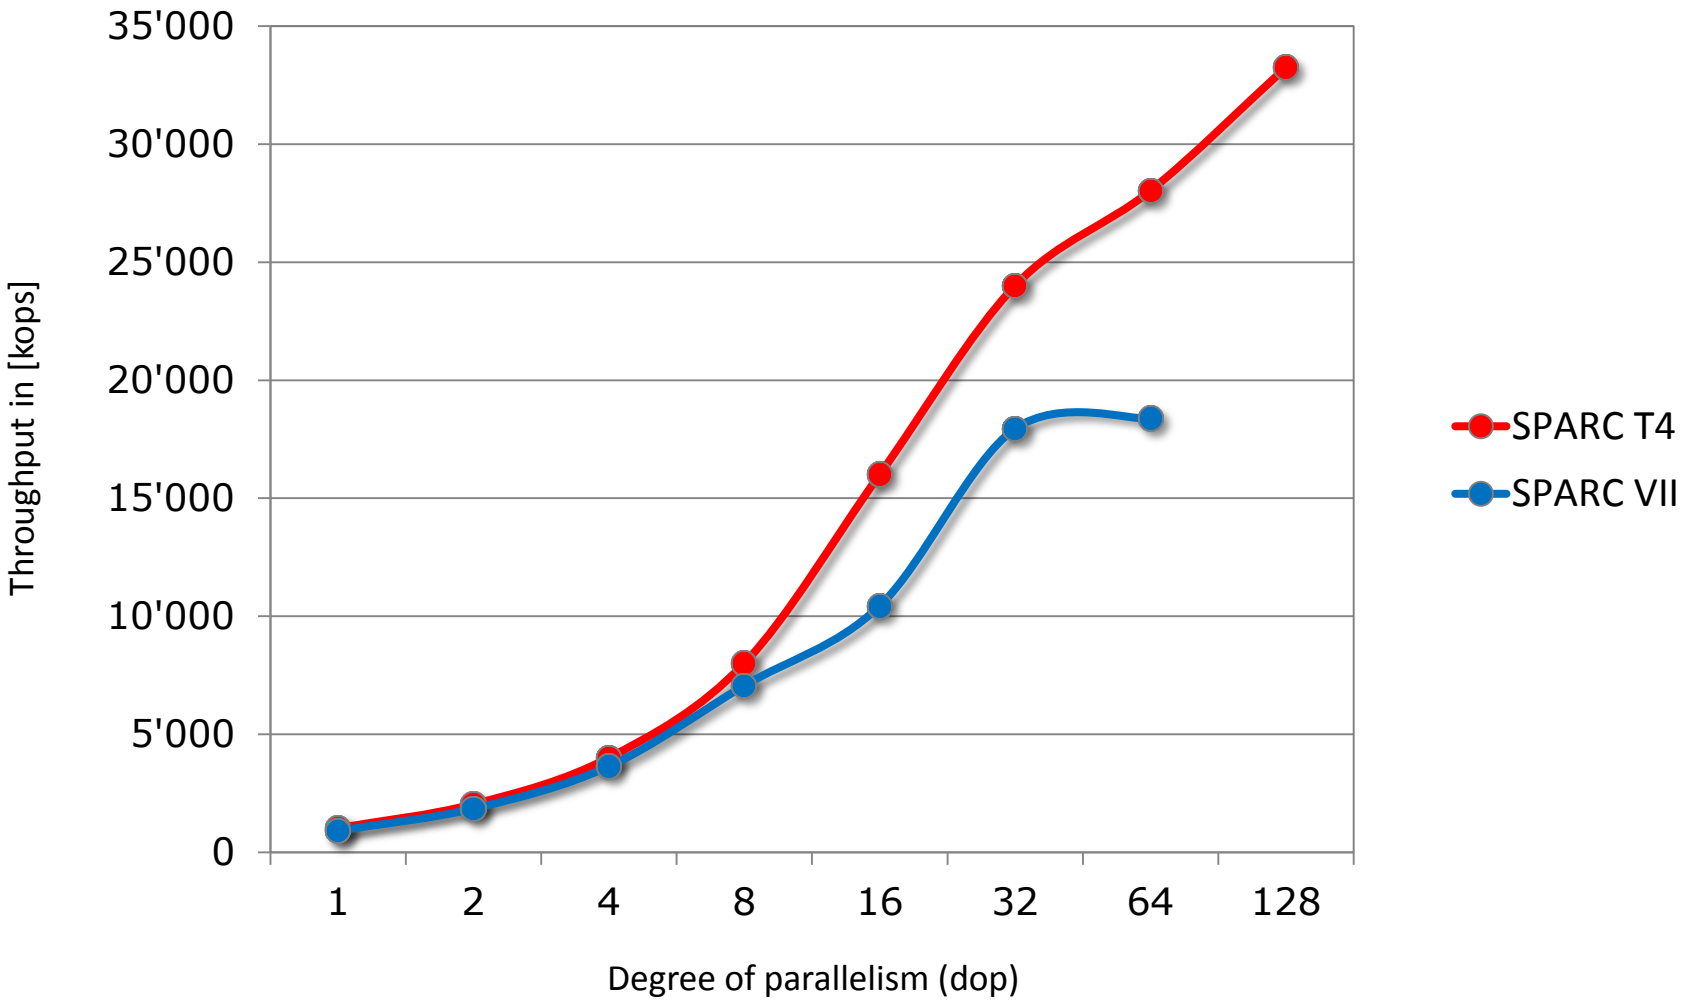

CPU Performance

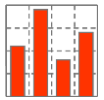

PL/SQL string processing (data type VARCHAR2)

