


peakmarks® Performance Study on Scalability of Intel Xeon servers

2-socket Server vs 4-socket Server
February 2022



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



Processor	Intel Xeon 3 rd gen scalable processors Intel Xeon 8358, 2.6 – 3.4 GHz	Intel Xeon 3 rd gen scalable processors Intel Xeon 8380HL, 2.9 – 4.3 GHz
Launch date	Q2/2021	Q2/2021
#cores per cpu	32	28
Multithreading	Enabled, 2-fold	Enabled, 2-fold
#threads per cpu	64	56
Memory type	DDR4, 3.2 GHz	DDR4 RDIMM, 3.2 GHz
PCI Express	Gen 4	Gen 4
SPECrate 2017 Integer per core	~ 7.4 peak	~ 7.2 peak



Server	Intel Xeon 3 rd gen scalable processors Intel Xeon 8358, 2.6 – 3.4 GHz	Intel Xeon 3 rd gen scalable processors Intel Xeon 8380HL, 2.9 – 4.3 GHz
Technology	Bare metal	Bare metal
#sockets, total	2	4
#cores, total	64	112
#threads, total	128	224
Operating System	OEL 7.9	RH 8.4



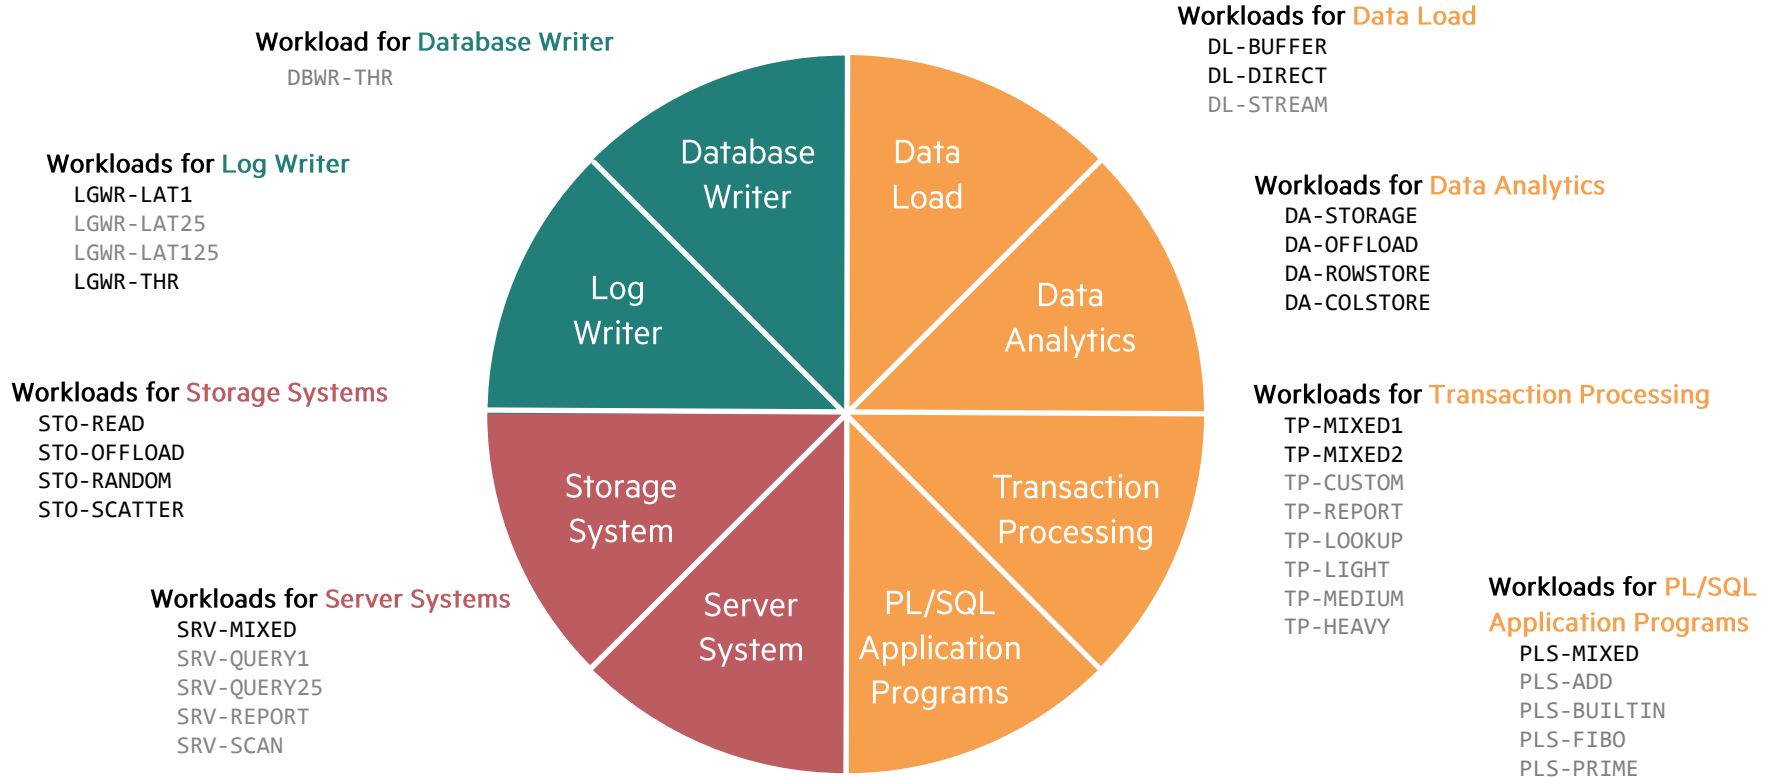
Database	Intel Xeon 3 rd gen scalable processors Intel Xeon 8358, 2.6 – 3.4 GHz	Intel Xeon 3 rd gen scalable processors Intel Xeon 8380HL, 2.9 – 4.3 GHz
Oracle version	19.3 Enterprise Edition	19.3 Enterprise Edition
peakmarks® Software	Version 9.4, Build 191130	Version 9.4, Build 191130



