peakmarks® Benchmark Software

August 2023

contact@peakmarks.com

Foreword

The peakmarks® Benchmark Software was introduced to a broader benchmark audience at the 11th TPC Technology Conference 2019 in Los Angeles[1]. Since then, the peakmarks® Benchmark Software has been adapted to new technologies and is now available in version 10 for all Oracle database platforms from Oracle 19c onwards, with and without CDB architecture. This publication is an updated version of the original TPC publication.

Introduction

The peakmarks® Benchmark Software enables the fast and reliable determination of generally understandable performance metrics for Oracle Database Services, both on-premises and in the cloud.

Performance metrics of Oracle Database Services are necessary for various tasks.

Evaluation. No comparable and comprehensible performance indicators can be found even for the well-known cloud services from Amazon, Google, and Microsoft. Making a price/performance comparison of different offerings requires the use of performance indicators. These indicators provide factual information to make informed decisions.

Capacity planning. When migrating databases to new platforms, whether on-premises or in the cloud, performance metrics help with sound capacity planning.

License cost optimization. License costs often far exceed infrastructure costs. License costs can often be significantly reduced by optimizing the infrastructure. The most suitable infrastructure components are quickly and clearly identified with the help of performance indicators.

Quality Assurance. Cloud services can change their infrastructure at any time and without notice. Providers' value propositions can be periodically verified by regularly reviewing performance - quickly and reliably.

[1] Drozd: *Benchmarking Database Cloud Services*. In: Nambiar, R., Poess, M. (eds.) Performance Evaluation and Benchmarking for the Era of Cloud(s), Lecture Notes in Computer Science (LNCS), vol. 12257, pp. 139-153. Springer, Switzerland (2020).

Requirements for benchmark tools

Huppler[2] describes the five most important characteristics of a good benchmark: relevant, repeatable, fair, traceable and economical. peakmarks® meets all these requirements.

But other features are also critical to customer acceptance of benchmark tools.

Simplicity. Installing the benchmark software, running the benchmark, and interpreting the results must be easy. peakmarks® is implemented with the database tools and without operating system scripts. Therefore, peakmarks® runs unchanged wherever the Oracle database software is available. Any database administrator can easily run the benchmark software without additional know-how.

Speed. The installation, the loading of the data, the processing of the different workloads, and the evaluation of the performance metrics should be fast. peakmarks® is installed in a few hours, including all adjustments to the database. The database loading time depends on the database size and the performance of the infrastructure. The scalable loading process automatically adapts to the performance of the database platform. On powerful systems, load times of 4 TByte per hour are measured. Complete benchmark runs with all workloads take between 12 and 24 hours; results are available immediately. A comprehensive benchmark project can be completed within one week.

Understandable performance metrics. Many benchmarks provide only a single performance metric. This simplifies the comparison of different systems. However, a single complexly created metric is difficult to interpret[3]. peakmarks® provides a set of representative and easy-to-understand metrics for different aspects. Concrete questions about performance can thus be answered more easily. Performance bottlenecks and malfunctions are identified more quickly.

Different load situations. Often, it is not the maximum value of a performance metric that is of interest but the optimal performance range where a sustained and predictable performance output occurs. peakmarks® analyzes the performance of a database service in different load situations. Workloads start with a low load. The load is continuously increased until system saturation. In this way, the optimal performance range of a Database Service can be determined.

Continuous further development. In contrast to many open-source benchmark tools, the peakmarks® Benchmark Software is continuously developed. The requirements of new hardware technologies such as *flash storage* and *persistent memory* are incorporated into further development, as are new database technologies (*smart scan*, *in-memory column store*, *memory-optimized tables,* etc.). The list of workloads is supplemented as far as they are relevant for measuring the infrastructure and choosing the solution architecture.

For more in-depth considerations on systematic benchmarking of cloud services, please refer to the literature by Bermbach, Wittern, and Tai[4].

---

[2] Huppler: *The Art of Building a Good Benchmark*. In: Nambiar, R., Poess, M. (eds.) Performance Evaluation and Benchmarking, LNCS, vol. 5895, pp. 18-30. Springer, Heidelberg (2009).

[3] Crolotte: *Issues in Benchmark Metric Selection*. In: Nambiar, R., Poess, M. (eds.) Performance Evaluation and Benchmarking 2009, LNCS, vol. 5895, pp. 146-152. Springer, Heidelberg (2009).

[4] Bermbach, Wittern, Tai.: *Cloud Service Benchmarking - Measuring Quality of Cloud Services from a Client Perspective*. Springer International Publishing 2017.

The architecture of the peakmarks® Benchmark Software

The peakmarks® Benchmark Software is stored in the database as stored procedures. Tools such as Putty are required to access the database server. For the transfer of the benchmark software and the exchange of the results, tools based on SCP are used.

The peakmarks® Benchmark Software consists of two components. A workload generator creates the database load, the so-called workloads. All workloads are generated within the database and executed by database jobs.

A performance monitor collects all relevant statistics before and after each single performance test and immediately displays the result in tabular form. Performance is measured at the interface to the application, i.e., the measured performance is directly available to the application.

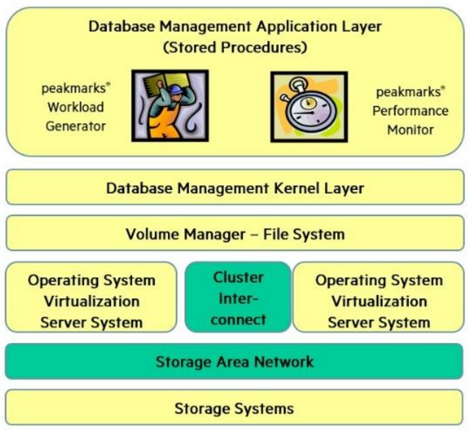

peakmarks® Benchmark Software
On-Premise

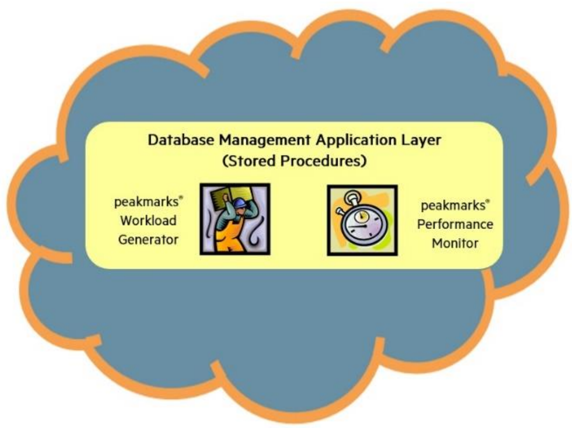

peakmarks® Benchmark Software
in der Cloud

The synthetic benchmark database can be adapted to customer requirements via various peakmarks® configuration parameters.

- Number of CPUs.
- Size database buffer cache.
- Database size between 64 GByte and 64 TByte per database instance.
- Optional use of Database Flash Cache for conventional database servers and Cell Flash Cache for engineered systems.

The runtime of performance tests can be set between 60 seconds and 12 hours.

Optionally, Oracle Data Guard can be configured. Oracle Real Application Cluster technology with multiple database servers is also supported.

Full 360-degree performance overview

With over 30 workloads, so-called micro-benchmarks, in 8 workload groups, the peakmarks® Benchmark Software provides a representative and complete overview of the performance of an Oracle Database Service in all load situations.

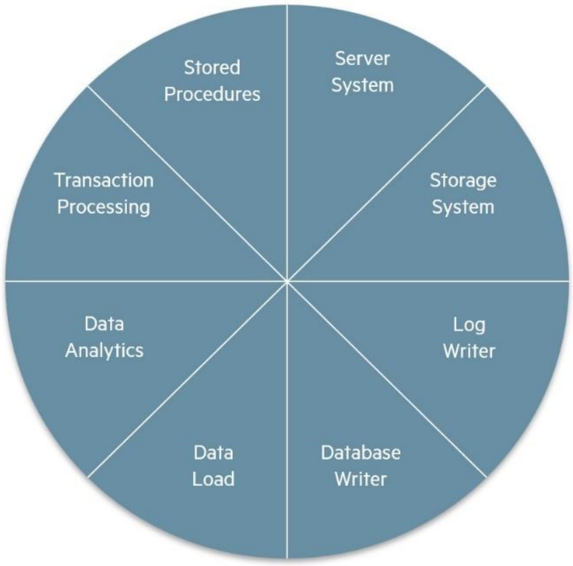

The two most important components of a platform - server system and storage system - have a decisive influence on the performance of all database operations and determine Oracle license and maintenance costs. Therefore, one group of workloads analyzes a server system's performance behavior in Oracle database operation. Another workload group determines the performance behavior of a storage system in Oracle database operation.

Two service processes are of great importance for the smooth operation of the database. Log writer processes are responsible for transaction management, database writer processes for buffer management. The peakmarks® Benchmark Software examines the performance of these two important database service processes with specially developed workloads.

Performance indicators for representative database operations, such as data load, data analytics, and transaction processing, are required for capacity planning. Different technologies can be used for this.

The Data Load Workloads determine the throughput for loading data using different methods such as *Buffered Load*, *Direct Load,* and *Stream Load* for IoT applications.

The data analytics workloads explore the performance of various technologies to accelerate the searching of non-indexed data, such as *Smart Scan* technology on Oracle Engineered Systems or *In-Memory* technology available on all platforms.