Server T4-2

Run	Tst	Code	#N	#J	#T	Rows/sec [rps]	Ops/sec [ops]	CPU [%]	Physical read [iops]	Physical write [iops]	Total [iops]	Physical read [MBps]	Physical write [MBps]	Total [MBps]	REDO [MBps]	Time [sec]
1	9	CP31	1	1	1	0.000E+00	1.017E+06	1	2	8	10	0	0	0	0	59
	10	CP31	1	2	1	0.000E+00	2.034E+06	2	1	6	7	0	0	0	0	59
	11	CP31	1	4	1	0.000E+00	4.000E+06	3	1	6	7	0	0	0	0	60
	12	CP31	1	8	1	0.000E+00	8.000E+06	6	1	6	7	0	0	0	0	60
	13	CP31	1	16	1	0.000E+00	1.600E+07	12	1	7	8	0	0	0	0	60
	14	CP31	1	32	1	0.000E+00	2.400E+07	25	1	5	6	0	0	0	0	80
	15	CP31	1	64	1	0.000E+00	2.803E+07	44	1	4	5	0	0	0	0	137
	16	CP31	1	128	1	0.000E+00	3.325E+07	99	3	4	7	0	0	0	0	231

CPU Performance

PL/SQL integer processing (data type NUMBER)

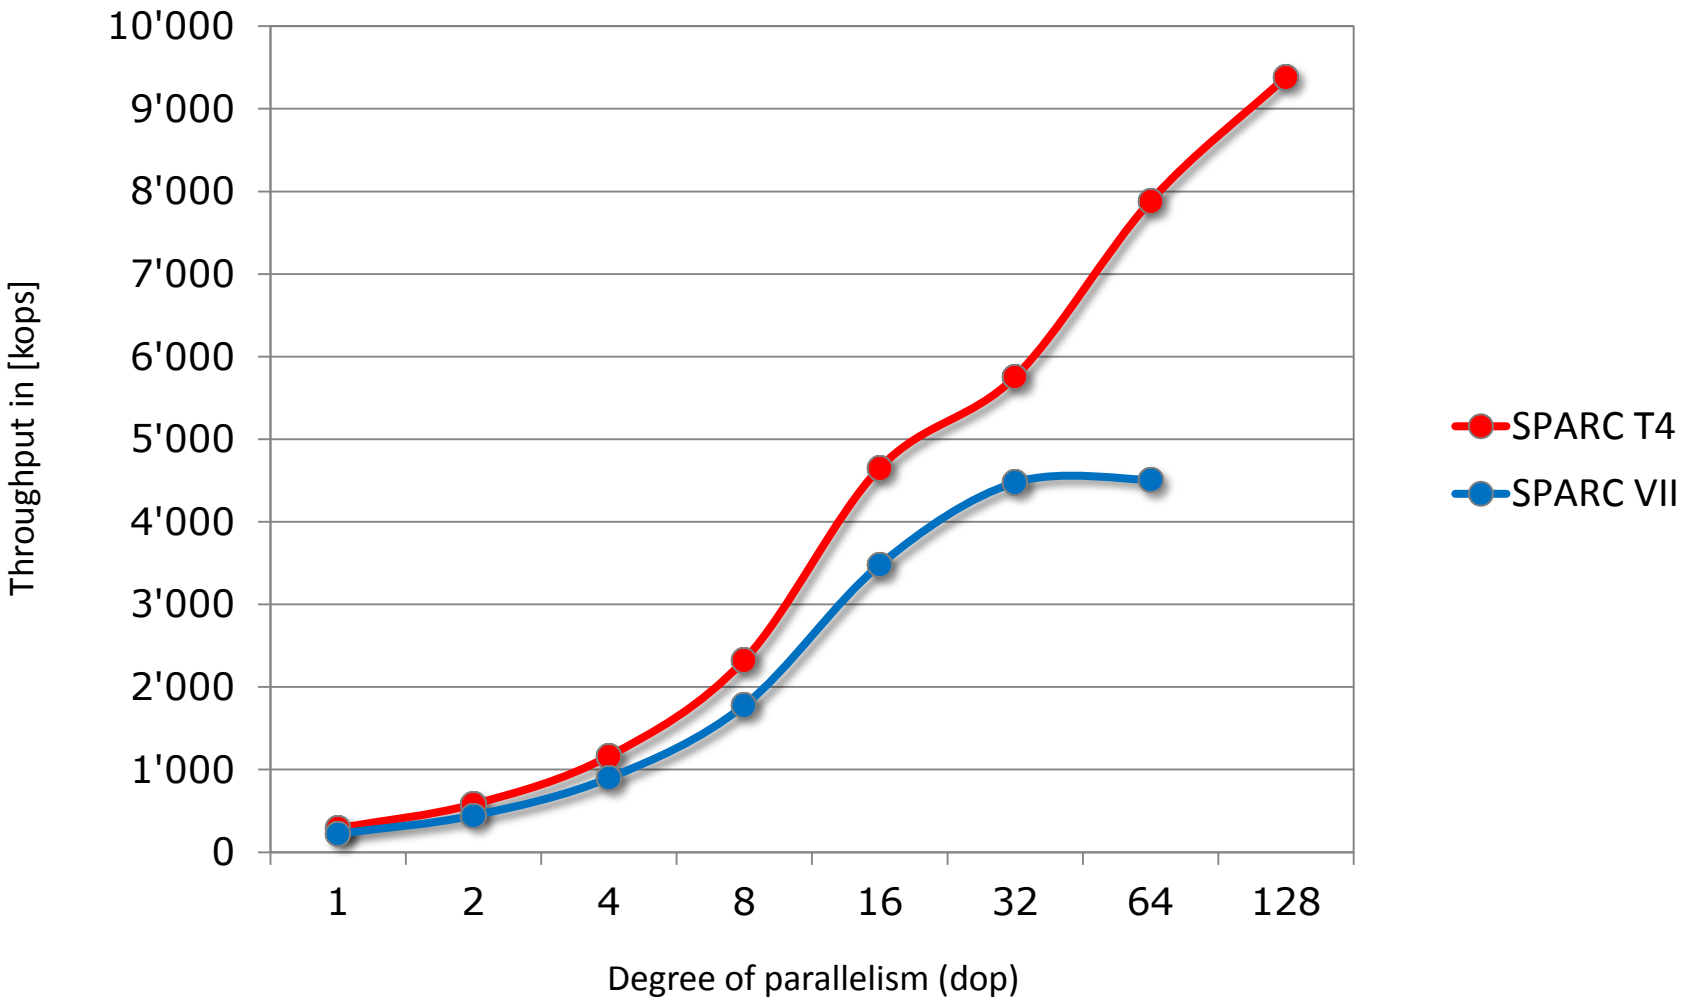

CPU Performance

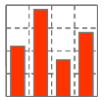

PL/SQL integer processing (data type NUMBER)

