


# peakmarks® Performance Study on Intel Hyperthreading

June 2020



peakmarks® showcased its software at the  
11<sup>th</sup> TPC Technology Conference 2019.



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

[ $\mu$ 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

## Server

	Database Server without Hyperthreading	Database Server with Hyperthreading
Processor	Intel Xeon 6234, 3.3 – 4.0 GHz	Intel Xeon 6234, 3.3 – 4.0 GHz
Launch date	2019	2019
#cpus, total	1	1
#cores, total	8	8
#threads, total	8	16
PCI Express	Gen 3	Gen 3
Memory type	DDR4	DDR4
DRAM capacity, total	384 GByte	384 GByte
DRAM capacity, per core	48	48
Operating System	Bare metal, OEL	Bare metal, OEL

## Database

	Database Server without Hyperthreading	Database Server with Hyperthreading
Oracle version	19.3 Enterprise Edition	19.3 Enterprise Edition
Database block size	8 kByte	8 kByte
Log Modus	NOARCHIVELOG	NOARCHIVELOG
DataGuard	No	No
REDO Log Files, per instance	4 x 4 GByte, non-multiplexed	4 x 4 GByte, non-multiplexed
SGA size	192 GByte	192 GByte
Database encryption	AES256	AES256
peakmarks® Software	Version 9.5, Build 200415	Version 9.5, Build 200415
peakmarks® Database size	1 TByte	1 TByte