Simple. Representative. Fast.

peakmarks[®] Workload Overview


More than 30 micro-benchmarks in 8 workload groups





Stop guessing. Start measuring.

Workloads to determine the Server Performance in Database Operations





Motivation

The server performance significantly impacts the performance of all database operations.

The goal is to

- Validate the performance capabilities (speed, throughput, scalability) of server components in database operation: processors, main memory, and internal memory channels
- Determine the impact of server virtualization, multithreading, NUMA effects, and encryption on server performance
- Optimize database license and maintenance costs for server system

Notes

- Some cloud service providers do not publish their server components and configurations, and these are subject to change without prior notice.
- Customers need to know the per-thread performance, which significantly impacts application process performance.
- Customers need to know the per-core performance, which significantly impacts Oracle license costs. In some cases, Oracle licensing costs exceed infrastructure costs.



Key Performance Metrics

- **SQL query throughput** in queries per second [qps]
- **SQL query response time** in milliseconds [ms]
- **Logical reads** in database blocks per second [dbps]
- **SQL buffer cache scan rate** in megabytes per second [MBps]



Description

Workload	Measurement Unit	Action
SRV-QUERY1	[qps] [ms]	Latency-oriented look-up query – select 1 row via index, e.g., select customer, account, product, order, invoice. This workload shows maximum query throughput and minimum response time for simple queries.
SRV-QUERY25	[qps] [ms]	Data volume-oriented look-up query – select Ø 25 rows via index, e.g., select last month's bank account bookings; select item list of order. This workload shows maximum query throughput and minimum response time for more complex queries.
SRV-REPORT	[dbps]	Online Report – select Ø 125 rows via index, e.g., select last month's cell phone call records. This workload shows maximum logical read throughput.
SRV-SCAN	[MBps]	Full table scan. Search for data without index support. This workload shows a maximum database buffer cache scan rate.

Notes

- All accessed data is completely stored in the database buffer cache. There are no I/O operations, and all SRV workloads are CPU-bound.
- These kinds of queries are generic to all applications in all industries.




