


Performance Baseline of Exadata X4-2

Part III: Storage (STO) Performance

Technical Presentation

June 2014



1 Introduction to Storage (STO) Performance Tests


2 Storage System Configuration

3 Storage Benchmark Results – Sequential I/O

4 Storage Benchmark Results – Random I/O


5 Reviewing Storage Benchmark Results

Introduction to Storage (STO) Performance Tests




- Why measure storage system performance?
- What is measured?
- How is storage system performance measured?
- Overview Benchware STO performance test for Oracle
- Monitoring Benchware STO performance tests

- Look at www.benchware.ch/methodology for detailed information



- 1 Introduction to Storage (STO) Performance Tests
- 2 Storage System Configuration**
- 3 Storage Benchmark Results – Sequential I/O
- 4 Storage Benchmark Results – Random I/O
- 5 Reviewing Storage Benchmark Results


Storage System Configuration



Server configuration for load generation

Server	X2-2	X3-2	X4-2
CPU Type	X5675	E5-2690	E5-2697 V2
#sockets	2	2	2
#cores	12	16	24
#threads	24	32	48
Memory capacity [GByte]	96	512	512
Performance Numbers from other Benchmarks			
SPECint_base_rate_2006 (throughput)	367	630	806
Cluster			
#server within cluster	4	8	8
Software			
Operating System	Oracle Lx	Oracle Lx	Oracle Lx
Oracle Database System	11.2	11.2	11.2
Benchware Performance Suite	8.6	8.6	8.6

Storage System Configuration




Oracle Capacities and Licensing on Exadata Storage Systems

Storage Capacities, raw, 1 Exadata Storage Server	X2-2	X3-2	X4-2
<ul style="list-style-type: none"> Flash in [GByte] High Performance HDD in [TByte] High Capacity HDD in [TByte] 	384 7.2@15K 36@7.2K	1'600 7.2@15K 36@7.2K	3'200 14.4@10K 48@7.2K
Ratio of Flash capacity in percentage of HDD capacity	1%	4.5%	6.7%

Exadata has:


- No flash storage, but flash cache
- Auto tiering technology
- Service time is not guaranteed

Oracle Storage Server License	X2-2	X3-2	X4-2
Oracle license cost <small>(list price 13th of February 2014)</small>			
<ul style="list-style-type: none"> Eighth Rack Quarter Rack Half Rack 	- 360'000 840'000	180'000 360'000 840'000	180'000 360'000 840'000
Software license cost in [USD] per 1 TByte Capacity	3'333	3'333	2'500

