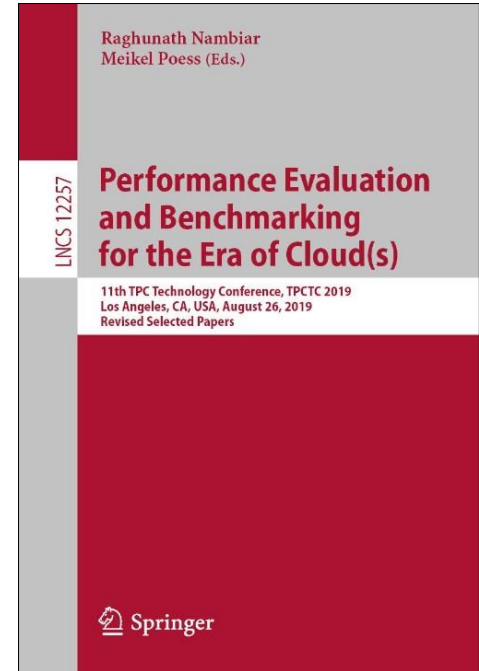


# peakmarks® Performance Study on IBM Database Server

IBM Power versus IBM LinuxONE  
October 2021



peakmarks® showcased its software at the  
11<sup>th</sup> 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

[ $\mu$ 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

	IBM LinuxONE II Rockhopper Database Server	IBM Power9 Database Server
Processor	IBM z14, 4.5 GHz	IBM Power9, 4.0 GHz
Launch date	2018	2018
#cores, total	24	32
#threads, total	48	256
PCI Express	Gen 3	Gen 3
Memory type	RAIM – redundant array of independent memory	DDR4
Operating System	RedHat Linux	AIX

Notes

- The IBM Power uses configurable simultaneous multithreading (SMT), in this case SMT-8 (8 threads per core).
- IBM introduced simultaneous multithreading (SMT) for the z13 processor. z processors use SMT-2 (2 threads per core).
- IBM introduced SIMD support for the z13 processor, which is used by the Oracle 19c in-memory technology.
- The peakmarks® performance reports show the performance per CPU (in terms of Oracle CPU\_COUNT) to enable comparability of the results.



## Database

	IBM LinuxONE II Rockhopper Database Server	IBM Power9 Database Server
Oracle version	19.3	19.3
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

Note

Both servers have the same core factor (1.0) according to Oracle licensing policy.