Description

Workload	Measurement Unit	Action
SRV-MIXED	[qps] [ms]	<p>Mixed queries and full table scans on cached data.</p> <p>This complex workload comprises the equally weighted simple workloads SRV-QUERY1, SRV-QUERY25, SRV-REPORT, and SRV-SCAN.</p> <p>It's important to note that SPEC numbers may not always be available or provide an accurate representation of Oracle database operations. To determine server system performance in Oracle database operations, we rely on the SRV-MIXED peakmarks workload as it is the most representative. This workload is used to compare different CPU architectures, such as ARM, Intel Xeon, AMD EPYC, IBM Power, and IBM Z.</p>

Notes

- All accessed data is completely stored in the database buffer cache. There are no I/O operations, and all SRV workloads are CPU-bound.
- These kinds of queries are generic to all applications in all industries.

Server System Performance





Workload SRV-QUERY1 – Simple look-up query, highest throughput, lowest response time

Intel Xeon 8358
2s, 64c, 128t

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Queries total [qps]	Queries per cpu [qps]	Response time [ms]	Log reads total [dbps]	Log reads per cpu [dbps]	BuCache read [%]	Elapsed time [s]
2	2	SRV-QUERY1	1	1	1	1	0	99	126,902	126,902	0.008	460,106	460,106	100.00	300
	3	SRV-QUERY1	1	32	26	25	0	74	3,368,485	105,265	0.009	10,171,543	317,861	100.00	301
	4	SRV-QUERY1	1	64	51	50	0	49	5,652,614	88,322	0.011	16,986,701	265,417	100.00	301
	5	SRV-QUERY1	1	96	76	75	0	24	6,098,959	63,531	0.016	18,283,906	190,457	100.00	301
	6	SRV-QUERY1	1	128	97	96	1	3	6,303,748	49,248	0.020	18,863,574	147,372	100.00	301

Peak single thread query throughput **126,902 qps**

Max query throughput per core **98,496 qps**

Query response time between **8 μs** and **20 μs**

Intel Xeon 8380HL
4s, 112c, 224t

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Queries total [qps]	Queries per cpu [qps]	Response time [ms]	Log reads total [dbps]	Log reads per cpu [dbps]	BuCache read [%]	Elapsed time [s]
3	2	SRV-QUERY1	1	1	1	1	0	99	112,509	112,509	0.009	337,848	337,848	100.00	300
	3	SRV-QUERY1	1	56	25	25	0	75	4,356,387	77,793	0.013	13,008,828	232,301	100.00	301
	4	SRV-QUERY1	1	112	50	49	0	50	5,235,239	46,743	0.021	15,578,179	139,091	100.00	301
	5	SRV-QUERY1	1	168	74	73	0	26	5,504,966	32,768	0.030	16,350,886	97,327	100.00	301
	6	SRV-QUERY1	1	224	96	96	0	4	5,649,405	25,221	0.040	16,752,951	74,790	100.00	301

Peak single thread query throughput **112,509 qps**

Max query throughput per core **50,441 qps**

Query response time between **9 μs** and **40 μs**



Workload SRV-QUERY25 – More complex query, highest throughput, lowest response time

Intel Xeon 8358
2s, 64c, 128t

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Queries total [qps]	Queries per cpu [qps]	Response time [ms]	Log reads total [dbps]	Log reads per cpu [dbps]	BuCache read [%]	Elapsed time [s]
2	7	SRV-QUERY25	1	1	1	1	0	99	36,798	36,798	0.027	1,072,999	1,072,999	99.98	301
	8	SRV-QUERY25	1	32	25	25	0	75	830,549	25,955	0.038	22,449,732	701,554	100.00	301
	9	SRV-QUERY25	1	64	51	50	0	49	1,475,883	23,061	0.043	39,673,583	619,900	100.00	301
	10	SRV-QUERY25	1	96	76	75	0	24	1,687,416	17,577	0.057	45,340,544	472,297	100.00	301
	11	SRV-QUERY25	1	128	98	97	0	2	1,849,399	14,448	0.069	49,682,958	388,148	100.00	301