The transaction processing workloads determine transaction throughput and response time behavior for transactions of varying complexity.

Some applications encapsulate centrally important business functions and transactions in stored procedures. The peakmarks® Benchmark Software offers workloads to analyze the performance behavior of PL/SQL application code on different processors.

All key performance indicators are presented in peakmarks® own reports and summarized in a performance certificate. Oracle AWR reports are automatically generated for a detailed analysis of all performance tests.

Configuration of peakmarks® Performance Tests

For the configuration of a single performance test, the peakmarks® Benchmark Software needs the following information:

- Workload to be performed.
- For some workloads there are optional workload parameters.
- Internal Oracle parallelism (DOP) can be selected for some workloads.
- Number of cluster nodes on which the workload is to be executed.
- Number of processes executing the workload. In a cluster, the processes are distributed equally across all cluster nodes in a *round-robin fashion.*
- Target runtime for the performance test.

Workloads are configured for different load situations to measure a database service fully. This requires a sequence of performance tests that can be configured in different ways.

Smart Benchmark Configuration. *Smart Benchmark Configuration* is fully automated. It is the most convenient and fastest way to get a complete overview of the performance of an Oracle Database Service. For each workload, a sequence of performance tests (based on the Oracle configuration parameter CPU_COUNT, which can be overridden manually) is generated automatically. Should the system reach saturation as the load increases, the test sequence will not continue.

Sample Benchmark Configuration. The *Sample Benchmark Configuration* also runs completely automatically but only performs a selected set of pre-configured performance tests. This benchmark configuration is often used with cloud services to periodically check within a maintenance window of 60 minutes whether the database service has still fully maintained its performance characteristics.

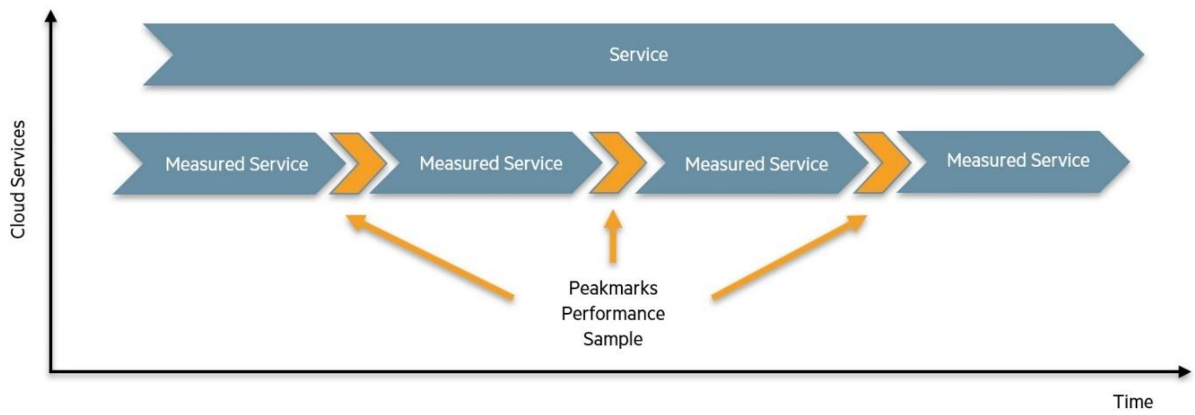

Manual Benchmark Configuration. The *Manual Benchmark Configuration* offers the highest flexibility for configuring benchmark tests. Each parameter of a performance test can be selected individually. Engineers prefer this method to analyze special load situations in detail.

Structure of the peakmarks® Benchmark Reports

Units of measurement and abbreviations

The following units of measurement and abbreviations are used in the benchmark reports; the columns of the respective key performance metrics are marked in red.

Abbreviation	Meaning	Abbreviation	Meaning
[MBps]	megabyte per second	[s]	seconds
[GBps]	gigabyte per second	[ms]	milliseconds
[TBph]	terabyte per hour	[µs]	microseconds
[rps]	rows per second	[dbps]	database blocks per second
[qps]	queries per second	[dbpt]	Database blocks per transaction
[tps]	transactions per second	[rbps]	redo blocks per second
		[kBpt]	kiloByte per transaction
[IOPS]	I/O operations per second	[Mops]	million operations per seconds
BuCache	Oracle buffer cache	FlCache	Flash Cache (Database, Exadata)

Examples

The following examples show two tabular peakmarks® benchmark reports[5].

The first example shows the performance behavior of the system when running the SRV-QUERY25 workload. In the 2-node cluster, each server has 48 cores and 96 threads.

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-QUERY25	2	2	2	1	0	98	53,726	26,863	0.037	1,506,933	753,467	99.96	180
	2	SRV-QUERY25	2	48	25	25	0	75	902,585	18,804	0.053	25,235,088	525,731	100.00	182
	3	SRV-QUERY25	2	96	51	50	0	49	1,577,747	16,435	0.061	43,988,376	458,212	100.00	182
	4	SRV-QUERY25	2	144	75	73	1	25	1,802,631	12,518	0.079	50,170,859	348,409	100.00	182
	5	SRV-QUERY25	2	192	97	96	1	3	2,000,281	10,418	0.096	55,576,848	289,463	100.00	182

[5] All of the following benchmark reports were determined on the current peakmarks® reference system (as of August 2023): an Exadata Quarter Rack system with two Exadata X7 Database Servers and three Exadata X8 High Capacity Storage Servers in an ASM "high redundancy" configuration. The technical details of this configuration can be found in the corresponding data sheets.

The second example shows the performance behavior of the system when executing the workload STO-RANDOM (100% read).

Run	Test	Workload	Wri [%]	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Phys reads total [IOPS]	IO time read [ms]	Phys reads total [MBps]	Phys reads FlCache [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
2	1	STO-RANDOM	0	2	2	1	2	1	1	98	0	67,955	0.214	544	544	100.00	-0.33	183
	2	STO-RANDOM	0	2	16	1	13	7	4	87	0	476,058	0.259	3,743	3,743	100.00	1.65	182
	3	STO-RANDOM	0	2	32	1	21	12	6	79	0	745,340	0.334	5,855	5,855	100.00	1.88	182
	4	STO-RANDOM	0	2	48	1	24	11	8	76	0	1,168,859	0.581	9,179	9,179	100.00	1.71	183
	5	STO-RANDOM	0	2	64	1	26	12	9	74	0	1,287,733	0.857	10,122	10,122	100.00	2.11	182
	6	STO-RANDOM	0	2	80	1	28	13	10	72	0	1,341,404	1.209	10,549	10,549	100.00	2.53	182
	7	STO-RANDOM	0	2	96	1	30	14	11	70	0	1,377,906	1.588	10,831	10,831	100.00	2.57	182
	8	STO-RANDOM	0	2	112	1	32	15	11	68	0	1,399,140	2.008	11,001	11,001	100.00	2.85	182
	9	STO-RANDOM	0	2	128	1	32	15	11	68	0	1,410,819	2.489	11,094	11,094	100.00	3.00	182
	10	STO-RANDOM	0	2	144	1	33	15	12	67	0	1,418,645	2.964	11,154	11,154	100.00	3.04	182
	11	STO-RANDOM	0	2	160	1	34	16	12	66	0	1,422,483	3.455	11,187	11,187	100.00	3.10	182
	12	STO-RANDOM	0	2	176	1	35	17	12	65	0	1,425,543	3.955	11,212	11,212	100.00	3.10	182
	13	STO-RANDOM	0	2	192	1	35	17	12	65	0	1,428,829	4.482	11,236	11,236	100.00	3.23	182

Description of the performance tests

The first column of each benchmark report identifies the benchmark run (*Run* column). Within a benchmark run, all performance tests are numbered (*Test* column). The third column describes the workload. Optionally, there are additional workload parameters depending on the workload. In the second example, this is the percentage of write operations (column *Wri [%]*).

The *Nodes column* indicates how many cluster nodes are involved in a test. The *Jobs column* documents the number of processes on all cluster nodes involved that generate the load. For cluster tests, the processes are distributed *round-robin to* all cluster nodes.

For some workloads, there is an additional column DOP (degree of parallelism), which describes Oracle's internal parallelism.

Choice of measurement points for CPU-bound performance tests

For *CPU-bound* workloads like SRV-QUERY25, the peakmarks® Benchmark Software automatically selects five measurement points to determine the performance of the server in different load situations: Single-thread, 25% CPU utilization, 50% CPU utilization, 75% CPU utilization, 100% CPU utilization.

Choice of measurement points for other workloads

The peakmarks® Benchmark Software selects up to thirty-two measurement points for all other workloads. Starting with one process per server, the load is automatically increased until system saturation occurs. The load increase is determined automatically depending on the Oracle configuration parameter CPU_COUNT value and can be manually overridden.

In the second example with the STO-RANDOM workload, the measurement series is stopped after the thirteenth test because there is no more performance increase. The system is saturated.

CPU utilization

The following four columns describe the CPU utilization of all involved servers. For workloads with a high I/O share, an additional column describes the percentage CPU I/O Wait share (*CPU iow [%]* column).

Key performance indicators

Now columns with the key performance indicators follow. Depending on the workload, different performance indicators are listed. In the first example, the number of queries (column *Queries total [qps]*), the average response time of the queries (column *Response Time [ms]*), and the number of logical database block accesses in the database buffer cache (column *Log reads total [dbps]*).