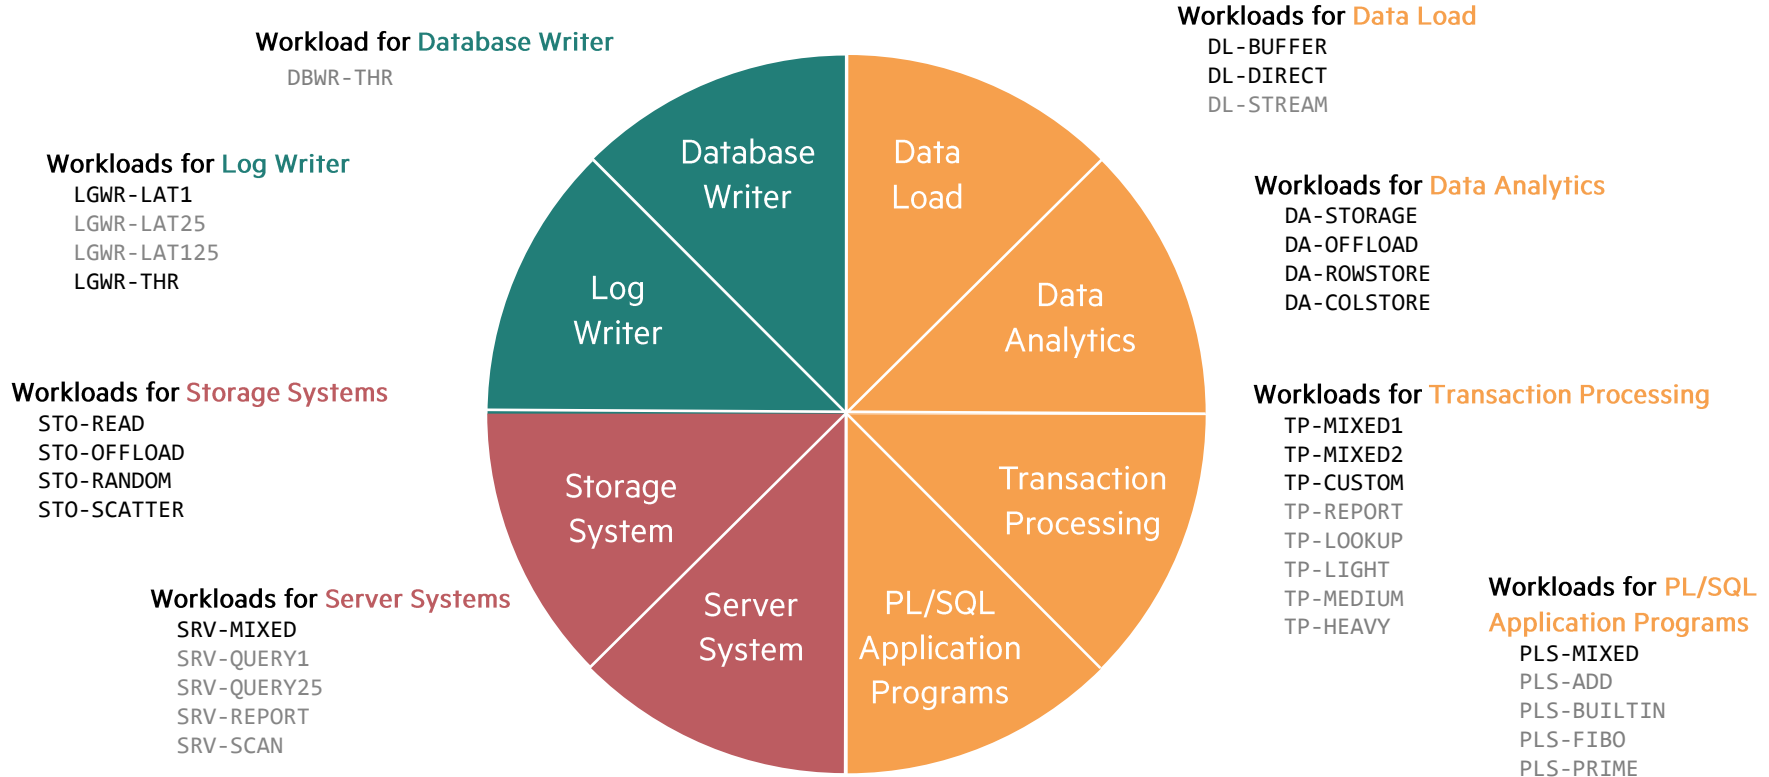




Simple. Representative. Fast.

## peakmarks<sup>®</sup> 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. <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]	<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

IBM Power9  
32c, 256t, Launch 2018

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	1	SRV-QUERY1	1	1	3	0	3	97	36,386	36,386	0.027	167,376	167,376	99.53	180
	2	SRV-QUERY1	1	32	12	12	0	88	1,184,757	37,024	0.027	3,611,232	112,851	100.00	182
	3	SRV-QUERY1	1	64	24	24	0	76	1,889,507	29,524	0.033	5,719,752	89,371	100.00	183
	5	SRV-QUERY1	1	128	48	48	0	52	2,433,129	19,009	0.052	7,342,610	57,364	100.00	182
	7	SRV-QUERY1	1	192	66	66	0	34	1,728,872	9,005	0.111	5,241,187	27,298	100.00	181
	9	SRV-QUERY1	1	256	85	85	0	15	2,490,612	9,729	0.101	7,496,094	29,282	100.00	183

IBM LinuxONE II  
24c, 48t, Launch 2018

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]
4	1	SRV-QUERY1	1	1	2	2	0	98	88,569	88,569	0.011	265,876	265,876	99.19	181
	2	SRV-QUERY1	1	24	49	45	4	51	1,928,976	80,374	0.014	5,787,410	241,142	100.00	183
	3	SRV-QUERY1	1	48	98	90	8	2	2,192,808	45,683	0.022	6,611,316	137,735	100.00	183

### Notes

- The Power9 server achieves 77,832 look-up queries per second per core (at full CPU utilization).
- The LinuxONE server achieves 91,367 look-up queries per second per core, providing significantly better response time (at full CPU utilization).
- The LinuxONE server provides better application process performance (thread performance), between factor 2.4 (low CPU utilization) and 4.7 (high CPU utilization).