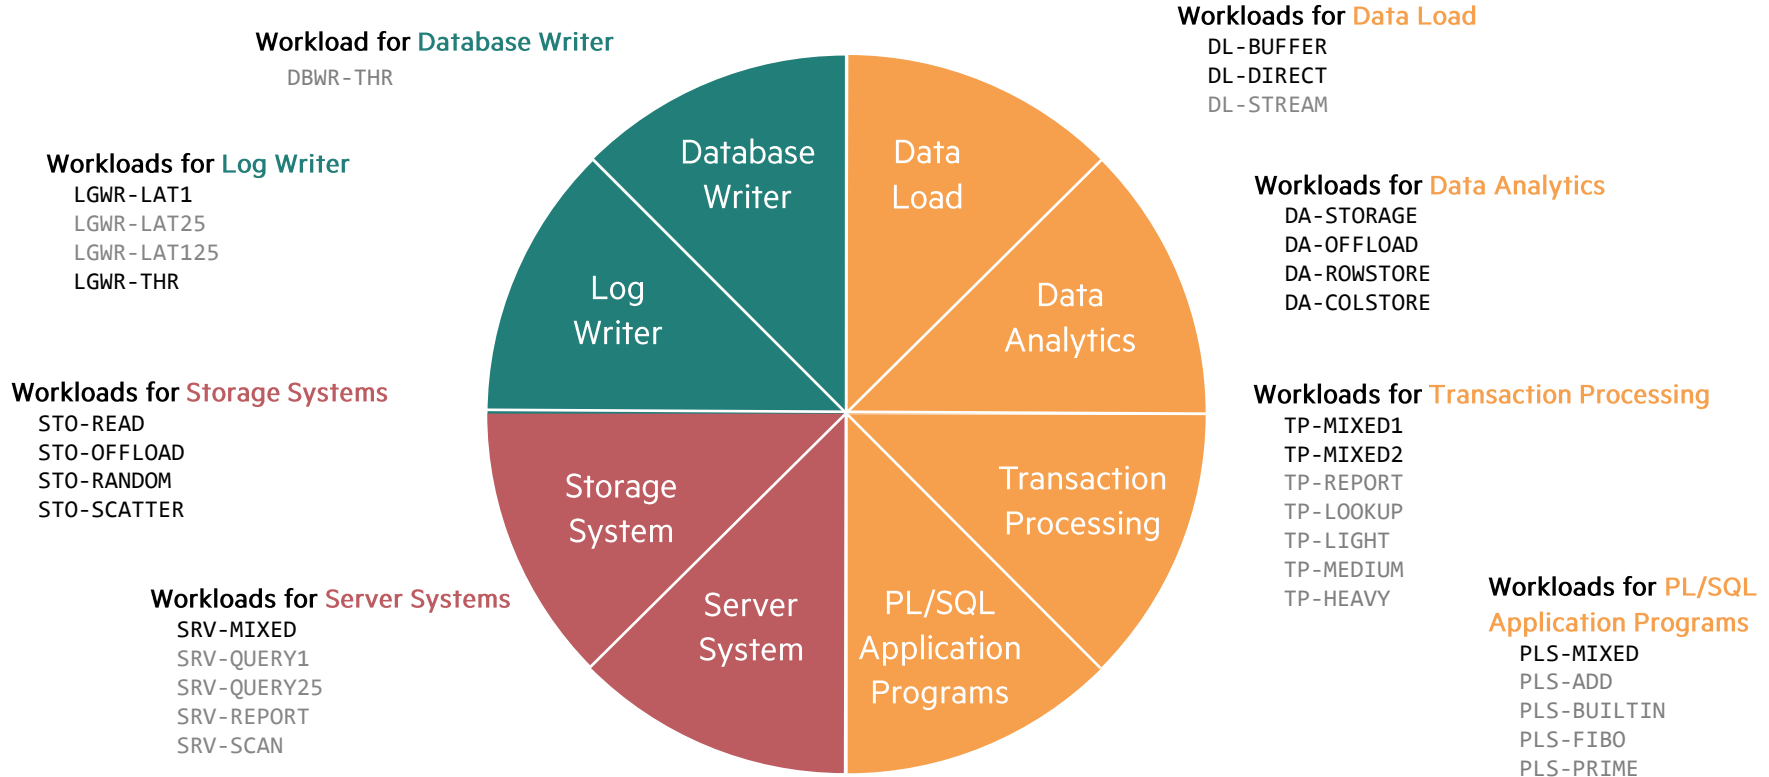


Simple. Representative. Fast.



## peakmarks<sup>®</sup> Workload Overview


More than 30 micro-benchmarks in 8 workload groups





Swiss precision in performance measurement.

## 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 (per-core performance matters)

### Notes

- Some cloud service providers do not publish their server components and configurations. Components and configurations of cloud services are subject to change without any prior notice.
- Customers need to know the per-core performance, significantly impacting application performance and the required Oracle licenses.
- In many cases, Oracle licensing costs far 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. <b>This workload shows maximum query throughput and minimum response time for simple queries.</b>
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. <b>This workload shows maximum query throughput and minimum response time for more complex queries.</b>
SRV-REPORT	[dbps]	Online Report – select Ø 125 rows via index, e.g., select last month's cell phone call records. <b>This workload shows maximum logical read throughput.</b>
SRV-SCAN	[MBps]	Full table scan. Search for data without index support. <b>This workload shows a maximum database buffer cache scan rate.</b>

### 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]	Mixed queries and full table scans on cached data.  This complex workload comprises the equally weighted simple workloads SRV-QUERY1, SRV-QUERY25, SRV-REPORT, and SRV-SCAN.  <b>SPEC numbers are not always representative of Oracle database operations. SRV-MIXED is the most representative peakmarks workload used to determine server system performance in Oracle database operations.</b>

### 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.
- Peakmarks uses the result of this workload to compare ARM, Intel Xeon, AMD EPYC, IBM POWER, and IBM Z processors.





Stop guessing. Start measuring.

## Server System Performance





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

Without Hyperthreading

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]
1	1	SRV-QUERY1	1	1	6	6	0	94	109,718	109,718	0.009	329,260	329,260	100.00	301
	2	SRV-QUERY1	1	4	25	25	0	75	465,308	116,327	0.009	1,396,062	349,016	100.00	302
	3	SRV-QUERY1	1	8	50	50	0	50	848,611	106,076	0.009	2,545,947	318,243	100.00	302
	4	SRV-QUERY1	1	12	75	75	0	25	1,254,599	104,550	0.010	3,763,946	313,662	100.00	201
	5	SRV-QUERY1	1	16	99	98	0	1	1,577,462	98,591	0.010	4,732,584	295,787	100.00	301