In the second example, the two key performance indicators are the number of I/O operations (*Phys reads total [IOPS]* column) and the I/O service time (*IO Time read [ms]* column).

Additional columns can often be found in the performance key figures, which show the performance per CPU used[6]. These metrics are important for price/performance comparisons to determine the performance per core. Oracle license and maintenance costs are usually based on the number of sockets used (Standard Edition) and cores (Enterprise Edition).

Other key figures

Columns with additional key figures follow this. These key figures often provide additional information to better understand the system behavior or check whether the workload has also run optimally.

In our first example, these are the Oracle buffer cache hit rate (*BuCache read [%]* column) and the workload processing time (*Elapsed time [s]* column). In the second example, we also find the Flash Cache Hitrate (column *FlCache read [%]*).

Workloads for Server Systems

peakmarks® offers various workloads for server systems. The selected server workloads occur across all industries in all database applications. They show the *real-world* performance of a server system in Oracle database operation.

All server workloads access tables via SQL with different access patterns. The affected tables are completely in the buffer cache. There are almost no I/O operations, so the performance of these workloads is completely limited by CPU power.

Server Performance Tests answer questions about the scalability of servers and the efficiency of technologies such as multithreading, virtualization, and data encryption. For on-premises platforms, they enable a simple price/performance comparison of processors and servers, taking Oracle licensing into account. They also provide comparable metrics on a cloud provider's price/performance ratio when evaluating cloud platforms, where the hardware components are not always known.

[6] The term CPU is based on the Oracle definition in the sense of the Oracle configuration parameter CPU_COUNT. For processors without multithreading, one CPU corresponds to one core; for systems with multithreading, one CPU corresponds to one thread.

Performance indicators for server systems in Oracle database operation

To measure the performance of server systems in Oracle database operation, the following performance indicators have proven to be useful:

- SQL *query throughput* in *queries per second* [qps].
- Average SQL query *response* time in *milliseconds* [ms].
- Number of logical database block accesses in the buffer cache (*logical reads throughput*) in *database blocks per second* [dbps].
- SQL Buffer Cache Scan Rate in *MegaByte per second* [MBps].

Workloads to determine the server performance in Oracle database operation

The following table shows the different workloads used to determine server performance in Oracle database operations.

Workload	Action	Performance indicator	Unit of measure
SRV-QUERY1	Query type A: Select 1 record via index. Example: Select account, product, order, invoice, etc. This workload shows the maximum query throughput of a server and the minimum query response time.	Query throughput Query response time	[qps] [ms]
SRV-QUERY25	Query type B: Select an average of 25 records via index. Example: Select account entries of the last week; list items of an order etc.	Query throughput Query response time	[qps] [ms]
SRV-REPORT	Online Report: Select an average of 125 data sets via index. This workload corresponds to a simple online report and shows the maximum value for *logical reads of* a server. Example: Online report mobile phone call records from last month, online report e-banking with account mutations from last month, etc.	Block accesses in the buffer cache	[dbps]
SRV-SCAN	Data search without index. This corresponds to a *full table scan*. Only one record is determined, but the entire data set of the table in the buffer cache must be searched.	Scan rate in buffer cache	[MBps]

Workload	Action	Performance indicator	Unit of measure
SRV-MIXED	Complex workload with a mix of equally weighted simple workloads running concurrently: ▪ SRV-QUERY1 ▪ SRV-QUERY25 ▪ SRV-REPORT ▪ SRV-SCAN SRV-MIXED displays server system performance in Oracle database operation for a representative cross-section of peakmarks® server workloads.	Query throughput Query response time	[qps] [ms]

Example: Workload SRV-QUERY1

The workload SRV-QUERY1 (1 record per query) shows the maximum possible number of queries and the shortest response time of queries for a server system.

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	2	2	2	1	0	98	233,345	116,673	0.009	933,557	466,779	99.95	180
	2	SRV-QUERY1	2	48	25	25	0	75	4,178,134	87,044	0.011	16,669,620	347,284	100.00	182
	3	SRV-QUERY1	2	96	51	50	1	49	6,520,208	67,919	0.015	25,944,413	270,254	100.00	182
	4	SRV-QUERY1	2	144	75	73	1	25	6,501,930	45,152	0.022	25,846,679	179,491	100.00	182
	5	SRV-QUERY1	2	192	98	96	1	2	6,795,993	35,396	0.028	26,987,811	140,562	100.00	182

Example: Workload SRV-REPORT

The workload SRV-REPORT (an online report with an average of 125 records per query) shows a database server's maximum number of logical reads.

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	11	SRV-REPORT	2	2	2	1	0	98	13,884	6,942	0.144	1,775,598	887,799	100.00	180
	12	SRV-REPORT	2	48	25	25	0	75	221,925	4,623	0.215	28,290,092	589,377	100.00	182
	13	SRV-REPORT	2	96	50	49	0	50	391,932	4,083	0.244	49,811,699	518,872	100.00	182
	14	SRV-REPORT	2	144	76	74	1	24	461,731	3,206	0.310	58,545,252	406,564	100.00	182
	15	SRV-REPORT	2	192	99	97	1	1	525,654	2,738	0.363	66,502,195	346,366	100.00	182

Example: Workload SRV-SCAN

The workload SRV-SCAN (query with a *full table scan*, where 1 hit is determined) shows the scan rate in the buffer cache.

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	21	SRV-SCAN	2	2	2	1	0	98	5,543	2,772	710,561	355,281	99.99	182
	22	SRV-SCAN	2	48	25	25	0	75	108,285	2,256	13,855,451	288,655	100.00	183
	23	SRV-SCAN	2	96	51	50	0	49	187,009	1,948	23,871,434	248,661	100.00	183
	24	SRV-SCAN	2	144	75	74	1	25	227,210	1,578	28,705,443	199,343	100.00	183
	25	SRV-SCAN	2	192	98	97	1	2	258,579	1,347	32,710,285	170,366	100.00	183

Example: Workload SRV-MIXED

The workload SRV-MIXED is often used instead of the other workloads for server systems to quickly obtain a representative statement about the performance of a server system in Oracle database operation.

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	31	SRV-MIXED	2	8	5	4	0	95	272,844	34,106	0.029	4,290,051	536,256	100.00	182
	32	SRV-MIXED	2	48	25	25	0	75	1,331,390	27,737	0.036	21,115,772	439,912	100.00	183
	33	SRV-MIXED	2	96	51	50	0	49	2,238,224	23,315	0.042	36,500,369	380,212	100.00	183
	34	SRV-MIXED	2	144	75	74	1	25	2,442,773	16,964	0.058	42,067,495	292,135	100.00	182
	35	SRV-MIXED	2	192	98	97	1	2	2,736,163	14,251	0.070	47,252,138	246,105	100.00	182

Notes

Main memory accesses[7] **enable ultimate database performance.** Extremely short response times between 9 and 28 µs are determined for the SRV-QUERY1 workload.

Misleading CPU utilization. If Intel Xeon processors are operated with multithreading, the CPU utilization information is not always meaningful. In the SQL-QUERY1 workload, all cores are occupied by processes in test 3. The CPU utilization shows 51%, and 6,520,208 queries are processed per second. In test 5, processes occupy all threads, CPU utilization increases to 98%, and 6,795,993 queries are processed per second. The server is now fully utilized. CPU utilization has almost doubled between the two tests, but throughput has only increased by less than 4%. I.e., already in test 3, the actual CPU utilization was 96% and not 51%!

Many operating systems value threads as resources when calculating CPU utilization, just like cores. But they are not. Threads are merely a concept for improving the parallel processing of instructions.

The multithreading of Intel Xeon processors can lead to a throughput improvement of up to 25% for Oracle workloads. Other values apply to processor architectures such as AMD EPYC, IBM POWER, or IBM z.

Scalability with increasing load. For each server workload, we observe a decreasing performance per CPU with increasing load. For example, the scan rate per CPU for the SRV-SCAN workload drops from 2.772 MBps (test 21, single thread performance) to 1.347 MBps (test 25, system saturation). The performance per CPU drops to about 50% of the peak value with increasing load.

We observe this behavior, especially in Intel Xeon processors with a large spread of the clock rate. In our case (Intel Xeon 8160), the maximum clock rate is 3.7 GHz. The clock rate is reduced to 2.1 GHz for thermal reasons when the load increases. This means a reduction of the power output by 43%.

As a rule, Intel Xeon processors with many cores also have a high clock rate spread. And vice versa, Intel Xeon processors with fewer cores also have a lower clock rate spread and can thus guarantee a more stable performance output.

[7] For orientation: the fastest processors (as of August 2023) require less than 0.1 nanoseconds for a DRAM access. The database needs about 1 microsecond for a block access in the Oracle buffer cache and between 20 and 500 microseconds (depending on the storage fabric and protocol) for a block access on flash storage.

Server evaluation for Oracle database operation. The use of Intel Xeon processors with a lower number of cores is recommended to achieve predictable and persistent CPU performance in all load ranges. Such processors also deliver the highest per-core performance and therefore improve a server's price/performance ratio, taking into account Oracle licensing and maintenance costs.

More sockets may then be necessary to achieve scalability requirements. This also leads to the ability to use significantly higher main memory capacities, which in turn can lead to a further improvement in performance and benefits the use of *In-Memory* technologies.