Peak single thread query throughput **36,798 qps**

Max query throughput per core **28,897 qps**

Query response time between **27 μs** and **69 μs**

Intel Xeon 8380HL
4s, 112c, 224t

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Queries total [qps]	Queries per cpu [qps]	Response time [ms]	Log reads total [dbps]	Log reads per cpu [dbps]	BuCache read [%]	Elapsed time [s]
3	7	SRV-QUERY25	1	1	1	1	0	99	26,260	26,260	0.038	711,817	711,817	99.92	301
	8	SRV-QUERY25	1	56	25	25	0	75	994,118	17,752	0.056	26,831,595	479,136	100.00	301
	9	SRV-QUERY25	1	112	49	49	0	51	1,650,921	14,740	0.068	44,498,370	397,307	100.00	301
	10	SRV-QUERY25	1	168	73	73	0	27	1,852,694	11,028	0.091	49,880,548	296,908	100.00	301
	11	SRV-QUERY25	1	224	96	95	0	4	1,920,436	8,573	0.117	51,640,046	230,536	100.00	301

Peak single thread query throughput **26,260 qps**

Max query throughput per core **17,147 qps**

Query response time between **38 μs** and **117 μs**



Workload SRV-REPORT – Online Report, max throughput of Logical Reads

Intel Xeon 8358
2s, 64c, 128t

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Queries total [qps]	Queries per cpu [qps]	Response time [ms]	Log reads total [dbps]	Log reads per cpu [dbps]	BuCache read [%]	Elapsed time [s]
2	12	SRV-REPORT	1	1	1	1	0	99	9,447	9,447	0.106	1,274,085	1,274,085	99.98	301
	13	SRV-REPORT	1	32	26	25	0	74	209,855	6,558	0.152	26,515,043	828,595	100.00	301
	14	SRV-REPORT	1	64	51	50	0	49	368,199	5,753	0.174	46,222,534	722,227	100.00	301
	15	SRV-REPORT	1	96	76	75	0	24	419,855	4,373	0.228	52,677,112	548,720	100.00	301
	16	SRV-REPORT	1	128	99	98	0	1	469,143	3,665	0.273	58,864,456	459,879	100.00	301

Peak single thread throughput logical reads **1,274,085 dbps**

Max throughput logical reads per core **919,757 dbps**

Intel Xeon 8380HL
4s, 112c, 224t

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Queries total [qps]	Queries per cpu [qps]	Response time [ms]	Log reads total [dbps]	Log reads per cpu [dbps]	BuCache read [%]	Elapsed time [s]
3	12	SRV-REPORT	1	1	1	1	0	99	6,498	6,498	0.154	822,277	822,277	99.93	301
	13	SRV-REPORT	1	56	25	25	0	75	238,734	4,263	0.234	30,156,852	538,515	100.00	301
	14	SRV-REPORT	1	112	49	49	0	51	386,227	3,448	0.290	48,717,953	434,982	100.00	301
	15	SRV-REPORT	1	168	73	72	0	27	425,437	2,532	0.394	53,595,456	319,021	100.00	301
	16	SRV-REPORT	1	224	96	95	0	4	438,646	1,958	0.510	55,187,167	246,371	100.00	301

Peak single thread throughput logical reads **822,277 dbps**

Max throughput logical reads per core **492,742 dbps**



Workload SRV-SCAN – Scan-Rate in Oracle Buffer Cache

Intel Xeon 8358
2s, 64c, 128t

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Scan rate total [MBps]	Scan rate per cpu [MBps]	Log reads total [dbps]	Log reads per cpu [dbps]	BuCache read [%]	Elapsed time [s]
2	17	SRV-SCAN	1	1	1	1	0	99	4,298	4,298	550,096	550,096	100.00	300
	18	SRV-SCAN	1	32	26	25	0	74	89,935	2,810	11,511,621	359,738	100.00	301
	19	SRV-SCAN	1	64	51	50	0	49	143,916	2,249	18,421,249	287,832	100.00	301
	20	SRV-SCAN	1	96	76	75	0	24	164,124	1,710	21,007,822	218,831	100.00	301
	21	SRV-SCAN	1	128	99	98	1	1	168,839	1,319	21,611,373	168,839	100.00	301