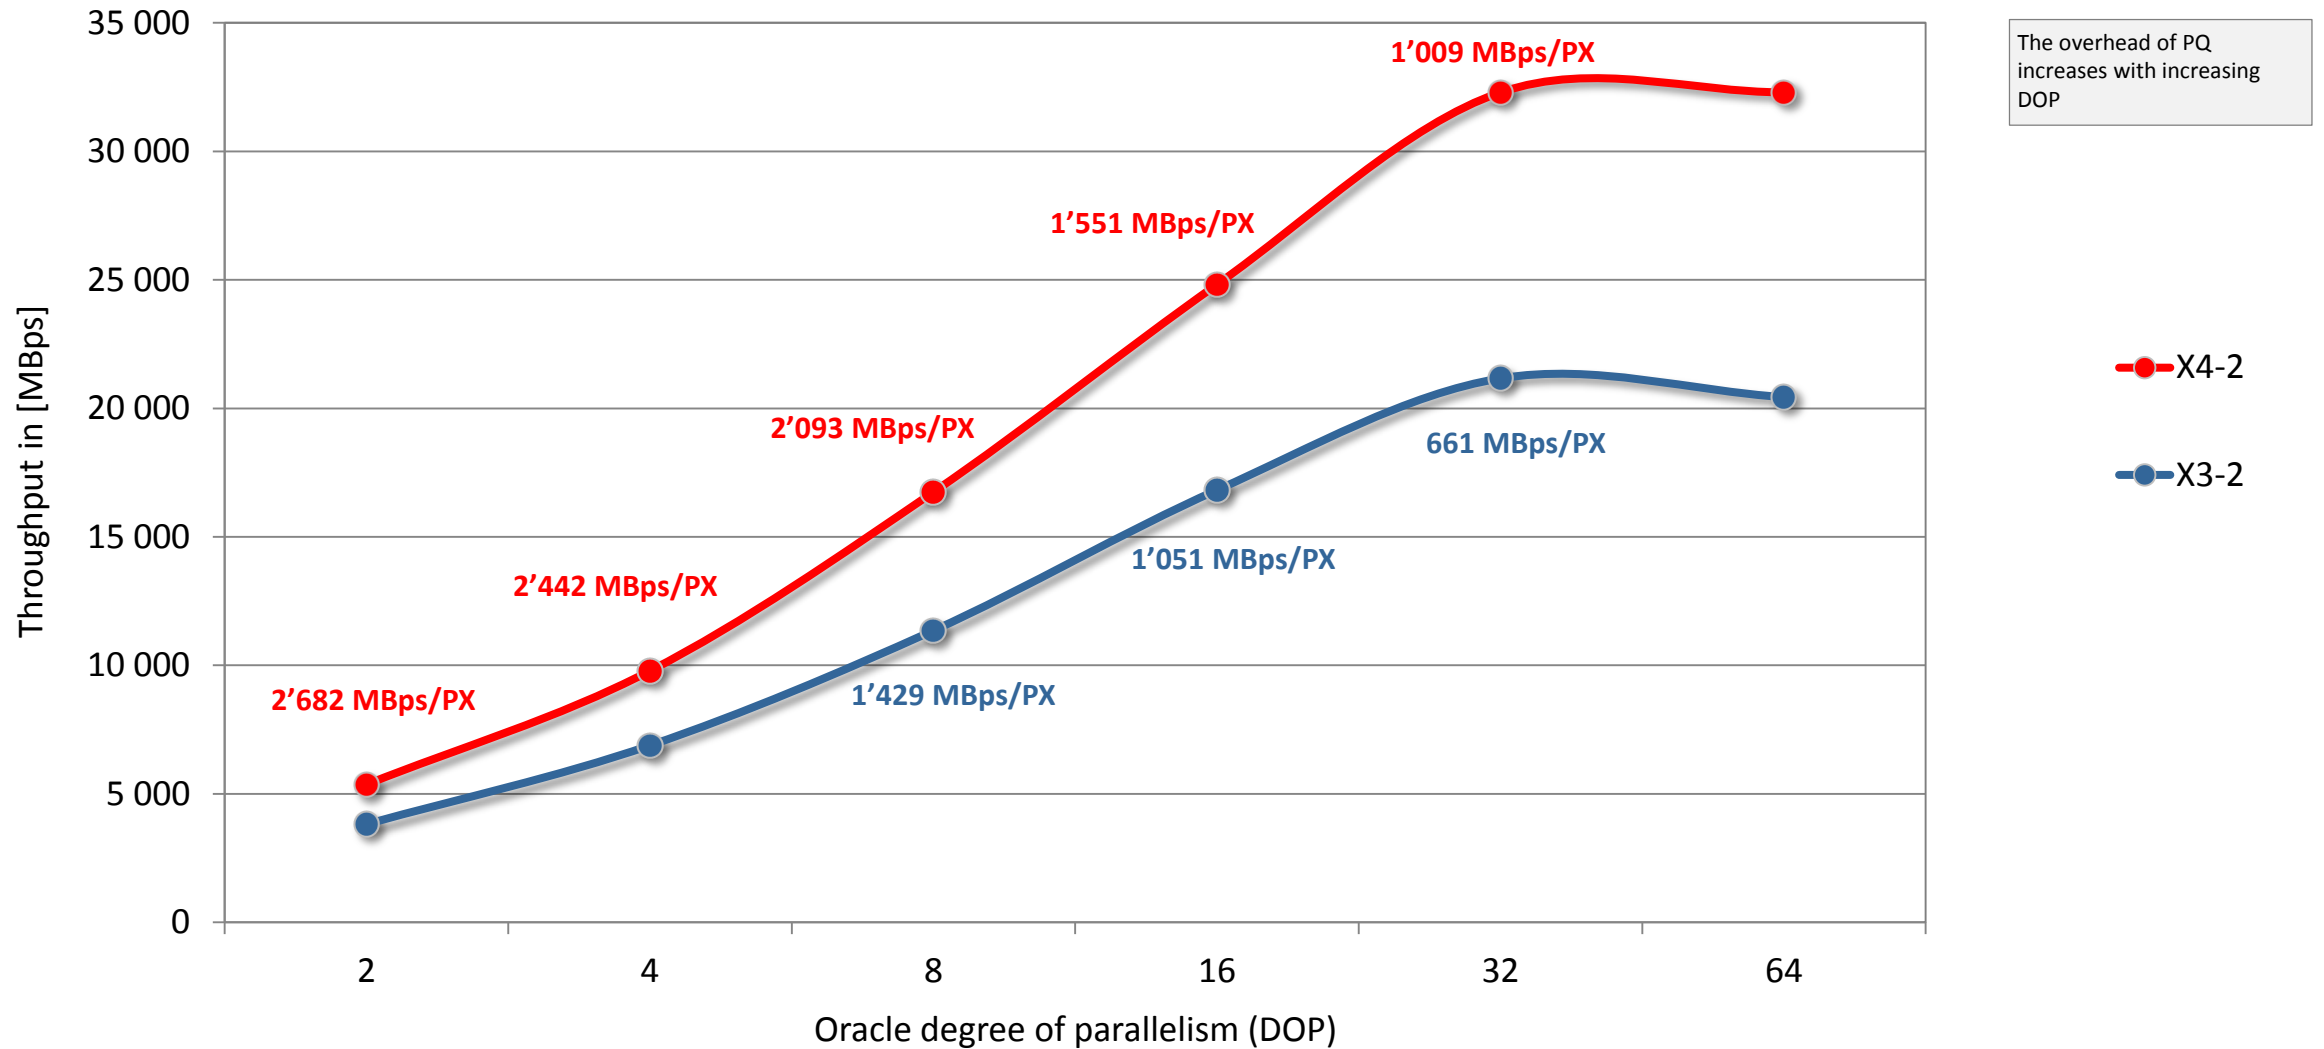


- 1 Introduction to Storage (STO) Performance Tests
- 2 Storage System Configuration
- 3 Storage Benchmark Results – Sequential I/O**
- 4 Storage Benchmark Results – Random I/O
- 5 Reviewing Storage Benchmark Results