Workloads for Storage Systems

Conventional I/O benchmark tools such as fio, vdbench, iometer, and Orion often show performance values not achieved in real database operation. The reason for this is the complexity of database I/O operations, which is not taken into account by these tools.

When a data block is read, the buffer cache management of the database has to perform many tasks:

- find a free place for the block.
- if there is no free space, replace older blocks.
- synchronize all database processes simultaneously trying to occupy free spaces in the buffer cache.
- when using a shared disk cluster architecture (Oracle Real Application Cluster), synchronization must be performed cluster-wide, which in the best case requires exchanging information (messages) between cluster nodes. In the worst case, additionally, the exchange of database blocks.
- Finally, blocks are checked for integrity and consistency during I/O transfer. Oracle configuration parameters can define the scope of integrity and consistency tests. These tests can lead to an additional load on processors and I/O system.

The peakmarks® Benchmark Software generates I/O load with so-called SQL-generated I/O operations to determine representative performance indicators for the storage system in Oracle database operation. In the case of sequential operations, the intra-SQL parallelism can also be tested to find the optimum SQL parallelism (Oracle DOP).

Storage performance tests provide answers to questions about the efficiency of protocols such as NVMe *over fabrics* (IP, FC), RoCE (*RDMA over Converged Ethernet*), and technologies such as flash storage, persistent memory, storage tiering, data encryption, offload functions of intelligent storage systems and data reduction through de-duplication and compression. They provide well-founded key figures for price/performance comparisons of storage systems for Oracle databases and the evaluation of suitable storage replication methods.

Performance indicators for storage systems in Oracle database operation

The following performance indicators are used to measure the performance of storage systems in database operation:

- *Sequential* read throughput of storage systems in *megabytes per second* [MBps]. Large I/O units of 1 MByte are typically used.
- *Random* access of storage systems (*random I/O*) in *I/O operations per second* [IOPS]. Typically, smaller I/O units like Oracle database blocks are used. The size of Oracle database blocks is configurable: 2 KByte, 4 KByte, 8 KByte, 16 KByte, 32 KByte). The default value is 8 Kbytes.
- I/O service time during random access in *milliseconds* [ms].

Workloads for determining storage performance in database operation

The following table shows the available workloads of the peakmarks® Benchmark Software for measuring the performance of storage systems in Oracle database operation.

Workload	Action	Performance indicator	Unit of measure
STO-READ	Sequential read operation generated by SQL statements.	Throughput *sequential* I/O	[MBps]
STO-OFFLOAD	Sequential read operation generated by SQL statements when using intelligent storage servers with offload functionality (*smart scan*).	Throughput *sequential* I/O	[MBps]
STO-RANDOM	Random read and write operations generated by SQL statements. In this workload, read operations are performed by foreground processes and write operations by background processes. The ratio of database blocks that are read or written is configurable[8].	Throughput *random* I/O I/O service time	[IOPS] [ms]
STO-SCATTER	Random write operations generated by SQL statements. In this workload, write operations are performed by foreground processes.	Throughput *random* I/O	[dbps]

[8] Oracle often performs write operations as multiblock I/O operations; a list of blocks is written with one I/O operation.

Example: Workload STO-READ

The workload STO-READ shows the I/O throughput between the database server and storage system during sequential read operations. In this example, an Oracle DOP of 4 is used. The data volume is loaded from the flash cache (*FlCache read* column).

Run	Test	Workload	Wri [%]	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Phys reads total [IOPS]	IO time read [ms]	Phys reads total [MBps]	Phys reads FlCache [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
4	1	STO-READ	0	2	2	4	2	1	1	98	0	6,628	1.015	6,531	6,531	100.00	0.00	175
	2	STO-READ	0	2	8	4	2	1	1	98	0	11,147	2.724	11,078	11,078	100.00	0.00	175
	3	STO-READ	0	2	16	4	3	2	1	97	0	11,682	5.297	11,631	11,631	100.00	0.00	183
	4	STO-READ	0	2	24	4	3	2	1	97	0	12,024	7.740	11,981	11,981	100.00	0.00	188
	5	STO-READ	0	2	32	4	3	2	1	97	0	12,118	10.107	12,067	12,067	100.00	0.00	193
	6	STO-READ	0	2	40	4	3	2	1	97	0	12,174	12.646	12,132	12,132	100.00	0.00	204
	7	STO-READ	0	2	48	4	3	2	1	97	0	12,108	15.116	12,061	12,061	100.00	0.00	194

Example: Workload STO-OFFLOAD

The STO-OFFLOAD workload is the same as STO-READ but uses the storage systems' intelligence and *Smart Scan* technology. When using Smart Scan technology, it is unnecessary to specify a DOP because the storage servers optimize access independently.

Run	Test	Workload	Wri [%]	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Phys reads total [IOPS]	IO time read [ms]	Phys reads total [MBps]	Phys reads FlCache [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
4	11	STO-OFFLOAD	0	2	2	1	2	1	1	98	0	45,224	0.478	44,931	44,931	100.00	0.00	171
	12	STO-OFFLOAD	0	2	8	1	1	1	1	99	0	64,711	0.668	64,303	64,303	100.00	0.00	172
	13	STO-OFFLOAD	0	2	16	1	2	1	1	98	0	69,557	1.185	69,122	69,122	100.00	0.00	173
	14	STO-OFFLOAD	0	2	24	1	2	1	1	98	0	72,848	1.749	72,394	72,394	100.00	0.00	174
	15	STO-OFFLOAD	0	2	32	1	1	1	1	99	0	72,396	2.372	71,943	71,943	100.00	0.00	175

Example: Workload STO-RANDOM

The workload STO-RANDOM shows the I/O throughput between the database server and storage system with *random* access to individual database blocks. In this example, the size of the Oracle database blocks is 8 KByte.

Only read operations are measured in the first benchmark report (tests 21 to 27). In the second benchmark report (tests 31 - 37), the ratio of read-to-write operations is 80:20 (column *Wri [%]*).

Run	Test	Workload	Wri [%]	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Phys reads total [IOPS]	IO time read [ms]	Phys reads total [MBps]	Phys reads FlCache [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
4	21	STO-RANDOM	0	2	2	1	2	1	1	98	0	67,955	0.214	544	544	100.00	-0.33	183
	22	STO-RANDOM	0	2	32	1	21	12	6	79	0	745,340	0.334	5,855	5,855	100.00	1.88	182
	23	STO-RANDOM	0	2	64	1	26	12	9	74	0	1,287,733	0.857	10,122	10,122	100.00	2.11	182
	24	STO-RANDOM	0	2	96	1	30	14	11	70	0	1,377,906	1.588	10,831	10,831	100.00	2.57	182
	25	STO-RANDOM	0	2	128	1	32	15	11	68	0	1,410,819	2.489	11,094	11,094	100.00	3.00	182
	26	STO-RANDOM	0	2	160	1	34	16	12	66	0	1,422,483	3.455	11,187	11,187	100.00	3.10	182
	27	STO-RANDOM	0	2	192	1	35	17	12	65	0	1,428,829	4.482	11,236	11,236	100.00	3.23	182

Run	Test	Workload	Wri [%]	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Phys reads total [IOPS]	IO time read [ms]	Phys reads total [MBps]	Phys reads FlCache [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
4	31	STO-RANDOM	20	2	2	1	2	1	1	98	0	66,545	0.227	524	524	100.00	17.99	182
	32	STO-RANDOM	20	2	16	1	10	5	3	90	0	423,127	0.336	3,319	3,319	100.00	15.93	183
	33	STO-RANDOM	20	2	32	1	17	9	5	83	0	662,583	0.534	5,208	5,208	100.00	16.78	183
	34	STO-RANDOM	20	2	48	1	21	11	6	79	0	781,877	0.817	6,153	6,153	100.00	17.56	183
	35	STO-RANDOM	20	2	64	1	24	13	7	76	0	845,871	1.105	6,651	6,651	100.00	18.18	184
	36	STO-RANDOM	20	2	80	1	26	14	7	74	0	883,718	1.471	6,951	6,951	100.00	18.69	184
	37	STO-RANDOM	20	2	96	1	27	15	8	73	0	904,607	1.862	7,116	7,116	100.00	18.91	183

Notes

I/O bandwidth. With workload STO-READ, the amount of data transferred is usually limited by the I/O bandwidth between the server and storage systems.

I/O service time. With the workload STO-RANDOM, the I/O service time is influenced by the efficiency of the I/O stack. Service times of less than 500 microseconds are expected when using all-flash arrays. With new protocols such as RoCE (*RDMA over converged Ethernet*), I/O service times of less than 20 microseconds are now possible if data is found in the storage system's cache.

Log Writer Workloads

Log Writer processes are primarily responsible for transaction logging and database recovery from system failures. These processes are critical to Oracle's overall performance when processing transactions. Transaction logging latency can have a significant impact on user transaction response time.

Optionally, the Log Writer processes are also used for database replication to synchronize standby databases. This technology is very popular for disaster recovery solutions. Replication can be performed in synchronous or asynchronous mode. The data transfer between the primary and standby databases can be optionally encrypted and/or compressed. With synchronous replication, local transactions must wait until the standby databases have received the transaction log. This can significantly delay local transaction processing[9].