Server T4-2

Run	Tst	Code	#N	#J	#T	Rows/sec [rps]	Ops/sec [ops]	CPU [%]	Physical read [iops]	Physical write [iops]	Total [iops]	Physical read [MBps]	Physical write [MBps]	Total [MBps]	REDO [MBps]	Time [sec]
1	17	CP32	1	1	1	0.000E+00	2.907E+05	1	1	5	6	0	0	0	0	86
	18	CP32	1	2	1	0.000E+00	5.814E+05	2	1	4	5	0	0	0	0	86
	19	CP32	1	4	1	0.000E+00	1.163E+06	3	1	4	5	0	0	0	0	86
	20	CP32	1	8	1	0.000E+00	2.326E+06	6	1	5	6	0	0	0	0	86
	21	CP32	1	16	1	0.000E+00	4.651E+06	12	1	4	5	0	0	0	0	86
	22	CP32	1	32	1	0.000E+00	5.755E+06	21	1	3	4	0	0	0	0	139
	23	CP32	1	64	1	0.000E+00	7.882E+06	45	1	2	4	0	0	0	0	203
	24	CP32	1	128	1	0.000E+00	9.384E+06	99	1	2	3	0	0	0	0	341

CPU Performance

PL/SQL floating point processing (data type FLOAT)