Peak single thread scan rate **4,298 MBps**

Max scan rate throughput per core **2,638 MBps**

Intel Xeon 8380HL
4s, 112c, 224t

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Scan rate total [MBps]	Scan rate per cpu [MBps]	Log reads total [dbps]	Log reads per cpu [dbps]	BuCache read [%]	Elapsed time [s]
3	17	SRV-SCAN	1	1	1	1	0	99	1,948	1,948	249,332	249,332	99.99	300
	18	SRV-SCAN	1	56	25	25	0	75	81,267	1,451	10,402,235	185,754	100.00	301
	19	SRV-SCAN	1	112	49	49	0	51	100,280	895	12,835,811	114,605	100.00	301
	20	SRV-SCAN	1	168	73	72	0	27	102,917	613	13,173,314	78,413	100.00	301
	21	SRV-SCAN	1	224	96	95	0	4	103,984	464	13,309,984	59,420	100.00	301

Peak single thread scan rate **1,948 MBps**

Max scan rate throughput per core **928 MBps**



Workload SRV-MIXED – Mixed queries and full table scans

Intel Xeon 8358
2s, 64c, 128t

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Queries total [qps]	Queries per cpu [qps]	Response time [ms]	Log reads total [dbps]	Log reads per cpu [dbps]	BuCache read [%]	Elapsed time [s]
2	22	SRV-MIXED	1	4	4	3	0	96	162,648	40,662	0.025	2,821,463	705,366	100.00	301
	23	SRV-MIXED	1	32	26	25	0	74	1,109,740	34,679	0.029	17,593,398	549,794	100.00	301
	24	SRV-MIXED	1	64	51	50	0	49	2,002,292	31,286	0.032	30,099,887	470,311	100.00	301
	25	SRV-MIXED	1	96	76	75	0	24	2,244,089	23,376	0.043	34,482,989	359,198	100.00	301
	26	SRV-MIXED	1	128	99	98	0	1	2,467,971	19,281	0.052	38,356,911	299,663	100.00	301

Single thread query throughput **40,662 qps**

Max query throughput per core **38,562 qps**

Query response time between **25 µs** and **52 µs**

Intel Xeon 8380HL
4s, 112c, 224t

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Queries total [qps]	Queries per cpu [qps]	Response time [ms]	Log reads total [dbps]	Log reads per cpu [dbps]	BuCache read [%]	Elapsed time [s]
3	22	SRV-MIXED	1	4	2	2	0	98	103,102	25,776	0.039	1,893,618	473,405	100.00	301
	23	SRV-MIXED	1	56	25	25	0	75	1,521,356	27,167	0.037	20,321,813	362,890	100.00	301
	24	SRV-MIXED	1	112	49	49	0	51	2,579,228	23,029	0.043	32,331,062	288,670	100.00	301
	25	SRV-MIXED	1	168	73	72	0	27	2,939,100	17,495	0.057	35,108,707	208,980	100.00	301
	26	SRV-MIXED	1	224	96	95	0	4	2,989,363	13,345	0.075	35,984,158	160,644	100.00	301

Single thread query throughput **25,776 qps**

Max query throughput per core **26,690 qps**

Query response time between **39 µs** and **75 µs**



Swiss precision in performance measurement.



The per-core performance determines the number of cores required to run an application and thus significantly influences Oracle license and maintenance costs.

The per-thread performance has a significant impact on

- the response time of queries and transactions in online applications
- the degree of parallelism for batch-oriented applications, e.g., end-of-day processing



The 2-socket server is more efficient: higher per-thread performance and better response times for all server workloads.

The 4-socket server requires 75% higher license and maintenance costs but delivers barely higher throughput in absolute terms (workload SRV-QUERY25 +4%, workload SRV-MIXED +21%).

The higher total CPU performance of the 4-socket server has only a limited effect.

We recommend using 4-socket servers only if a higher main memory capacity is required.



peakmarks Mission

Identify Key Performance Metrics for Oracle Database Platforms.

On-Premises and in the Cloud.

For Quality Assurance, Evaluations, and Capacity Planning.