Performance metrics for Log Writer processes

The performance metrics for Log Writer processes are:

- Number of SQL commit operations in *transactions per second* [tps].
- Average latency for SQL commit operations in *milliseconds* [ms].
- REDO throughput in *megabytes per second* [MBps].

[9] When using Oracle Data Guard as a BCP solution, synchronous replication is used for *mission-critical* systems. This means an additional latency of at least 1 ms for the log writers.

Workloads for Log Writer processes

To analyze the performance behavior of the Log Writer processes for different transaction sizes, the peakmarks® Benchmark Software offers different workloads.

Workload	Action	Performance indicator	Unit of measure
LGWR-LAT1 LGWT-LAT25 LGWR-LAT125	Transactions with approx. 1, 25, or 125 KByte REDO-data volume per transaction. All transactions conclude with a *COMMIT WAIT IM-MEDIATE.*	SQL Commit Operations SQL Commit Latency	[tps] [ms]
LGWR-THR	Transactions that generated a large amount of REDO data. All transactions conclude with a *COMMIT NOWAIT BATCH*.	Log Writer Throughput	[MBps]

Example: Workload LGWR-LAT1

The workload LGWR-LAT1 shows the maximum number of commit operations the system can process. The average transaction size, measured by the REDO data volume, is just under 1.4 kByte (column *REDO data [kBpt]*).

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	#prc LGWR	Commit thrput [tps]	Commit latency [ms]	LogFile sync [ms]	REDO data [rbps]	REDO data [IOPS]	REDO data [MBps]	REDO data [kBpt]	FlCache write [%]	Elapsed time [s]
5	1	LGWR-LAT1	2	2	2	1	1	98	0	18	6,165	0.323	0.234	19,487	6,258	8	1.396	99.98	182
	2	LGWR-LAT1	2	32	6	4	1	94	0	18	55,534	0.572	0.461	197,131	10,794	77	1.417	100.00	183
	3	LGWR-LAT1	2	64	8	6	2	92	0	18	88,145	0.722	0.584	300,785	9,020	121	1.404	100.00	182
	4	LGWR-LAT1	2	96	11	9	2	89	0	18	110,529	0.863	0.695	364,364	7,960	151	1.401	100.00	183
	5	LGWR-LAT1	2	128	14	11	2	86	0	18	131,324	0.969	0.781	421,304	7,288	179	1.395	100.00	182
	6	LGWR-LAT1	2	160	16	13	2	84	0	18	153,826	1.031	0.833	478,342	5,434	209	1.388	87.62	183
	7	LGWR-LAT1	2	192	18	14	2	82	0	18	172,065	1.104	0.882	526,155	3,995	234	1.390	100.00	184
	8	LGWR-LAT1	2	224	20	16	3	80	0	18	183,894	1.187	0.937	558,979	3,749	250	1.393	90.58	186
	9	LGWR-LAT1	2	256	23	19	3	77	0	18	195,328	1.276	1.002	590,846	3,570	266	1.394	77.85	186
	10	LGWR-LAT1	2	288	25	21	3	75	0	18	203,230	1.356	1.073	610,915	3,350	276	1.392	100.00	189
	11	LGWR-LAT1	2	320	27	22	3	73	0	18	209,735	1.443	1.152	627,172	3,097	285	1.391	100.00	192

Example: Workload LGWR-THR

The workload LGWR-THR shows the maximum REDO data volume that the Log Writer can process. The number of Log Writer processes is not configurable and is determined automatically by Oracle.

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	#prc LGWR	Commit thrput [tps]	Commit latency [ms]	LogFile sync [ms]	REDO data [rbps]	REDO data [IOPS]	REDO data [MBps]	REDO data [kBpt]	FlCache write [%]	Elapsed time [s]
5	12	LGWR-THR	2	2	2	1	1	98	0	18	488	4.092	1.348	270,752	468	128	268.396	100.00	181
	13	LGWR-THR	2	16	9	7	1	91	0	18	3,026	5.261	1.693	1,614,870	2,279	760	257.316	100.00	181
	14	LGWR-THR	2	32	14	11	2	86	0	18	4,501	7.090	3.784	2,302,606	1,883	1,085	246.868	100.00	181
	15	LGWR-THR	2	48	16	13	2	84	0	18	4,986	9.615	6.496	2,437,707	1,578	1,150	236.091	100.00	181
	16	LGWR-THR	2	64	18	14	2	82	0	18	5,154	12.388	8.322	2,491,403	1,446	1,175	233.485	100.00	182
	17	LGWR-THR	2	80	18	14	3	82	0	18	5,281	15.100	10.556	2,497,715	1,366	1,178	228.475	100.00	181
	18	LGWR-THR	2	96	19	15	3	81	0	18	5,502	17.362	12.947	2,556,880	1,401	1,206	224.497	95.20	182
	19	LGWR-THR	2	112	20	16	3	80	0	18	5,558	20.076	14.167	2,576,229	1,313	1,215	223.921	100.00	181

Note

Data reduction technologies on Storage Systems. Some cloud service providers use storage systems with data reduction processes (deduplication, compression). These technologies reduce the required storage capacity but can harm the latency of log writer processes and the transaction throughput of OLTP systems.

Database Writer Workload

The database writers are responsible for managing the buffer cache. These processes are critical to Oracle's overall performance when many blocks are modified, such as during conventional buffer cache data loading or intensive transaction processing.

The number of Database Writer processes is determined automatically by Oracle. If this value is not sufficient, it can be adjusted manually.

Performance metrics for Database Writer processes

The performance metric for Database Writer processes is

- Throughput of database blocks written back in *database blocks per second* [dbps].

Workloads to determine the Database Writer performance

The DBWR-THR workload is available to determine Database Writer performance.

Workload	Action	Performance indicator	Unit of measure
DBWR-THR	Massive block changes in the buffer cache	Database Writer Throughput	[MBps]

Example: Workload DBWR-THR

The DBWR-THR workload indicates the maximum number of blocks database writer processes can write back.

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	#prc DBWR	Phys writes total [dbps]	Phys writes total [IOPS]	Phys writes total [MBps]	REDO data [MBps]	Phys writes FlCache [MBps]	FlCache write [%]	Elapsed time [s]
6	1	DBWR-THR	2	2	4	2	1	96	0	24	39,488	27,076	373	62	373	100.00	181
	2	DBWR-THR	2	8	9	6	2	91	0	24	197,333	34,425	1,778	227	1,778	100.00	182
	3	DBWR-THR	2	16	16	11	3	84	0	24	300,089	83,328	2,928	366	2,643	90.27	182
	4	DBWR-THR	2	24	23	16	5	77	0	24	328,586	182,857	3,047	462	2,735	89.77	183
	5	DBWR-THR	2	32	25	16	6	75	0	24	304,372	241,604	2,858	462	2,702	94.54	182

Notes

ASM redundancy. When configuring ASM, different redundancy levels can be defined (*external*, *normal redundancy*, *high redundancy*). This has a significant impact on the performance of writing processes.

Our example uses the "high redundancy" level common to Exadata systems. This means that a total of three times the amount of data must be written, e.g., in test 4, a total of 3 x 328.586 database blocks!

Data Load Workloads

System architects and capacity planners need performance metrics from database services regarding their ability to load data. This is especially important for data warehouse and data analytics systems, where data volumes are increasing and the time windows for loading are decreasing.

Oracle provides several technologies for loading data:

- conventional loading via the buffer cache is preferred in *transaction* processing *systems.*
- direct loading by bypassing the buffer cache; this method is mainly used in the *data warehouse* and *data analytics* environment to load large amounts of data as quickly as possible.
- as of Oracle 19c, a new loading procedure is available for IOT (*Internet of Things*) applications. Mass data can be loaded quickly, but transaction integrity is not always guaranteed. If aggregates such as average values, etc., do not depend on each data set, this loading procedure is acceptable.

Key figures for data load performance

There is only one performance metric for data load workloads:

- Throughput of loaded data (*data load rate*) in *megabyte per second* [MBps] or *terabyte per hour* ([TBph].

Workloads for determining the data load performance

The peakmarks® Benchmark Software offers workloads for all three charging techniques.

Workload	Action	Performance indicator	Unit of measure
DL-BUFFER	The data is program-generated and conventionally loaded via the buffer cache. 3 additional indexes are updated during loading. Typically, only a few records are loaded per transaction. The load process uses *COMMIT WRITE WAIT IMMEDIATE* at transaction completion.	Data Load Throughput	[MBps]
DL-DIRECT	Data is loaded from a data source directly, bypassing the buffer cache. Typically, large amounts of data are loaded per transaction. Only one additional index is updated during loading. The loading process uses the *NOLOGGING* option, and at transaction completion, *COMMIT WRITE NOWAIT BATCH*.	Data Load Throughput	[MBps]

Workload	Action	Performance indicator	Unit of measure
DL-STREAM	The data is program-generated and uses the *memory-optimized row store for write* for IOT applications. Only one additional index is updated during loading. The load process does not use a commit but only a flush operation of the memory-optimized pool.	Data Load Throughput	[MBps]

Example: Workload DL-BUFFER

The workload DL-BUFFER shows the maximum load rate achieved with the *buffered data load.* This method of loading is preferred in transaction-oriented systems.