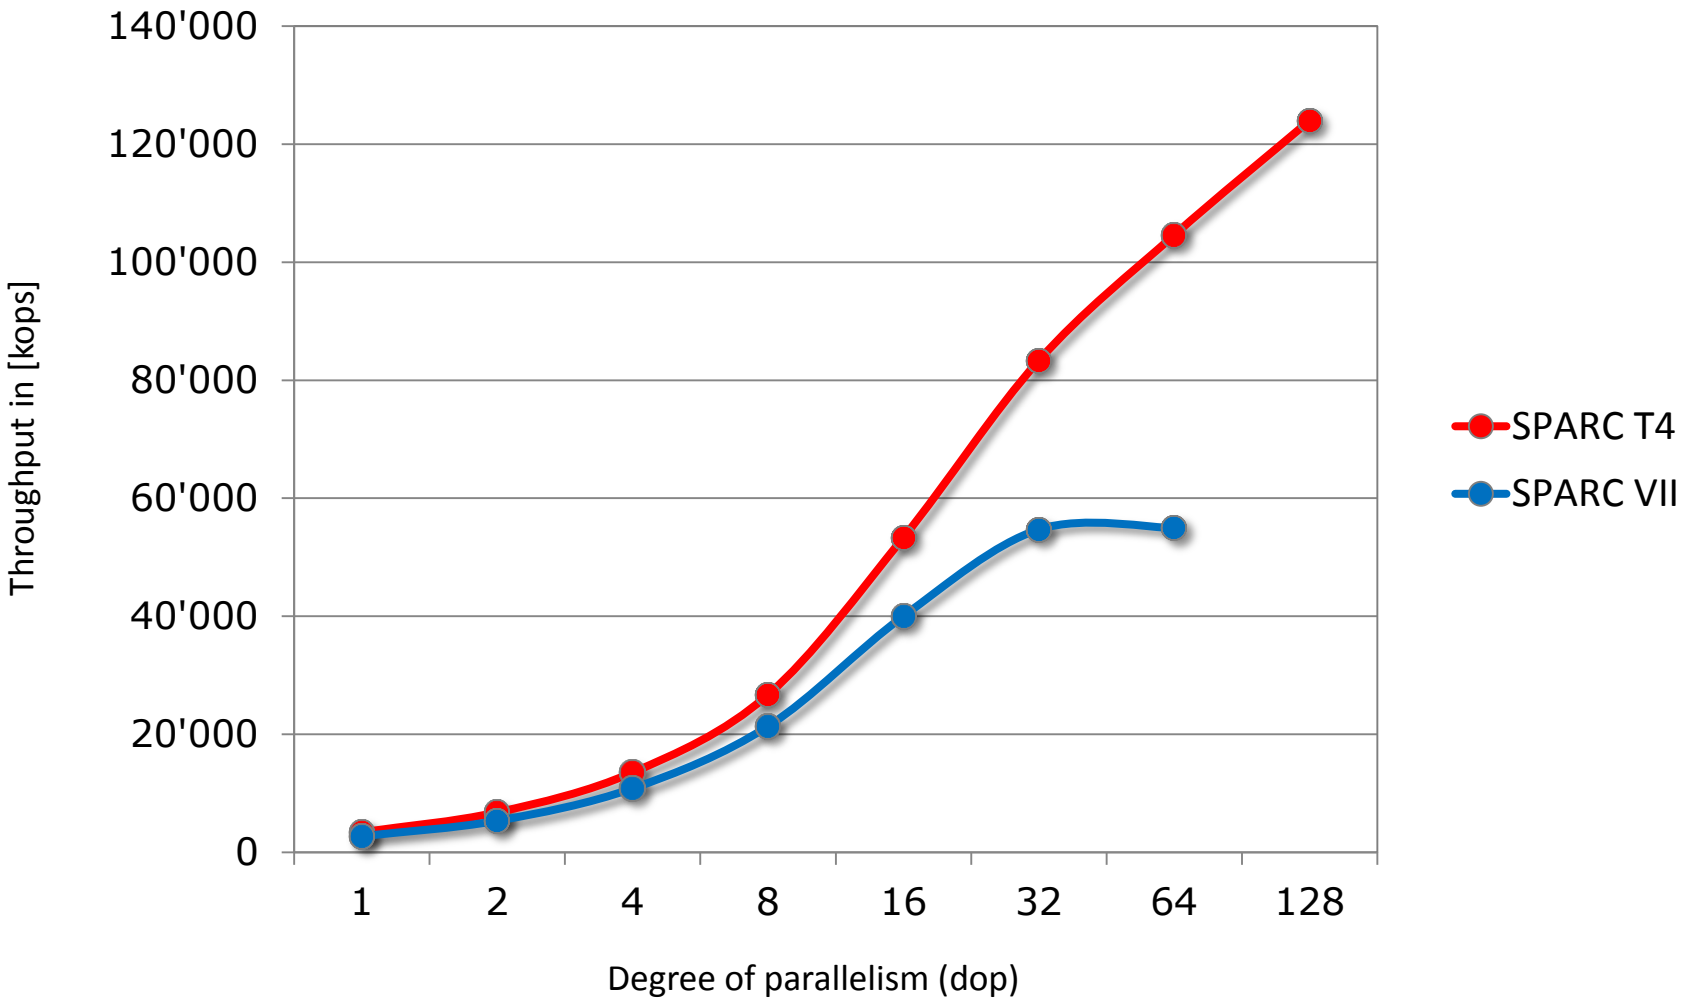

CPU Performance

PL/SQL floating point processing (data type FLOAT)

Server T4-2

Run	Tst	Code	#N	#J	#T	Rows/sec [rps]	Ops/sec [ops]	CPU [%]	Physical read [iops]	Physical write [iops]	Total [iops]	Physical read [MBps]	Physical write [MBps]	Total [MBps]	REDO [MBps]	Time [sec]
1	25	CP33	1	1	1	0.000E+00	3.378E+06	1	1	6	7	0	0	0	0	74
	26	CP33	1	2	1	0.000E+00	6.757E+06	2	1	5	6	0	0	0	0	74
	27	CP33	1	4	1	0.000E+00	1.351E+07	3	1	5	6	0	0	0	0	74
	28	CP33	1	8	1	0.000E+00	2.667E+07	6	1	5	7	0	0	0	0	75
	29	CP33	1	16	1	0.000E+00	5.333E+07	12	2	6	8	0	0	0	0	75
	30	CP33	1	32	1	0.000E+00	8.333E+07	25	1	4	6	0	0	0	0	96
	31	CP33	1	64	1	0.000E+00	1.046E+08	46	1	4	5	0	0	0	0	153
	32	CP33	1	128	1	0.000E+00	1.245E+08	99	1	2	3	0	0	0	0	257

CPU Performance

PL/SQL algorithm interpreted (fibonacci, recursive, n=39)

