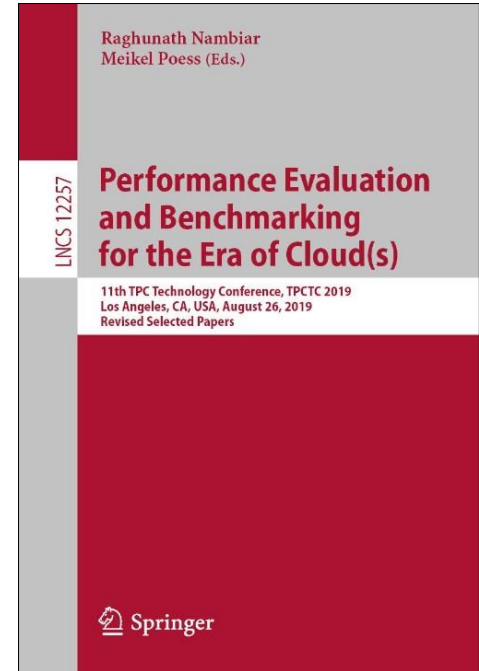


peakmarks® Performance Study on Exadata PMEM for Log Writer

December 2021



peakmarks® showcased its software at the
11th TPC Technology Conference 2019.



peakmarks® Software and its documentation are protected under intellectual property laws. The license agreement states that explicit permission is mandatory for any use, display, modification, distribution, transmission, licensing, transfer, publication, or demonstration of the peakmarks® Software and its documentation. The license also strictly prohibits reengineering, disassembling, or decompiling of the software.

peakmarks® is a registered trademark. Other names may be trademarks of their respective owners.



All performance data were determined with the peakmarks® software under certain conditions and do not necessarily correspond to the manufacturer's specifications.



[MBps] megabyte per second

[GBps] gigabyte per second

[dbps] database blocks per second

[rbps] redo blocks per second

[dbpt] database blocks per transaction

[kBpt] kilobyte per transaction

[s] seconds

[ms] milliseconds

[μ s] microseconds

[IOPS] I/O operations per second

[qps] queries per second

[rps] rows per second

[tps] transactions per second

[Mops] million operations per second

Nodes number of cluster nodes

Jobs number of workload processes

BuCache Database Buffer Cache

FlCache Database or Exadata Flash Cache



Performance is not everything.
But without performance, everything is worth nothing.

Platform description



Storage	Exadata X5-2 Quarter Rack High Capacity	Exadata X9M-2 Quarter Rack High Capacity	2-node x86 Cluster with All-Flash-Array
Storage system	3 x Scale-Out Exadata X5 Storage Server High Capacity	3 x Scale-Out Exadata X9 Storage Server High Capacity	All-Flash-Array
Launch date	2015	2021	2021
DRAM capacity, total	288 GByte	756 GByte	1,280 GByte
PMEM capacity, total	-	4.5 TByte	-
Flash capacity, total raw	19.2 TByte	76.8 TByte	274 TByte
Disk capacity, total raw	144 TByte	648 TByte	-
Connectivity	Infiniband, 2 x 40 Gbps	RDMA over Converged Ethernet, 2 x 100 Gbps	NVME over Fibre Channel, 8 x 32 Gbps
File system	ASM, normal redundancy ASM allocation unit 4 MByte	ASM, normal redundancy ASM allocation unit 4 MByte	ASM, external redundancy ASM allocation unit 4 MByte
Deduplication	No	No	Yes
Compression	No	No	Yes



peakmarks® Software	Exadata X5-2 Quarter Rack High Capacity	Exadata X9M-2 Quarter Rack High Capacity	2-node x86 Cluster with All-Flash-Array
Version	9.6	9.6	9.6
Build	210601	211201	211201
Database size	8 TByte	8 TByte	8 TByte

Note

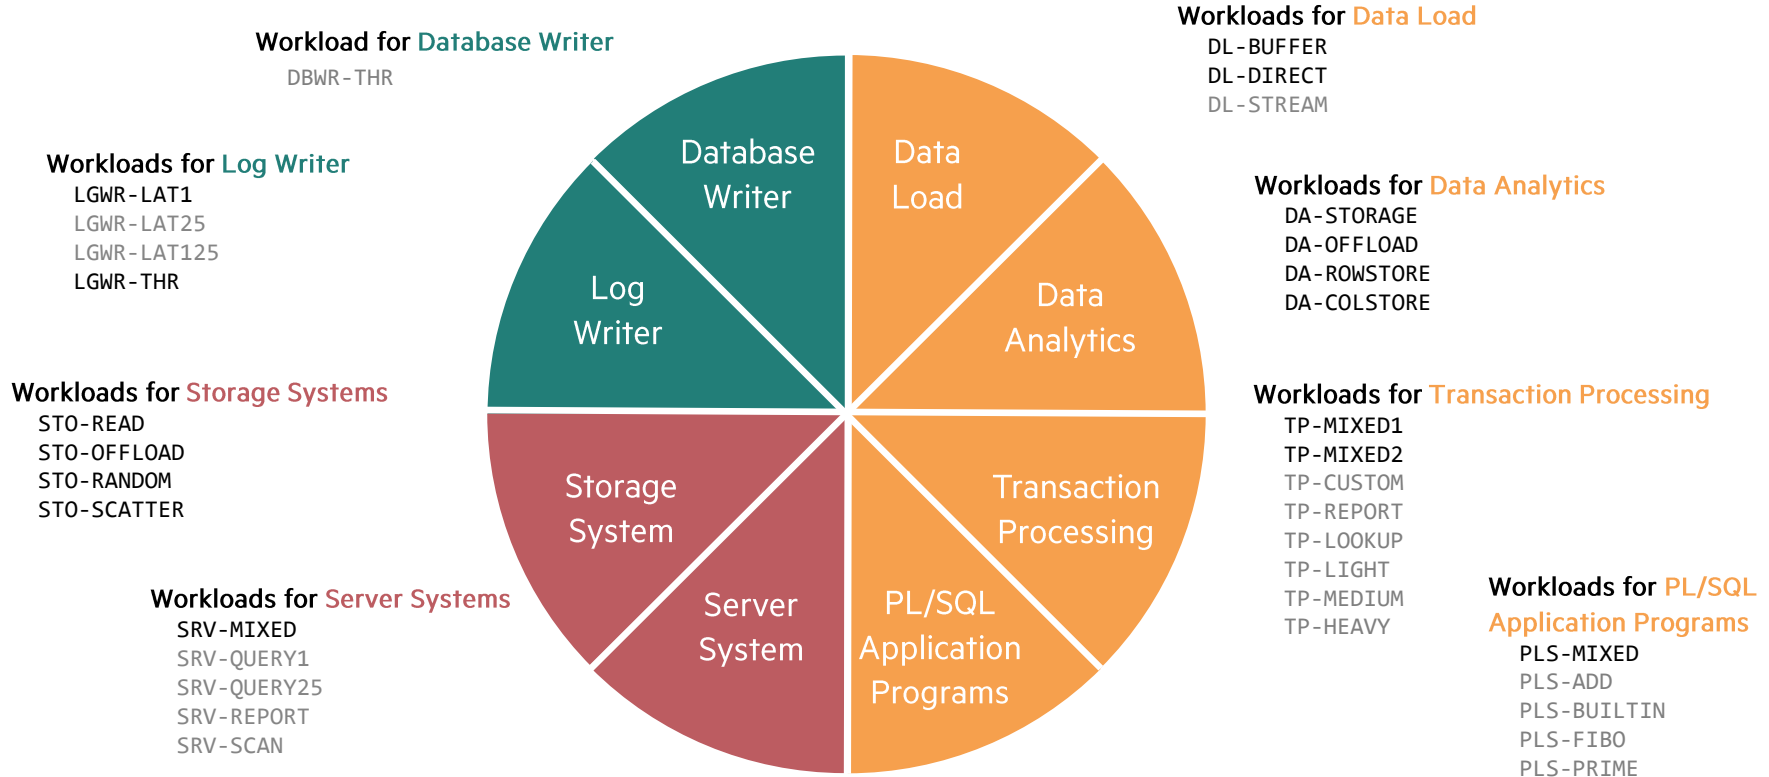
- To ensure full transparency, the peakmarks® software generates individual Oracle AWR reports for each performance test.