Run	Test	Workload	TX size [rpt]	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Load rate total [rps]	Load rate total [MBps]	Phys writes total [IOPS]	Phys writes total [MBps]	REDO data [MBps]	FlCache write [%]	Elapsed time [s]
7	1	DL-BUFFER	5	2	2	2	1	1	98	0	25,151	7	5,796	46	24	100.00	182
	2	DL-BUFFER	5	2	32	7	5	1	93	0	204,245	58	14,866	398	193	100.00	183
	3	DL-BUFFER	5	2	64	11	8	2	89	0	303,210	87	18,297	617	275	100.00	182
	4	DL-BUFFER	5	2	96	13	10	2	87	0	377,273	108	19,309	828	347	100.00	182
	5	DL-BUFFER	5	2	128	17	13	2	83	0	420,796	120	23,709	988	392	100.00	182
	6	DL-BUFFER	5	2	160	20	16	3	80	0	452,938	130	28,397	1,123	416	100.00	182
	7	DL-BUFFER	5	2	192	23	18	3	77	0	476,194	136	32,492	1,248	439	100.00	184

Example: Workload DL-DIRECT

The workload DL-DIRECT shows the maximum load rate that is achieved with the *direct data load* bypassing the buffer cache. This loading method is preferably used in data warehouse systems.

Run	Test	Workload	TX size [rpt]	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Load rate total [rps]	Load rate total [MBps]	Phys writes total [IOPS]	Phys writes total [MBps]	REDO data [MBps]	FlCache write [%]	Elapsed time [s]
7	14	DL-DIRECT	125,000	2	2	2	1	0	98	0	940,670	269	1,917	411	62	100.00	181
	15	DL-DIRECT	125,000	2	16	8	7	1	92	0	5,540,540	1,585	12,476	2,437	362	100.00	182
	16	DL-DIRECT	125,000	2	32	11	9	1	89	0	6,851,879	1,960	18,829	3,021	448	100.00	182
	17	DL-DIRECT	125,000	2	48	13	11	2	87	0	6,923,907	1,981	28,546	3,085	454	100.00	182
	18	DL-DIRECT	125,000	2	64	13	10	2	87	0	6,742,719	1,929	36,474	3,025	443	100.00	182

Data Analytics Workloads

System architects and capacity planners require performance metrics from Database Services regarding their ability to search large amounts of data. These applications are typically based on *full table scan* operations and are not supported by index structures.

Full table scans' performance depends on the data's position in the storage hierarchy (storage, main memory) and the technology used to increase scan performance (*Smart Scan*, *In-Memory Column Store*).

The *Smart Scan* technology is only available on Oracle Engineered Systems and is additionally subject to licensing. The *In-Memory* technology is available on all platforms and is license-free from Oracle 21c up to a certain capacity.

Data analytics workloads answer questions about both technologies' efficiency and enable price/performance comparisons.

Key figures for data analytics performance

For data analytics workloads, there is only one performance metric:

- Throughput of scanned data (*data scan rate*) in *megabyte per second* [MBps] or *rows per second* [rps].

Workloads for determining the data load performance

peakmarks® offers workloads for testing different data locations (storage, main memory) and different boost technologies (*Smart Scan*, *In-Memory Column Store*). The intra-SQL parallelism can also be changed to find the optimal database parallelism (Oracle DOP) during Data Scan.

Workload	Action	Performance indicator	Unit of measure
DA-STORAGE	*Full table scan* with an aggregate of low complexity. The data is read from the storage system.	Data Scan Throughput	[MBps]
DA-OFFLOAD	*Full table scan* with an aggregate of low complexity. The intelligent storage system with offload functionality (smart scan) reads the data.	Data Scan Throughput	[MBps]
DA-ROWSTORE	*Full table scan* with an aggregate of low complexity. The data is read from the *Row Store* in the buffer cache.	Data Scan Throughput	[MBps]
DA-COLSTORE	*Full table scan* with an aggregate of low complexity. The column store in the buffer cache processes the data.	Data Scan Throughput	[rps]

Example: Workload DA-STORAGE

The DA-STORAGE workload shows the maximum scan rate when all data must be read from the storage. The data scan (filter operations) is processed on the database server.

The Scanned user data [MBps] metric corresponds to the actual amount of data scanned. The Scanned *user data [rps]* metric depends on the data model and applies to a record length of 300 bytes (record length of the synthetic tables of the peakmarks® benchmark software).

Run	Test	Workload	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Scan rate total [rps]	Scan rate total [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
8	1	DA-STORAGE	2	2	4	2	1	1	98	0	16,450,340	5,601	100.00	0.00	171
	2	DA-STORAGE	2	16	4	3	2	1	97	0	34,464,955	11,734	100.00	0.00	172
	3	DA-STORAGE	2	32	4	4	2	1	96	0	35,703,125	12,155	100.00	0.00	172
	4	DA-STORAGE	2	48	4	4	2	1	96	0	36,270,748	12,348	100.00	0.00	172
	5	DA-STORAGE	2	64	4	4	3	1	96	0	36,527,567	12,436	100.00	0.00	173
	6	DA-STORAGE	2	80	4	4	3	1	96	0	36,640,231	12,474	100.00	0.00	174
	7	DA-STORAGE	2	96	4	4	3	1	96	0	36,765,045	12,516	100.00	0.00	174
	8	DA-STORAGE	2	112	4	4	3	1	96	0	36,773,204	12,519	100.00	0.00	174

Example: Workload DA-OFFLOAD

The DA-OFFLOAD workload shows the maximum scan rate when all data is processed by the intelligent storage system (filter operations). The database server only receives the amount of hits transferred.

This offload technology is only available on Oracle Engineered Systems and requires additional licenses for their storage servers.

Run	Test	Workload	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Scan rate total [rps]	Scan rate total [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
8	11	DA-OFFLOAD	2	2	1	1	1	0	99	0	58,173,231	19,819	100.00	0.00	171
	12	DA-OFFLOAD	2	16	1	2	1	1	98	0	164,581,594	56,070	100.00	0.00	171
	13	DA-OFFLOAD	2	32	1	2	1	1	98	0	187,076,316	63,733	100.00	0.00	171
	14	DA-OFFLOAD	2	48	1	2	1	1	98	0	195,758,609	66,691	100.00	0.00	171
	15	DA-OFFLOAD	2	64	1	2	1	1	98	0	196,795,769	67,044	100.00	0.00	172
	16	DA-OFFLOAD	2	80	1	2	1	1	98	0	199,305,704	67,899	100.00	0.00	172
	17	DA-OFFLOAD	2	96	1	2	1	1	98	0	201,142,122	68,525	100.00	0.00	172
	18	DA-OFFLOAD	2	104	1	3	1	1	97	0	201,183,310	68,539	100.00	0.00	172
	19	DA-OFFLOAD	2	112	1	2	1	1	98	0	200,329,561	68,248	100.00	0.00	172

Example: Workload DA-ROWSTORE

The workload DA-ROWSTORE shows the maximum scan rate when all data is processed in the buffer cache of the Database Server.

This form of processing is also available for Oracle Standard Edition. A large main memory capacity has an advantageous effect in storing as much data as possible in the buffer cache.

Run	Test	Workload	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Scan rate total [rps]	Scan rate total [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
8	21	DA-ROWSTORE	2	2	1	2	1	0	98	0	15,223,603	4,600	0.00	99.74	181
	22	DA-ROWSTORE	2	48	1	25	25	0	75	0	362,635,020	109,529	0.00	100.00	183
	23	DA-ROWSTORE	2	96	1	50	49	0	50	0	617,217,180	186,421	0.00	100.00	183
	24	DA-ROWSTORE	2	144	1	75	73	1	25	0	754,465,316	227,875	0.00	100.00	183
	25	DA-ROWSTORE	2	192	1	96	94	1	4	0	845,157,690	255,266	0.00	100.00	183

Example: Workload DA-COLSTORE

The DA-COLSTORE workload shows the maximum scan rate when all data is processed in the *column store of* the Database Server. The use of this *In-Memory* technology requires appropriate licenses on the database server.

It is unnecessary to store entire tables in the column store, but only individual selected columns frequently used as search criteria. Compression procedures can also increase the amount of data stored in the column store. Here, too, a large main memory capacity is advantageous to store as much data as possible in the *column store.*

Run	Test	Workload	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Scan rate total [rps]	Scan rate total [MBps]	FlCache read [%]	BuCache read [%]	Elapsed time [s]
8	31	DA-COLSTORE	2	2	1	2	1	0	98	0	10,861,614,976	3,287,494	0.00	99.99	182
	32	DA-COLSTORE	2	48	1	26	25	0	74	0	141,771,840,591	42,916,756	0.00	100.00	182
	33	DA-COLSTORE	2	96	1	52	50	1	48	0	116,215,766,166	35,198,349	0.00	100.00	181
	34	DA-COLSTORE	2	144	1	75	72	1	25	0	125,054,889,772	37,885,745	0.00	100.00	181
	35	DA-COLSTORE	2	192	1	96	93	1	4	0	131,717,969,993	39,916,054	0.00	100.00	182

The Scanned user data [rps] metric corresponds to the actual amount of data scanned. The Scanned *user data [MBps]* metric depends on the data model and applies to a record length of 300 bytes (record length of the synthetic tables of the peakmarks® Benchmark Software).