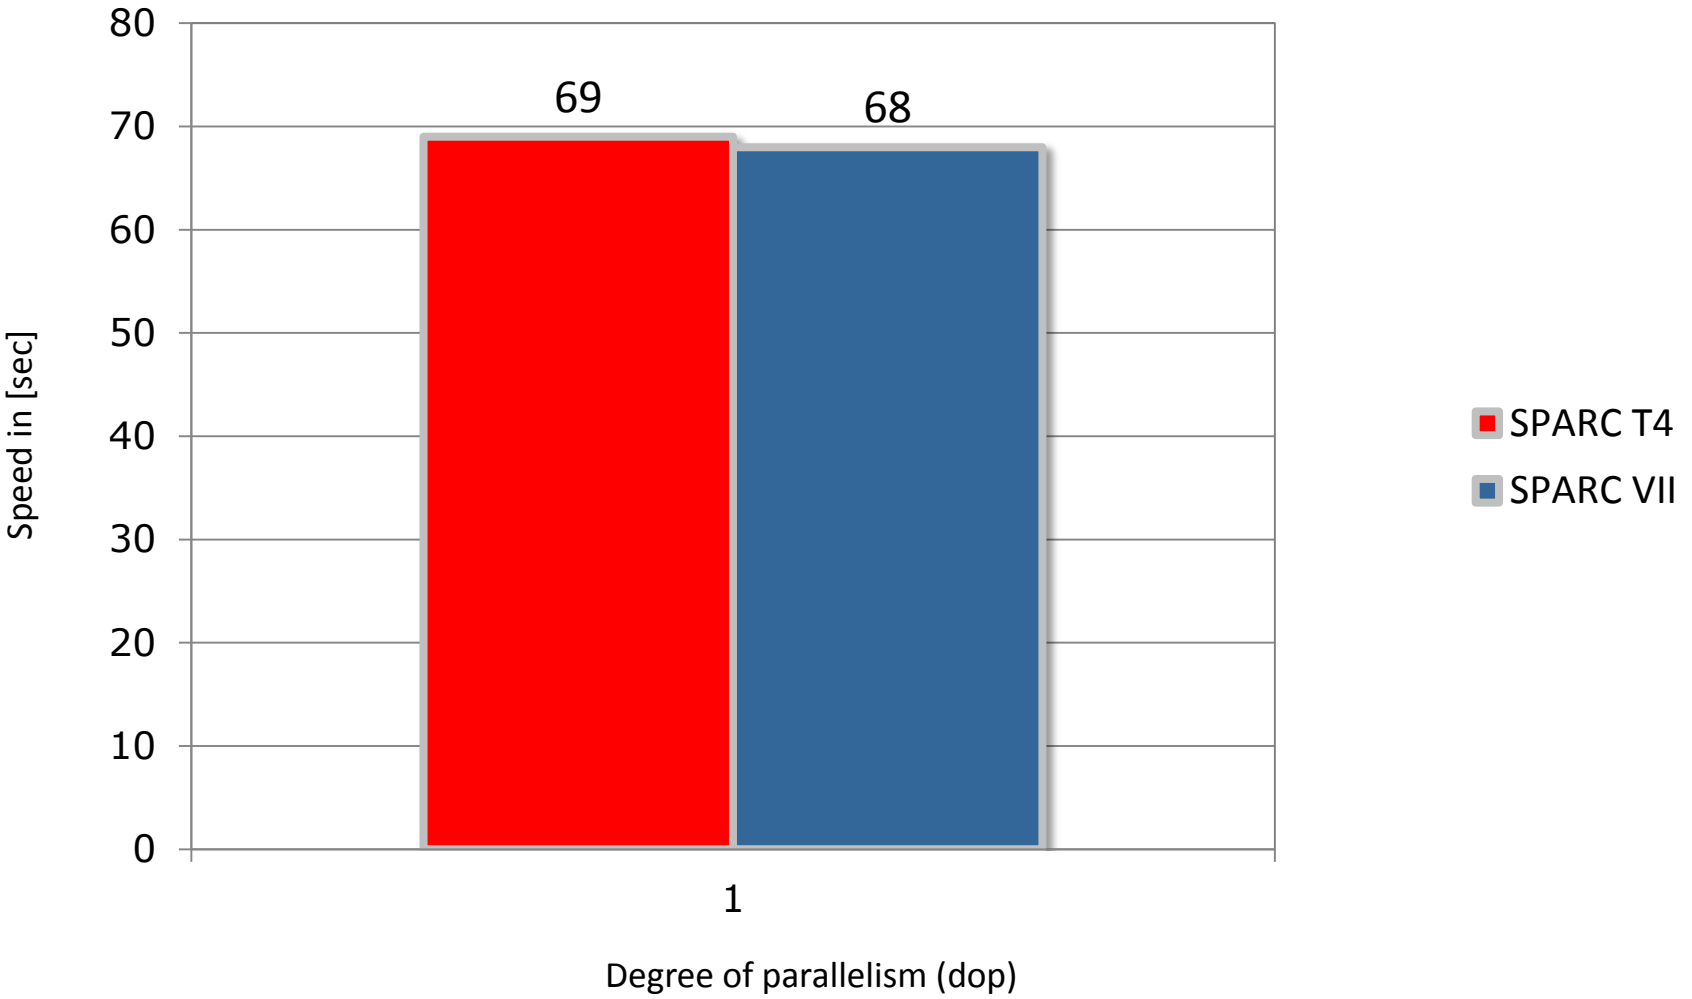

CPU Performance

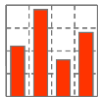

Summary CPU Performance

	Metric	M5000	T4-2
#cores		16	16
#threads		32	128
PL/SQL operations			
String processing • Speed (single thread) • Throughput	 [kops] [kops]	 909 18’373	 1’017 33’250
NUMBER processing • Speed (single thread) • Throughput	 [kops] [kops]	 224 4’507	 290 9’384
Floating point processing • Speed (single thread • Throughput	 [kops] [kops]	 2’702 54’935	 3’378 124’500
Algorithms • Speed fibonacci recursive (n=39)	 [s]	 68	 69

Contents

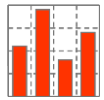

Oracle Database Platform

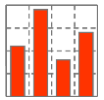

Server

Server	M5000	T4-2
#sockets	4	2
#cores	16	16
#threads (CPU_COUNT)	32	128
Oracle licensing (Oracle processors)	12	8
Main memory [GByte]	128	64
Host-Bus-Adapter (type, quantity, throughput)	-	-
Operating System	Solaris 10	Solaris 10
Cluster		
#server	-	-

Most OLTP applications avoid I/O operations as much as possible and work predominately in RAM – server performance is essential for these kind of OLTP applications!