Simple. Representative. Fast.

peakmarks[®] Workload Overview

More than 30 micro-benchmarks in 8 workload groups





Stop guessing. Start measuring.

Workloads to determine the Log Writer Performance





Motivation

Log Writer processes are primarily responsible for transaction management and database consistency in case of failures.

These processes are critical to overall Oracle performance, especially with a high transaction load.

The goal is to

- Optimize throughput and latency of REDO log writers
- Validate the impact of Oracle Data Guard on local transaction performance
- Validate the impact of several other factors on log writer performance, such as data deduplication, data compression, usage of ASM redundancy level, etc.



Key Performance Metrics

- **SQL commit throughput** in transactions per second [tps]
- **SQL commit latency** in milliseconds [ms]
- **Log writer throughput** in megabyte per second [MBps]



Description

Workload	Measurement Unit	Action
LGWR-LAT1	[tps]	Small-sized transaction; workload uses COMMIT WRITE WAIT IMMEDIATE.
	[ms]	This workload shows the maximum commit rate and the minimum commit latency for <u>small-sized</u> transactions generating between 1 kByte and 2 kByte redo data per transaction.
LGWR-LAT25	[tps]	Medium-sized transaction; workload uses COMMIT WRITE WAIT IMMEDIATE.
	[ms]	This workload shows the maximum commit rate and the minimum commit latency for <u>medium-sized</u> transactions generating around 25 kByte redo data per transaction.
LGWR-LAT125	[tps]	Large-sized transaction size; workload uses COMMIT WRITE WAIT IMMEDIATE.
	[ms]	This workload shows the maximum commit rate and the minimum commit latency for <u>larger-sized</u> transactions generating around 125 kByte redo data per transaction. This transaction size is typical for payment systems in banking applications.
LGWR-THR	[MBps]	Huge transaction; workload uses COMMIT WRITE NOWAIT BATCH. This workload shows the maximum redo data written by log writer processes.

Notes

- These kinds of transactions are generic to all applications in all industries.
- The actual REDO data volume per transaction can vary depending on the database size and the usage of RAC technology. Therefore, the exact value of this metric is explicitly shown in the peakmarks reports.

Log Writer Performance





Workload LGWR-LAT1: maximum commit throughput for small-sized transactions

Exadata X5-2
2015

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
5	34	LGWR-LAT1	2	2	2	1	1	98	0	3,949	0.506	17,191	3,967	3,949	8	2.07	0.363	92.63	301
	35	LGWR-LAT1	2	12	6	4	2	94	0	19,610	0.610	86,907	9,415	19,611	37	1.93	0.450	93.48	302
	36	LGWR-LAT1	2	24	8	6	2	92	0	33,804	0.708	150,907	8,766	33,806	64	1.94	0.544	93.24	302
	37	LGWR-LAT1	2	36	10	8	2	90	0	45,805	0.784	207,555	7,976	45,807	87	1.94	0.632	92.01	302
	38	LGWR-LAT1	2	48	12	10	2	88	0	57,159	0.837	258,902	7,615	57,163	109	1.95	0.701	91.01	302
	39	LGWR-LAT1	2	60	14	12	2	86	0	67,613	0.885	303,847	7,272	67,617	129	1.95	0.765	88.95	302
	40	LGWR-LAT1	2	72	16	13	2	84	0	77,783	0.922	345,624	6,813	77,788	148	1.95	0.812	87.06	302

Exadata X9M-2
With PMEM 2021

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
4	192	LGWR-LAT1	2	2	2	1	0	98	0	17,085	0.117	74,132	17,134	17,085	32	1.92	0.047	24.78	302
	194	LGWR-LAT1	2	16	7	6	1	93	0	118,158	0.135	549,506	49,541	118,160	226	1.96	0.058	29.76	302
	196	LGWR-LAT1	2	32	11	10	1	89	0	211,878	0.151	1,000,856	43,353	211,881	406	1.96	0.071	21.71	302
	198	LGWR-LAT1	2	48	16	14	1	84	0	283,116	0.169	1,317,684	36,121	283,119	541	1.96	0.085	22.98	302
	200	LGWR-LAT1	2	64	20	18	1	80	0	332,211	0.192	1,513,514	30,699	332,215	633	1.95	0.102	23.99	302
	201	LGWR-LAT1	2	72	22	20	1	78	0	350,942	0.204	1,582,717	28,571	350,947	668	1.95	0.111	24.28	302
	202	LGWR-LAT1	2	80	24	21	1	76	0	367,848	0.217	1,644,659	26,908	367,853	699	1.95	0.120	24.88	302

Note

- Small-sized transactions generate between 1 kByte and 2 kByte redo data.