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

IBM Power9  
32c, 256t, Launch 2018

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	10	SRV-QUERY25	1	1	3	0	3	97	10,393	10,393	0.096	339,923	339,923	99.82	180
	11	SRV-QUERY25	1	32	12	12	0	88	352,441	11,014	0.090	9,563,939	298,873	100.00	182
	12	SRV-QUERY25	1	64	24	24	0	76	564,658	8,823	0.113	15,257,500	238,398	100.00	182
	14	SRV-QUERY25	1	128	47	47	0	53	700,784	5,475	0.181	18,900,270	147,658	100.00	183
	16	SRV-QUERY25	1	192	70	70	0	30	780,072	4,063	0.243	21,056,294	109,668	100.00	183
	18	SRV-QUERY25	1	256	90	89	0	10	830,956	3,246	0.304	22,431,929	87,625	100.00	183

IBM LinuxONE II  
24c, 48t, Launch 2018

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]
4	4	SRV-QUERY25	1	1	3	3	0	97	41,213	41,213	0.023	1,110,690	1,110,690	99.94	181
	5	SRV-QUERY25	1	24	51	49	2	49	590,719	24,613	0.037	15,923,421	331,737	100.00	183
	6	SRV-QUERY25	1	48	99	96	3	1	711,527	14,823	0.065	19,107,986	398,093	100.00	183

### Notes

- The Power9 server achieves 25,969 queries per second per core (at full CPU utilization).
- The LinuxONE server achieves 29,647 queries per second per core, providing significantly better response time (at full CPU utilization).
- The LinuxONE server provides better application process performance (thread performance), between factor 3.9 (low CPU utilization) and 4.5 (high CPU utilization).



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

IBM Power9  
32c, 256t, Launch 2018

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	19	SRV-REPORT	1	1	3	0	3	97	2,697	2,697	0.371	400,187	400,187	99.85	181
	20	SRV-REPORT	1	32	12	12	0	88	95,448	2,983	0.334	12,107,278	378,352	100.00	182
	21	SRV-REPORT	1	64	24	24	0	76	150,523	2,352	0.422	19,011,887	297,061	100.00	182
	23	SRV-REPORT	1	128	48	48	0	52	195,173	1,525	0.648	24,623,366	192,370	100.00	183
	25	SRV-REPORT	1	192	70	70	0	30	205,767	1,072	0.922	25,971,098	135,266	100.00	183
	27	SRV-REPORT	1	256	90	90	0	10	209,177	817	1.211	26,416,586	103,190	100.00	183

IBM LinuxONE II  
24c, 48t, Launch 2018

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]
4	7	SRV-REPORT	1	1	3	3	0	97	12,667	12,667	0.079	1,595,433	1,595,433	99.96	181
	8	SRV-REPORT	1	24	51	50	1	49	161,154	6,714	0.140	20,297,359	845,723	100.00	182
	9	SRV-REPORT	1	48	99	98	1	1	183,514	3,823	0.241	23,113,691	481,535	100.00	183

### Notes

- The Power9 server achieves 825,518 logical reads per second per core (at full CPU utilization).
- The LinuxONE server achieves 963,070 logical reads per second per core (at full CPU utilization).



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

IBM Power9  
32c, 256t, Launch 2018

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	28	SRV-SCAN	1	1	3	0	3	97	2,195	2,195	280,958	280,958	100.00	180
	29	SRV-SCAN	1	32	12	12	0	88	61,996	1,937	7,935,456	247,983	100.00	182
	30	SRV-SCAN	1	64	24	24	0	76	87,036	1,360	11,140,555	174,071	100.00	181
	32	SRV-SCAN	1	128	47	47	0	53	116,111	907	14,862,227	116,111	100.00	181
	34	SRV-SCAN	1	192	70	69	0	30	104,714	545	13,403,338	69,809	100.00	183
	36	SRV-SCAN	1	256	89	88	0	11	115,688	452	14,808,023	57,844	100.00	183