Notes

The efficiency of the various processes. It's clear that *In-Memory* technology with a column store performs better than *Smart Scan* technology. *In-Memory* technology has recently become the preferred technology for data analytics applications.

Price/performance comparison of the different processes. Additional licenses are required for both *Smart Scan* technology and *In-Memory* technology. It is noticeable that the *In-Memory* technology has a significantly better price/performance ratio than the *Smart Scan* technology if the scan rates are considered concerning the license costs.

However, *Smart Scan* technology works largely transparently for applications, and code changes are rarely required. *Smart Scan* technology can be combined with *In-Memory* technology on storage servers and can manage immense data capacities. However, *Smart Scan* technology can only be used on Oracle Engineered Systems.

On the other hand, In-Memory technology is available on all platforms, requires no special hardware, and delivers ultimate performance. However, the database server's main memory capacity (DRAM) limits the amount of data. Also, intervention in the application is required to specify which table columns should be included in the *column store.* Compression methods help to use the capacity of the *column store* optimally.

Transaction Processing Workloads

System architects and capacity planners need performance metrics from database services regarding their ability to execute typical transactions.

Transaction processing is an extremely complex form of data processing. All platform components and database service processes are significantly involved and must be precisely balanced. Transaction processing is the supreme discipline of database platforms.

Transaction processing performance depends on very many factors. The most important are:

- Ratio of database size and buffer cache size expressed in the buffer cache hit rate.
- I/O performance when reading data.
- I/O performance when writing the transaction log by REDO Log Writer processes.
- I/O performance when writing data through Database Writer processes.

The I/O performance when reading data can in turn be influenced by storage tiering and many other factors. The I/O performance of the log writer processes can be influenced by the data guard performance or other technologies for mirroring data between different data centers.

Key figures for transaction processing performance

For transaction processing workloads, the following performance metrics are of interest:

- SQL *transaction throughput* in *transactions per second* [tps].
- Average SQL *transaction response* time in *milliseconds* [ms].

Transaction Processing Performance Workload Overview

The peakmarks® Benchmark Software provides different types of workloads:

- Reading workloads TP-LOOKUP and TP-REPORT.
- Mutating workloads with different complexity TP-LIGHT, TP-MEDIUM, and TP-HEAVY. The ratio between reading and modifying operations can be configured for these workloads.
- Mix of workloads with SELECT, UPDATE, and INSERT statements, running multiple simple TP workloads concurrently and equally weighted, with the addition of a DL-BUFFER workload.

  TP-MIXED1 is a read-intensive transaction mix, and TP-MIXED2 is a write-intensive workload.

  These two workloads are preferred to quickly determine a representative statement of the transaction processing performance of an Oracle database platform.

Workload	Action	Performance indicator	Unit of Measure
TP-LOOKUP	Fast look-up query that checks the existence of a single record. Example: Check credit card, car registration, phone number, etc. If available on the platform, this workload uses *memory-optimized tables* for fast access via hash key. Otherwise, conventional tables with a *unique b-tree index* are used.	Query throughput Query response time	[tps] [ms]
TP-REPORT	Online report that selects an average of 25 data records per query.	Query throughput Query response time	[tps] [ms]
TP-LIGHT	Select/Mute single record via index. Example: Select/Mute account, product, order, invoice, etc. Transactions are always completed with a *COM-MIT WRITE WAIT IMMEDIATE.*	Transaction throughput Transaction response time	[tps] [ms]
TP-MEDIUM	Select/Mut Ø 25 records via index. Example: Select/Mutate account entries last week; item list of an order etc. Transactions are always completed with a *COM-MIT WRITE WAIT IMMEDIATE.*	Transaction throughput Transaction response time	[tps] [ms]
TP-HEAVY	Select/Mut Ø 125 records via index. Transactions are always completed with a *COM-MIT WRITE WAIT IMMEDIATE.* Example: select/mute mobile phone call records from last month, etc.	Transaction throughput Transaction response time	[tps] [ms]
TP-MIXED1	Balanced read-intensive workload mix: ▪ TP-REPORT ▪ TP-LOOKUP ▪ TP-MEDIUM (40% UPDATE) ▪ DL-BUFFER (2 rpt)	Transaction throughput Transaction response time	[tps] [ms]
TP-MIXED2	Balanced write-intensive workload mix: ▪ TP-LIGHT (40% UPDATE) ▪ TP-MEDIUM (30% UPDATE) ▪ TP-HEAVY (20% UPDATE) ▪ DL-BUFFER (3 rpt).	Transaction throughput Transaction response time	[tps] [ms]

The benchmark reports for transaction processing workloads include both the primary performance metrics such as throughput and transaction response time as well as important influencing factors: Buffer Cache Hit Rate (*BuCache read [%]* column), Flash Cache Hit Rate - if any (*FlCache read [%] column*), Random Read Service Time (*IO time read [ms] column),* and Latency from Log Writer at Transaction Completion (*Log File syn [ms]* column).

Example: Workload TP-MIXED1

Read-intensive workload mix TP-MIXED1.

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Transactions total [tps]	Response time [ms]	IO time read [ms]	LogFile sync [ms]	REDO data [kBpt]	BuCache read [%]	FlCache read [%]	Elapsed time [s]
9	1	TP-MIXED1	2	8	4	2	1	96	0	13,979	0.570	0.240	0.228	2.152	70.11	100.00	181
	2	TP-MIXED1	2	24	14	9	3	86	0	310,522	0.077	0.271	0.412	0.200	91.53	100.00	183
	3	TP-MIXED1	2	40	19	13	4	81	0	685,432	0.058	0.298	0.542	0.127	93.91	100.00	182
	4	TP-MIXED1	2	56	26	17	5	74	0	943,014	0.059	0.341	0.721	0.105	94.38	100.00	182
	5	TP-MIXED1	2	72	30	21	6	70	0	1,179,671	0.061	0.378	0.877	0.094	94.61	100.00	182
	6	TP-MIXED1	2	88	35	24	7	65	0	1,403,307	0.062	0.412	1.001	0.087	94.91	100.00	182
	7	TP-MIXED1	2	104	39	27	7	61	0	1,601,678	0.065	0.446	1.113	0.080	95.14	100.00	182
	8	TP-MIXED1	2	120	43	31	8	57	0	1,791,183	0.067	0.485	1.264	0.076	95.33	100.00	182
	9	TP-MIXED1	2	136	47	34	9	53	0	1,919,583	0.070	0.518	1.387	0.073	95.36	100.00	182
	10	TP-MIXED1	2	152	51	37	9	49	0	2,013,175	0.075	0.546	1.502	0.072	95.32	100.00	182
	11	TP-MIXED1	2	168	55	40	10	45	0	2,108,923	0.079	0.573	1.616	0.072	95.32	100.00	182
	12	TP-MIXED1	2	184	58	43	10	42	0	2,171,213	0.084	0.605	1.739	0.072	95.28	100.00	182
	13	TP-MIXED1	2	200	61	45	10	39	0	2,262,549	0.087	0.624	1.829	0.071	95.27	100.00	183
	14	TP-MIXED1	2	208	63	47	11	37	0	2,281,498	0.090	0.641	1.881	0.071	95.28	100.00	184
	15	TP-MIXED1	2	216	64	48	11	36	0	2,335,079	0.091	0.653	1.940	0.070	95.30	100.00	185

Example: Workload TP-MIXED2

Schreib-intensiver Workload Mix TP-MIXED2.

Run	Test	Workload	Upd [%]	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Transactions total [tps]	Response time [ms]	IO time read [ms]	LogFile sync [ms]	REDO data [kBpt]	BuCache read [%]	FlCache read [%]	Elapsed time [s]
18	2	TP-MIXED2	N/A	2	8	6	3	2	94	0	10,444	0.762	0.250	0.267	2.827	65.73	100.00	183
	3	TP-MIXED2	N/A	2	16	10	5	3	90	0	17,037	0.933	0.274	0.387	2.917	63.63	100.00	182
	4	TP-MIXED2	N/A	2	24	13	7	4	87	0	23,605	1.011	0.293	0.473	2.749	62.96	100.00	182
	5	TP-MIXED2	N/A	2	32	16	9	4	84	0	26,667	1.192	0.327	0.594	2.932	62.20	100.00	182
	6	TP-MIXED2	N/A	2	40	18	10	5	82	0	30,199	1.318	0.353	0.699	2.893	60.82	100.00	182
	7	TP-MIXED2	N/A	2	48	20	11	6	80	0	32,956	1.450	0.382	0.810	2.895	59.73	100.00	182
	8	TP-MIXED2	N/A	2	56	21	12	6	79	0	35,255	1.578	0.411	0.922	2.912	59.27	100.00	182
	9	TP-MIXED2	N/A	2	64	23	13	7	77	0	37,798	1.685	0.437	1.021	2.880	58.63	100.00	182
	10	TP-MIXED2	N/A	2	72	24	13	7	76	0	39,469	1.814	0.468	1.145	2.923	58.60	100.00	182
	11	TP-MIXED2	N/A	2	80	26	14	7	74	0	41,402	1.922	0.496	1.223	2.942	58.27	100.00	182
	12	TP-MIXED2	N/A	2	88	27	15	8	73	0	43,347	2.017	0.521	1.304	2.924	57.94	100.00	182
	13	TP-MIXED2	N/A	2	96	28	16	8	72	0	45,473	2.098	0.550	1.414	2.908	58.41	100.00	182
	14	TP-MIXED2	N/A	2	104	29	16	8	71	0	48,129	2.148	0.572	1.479	2.849	58.52	100.00	182
	15	TP-MIXED2	N/A	2	112	30	17	8	70	0	49,139	2.266	0.605	1.605	2.865	58.46	100.00	182
	16	TP-MIXED2	N/A	2	120	31	18	9	69	0	49,413	2.413	0.636	1.693	2.894	57.91	100.00	182
	17	TP-MIXED2	N/A	2	128	32	18	9	68	0	51,392	2.478	0.660	1.770	2.854	57.81	100.00	182

An additional report provides supplementary information on I/O behavior.