Workload LGWR-LAT1: maximum commit throughput for small-sized transactions

x86 Server with AFA
2021

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
2	155	LGWR-LAT1	2	2	1	0	0	99	0	6,835	0.292	7,179	6,874	6,835	13	1.95	0.205	0.00	301
	156	LGWR-LAT1	2	8	1	1	0	99	0	12,509	0.639	12,537	6,431	12,509	24	1.96	0.531	0.00	301
	157	LGWR-LAT1	2	16	1	1	0	99	0	14,125	1.131	13,898	3,623	14,126	27	1.96	1.001	0.00	301
	158	LGWR-LAT1	2	24	1	1	0	99	0	14,682	1.630	15,165	2,506	14,683	28	1.95	1.472	0.00	302
	159	LGWR-LAT1	2	32	1	1	0	99	0	20,049	1.593	21,257	2,570	20,052	39	1.99	1.422	0.00	302

Exadata X9M-2
With PMEM 2021

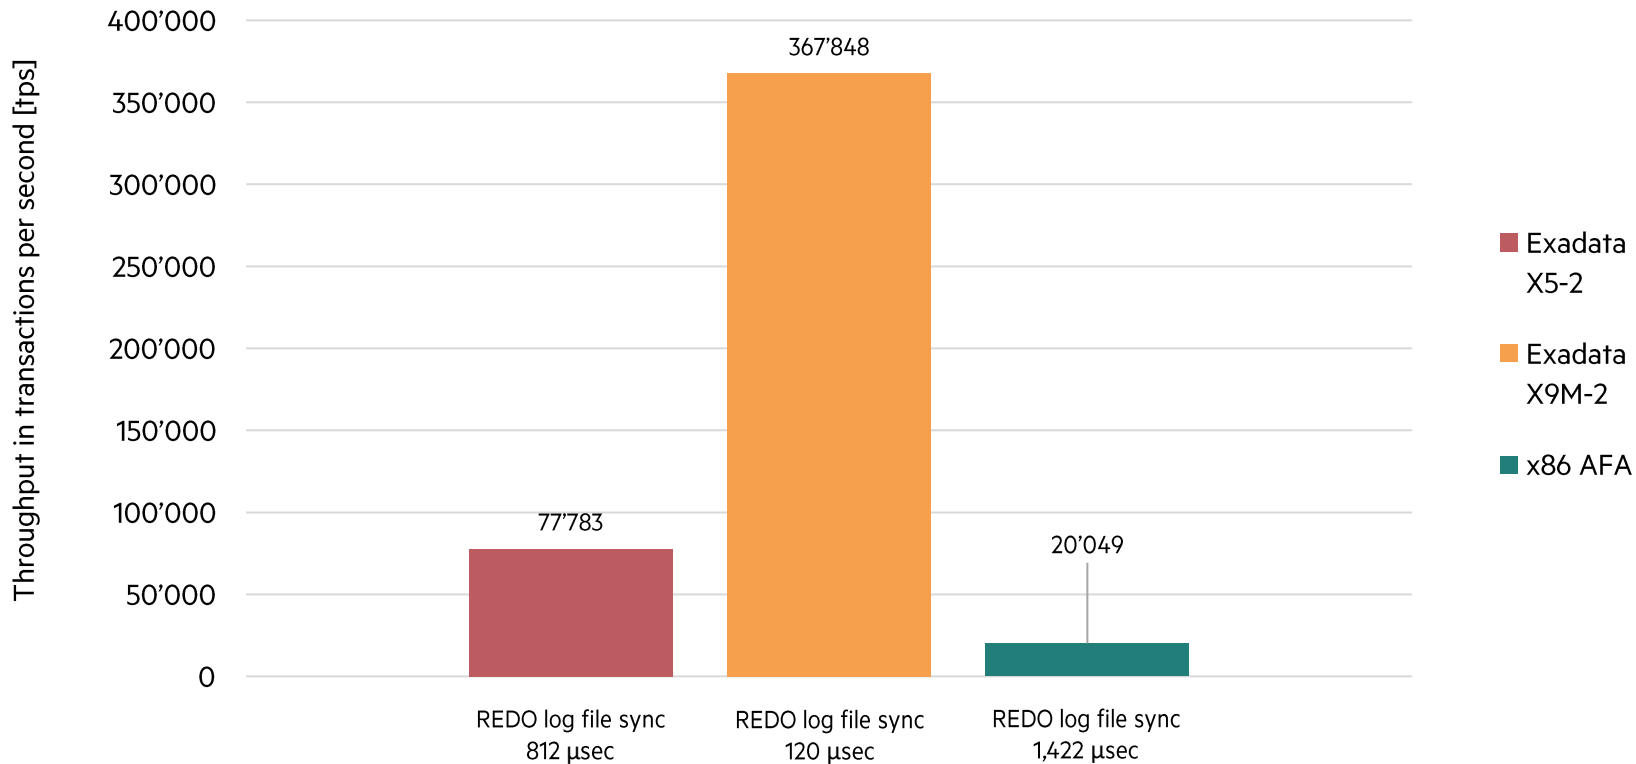
Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
4	192	LGWR-LAT1	2	2	2	1	0	98	0	17,085	0.117	74,132	17,134	17,085	32	1.92	0.047	24.78	302
	194	LGWR-LAT1	2	16	7	6	1	93	0	118,158	0.135	549,506	49,541	118,160	226	1.96	0.058	29.76	302
	196	LGWR-LAT1	2	32	11	10	1	89	0	211,878	0.151	1,000,856	43,353	211,881	406	1.96	0.071	21.71	302
	198	LGWR-LAT1	2	48	16	14	1	84	0	283,116	0.169	1,317,684	36,121	283,119	541	1.96	0.085	22.98	302
	200	LGWR-LAT1	2	64	20	18	1	80	0	332,211	0.192	1,513,514	30,699	332,215	633	1.95	0.102	23.99	302
	201	LGWR-LAT1	2	72	22	20	1	78	0	350,942	0.204	1,582,717	28,571	350,947	668	1.95	0.111	24.28	302
	202	LGWR-LAT1	2	80	24	21	1	76	0	367,848	0.217	1,644,659	26,908	367,853	699	1.95	0.120	24.88	302

Note

- Small-sized transactions generate between 1 kByte and 2 kByte redo data.