Server Performance

In-memory SQL, primary key access

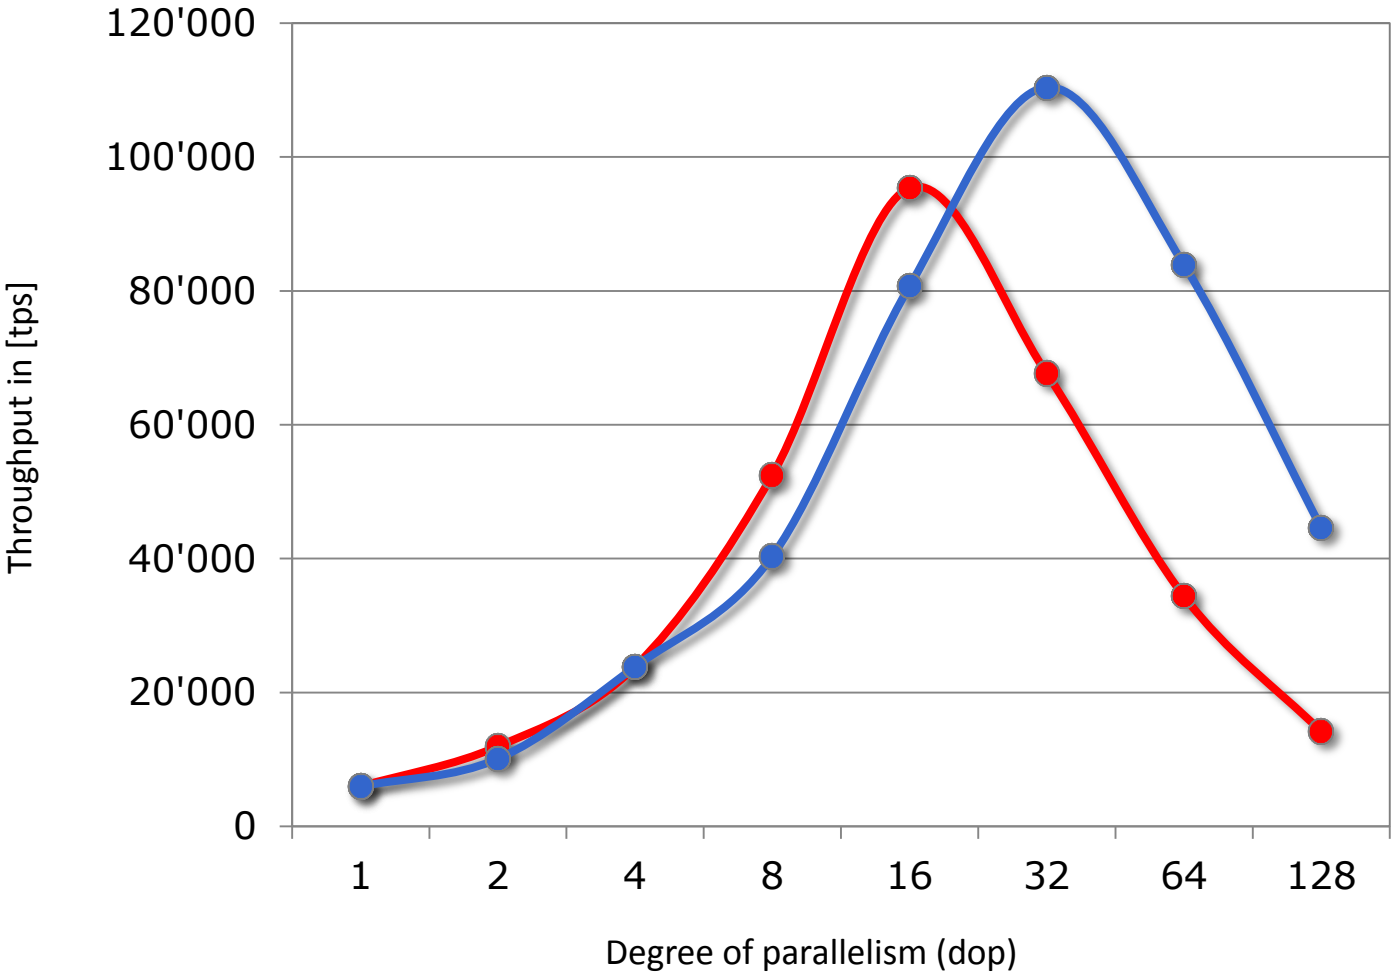

T4 does not scale better because of concurrency conflicts in smaller SGA.

On Solaris Oracle uses spinning CPU resources in some wait situations.

In OLTP systems on Solaris don't exceed a specific CPU utilization threshold.

Control CPU usage in these situations with parameter _spin_count

Server Performance

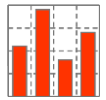

In-memory SQL, primary key access

Server T4-2

Run	Tst	Code	#N	#J	#T	Rows/sec [rps]	Ops/sec [ops]	CPU [%]	Physical read [iops]	Physical write [iops]	Total [iops]	Physical read [MBps]	Physical write [MBps]	Total [MBps]	REDO [MBps]	Time [sec]
5	1	CS12	1	1	1	5.958E+03	5.958E+03	0	3	41	44	0	1	1	0	11
	2	CS12	1	2	1	1.192E+04	1.192E+04	1	3	29	32	0	0	0	0	11
	3	CS12	1	4	1	2.383E+04	2.383E+04	1	2	31	33	0	0	0	0	11
	4	CS12	1	8	1	5.243E+04	5.243E+04	3	3	35	38	0	0	0	0	10
	5	CS12	1	16	1	9.533E+04	9.533E+04	5	2	29	32	0	0	0	0	11
	6	CS12	1	32	1	6.765E+04	6.765E+04	24	2	14	15	0	0	0	0	31
	7	CS12	1	64	1	3.438E+04	3.438E+04	49	1	4	5	0	0	0	0	122
	8	CS12	1	128	1	1.417E+04	1.417E+04	99	1	2	3	0	0	0	0	592

DOP = 32

```
Top 5 Timed Foreground Events
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

                                                    Avg
                                                   wait   % DB
Event                                 Waits     Time(s)   (ms)   time Wait Class
------------------------------ ------------ ----------- ------ ------ ----------
DB CPU                                            1,878          82.2
cursor: pin S                       33,098           6      0     .3 Concurrenc
```

DOP = 128

```
Top 5 Timed Foreground Events
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

                                                    Avg
                                                   wait   % DB
Event                                 Waits     Time(s)   (ms)   time Wait Class
------------------------------ ------------ ----------- ------ ------ ----------
DB CPU                                          150,266          95.3
cursor: pin S                    6,983,021      31,419      4   19.9 Concurrenc
```

Oracle Reference Manual:
"A session waits on this event when it wants to update a shared mutex pin and another session is currently in the process of updating a shared mutex pin for the same cursor object. This wait event should rarely be seen because a shared mutex pin update is very fast."