With Hyperthreading

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	1	SRV-QUERY1	1	1	3	3	0	97	124,828	124,828	0.008	374,600	374,600	100.00	300
	2	SRV-QUERY1	1	4	13	13	0	87	464,251	116,063	0.009	1,392,930	348,233	100.00	300
	3	SRV-QUERY1	1	8	25	25	0	75	892,126	111,516	0.009	2,676,498	334,562	100.00	298
	4	SRV-QUERY1	1	12	37	37	0	63	1,302,515	108,543	0.009	3,907,693	325,641	100.00	297
	5	SRV-QUERY1	1	16	50	50	0	50	1,603,485	100,218	0.010	4,810,583	300,661	100.00	297
	6	SRV-QUERY1	1	24	75	74	0	25	1,659,727	69,155	0.014	4,979,460	207,478	100.00	298
	7	SRV-QUERY1	1	32	99	99	0	1	1,811,665	56,615	0.018	5,435,147	169,848	100.00	297

### Notes

- With hyperthreading, the CPU utilization does not show correct values; from test 5 to test 7, the CPU utilization increases by 49%, but throughput rises only 13%. In test 5, the server utilization shows 50%, but it has already reached 89% of its maximum throughput.
- With hyperthreading workload SRV-QUERY1 increases by about 15%.
- Without hyperthreading, the server delivers stable query response times across all utilization ranges.



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

Without Hyperthreading

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]
1	6	SRV-QUERY25	1	1	6	6	0	94	31,213	31,213	0.032	843,741	843,741	100.00	299
	7	SRV-QUERY25	1	4	25	25	0	75	104,664	26,166	0.038	2,828,197	707,049	100.00	298
	8	SRV-QUERY25	1	8	50	50	0	50	196,185	24,523	0.041	5,301,440	662,680	100.00	296
	9	SRV-QUERY25	1	12	75	75	0	25	282,532	23,544	0.042	7,635,016	636,251	100.00	297
	10	SRV-QUERY25	1	16	100	100	0	0	365,222	22,826	0.044	9,869,366	616,835	100.00	296

With Hyperthreading

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	8	SRV-QUERY25	1	1	3	3	0	97	31,713	31,713	0.032	857,071	857,071	100.00	298
	9	SRV-QUERY25	1	4	13	13	0	87	108,493	27,123	0.037	2,932,037	733,009	100.00	298
	10	SRV-QUERY25	1	8	25	25	0	75	196,517	24,565	0.040	5,310,365	663,796	100.00	297
	11	SRV-QUERY25	1	12	38	37	0	62	284,304	23,692	0.042	7,682,134	640,178	100.00	296
	12	SRV-QUERY25	1	16	50	50	0	50	364,084	22,755	0.044	9,838,171	614,886	100.00	297
	13	SRV-QUERY25	1	24	75	75	0	25	410,172	17,090	0.058	11,083,796	461,825	100.00	298
	14	SRV-QUERY25	1	32	99	99	0	1	472,530	14,767	0.067	12,768,692	399,022	100.00	297

### Notes on Hyperthreading

- With hyperthreading workload SRV-QUERY25 increases by about 30%.
- On the other hand, the query response time doubles from 32 μs to 67 μs.



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

Without Hyperthreading

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]
1	11	SRV-REPORT	1	1	6	6	0	94	7,737	7,737	0.129	978,901	978,901	100.00	298
	12	SRV-REPORT	1	4	25	25	0	75	27,026	6,757	0.148	3,419,306	854,827	100.00	298
	13	SRV-REPORT	1	8	50	50	0	50	50,256	6,282	0.159	6,358,107	794,763	100.00	295
	14	SRV-REPORT	1	12	75	75	0	25	70,906	5,909	0.169	8,970,439	747,537	100.00	296
	15	SRV-REPORT	1	16	100	99	0	0	90,787	5,674	0.175	11,486,128	717,883	100.00	297

With Hyperthreading

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	15	SRV-REPORT	1	1	3	3	0	97	8,185	8,185	0.122	1,035,666	1,035,666	100.00	297
	16	SRV-REPORT	1	4	13	13	0	87	27,300	6,825	0.146	3,453,929	863,482	100.00	297
	17	SRV-REPORT	1	8	25	25	0	75	50,478	6,310	0.158	6,386,131	798,266	100.00	296
	18	SRV-REPORT	1	12	37	37	0	63	70,771	5,898	0.168	8,953,392	746,116	100.00	298
	19	SRV-REPORT	1	16	50	50	0	50	90,485	5,655	0.176	11,447,349	715,459	100.00	296
	20	SRV-REPORT	1	24	75	75	0	25	105,016	4,376	0.227	13,285,369	553,557	100.00	297
	21	SRV-REPORT	1	32	99	99	0	1	122,346	3,823	0.260	15,478,112	483,691	100.00	297

### Notes on Hyperthreading

- Hyperthreading increases the throughput of workload SRV-REPORT by about 35%.