Workload LGWR-LAT1: maximum commit throughput for small-sized transactions





Workload LGWR-LAT25: maximum commit throughput for medium-sized transactions

Exadata X5-2
2015

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
5	41	LGWR-LAT25	2	2	2	2	1	98	0	1,924	1.038	105,283	1,952	1,924	49	26.08	0.455	90.97	301
	42	LGWR-LAT25	2	12	8	6	2	92	0	8,625	1.387	477,167	6,086	8,626	222	26.36	0.670	84.37	302
	43	LGWR-LAT25	2	24	11	9	2	89	0	13,629	1.755	754,477	4,697	13,631	350	26.30	1.004	70.38	302
	44	LGWR-LAT25	2	36	14	12	2	86	0	17,242	2.083	953,326	4,004	17,245	442	26.25	1.294	61.92	302
	45	LGWR-LAT25	2	48	17	14	2	83	0	19,895	2.406	1,095,184	3,264	19,898	510	26.25	1.583	54.03	302
	46	LGWR-LAT25	2	60	19	16	2	81	0	22,391	2.672	1,231,362	3,164	22,395	573	26.20	1.807	50.17	302
	47	LGWR-LAT25	2	72	21	18	2	79	0	24,124	2.977	1,327,739	3,156	24,129	619	26.27	2.087	49.20	302

Exadata X9M-2
With PMEM 2021

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
4	224	LGWR-LAT25	2	2	2	1	0	98	0	2,927	0.683	160,312	2,996	2,927	75	26.24	0.078	27.45	300
	225	LGWR-LAT25	2	8	4	4	0	96	0	11,312	0.706	625,249	10,226	11,313	290	26.25	0.079	30.33	302
	226	LGWR-LAT25	2	16	8	7	1	92	0	21,187	0.753	1,180,312	14,207	21,188	543	26.24	0.112	31.24	302
	227	LGWR-LAT25	2	24	11	9	1	89	0	29,580	0.809	1,654,902	14,700	29,582	760	26.31	0.158	32.13	302
	228	LGWR-LAT25	2	32	13	12	1	87	0	36,313	0.878	2,031,070	13,548	36,316	932	26.28	0.213	32.93	302
	229	LGWR-LAT25	2	40	16	14	1	84	0	41,324	0.965	2,303,633	11,738	41,327	1,060	26.27	0.282	33.50	302
	230	LGWR-LAT25	2	48	17	16	1	83	0	44,681	1.071	2,481,798	10,026	44,684	1,146	26.26	0.361	33.42	302
	231	LGWR-LAT25	2	56	19	17	1	81	0	47,053	1.186	2,604,509	8,591	47,057	1,206	26.25	0.445	34.09	302
	232	LGWR-LAT25	2	64	20	18	1	80	0	48,589	1.314	2,682,045	7,326	48,593	1,245	26.24	0.537	34.96	302

Note

- Medium-sized transactions generate around 25 kByte redo data.



Workload LGWR-LAT25: maximum commit throughput for medium-sized transactions

x86 Server with AFA
2021

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn [IOPS]	REDO writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
2	187	LGWR-LAT25	2	2	1	1	0	99	0	2,045	0.976	14,701	2,177	2,045	53	26.54	0.291	0.00	301	
	188	LGWR-LAT25	2	8	1	1	0	99	0	4,923	1.622	37,714	3,359	4,924	127	26.42	0.847	0.00	301	
	189	LGWR-LAT25	2	16	2	2	0	98	0	6,739	2.372	52,276	2,208	6,740	174	26.44	1.537	0.00	301	
	190	LGWR-LAT25	2	24	3	3	0	97	0	8,426	2.843	64,249	1,697	8,428	219	26.61	1.903	0.00	301	
	191	LGWR-LAT25	2	32	4	4	0	96	0	9,790	3.262	72,933	1,442	9,792	254	26.57	2.185	0.00	302	
	192	LGWR-LAT25	2	40	5	4	0	95	0	10,569	3.777	77,834	1,237	10,572	275	26.64	2.565	0.00	302	
	193	LGWR-LAT25	2	48	6	5	0	94	0	11,070	4.326	80,808	1,088	11,073	288	26.64	2.931	0.00	302	
	194	LGWR-LAT25	2	56	7	6	0	93	0	11,885	4.702	85,823	976	11,889	309	26.62	3.284	0.00	302	
	195	LGWR-LAT25	2	64	8	7	0	92	0	12,368	5.163	88,710	879	12,372	322	26.66	3.650	0.00	302	
	196	LGWR-LAT25	2	72	9	8	1	91	0	12,932	5.557	91,978	805	12,937	336	26.61	3.963	0.00	301	
	197	LGWR-LAT25	2	80	9	8	1	91	0	13,790	5.790	97,759	772	13,795	359	26.66	4.161	0.00	301	
	198	LGWR-LAT25	2	88	10	9	1	90	0	13,962	6.284	98,878	707	13,968	365	26.77	4.496	0.00	302	