Server Performance

In-memory SQL, secondary key access

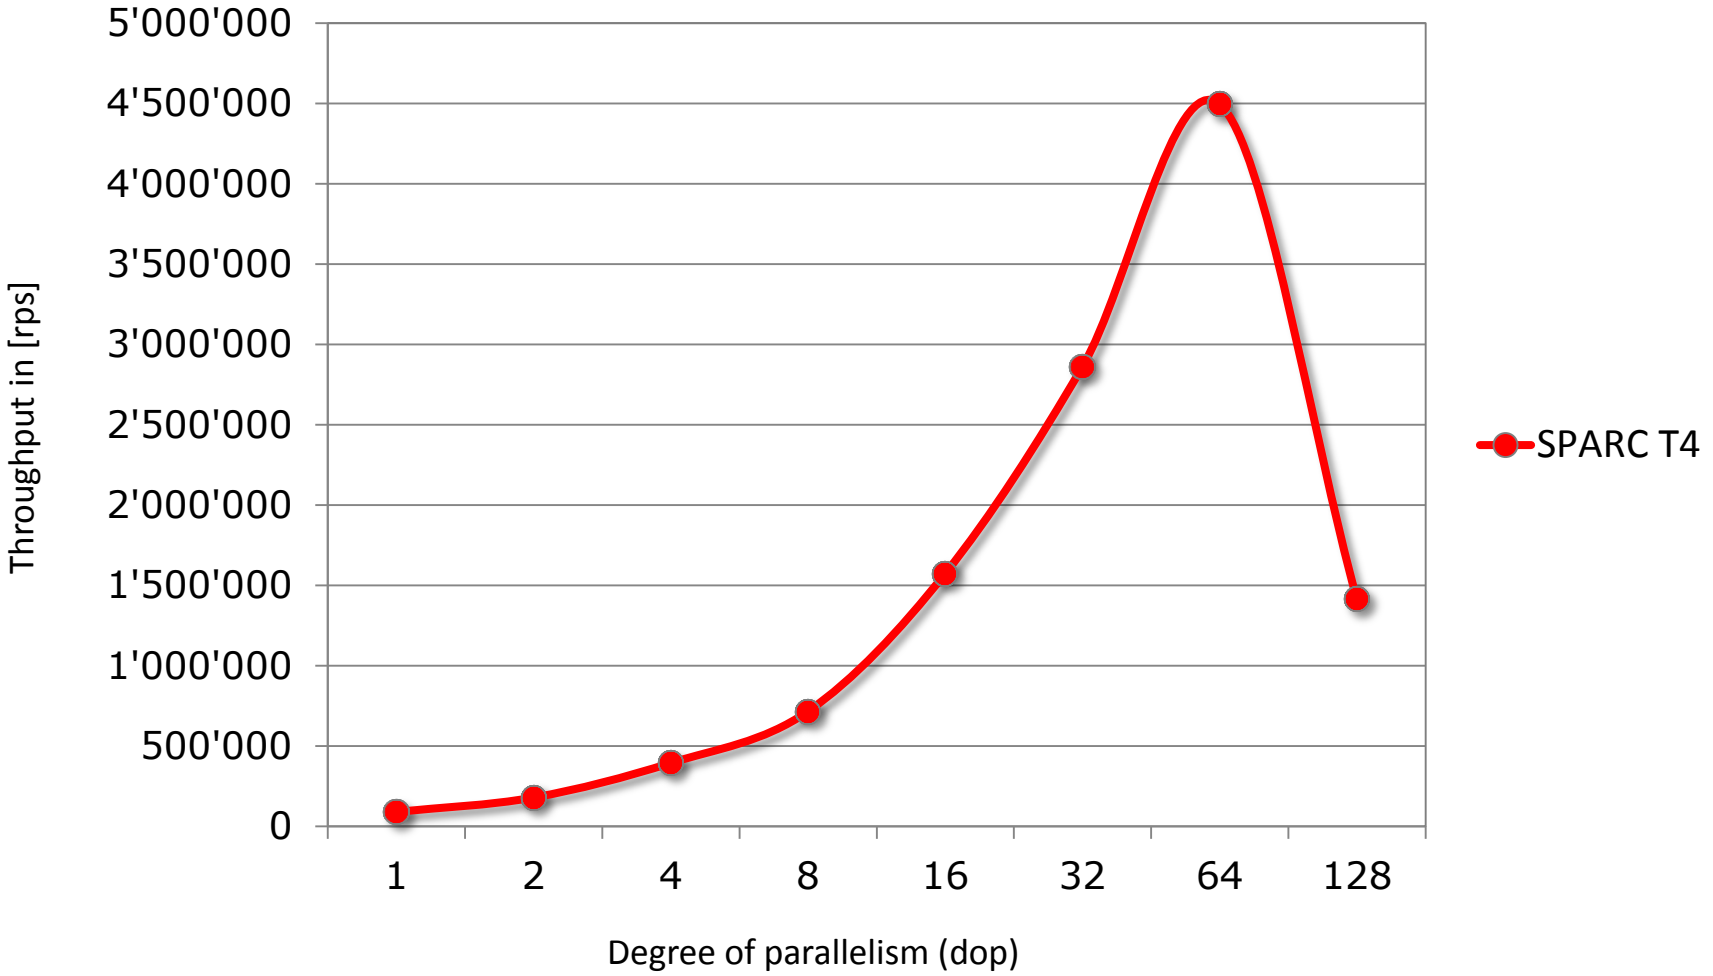

Server Performance

In-memory SQL, secondary key access

Server T4-2

Run	Tst	Code	#N	#J	#T	Rows/sec [rps]	Ops/sec [ops]	CPU [%]	Physical read [iops]	Physical write [iops]	Total [iops]	Physical read [MBps]	Physical write [MBps]	Total [MBps]	REDO [MBps]	Time [sec]
5	9	CS13	1	1	1	8.930E+04	1.490E+03	1	3	39	42	0	1	1	0	11
	10	CS13	1	2	1	1.788E+05	2.979E+03	1	2	32	34	0	0	0	0	11
	11	CS13	1	4	1	3.930E+05	6.554E+03	2	3	35	38	0	0	1	0	10
	12	CS13	1	8	1	7.153E+05	1.192E+04	4	3	31	34	0	0	0	0	11
	13	CS13	1	16	1	1.573E+06	2.621E+04	8	3	34	37	0	0	0	0	10
	14	CS13	1	32	1	2.860E+06	4.766E+04	21	2	37	39	0	0	0	0	11
	15	CS13	1	64	1	4.494E+06	7.490E+04	45	2	31	34	0	0	0	0	14
	16	CS13	1	128	1	1.414E+06	2.356E+04	98	1	6	7	0	0	0	0	89

Server Performance

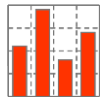

Summary Server Performance

	Metric	M5000	T4-2
#cores		16	16
#threads		32	128
Main memory capacity [GByte]		128	64
In-memory SQL operations			
Full table scan • Throughput	[rps]	-	-
Random table access via primary key • Throughput for DOP = 1 • Throughput max	[tps] [tps]	5'960 110'380	5'958 95'330
Random table access via secondary key • Throughput	[rps]	-	4'500'000

Contents

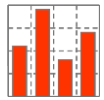

Conclusion

SPARC T4 versus SPARC VII

All prices are list prices (spring 2012)