IBM LinuxONE II  
24c, 48t, Launch 2018

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]
4	10	SRV-SCAN	1	1	2	2	0	98	8,197	8,197	1,000,067	1,000,067	100.00	181
	11	SRV-SCAN	1	24	50	50	0	50	47,856	1,994	5,838,769	243,282	100.00	183
	12	SRV-SCAN	1	48	99	99	0	1	51,745	1,078	6,313,356	131,528	100.00	184

### Notes

- The Power9 server achieves a scan rate of 3,615 MBps per second per core (at full CPU utilization).
- The LinuxONE server achieves a scan rate of 2,156 MBps per second per core (at full CPU utilization).



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

IBM Power9  
32c, 256t, Launch 2018

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	37	SRV-MIXED	1	4	4	2	3	96	52,620	13,155	0.076	1,059,105	264,776	100.00	181
	38	SRV-MIXED	1	128	48	48	0	52	855,659	6,685	0.149	16,169,406	126,323	100.00	182
	39	SRV-MIXED	1	256	90	90	0	10	1,064,694	4,159	0.239	19,250,904	75,199	100.00	182


IBM LinuxONE II  
24c, 48t, Launch 2018

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]
4	13	SRV-MIXED	1	4	11	10	1	89	153,056	38,264	0.013	3,290,704	822,676	100.00	182
	14	SRV-MIXED	1	24	49	47	2	51	726,598	30,274	0.019	15,523,759	646,823	100.00	183
	15	SRV-MIXED	1	48	97	94	3	3	784,725	16,348	0.036	16,489,034	343,521	100.00	184

### Notes

- The Power9 server achieves 33,271 queries and scans per second per core (at full CPU utilization).
- The LinuxONE server achieves 32,697 queries and scans per second per core (at full CPU utilization).
- Overall, the LinuxONE server provides better application process performance (thread performance), between factor 2.9 (low CPU utilization) and 3.9 (high CPU utilization).

## Workloads to determine the Data Analytics Performance







## Motivation

In general, data analytics operations cause full table scans. The performance of full table scan operation depends on the location of data in the storage hierarchy (storage, memory) and the technology used to boost scan performance (smart scan for the data location storage system, in-memory column store for the data location memory).

The goal is to

- Optimize data scan throughput
- Validate the impact of several factors on data analytics performance
  - » Smart scan offload technology
  - » In-memory column store technology
- Optimize Oracle license and maintenance costs



## Key Performance Metrics

- **SQL data scan rate** in megabytes per second [MBps] or rows per second [rps]




## Description

Workload	Measurement Unit	Action
DA-STORAGE	[MBps] [rps]	Simple aggregate after full table scan, using <b>conventional storage</b> .
DA-OFFLOAD	[MBps] [rps]	Simple aggregate after full table scan, using <b>smart-scan offload technology</b> .
DA-ROWSTORE	[MBps] [rps]	Simple aggregate after full table scan, using <b>row-store</b> .
DA-COLSTORE	[MBps] [rps]	Simple aggregate after full table scan, using <b>column-store</b> . <b>This workload uses in-memory optimized arithmetic operations introduced in 18c if supported by the underlying processor.</b>

### Notes:

- These analytic workloads are generic to all applications in all industries.
- The workloads DA-ROWSTORE and DA-COLSTORE are CPU-bound workloads.

## Data Analytics Performance





## Workload DA-ROWSTORE – Data scan using buffer cache row store

IBM Power9  
32c, 256f, Launch 2018

Run	Test	Workload	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Scan rate total [rps]	Scan rate per cpu [rps]	Scan rate total [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
4	1	DA-ROWSTORE	1	1	1	3	0	3	97	0	626,669	626,669	663	0.00	92.43	185
	2	DA-ROWSTORE	1	64	1	24	24	0	76	0	220,664,899	3,447,889	67,020	0.00	100.00	183
	3	DA-ROWSTORE	1	128	1	48	47	0	52	0	210,554,588	1,644,958	63,896	0.00	100.00	182
	4	DA-ROWSTORE	1	192	1	70	70	0	30	0	208,192,108	1,084,334	63,271	0.00	100.00	182
	5	DA-ROWSTORE	1	256	1	90	90	0	10	0	221,183,828	863,999	67,076	0.00	100.00	184

### Note

Run 4 test 1 did not deliver an optimal result, as the buffer cache was not sufficiently warmed up.