Storage Benchmark Results




Oracle sequential read: 1 user process with different DOP's




Storage Benchmark Results



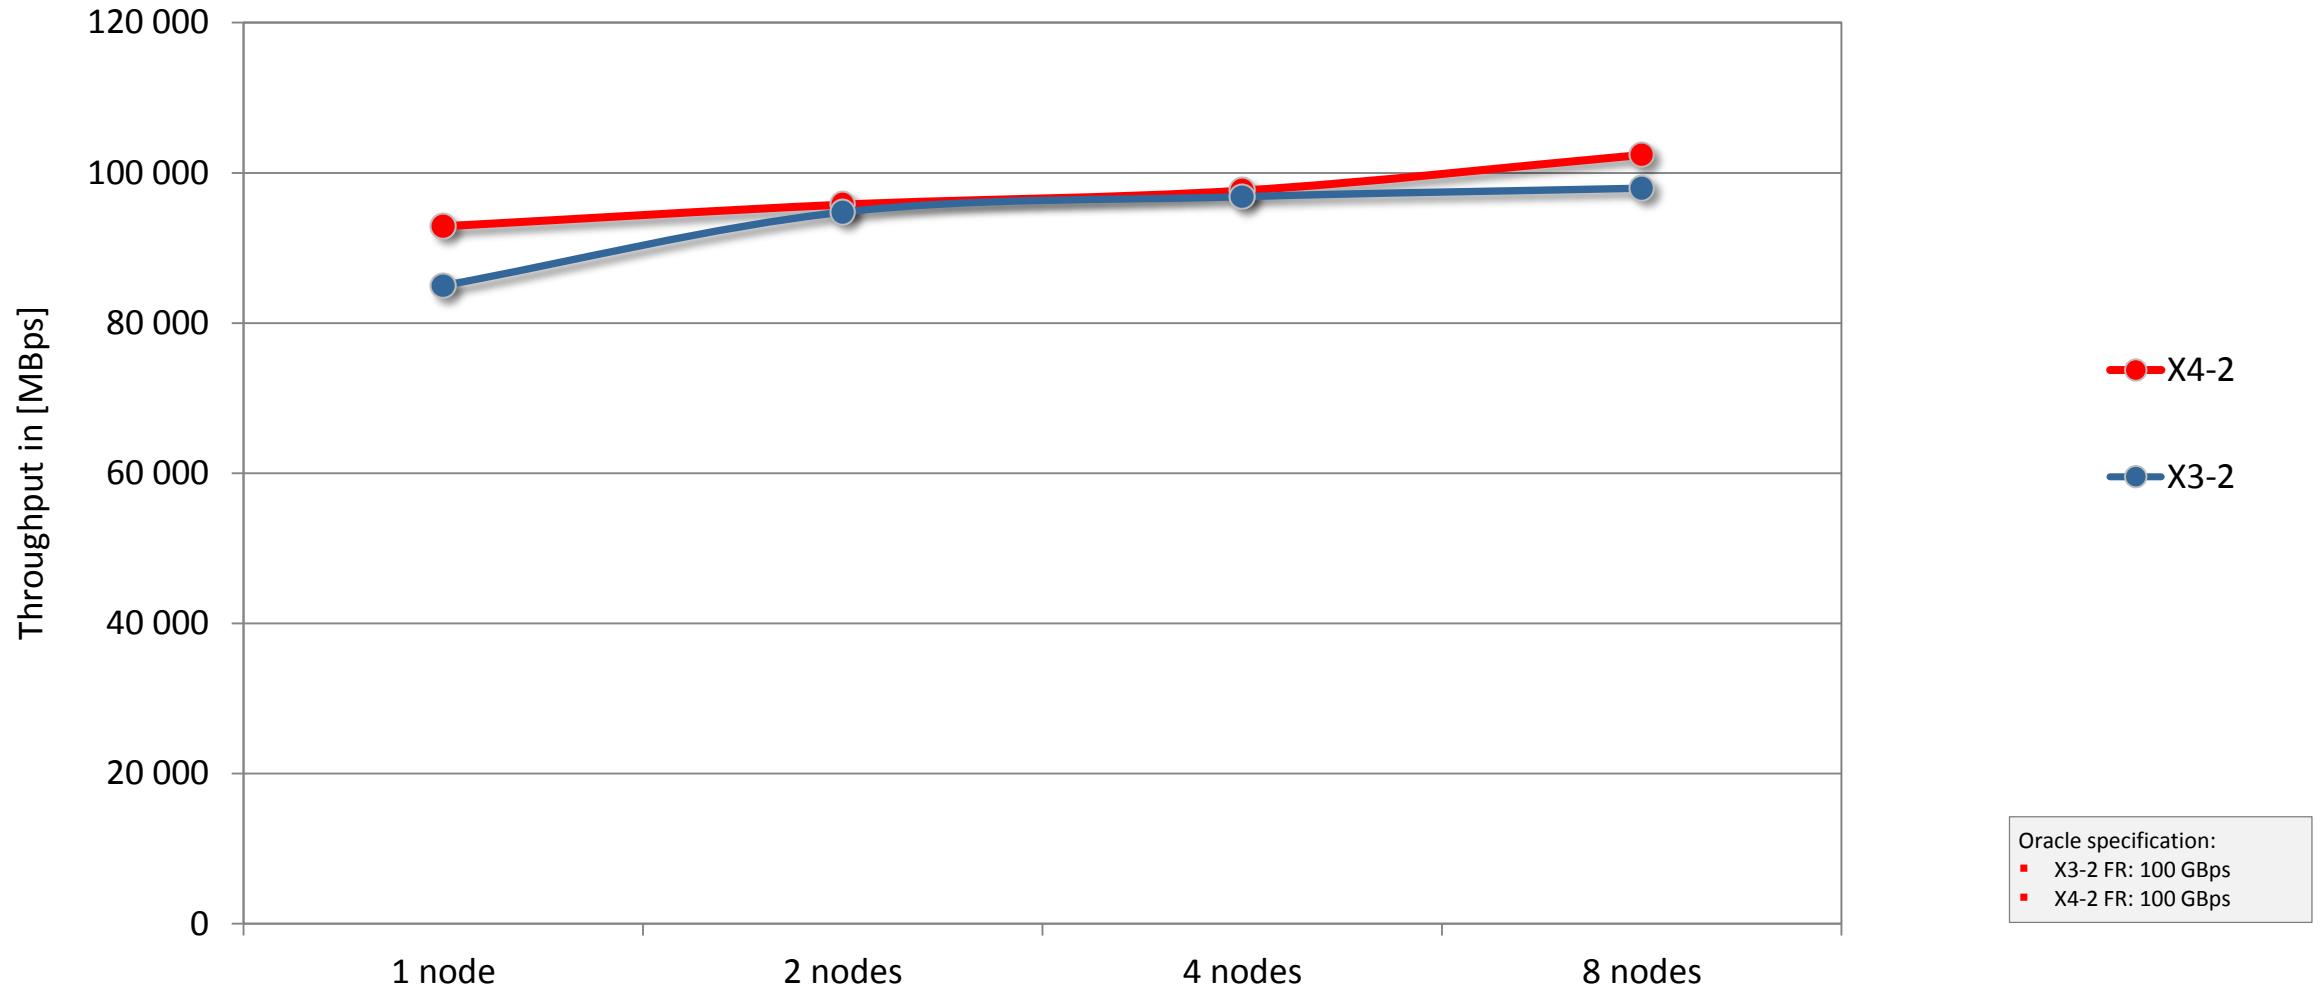
Oracle sequential read: 1 DB Server




Storage Benchmark Results



Oracle sequential read: Cluster



Storage Benchmark Results




Oracle sequential read: Cluster

Exadata X4-2 FR HC

Run	Tst	Code	#N	#J	#T	CPU busy [%]	CPU sys [%]	Physical read [iops]	Physical read [dbps]	Physical read [MBps]	Physical write [iops]	Physical write [dbps]	Physical write [MBps]	REDO write [iops]	Hitrate db flash [%]	Hitrate exa flash [%]	Elap time [s]
2	35	STO-14	1	1	4	1	1	9695	1233359	9636	19	26	0	1	0	100	300
	36	STO-14	1	2	4	2	1	17442	2223360	17370	19	25	0	1	0	100	300
	37	STO-14	1	4	4	3	1	31368	4003349	31276	19	25	0	2	0	100	301
	38	STO-14	1	8	4	4	1	51409	6564811	51288	20	25	0	2	0	100	301
	39	STO-14	1	16	4	5	1	71893	9182751	71741	22	26	0	3	0	100	301
	40	STO-14	1	32	4	6	2	85902	10973449	85730	23	28	0	3	0	98	301
	41	STO-14	1	64	4	7	2	93083	11891116	92900	26	29	0	6	0	95	303
	42	STO-14	2	128	4	7	2	95992	12259792	95780	26	28	1	6	0	90	308
	43	STO-14	4	256	4	7	2	97271	12416499	97004	35	71	1	10	0	87	305
	44	STO-14	8	512	4	8	2	98000	12500180	97658	41	33	1	16	0	84	304
	45	STO-14	8	512	4	8	2	102774	13107346	102401	43	35	1	16	0	87	308


Legend:

Run	benchmark run id	#N	number of RAC nodes	[rps]	rows per second	[iops]	i/o operations per second	[s]	time in seconds
Tst	benchmark test id	#J	number of load generators (jobs)	[tps]	transactions per second	[dbps]	database blocks per second	[ms]	time in milli seconds
Code	benchmark test code	#T	number of threads (PX)	[ops]	operations per second	[MBps]	mega byte per second	[μs]	time in micro seconds




- 1 Introduction to Storage (STO) Performance Tests
- 2 Storage System Configuration
- 3 Storage Benchmark Results – Sequential I/O
- 4 Storage Benchmark Results – Random I/O**
- 5 Reviewing Storage Benchmark Results