Exadata X9M-2
With PMEM 2021

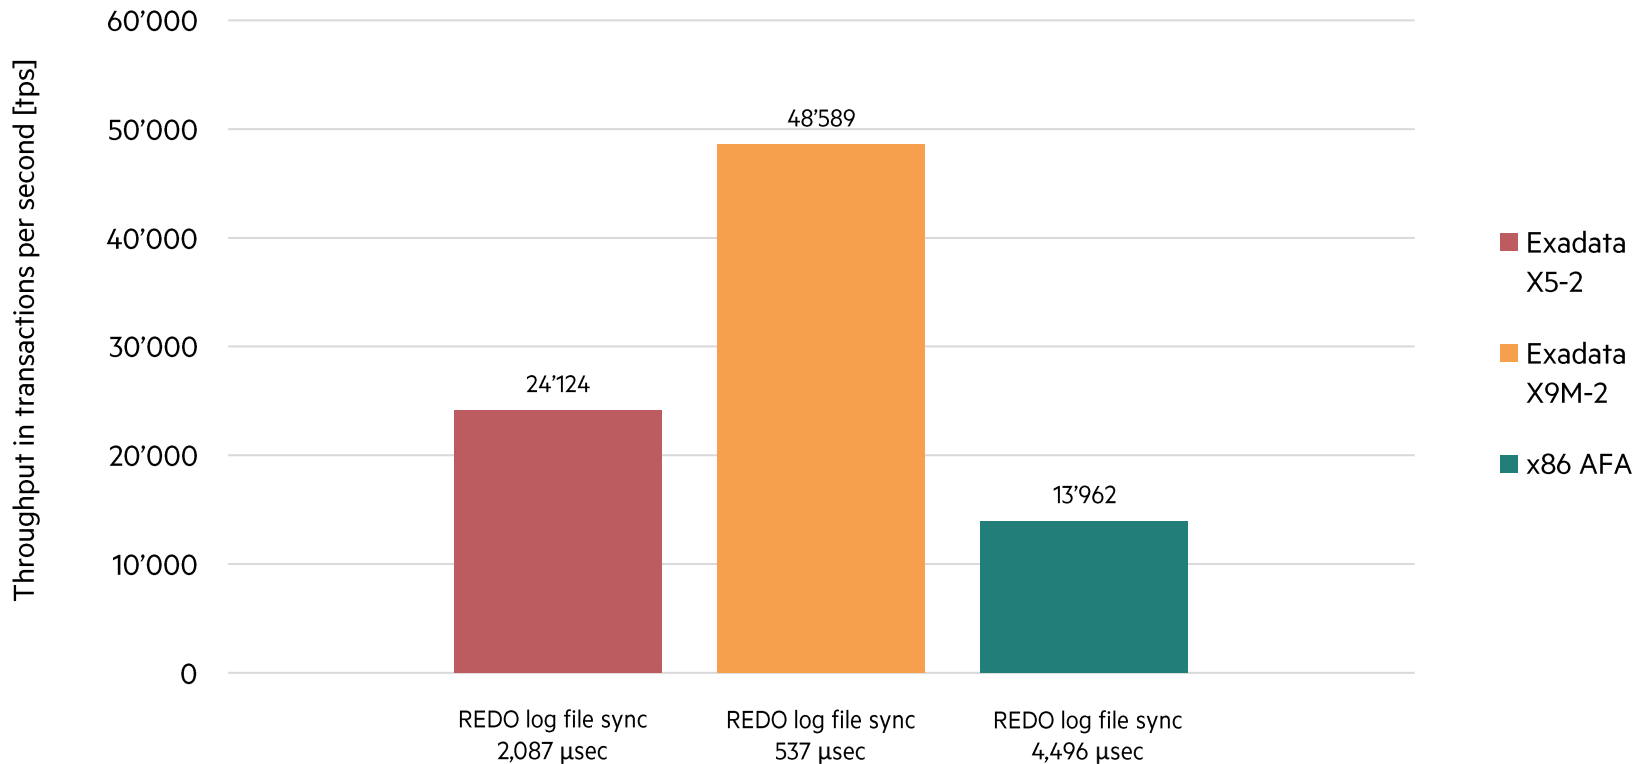
Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn [IOPS]	REDO writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
4	224	LGWR-LAT25	2	2	2	1	0	98	0	2,927	0.683	160,312	2,996	2,927	75	26.24	0.078	27.45	300	
	225	LGWR-LAT25	2	8	4	4	0	96	0	11,312	0.706	625,249	10,226	11,313	290	26.25	0.079	30.33	302	
	226	LGWR-LAT25	2	16	8	7	1	92	0	21,187	0.753	1,180,312	14,207	21,188	543	26.24	0.112	31.24	302	
	227	LGWR-LAT25	2	24	11	9	1	89	0	29,580	0.809	1,654,902	14,700	29,582	760	26.31	0.158	32.13	302	
	228	LGWR-LAT25	2	32	13	12	1	87	0	36,313	0.878	2,031,070	13,548	36,316	932	26.28	0.213	32.93	302	
	229	LGWR-LAT25	2	40	16	14	1	84	0	41,324	0.965	2,303,633	11,738	41,327	1,060	26.27	0.282	33.50	302	
	230	LGWR-LAT25	2	48	17	16	1	83	0	44,681	1.071	2,481,798	10,026	44,684	1,146	26.26	0.361	33.42	302	
	231	LGWR-LAT25	2	56	19	17	1	81	0	47,053	1.186	2,604,509	8,591	47,057	1,206	26.25	0.445	34.09	302	
	232	LGWR-LAT25	2	64	20	18	1	80	0	48,589	1.314	2,682,045	7,326	48,593	1,245	26.24	0.537	34.96	302	

Note

- Medium-sized transactions generate around 25 kByte redo data.



Workload LGWR-LAT25: maximum commit throughput for medium-sized transactions





Workload LGWR-LAT125: maximum commit throughput for large-sized transactions

Exadata X5-2
2015

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
5	48	LGWR-LAT125	2	2	2	2	1	98	0	669	2.985	171,583	703	669	81	123.98	0.692	68.13	301
	49	LGWR-LAT125	2	12	9	7	1	91	0	3,039	3.935	783,628	2,760	3,039	368	124.00	1.129	52.33	302
	50	LGWR-LAT125	2	24	15	13	2	85	0	4,994	4.790	1,290,444	3,404	4,996	604	123.85	1.718	52.43	302
	51	LGWR-LAT125	2	36	19	17	2	81	0	6,066	5.915	1,565,086	2,650	6,069	734	123.91	2.527	46.34	302
	52	LGWR-LAT125	2	48	23	20	2	77	0	6,967	6.870	1,795,180	2,104	6,970	843	123.90	3.215	43.89	302
	53	LGWR-LAT125	2	60	27	23	3	73	0	7,936	7.534	2,042,004	1,801	7,940	960	123.87	3.723	44.39	302
	54	LGWR-LAT125	2	72	29	25	3	71	0	8,455	8.482	2,174,396	1,549	8,459	1,023	123.90	4.413	45.74	302