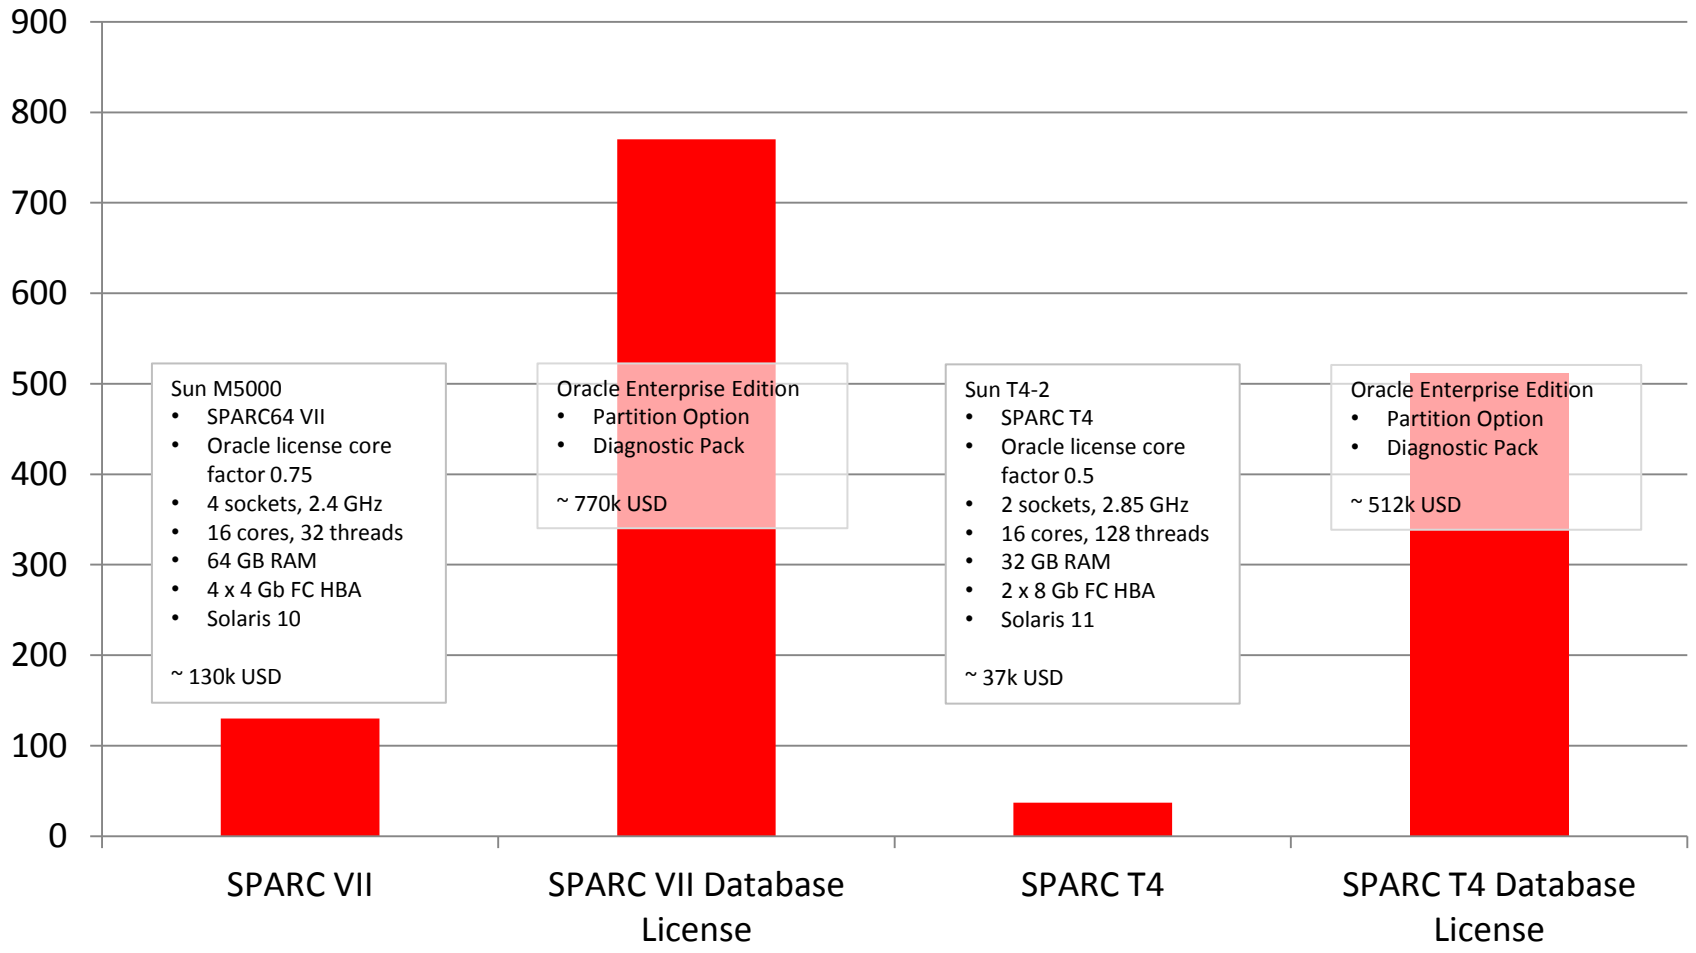

Conclusion

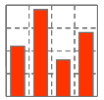

SPARC T4 versus SPARC VII

- Performance
  - Speed of SPARC T4 core is very similar to SPARC VII
  - Throughput of SPARC T4 core is up to factor 2 higher than SPARC VII
- Cost-efficient hardware
  - SPARC T4 is more cost-efficient than SPARC VII
  - Less Server investment
  - Less Oracle license fee
  - Less Oracle maintenance fee

Conclusion

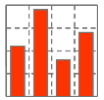

SPARC T4 versus SPARC VII

- Functionality
  - SPARC T4 support new technologies
  - Embedded cryptographic instruction set
  - PCI-based SSD technology, e.g. for Oracle Flash Cache as second-level Oracle buffer cache (available only on Solaris and OEL)

- SPARC T4 – perfect replacement for systems like
  - V-Series (V440, V480, V490)
  - Smaller M-Series with older SPARC chips: III, IV, V, VI and VII

Conclusion

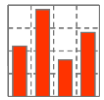

SPARC T4 versus SPARC VII

- Benchware uses fair, reproducible and representative benchmark tests delivering understandable *key performance metrics* (KPM)
- Benchware uses a list of defined *price performance ratios* (PPR) to evaluate platform cost
- Benchware publishes *price performance ratios* (PPR) to its customers only

swiss precision in performance measurement

www.benchware.ch

info@benchware.ch