IBM LinuxONE II  
24c, 48f, Launch 2018

Run	Test	Workload	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Scan rate total [rps]	Scan rate per cpu [rps]	Scan rate total [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
5	1	DA-ROWSTORE	1	1	1	2	2	0	98	0	28,587,677	28,587,677	8,179	0.00	100.00	182
	2	DA-ROWSTORE	1	12	1	25	25	0	75	0	95,354,006	7,946,167	27,281	0.00	100.00	181
	3	DA-ROWSTORE	1	24	1	50	50	0	50	0	167,260,065	6,969,169	48,013	0.00	100.00	182
	4	DA-ROWSTORE	1	36	1	75	75	0	25	0	170,942,192	4,748,394	49,233	0.00	100.00	182
	5	DA-ROWSTORE	1	48	1	99	99	0	1	0	181,122,407	3,773,383	52,165	0.00	100.00	183

### Notes

- The scan rate in [rps] is derived and depends on the data model. The peakmarks data model uses a row length of approx. 300 bytes.
- The Power9 server achieves a scan rate of 6,911,994 MBps per core (at full CPU utilization).
- The LinuxONE server achieves a scan rate of 6,655,902 MBs per core (at full CPU utilization).
- Overall, the LinuxONE server provides better application process performance (thread performance), between factor 2.9 (low CPU utilization) and 3.9 (high CPU utilization).



## Workload DA-COLSTORE – Data scan using column store

IBM Power9  
32c, 256f, Launch 2018

Run	Test	Workload	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Scan rate total [rps]	Scan rate per cpu [rps]	Scan rate total [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
4	6	DA-COLSTORE	1	1	1	3	0	3	97	0	32,755,468	32,755,468	12,426	0.00	0.00	181
	7	DA-COLSTORE	1	64	1	24	24	0	76	0	1,733,662,193	27,088,472	660,366	0.00	0.00	182
	8	DA-COLSTORE	1	128	1	48	47	0	52	0	1,982,154,002	15,485,578	755,025	0.00	0.00	182
	9	DA-COLSTORE	1	192	1	69	69	0	31	0	2,021,287,954	10,527,541	769,933	0.00	0.00	183
	10	DA-COLSTORE	1	256	1	91	90	0	9	0	2,048,217,223	8,000,848	780,191	0.00	0.00	182

IBM LinuxONE II  
24c, 48f, Launch 2018

Run	Test	Workload	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Scan rate total [rps]	Scan rate per cpu [rps]	Scan rate total [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
5	7	DA-COLSTORE	1	1	1	2	2	0	98	0	161,745,059	161,745,059	46,275	0.00	0.00	181
	8	DA-COLSTORE	1	12	1	25	25	0	75	0	1,799,223,637	149,935,303	514,762	0.00	0.00	181
	9	DA-COLSTORE	1	24	1	50	50	0	50	0	3,475,882,651	144,828,444	994,458	0.00	0.00	182
	10	DA-COLSTORE	1	36	1	75	75	0	25	0	4,471,229,780	124,200,827	1,279,229	0.00	0.00	183
	11	DA-COLSTORE	1	48	1	99	99	0	1	0	4,729,986,524	98,541,386	1,353,260	0.00	0.00	183

### Notes

- The scan rate in [MBps] is derived and depends on the data model. The peakmarks data model uses a row length of approx. 300 bytes.
- The Power9 server achieves a scan rate of 64,006,788 rps per core (at full CPU utilization).
- The LinuxONE server achieves a scan rate of 100,637,186 rps per core (at full CPU utilization).



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.