Storage Benchmark Results



Oracle random read: 1 DB server



Storage Benchmark Results



Oracle random read: 1 DB server

Exadata X4-2 FR HC


Run	Tst	Code	#N	#J	#T	CPU busy [%]	CPU sys [%]	Physical read [iops]	Physical read [dbps]	Physical read [MBps]	Physical write [iops]	Physical write [dbps]	Physical write [MBps]	REDO write [iops]	Hitrate db flash [%]	Hitrate exa flash [%]	Elap time [s]
2	52	STO-64	1	1	1	1	0	79092	79073	618	10	12	0	1	0	96	298
	53	STO-64	1	2	1	1	0	128788	128770	1006	7	8	0	1	0	96	296
	54	STO-64	1	4	1	1	0	228344	228325	1784	16	17	0	2	0	96	308
	55	STO-64	1	8	1	1	0	355593	355573	2778	15	19	0	1	0	95	304
	56	STO-64	1	16	1	2	1	433914	433895	3390	18	19	0	3	0	95	307
	57	STO-64	1	32	1	3	1	496501	496482	3879	19	22	0	2	0	93	305
	58	STO-64	1	64	1	5	1	567283	567264	4432	24	25	0	4	0	93	307
	59	STO-64	1	128	1	9	2	660586	660568	5161	27	27	0	3	0	93	316
	60	STO-64	1	256	1	12	3	693597	693577	5419	31	31	0	6	0	92	333
	61	STO-64	1	512	1	11	3	598168	598148	4673	33	29	0	9	0	91	362

Legend:

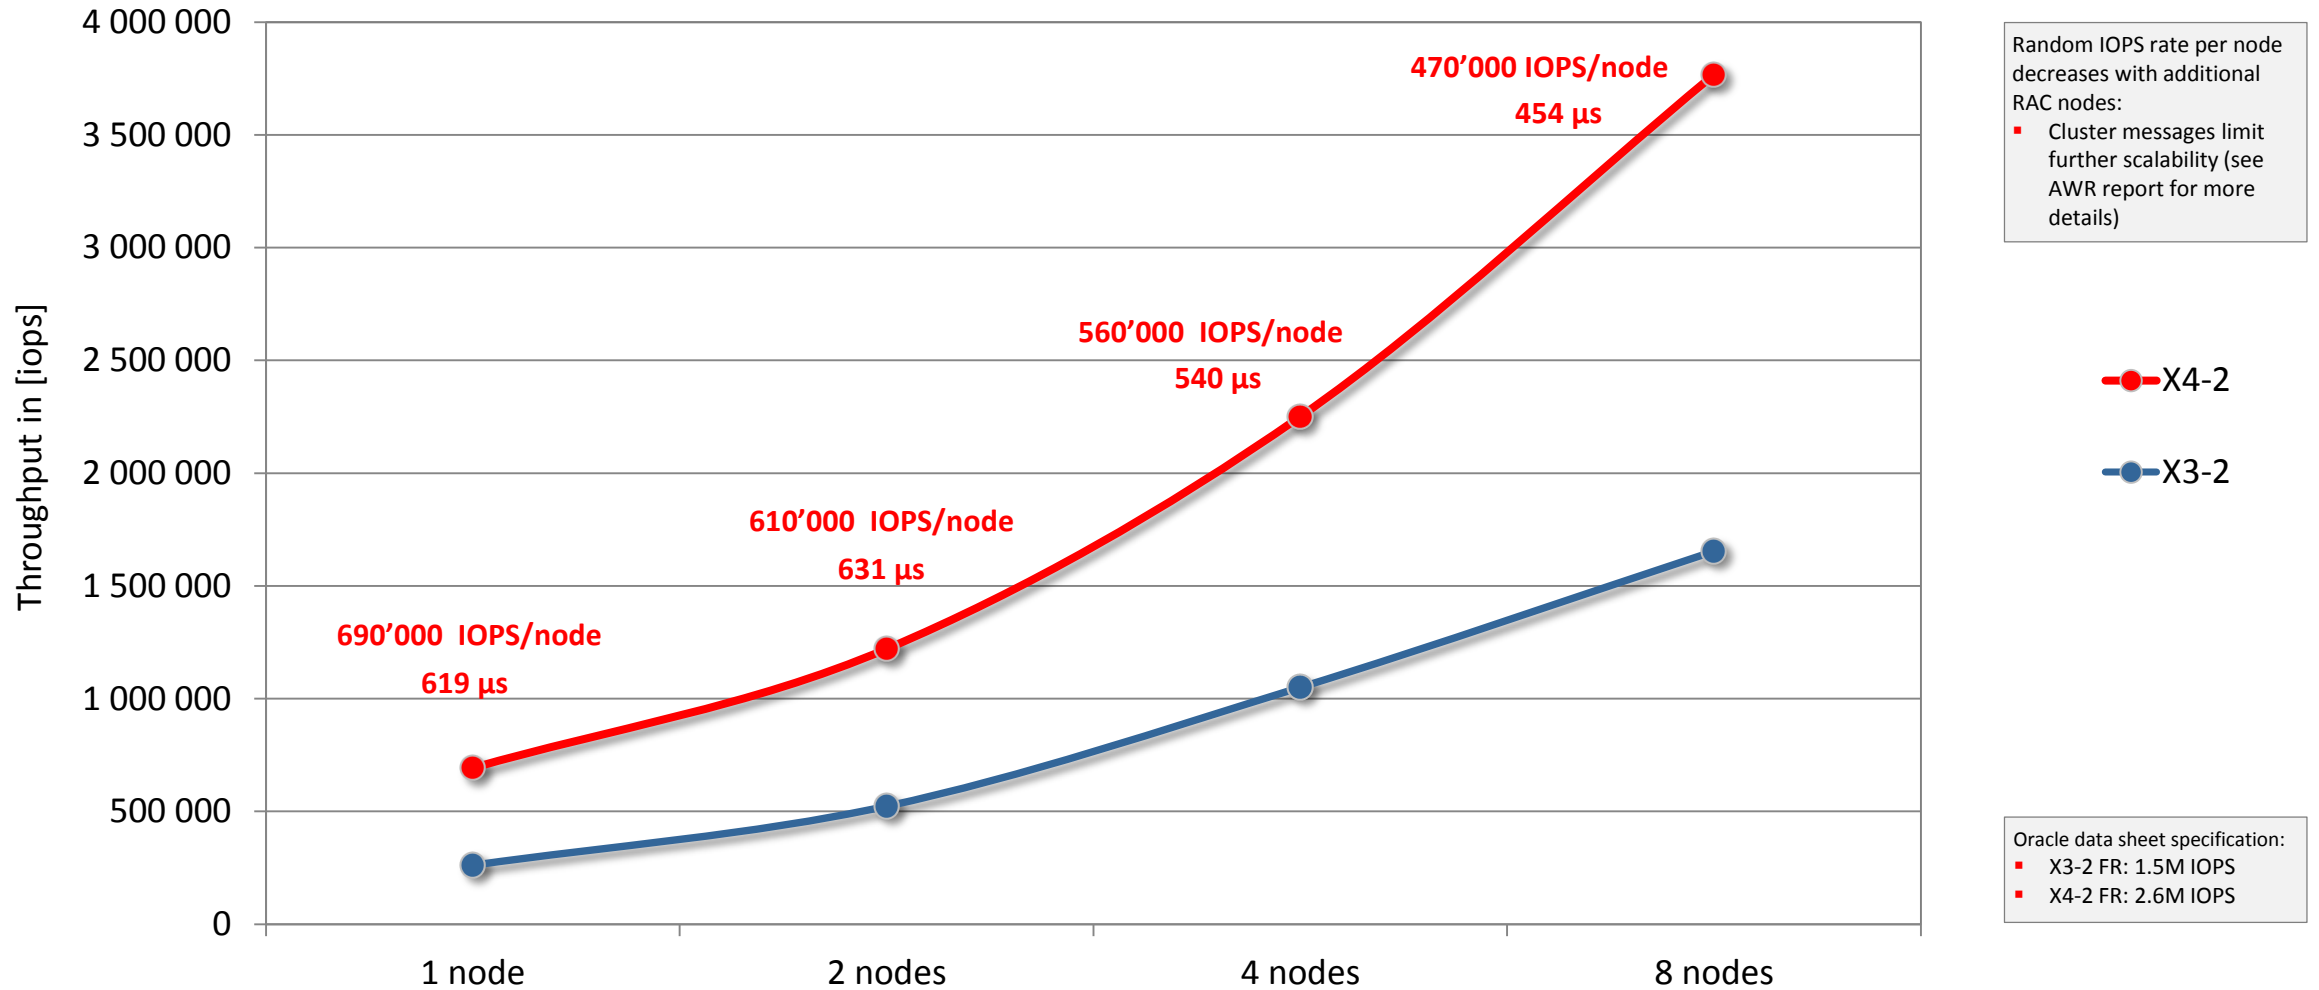
Run	benchmark run id	#N	number of RAC nodes	[rps]	rows per second	[iops]	i/o operations per second	[s]	time in seconds
Tst	benchmark test id	#J	number of load generators (jobs)	[tps]	transactions per second	[dbps]	database blocks per second	[ms]	time in milli seconds
Code	benchmark test code	#T	number of threads (PX)	[ops]	operations per second	[MBps]	mega byte per second	[μs]	time in micro seconds