Exadata X9M-2
With 2021

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
4	256	LGWR-LAT125	2	2	1	1	0	99	0	688	2.901	175,535	754	688	83	123.53	0.180	29.17	301
	257	LGWR-LAT125	2	8	4	4	0	96	0	2,630	3.039	674,278	2,549	2,631	317	123.43	0.231	31.65	301
	258	LGWR-LAT125	2	16	7	7	1	93	0	4,935	3.231	1,268,571	3,817	4,936	595	123.46	0.352	33.72	302
	259	LGWR-LAT125	2	24	11	9	1	89	0	6,940	3.444	1,786,436	4,297	6,942	837	123.50	0.504	34.16	302
	260	LGWR-LAT125	2	32	13	12	1	87	0	8,571	3.724	2,208,063	4,132	8,573	1,035	123.65	0.715	35.30	302
	261	LGWR-LAT125	2	40	15	14	1	85	0	9,766	4.085	2,513,780	3,680	9,768	1,178	123.52	0.979	35.55	302
	262	LGWR-LAT125	2	48	17	15	1	83	0	10,522	4.551	2,704,314	3,018	10,525	1,269	123.50	1.324	35.87	302
	263	LGWR-LAT125	2	56	18	16	1	82	0	10,789	5.177	2,767,696	2,286	10,793	1,301	123.48	1.798	36.40	302

Note

- Large-sized transactions generate around 125 kByte redo data.



Workload LGWR-LAT125: maximum commit throughput for large-sized transactions

x86 Server with AFA
2021

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn [IOPS]	REDO writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
2	219	LGWR-LAT125	2	2	1	1	0	99	0	558	3.580	17,968	742	558	68	124.79	0.477	0.00	301	
	221	LGWR-LAT125	2	16	4	3	0	96	0	2,522	6.327	86,463	1,286	2,523	307	124.65	2.610	0.00	301	
	223	LGWR-LAT125	2	32	8	7	0	92	0	3,815	8.370	125,392	878	3,817	466	125.08	3.795	0.00	302	
	225	LGWR-LAT125	2	48	11	10	1	89	0	4,760	10.054	154,058	606	4,763	583	125.42	5.121	0.00	302	
	227	LGWR-LAT125	2	64	12	11	1	88	0	5,431	11.745	174,267	484	5,435	665	125.38	6.292	0.00	302	
	228	LGWR-LAT125	2	72	13	12	1	87	0	5,695	12.576	182,089	442	5,699	696	125.15	6.790	0.00	302	
	229	LGWR-LAT125	2	80	12	11	1	88	0	6,592	12.085	209,507	470	6,597	802	124.58	6.414	0.00	302	
	230	LGWR-LAT125	2	88	13	12	1	87	0	7,214	12.143	228,178	454	7,220	875	124.20	6.538	0.00	302	
	231	LGWR-LAT125	2	96	13	12	1	87	0	7,575	12.619	239,596	423	7,581	921	124.50	6.910	0.00	302	
	232	LGWR-LAT125	2	104	14	13	1	86	0	7,911	13.087	249,574	402	7,918	960	124.26	7.279	0.00	302	
	233	LGWR-LAT125	2	112	14	13	1	86	0	8,027	13.877	253,506	385	8,035	976	124.51	7.756	0.00	303	

Exadata X9M-2
With 2021

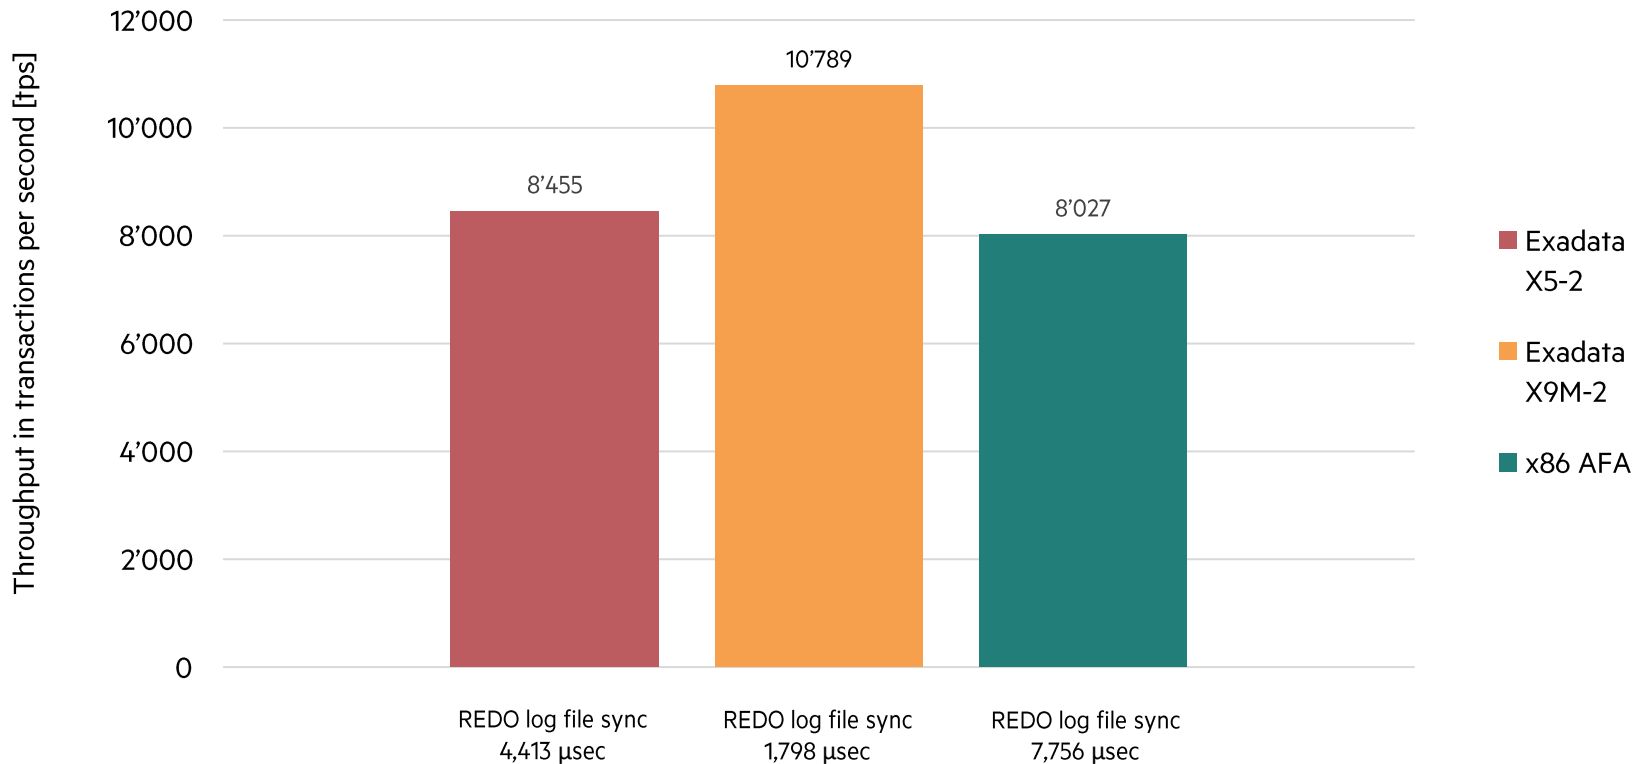
Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn [IOPS]	REDO writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
4	256	LGWR-LAT125	2	2	1	1	0	99	0	688	2.901	175,535	754	688	83	123.53	0.180	29.17	301	
	257	LGWR-LAT125	2	8	4	4	0	96	0	2,630	3.039	674,278	2,549	2,631	317	123.43	0.231	31.65	301	
	258	LGWR-LAT125	2	16	7	7	1	93	0	4,935	3.231	1,268,571	3,817	4,936	595	123.46	0.352	33.72	302	
	259	LGWR-LAT125	2	24	11	9	1	89	0	6,940	3.444	1,786,436	4,297	6,942	837	123.50	0.504	34.16	302	
	260	LGWR-LAT125	2	32	13	12	1	87	0	8,571	3.724	2,208,063	4,132	8,573	1,035	123.65	0.715	35.30	302	
	261	LGWR-LAT125	2	40	15	14	1	85	0	9,766	4.085	2,513,780	3,680	9,768	1,178	123.52	0.979	35.55	302	
	262	LGWR-LAT125	2	48	17	15	1	83	0	10,522	4.551	2,704,314	3,018	10,525	1,269	123.50	1.324	35.87	302	
	263	LGWR-LAT125	2	56	18	16	1	82	0	10,789	5.177	2,767,696	2,286	10,793	1,301	123.48	1.798	36.40	302	