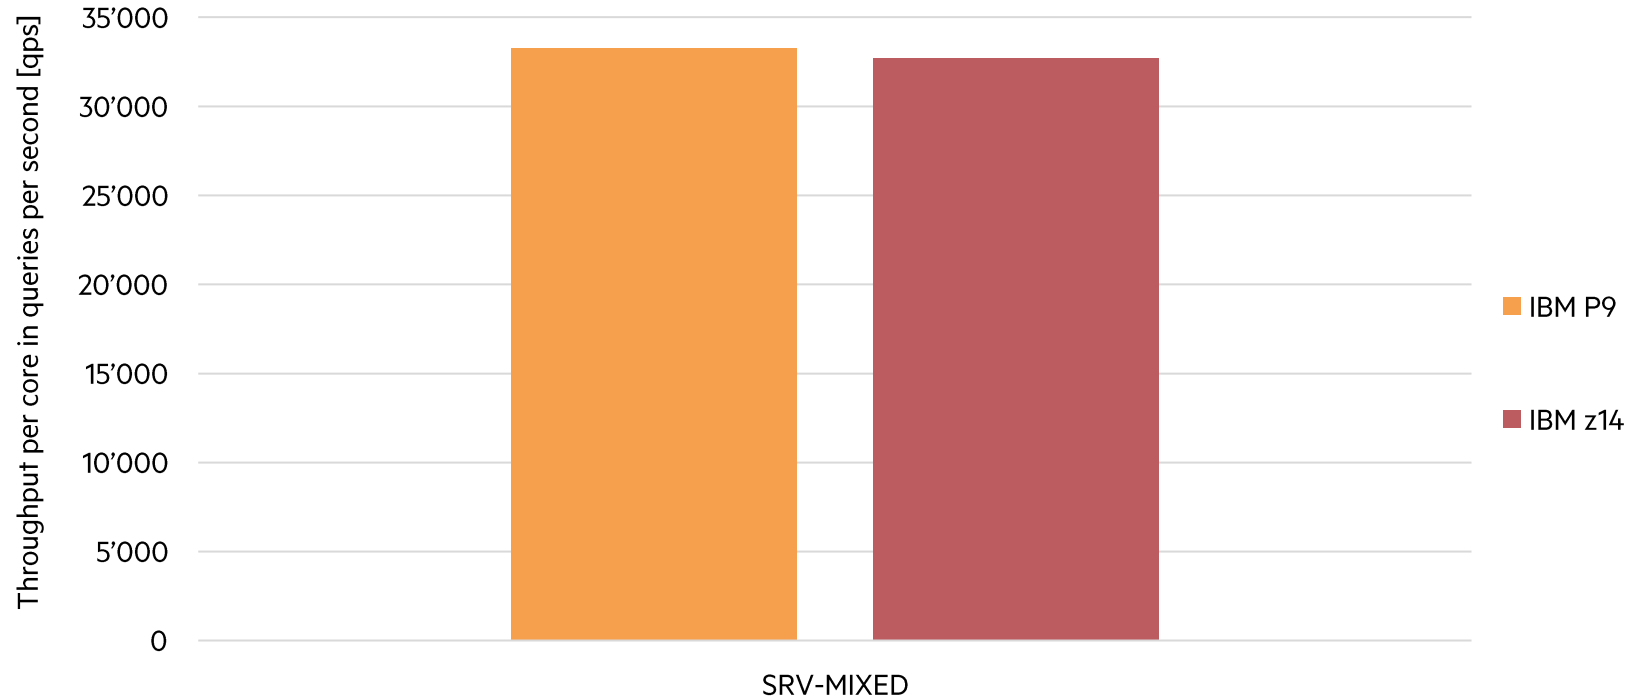
The per-thread performance has a significant impact on