- CPU of one server completely utilized
- Ø 28'900 IOPS per x86 core
- Its hard to get 100% cell flash cache hit rate – even if database is much smaller than flash capacity


Storage Benchmark Results



Oracle random read: Cluster



Storage Benchmark Results




Oracle random read: Cluster

Exadata X4-2 FR HC

Run	Tst	Code	#N	#J	#T	CPU busy [%]	CPU sys [%]	Physical read [iops]	Physical read [dbps]	Physical read [MBps]	Physical write [iops]	Physical write [dbps]	Physical write [MBps]	REDO write [iops]	Hitrate db flash [%]	Hitrate exa flash [%]	Elap time [s]
2	46	STO-64	8	4096	1	30	7	3581770	3581749	27983	150	75	1	96	0	86	310
	47	STO-64	8	2048	1	31	8	3767008	3766988	29430	99	51	1	53	0	88	307
	48	STO-64	4	2048	1	36	9	2106834	2106814	16460	96	62	1	47	0	89	332
	49	STO-64	4	1024	1	37	10	2251293	2251273	17588	69	48	1	25	0	91	322
	50	STO-64	2	1024	1	23	6	1064359	1064339	8315	50	37	1	22	0	90	336
	51	STO-64	2	512	1	24	6	1221211	1221192	9541	45	37	0	13	0	91	329
	52	STO-64	1	1	1	1	0	79092	79073	618	10	12	0	1	0	96	298
	53	STO-64	1	2	1	1	0	128788	128770	1006	7	8	0	1	0	96	296
	54	STO-64	1	4	1	1	0	228344	228325	1784	16	17	0	2	0	96	308
	55	STO-64	1	8	1	1	0	355593	355573	2778	15	19	0	1	0	95	304
	56	STO-64	1	16	1	2	1	433914	433895	3390	18	19	0	3	0	95	307
	57	STO-64	1	32	1	3	1	496501	496482	3879	19	22	0	2	0	93	305
	58	STO-64	1	64	1	5	1	567283	567264	4432	24	25	0	4	0	93	307
	59	STO-64	1	128	1	9	2	660586	660568	5161	27	27	0	3	0	93	316
	60	STO-64	1	256	1	12	3	693597	693577	5419	31	31	0	6	0	92	333
	61	STO-64	1	512	1	11	3	598168	598148	4673	33	29	0	9	0	91	362


Legend:

Run	benchmark run id	#N	number of RAC nodes	[rps]	rows per second	[iops]	i/o operations per second	[s]	time in seconds
Tst	benchmark test id	#J	number of load generators (jobs)	[tps]	transactions per second	[dbps]	database blocks per second	[ms]	time in milli seconds
Code	benchmark test code	#T	number of threads (PX)	[ops]	operations per second	[MBps]	mega byte per second	[μs]	time in micro seconds



- 1 Introduction to Storage (STO) Performance Tests
- 2 Storage System Configuration
- 3 Storage Benchmark Results – Sequential I/O
- 4 Storage Benchmark Results – Random I/O
- 5 Reviewing Storage Benchmark Results**

Reviewing Storage Benchmark Results



Summary

	Metric		X3-2	X4-2
Sequential Read I/O				
▪ 1 User Process	[GBps]		24.7	32.3
▪ 1 Database Server	[GBps]		85.0 (87%)	95.8 (94%)
▪ Cluster	[GBps]		98.0	102.4
Random Read I/O				
▪ Throughput 1 Database Server	[IOPS]		261'000	693'000
▪ Service time	[μs]		-	619
Random Read I/O				
▪ Throughput Cluster	[IOPS]		1'658'000	3'767'000
▪ Service time	[μs]		-	454

BENCHWARE

swiss precision in performance measurement

www.benchmarkware.ch

info@benchmarkware.ch