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

Without Hyperthreading

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]
1	16	SRV-SCAN	1	1	7	6	0	93	2,923	2,923	356,803	356,803	100.00	296
	17	SRV-SCAN	1	4	25	25	0	75	11,229	2,807	1,370,782	342,696	100.00	296
	18	SRV-SCAN	1	8	50	50	0	50	22,116	2,765	2,699,685	337,461	100.00	298
	19	SRV-SCAN	1	12	74	74	0	26	32,543	2,712	3,972,561	331,047	100.00	299
	20	SRV-SCAN	1	16	99	99	0	1	42,329	2,646	5,167,078	322,942	100.00	299

With Hyperthreading

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	22	SRV-SCAN	1	1	3	3	0	97	2,837	2,837	346,306	346,306	100.00	295
	23	SRV-SCAN	1	4	13	12	0	87	10,578	2,645	1,291,199	322,800	100.00	298
	24	SRV-SCAN	1	8	25	25	0	75	21,261	2,658	2,595,392	324,424	100.00	298
	25	SRV-SCAN	1	12	38	37	0	62	31,575	2,631	3,854,346	321,196	100.00	296
	26	SRV-SCAN	1	16	50	50	0	50	41,492	2,593	5,064,906	316,557	100.00	296
	27	SRV-SCAN	1	24	75	75	0	25	48,719	2,030	5,947,093	247,796	100.00	298
	28	SRV-SCAN	1	32	99	99	0	1	57,668	1,802	7,039,579	219,987	100.00	297

Notes:

- The throughput of the workload SRV-SCAN increases by about 36% with hyperthreading.
- Without hyperthreading the per-thread performance delivers stable scan rates across all utilization ranges



## Workload SRV-MIXED – Mixed queries and full table scans

Without Hyperthreading

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]
1	21	SRV-MIXED	1	4	25	25	0	75	145,644	36,411	0.027	2,259,273	564,818	100.00	296
	22	SRV-MIXED	1	8	50	50	0	50	270,027	33,753	0.030	4,222,234	527,779	100.00	295
	23	SRV-MIXED	1	16	99	99	0	1	507,808	31,738	0.031	7,614,683	475,918	100.00	299

With Hyperthreading

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	29	SRV-MIXED	1	4	13	13	0	87	149,844	37,461	0.027	2,252,890	563,223	100.00	296
	30	SRV-MIXED	1	8	25	25	0	75	271,045	33,881	0.029	4,232,125	529,016	100.00	296
	31	SRV-MIXED	1	16	50	50	0	50	522,758	32,672	0.030	7,662,755	478,922	100.00	297
	32	SRV-MIXED	1	24	75	75	0	25	591,556	24,648	0.040	8,884,085	370,170	100.00	297
	33	SRV-MIXED	1	32	99	99	0	1	660,865	20,652	0.048	10,126,183	316,443	100.00	298

**Notes:**

- The throughput of workload SRV-SCAN increases by about 30% with hyperthreading.
- Without hyperthreading, the server delivers stable query response times across all utilization ranges



Hyperthreading increases the server's throughput by 15% to 39%, depending on the workload (on average, around 30%).

However, the query service time increases at full server utilization up to factor 2.

For transaction-oriented applications that require predictable performance, consider Intel Xeon processors without Hyperthreading for consistent monitoring of CPU utilization and stable response time.

If high throughput is more important than relatively stable response time, such as in data warehouse systems, consider Intel Xeon processors with Hyperthreading.





# peakmarks Mission

Identify Key Performance Metrics for Oracle Database Platforms.

On-Premises and in the Cloud.

For Quality Assurance, Evaluations, and Capacity Planning.