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

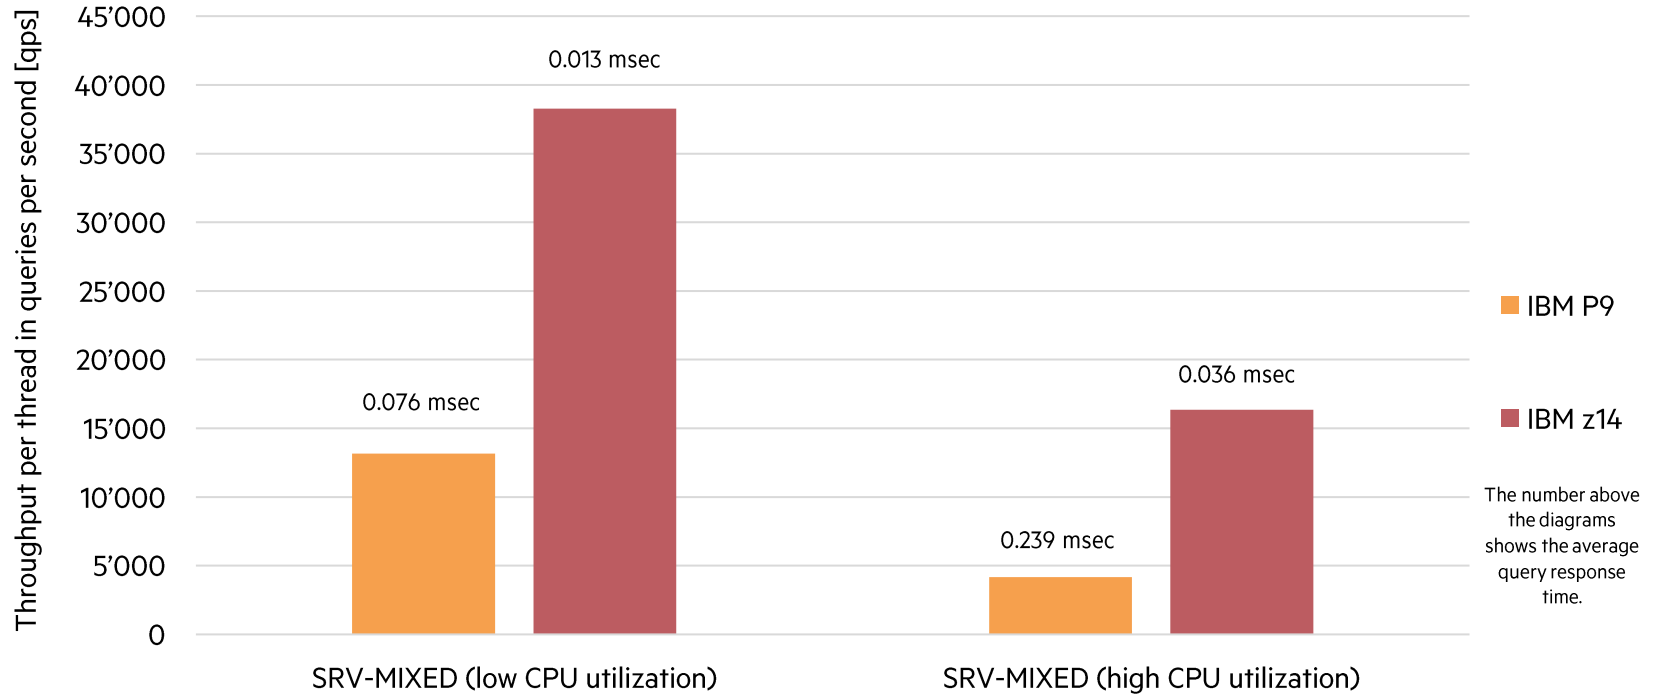




## Per-Core Performance based on peakmarks® workload SRV-MIXED



## Per-Thread Performance based on peakmarks® workload SRV-MIXED





## Summary

Both IBM servers offer top-of-the-line technology in terms of virtualization, availability [1], and scalability, even without the need for Oracle RAC.

Both IBM servers deliver similar per-core performance based on peakmarks® workload SRV-MIXED.

However, IBM LinuxONE servers deliver up to a factor of 4 better per-thread performance for individual application processes based on peakmarks® workload SRV-MIXED. Online transaction systems on LinuxONE Servers deliver considerably better response times, and batch applications can work with significantly lower parallelism and, thus, much more efficiently.

We recommend that IBM Power customers consider using IBM LinuxONE servers for Oracle databases for the following lifecycle.

[1] *Worldwide AL4 Server Market Shares, 2019: Fault-Tolerant Systems Become Digital Transformation Platforms*, IDC Report, July 2020



# peakmarks Mission

Identify Key Performance Metrics for Oracle Database Platforms.

On-Premises and in the Cloud.

For Quality Assurance, Evaluations, and Capacity Planning.