Note

- Large-sized transactions generate around 125 kByte redo data.



Workload LGWR-LAT125: maximum commit throughput for large-sized transactions





Workload LGWR-THR: maximum throughput of REDO data

Exadata X5-2
2015

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn [IOPS]	REDO writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
5	55	LGWR-THR	2	2	3	2	1	97	0	201	9.916	200,679	259	202	95	483.98	1.323	33.14	301	
	56	LGWR-THR	2	12	10	8	1	90	0	970	12.325	969,991	1,268	972	457	482.44	1.875	46.34	302	
	57	LGWR-THR	2	24	18	15	2	82	0	1,631	14.641	1,632,897	1,754	1,633	768	482.18	2.967	43.78	303	
	58	LGWR-THR	2	36	23	20	2	77	0	2,082	17.194	2,086,533	1,656	2,086	982	482.98	4.009	42.20	303	
	59	LGWR-THR	2	48	26	22	3	74	0	2,124	22.474	2,135,336	1,171	2,130	1,006	485.00	7.532	53.32	303	
	60	LGWR-THR	2	60	26	23	3	74	0	2,100	28.497	2,107,466	856	2,107	994	484.69	12.978	60.80	301	
	61	LGWR-THR	2	72	27	23	3	73	0	2,082	34.446	2,089,529	693	2,091	986	484.95	17.607	60.30	302	

Exadata X9M-2
With PMEM 2021

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn [IOPS]	REDO writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
4	288	LGWR-THR	2	2	1	1	0	99	0	180	11.118	179,031	269	180	85	483.56	0.492	27.91	301	
	289	LGWR-THR	2	8	4	3	0	96	0	691	11.531	689,476	875	692	325	481.62	0.762	30.04	302	
	290	LGWR-THR	2	16	7	6	1	93	0	1,275	12.518	1,273,946	1,259	1,277	601	482.69	1.459	31.61	302	
	291	LGWR-THR	2	24	10	9	1	90	0	1,773	13.506	1,772,599	1,483	1,776	835	482.26	2.136	32.39	302	
	292	LGWR-THR	2	32	12	11	1	88	0	2,167	14.724	2,167,300	1,489	2,171	1,021	482.47	2.955	33.20	302	
	293	LGWR-THR	2	40	14	13	1	86	0	2,500	15.943	2,503,623	1,439	2,505	1,180	483.33	3.662	34.24	302	
	294	LGWR-THR	2	48	17	15	1	83	0	2,790	17.141	2,789,493	1,258	2,796	1,315	482.64	4.337	35.02	302	
	295	LGWR-THR	2	56	18	17	1	82	0	3,015	18.490	3,018,031	1,029	3,022	1,423	483.30	4.976	35.67	303	
	296	LGWR-THR	2	64	20	18	1	80	0	3,140	20.289	3,140,976	796	3,147	1,482	483.30	5.680	36.48	302	



Workload LGWR-THR: maximum throughput of REDO data

x86 Server with AFA
2021

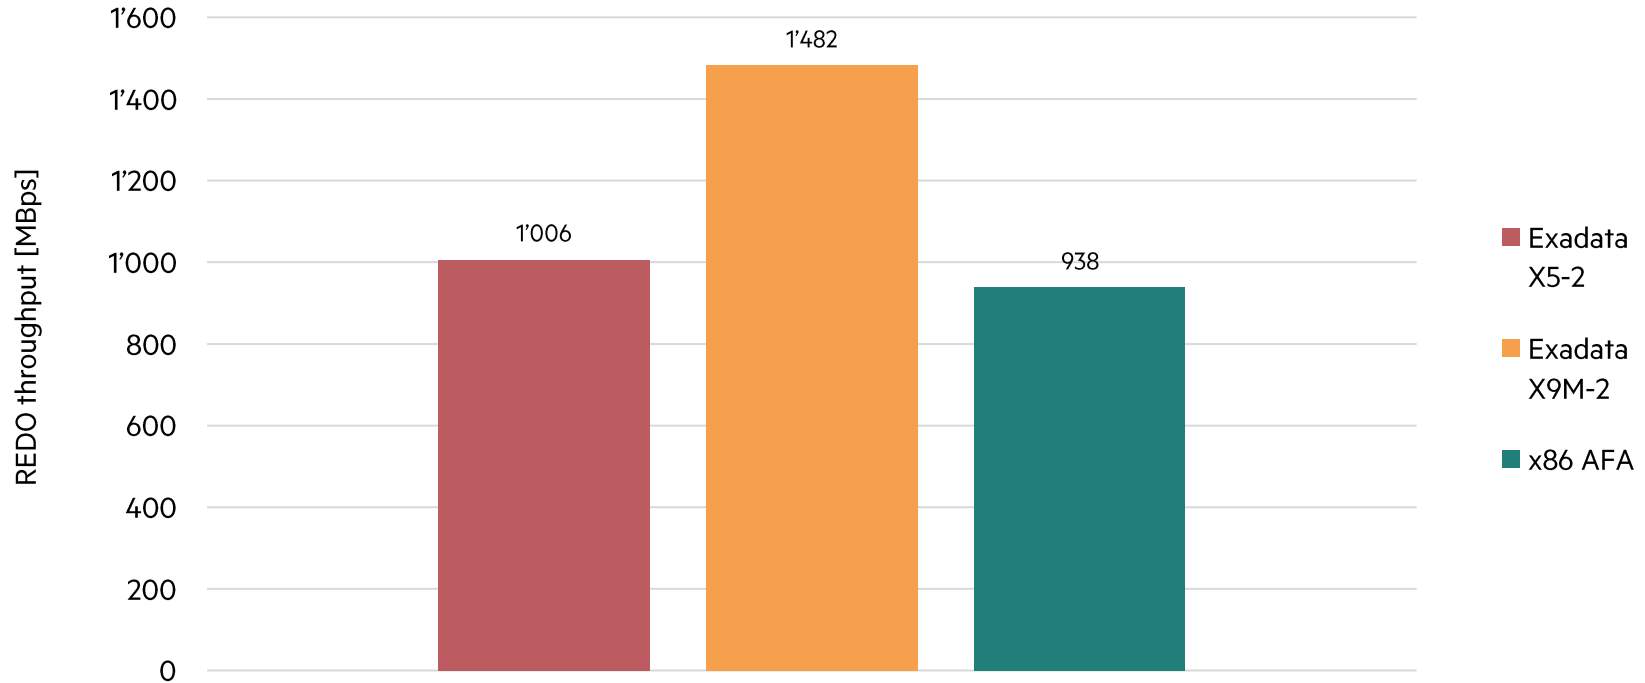
Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn [IOPS]	REDO writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
2	251	LGWR-THR	2	2	1	1	0	99	0	158	12.589	19,493	347	159	75	486.08	0.931	0.00	302	
	252	LGWR-THR	2	8	2	2	0	98	0	548	14.560	69,273	910	549	259	483.97	1.992	0.00	302	
	253	LGWR-THR	2	16	5	4	0	95	0	929	17.131	119,107	983	931	441	486.10	3.208	0.00	303	
	254	LGWR-THR	2	24	8	7	0	92	0	1,179	20.278	150,450	839	1,182	563	488.98	4.035	0.00	302	
	255	LGWR-THR	2	32	11	10	1	89	0	1,433	22.198	180,634	655	1,436	685	489.49	5.264	0.00	303	
	256	LGWR-THR	2	40	12	11	1	88	0	1,698	23.433	212,558	556	1,703	812	489.69	6.120	0.00	303	
	257	LGWR-THR	2	48	13	12	1	87	0	1,914	24.971	238,204	475	1,919	914	488.99	7.243	0.00	302	
	258	LGWR-THR	2	56	15	14	1	85	0	1,962	28.425	243,819	391	1,968	938	489.56	8.738	0.00	302	

Exadata X9M-2
with PMEM
2021

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Commit throughput [tps]	Commit latency [ms]	REDO blocks [rbps]	REDO writes [IOPS]	REDO syn [IOPS]	REDO writes [IOPS]	REDO data [MBps]	REDO data [kBpt]	LogFile sync [ms]	FlCache write [%]	Elapsed time [s]
4	288	LGWR-THR	2	2	1	1	0	99	0	180	11.118	179,031	269	180	85	483.56	0.492	27.91	301	
	289	LGWR-THR	2	8	4	3	0	96	0	691	11.531	689,476	875	692	325	481.62	0.762	30.04	302	
	290	LGWR-THR	2	16	7	6	1	93	0	1,275	12.518	1,273,946	1,259	1,277	601	482.69	1.459	31.61	302	
	291	LGWR-THR	2	24	10	9	1	90	0	1,773	13.506	1,772,599	1,483	1,776	835	482.26	2.136	32.39	302	
	292	LGWR-THR	2	32	12	11	1	88	0	2,167	14.724	2,167,300	1,489	2,171	1,021	482.47	2.955	33.20	302	
	293	LGWR-THR	2	40	14	13	1	86	0	2,500	15.943	2,503,623	1,439	2,505	1,180	483.33	3.662	34.24	302	
	294	LGWR-THR	2	48	17	15	1	83	0	2,790	17.141	2,789,493	1,258	2,796	1,315	482.64	4.337	35.02	302	
	295	LGWR-THR	2	56	18	17	1	82	0	3,015	18.490	3,018,031	1,029	3,022	1,423	483.30	4.976	35.67	303	
	296	LGWR-THR	2	64	20	18	1	80	0	3,140	20.289	3,140,976	796	3,147	1,482	483.30	5.680	36.48	302	



Workload LGWR-LAT125: maximum commit throughput for large-sized transactions





Swiss precision in performance measurement.

Summary



The new PMEM technology is particularly effective for short transactions.

When migrating from older Exadata generations to newer Exadata generations with PMEM technology

- REDO log writer latency decreases by up to a factor of 6
- REDO log writer commit throughput increases up to a factor of 4
- REDO log writer data throughput increases up to 47%



peakmarks Mission

Identify Key Performance Metrics for Oracle Database Platforms.

On-Premises and in the Cloud.

For Quality Assurance, Evaluations, and Capacity Planning.