Run	Test	Workload	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	CPU iow [%]	Transactions total [tps]	Logical reads [dbpt]	Logical writes [dbpt]	Physical reads [dbpt]	Physical writes [dbpt]	Phys reads total [IOPS]	Phys reads total [MBps]	Phys writes total [IOPS]	Phys writes total [MBps]	Elapsed time [s]
18	2	TP-MIXED2	2	8	6	3	2	94	0	10,444	21.204	11.376	7.267	2.561	75,680	616	30,931	242	183
	3	TP-MIXED2	2	16	10	5	3	90	0	17,037	23.482	11.774	8.541	2.918	145,085	1,181	53,268	443	182
	4	TP-MIXED2	2	24	13	7	4	87	0	23,605	22.884	11.230	8.476	2.874	199,417	1,627	69,387	600	182
	5	TP-MIXED2	2	32	16	9	4	84	0	26,667	25.522	12.018	9.648	3.060	256,338	2,080	80,782	722	182
	6	TP-MIXED2	2	40	18	10	5	82	0	30,199	25.502	11.891	9.993	3.049	300,517	2,450	89,758	813	182
	7	TP-MIXED2	2	48	20	11	6	80	0	32,956	26.087	11.935	10.505	3.026	344,581	2,804	96,018	882	182
	8	TP-MIXED2	2	56	21	12	6	79	0	35,255	26.570	12.018	10.823	3.001	379,739	3,090	101,000	937	182
	9	TP-MIXED2	2	64	23	13	7	77	0	37,798	26.592	11.902	11.001	2.925	413,558	3,361	104,637	981	182
	10	TP-MIXED2	2	72	24	13	7	76	0	39,469	27.422	12.118	11.352	2.926	445,688	3,617	108,548	1,026	182
	11	TP-MIXED2	2	80	26	14	7	74	0	41,402	27.513	12.188	11.482	2.872	472,848	3,831	111,366	1,060	182
	12	TP-MIXED2	2	88	27	15	8	73	0	43,347	27.535	12.134	11.582	2.783	499,383	4,047	112,232	1,078	182
	13	TP-MIXED2	2	96	28	16	8	72	0	45,473	27.787	12.121	11.557	2.735	522,949	4,238	115,367	1,113	182
	14	TP-MIXED2	2	104	29	16	8	71	0	48,129	27.249	11.882	11.304	2.625	541,540	4,378	116,818	1,134	182
	15	TP-MIXED2	2	112	30	17	8	70	0	49,139	27.756	11.967	11.531	2.630	564,340	4,559	119,305	1,159	182
	16	TP-MIXED2	2	120	31	18	9	69	0	49,413	28.158	12.108	11.852	2.658	583,253	4,710	120,994	1,178	182
	17	TP-MIXED2	2	128	32	18	9	68	0	51,392	27.674	11.934	11.676	2.606	597,772	4,831	123,293	1,202	182

Performance of PL/SQL application code

PL/SQL is the preferred programming language for complex transaction logic and algorithms. PL/SQL programs are stored and compiled in the database server. Some large applications, such as *Core Banking Systems*, are largely implemented in PL/SQL.

Key figures for PL/SQL performance

For PL/SQL workloads, the following performance metrics are of interest:

- Number of executed PL/SQL operations per time unit in *million operations per second* [Mops].
- Execution time of PL/SQL algorithms in *seconds* [s].

Workloads to determine PL/SQL performance

peakmarks® offer workloads to test the efficiency of PL/SQL programs on a specific processor. These workloads are completely limited by CPU performance.

Workload	Action	Performance indicator	Unit of measure
PLS-ADD	Addition of numbers	Throughput PL/SQL operations	[Mops]
PLS-BUILTIN	Data type specific operations including SQL functions, based on Core Banking Systems and Telco Billing Systems	Throughput PL/SQL operations	[Mops]
PLS-PRIME	Calculation of the first n prime numbers	Processing time of an algorithm	[s]
PLS-FIBO	Calculation of the first n Fibonacci numbers (recursive algorithm)	Processing time of an algorithm	[s]
PLS-MIXED	Complex workload with a mix of equally weighted simple workloads running concurrently: ▪ PLS_ADD with data type NUMBER ▪ PLS_ADD with data type PLS_INTEGER ▪ PLS_BUILTIN with data type NUMBER ▪ PLS_BUILTIN with data type VARCHAR2	Throughput PL/SQL operations	[Mops]

Example: Workload PLS-MIXED

The PLS-MIXED workload shows the throughput of different PL/SQL operations with different data types.

Run	Test	Workload	Parameter	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Operations total [Mops]	Operations per cpu [Mops]	Elapsed time [s]
20	1	PLS-MIXED	N/A	2	8	5	4	0	95	737.47	92.18	182
	2	PLS-MIXED	N/A	2	48	25	25	0	75	4,004.28	83.42	182
	3	PLS-MIXED	N/A	2	96	50	49	0	50	6,447.20	67.16	182
	4	PLS-MIXED	N/A	2	144	74	73	0	26	6,937.48	48.18	183
	5	PLS-MIXED	N/A	2	192	95	94	0	5	6,572.97	34.23	183

Example: Workload PLS-FIBO

The PLS-FIBO workload displays the processing time for an algorithm that calculates Fibonacci numbers (column *N*). The data type used can be configured (*Data type* column).

Run	Test	Workload	Parameter	Nodes	Jobs	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU idle [%]	Operations total [Mops]	Operations per cpu [Mops]	Elapsed time [s]
19	31	PLS-FIBO	42	2	1	1	1	0	99	0.00	0.00	69
	32	PLS-PRIME	8000	2	1	1	1	0	99	0.00	0.00	79

Conclusion

peakmarks® Benchmark Software provides system engineers and system architects with a robust and comprehensive benchmark framework for determining meaningful and understandable performance metrics of Oracle Database Services - **on-premise and** in the **cloud**.

The easy-to-understand performance metrics of peakmarks® make technologies, configurations, components and complete systems comparable across manufacturers and vendors.

The peakmarks® Benchmark Software covers all conceivable workloads for a performance test of the infrastructure. The software is constantly being further developed and adapted to new database versions.

Benefits for IT operations. Many IT organizations can reduce the cost of database services by choosing *best-in-class* infrastructure components or *best-in-class* cloud services and minimizing licensing and maintenance costs.

Bottlenecks, misconfigurations and malfunctions are detected before going live. Regular performance analyses of cloud services ensure performance quality. peakmarks® performance metrics are a solid basis for capacity planning when migrating to other platforms or cloud services.

The result is Oracle Database Services with *predictable* and *persistent* performance in all load situations.

Benefits for IT hardware vendors and service providers. peakmarks® benchmarks are fast and efficient and support the process of developing and marketing license-optimized Oracle database solutions with the best price/performance ratio.

Customers and prospects benefit from understandable, user-friendly and comparable performance metrics that can reduce time-consuming and expensive proofs-of-concepts.

peakmarks® performance indicators help to strengthen the positioning of a provider compared to the competition.

Appendix: Performance Certification for Reference System

Oracle Database Service

Platform	Database Server, each	Storage Server High Capacity, each	Database	peakmarks® configuration
▪ Exadata Flex ▪ 2 database servers Exadata X7 ▪ 3 storage server Exadata X8 high capacity	▪ Intel Xeon 8160 ▪ 2 sockets, 2.1 - 3.7 GHz ▪ 48 cores ▪ 96 threads ▪ 768 GByte DDR4-2666, PCI Gen 3	▪ 12 x 14 = 168 TByte HDD (raw) ▪ 4 x 6.4 = 25.6 TByte flash cache ▪ ASM high redundancy	▪ Oracle 19.20 EE ▪ db block size 8 kByte ▪ RAC ▪ data guard off ▪ archiving off	▪ Version 10.1 ▪ Build 230801 ▪ database size 2 x 2 TByte

peakmarks® Key Performance Metrics for platform components in Oracle database operations

Category	Key Performance Metric	Response time		Throughput		peakmarks® Workload
Server System	Mix of queries and scans	0.029 - 0.070	ms	2,736,163	qps	SRV-MIXED
	Buffer cache logical read throughput	-		66,502,195	dbps	SRV REPORT
	Buffer cache scan rate	-		258,579	MBps	SRV-SCAN

Category	Key Performance Metric	Throughput		peakmarks® Workload
Storage System	SQL sequential I/O	12,132	MBps	STO-READ
	SQL sequential I/O - using smart scan	72,394	MBps	STO-OFFLOAD
	SQL random read, 100% read, < 500 µs	745,340	IOPS	STO-RANDOM
	SQL random read, 80% read, < 500 µs	423,127	IOPS	STO-RANDOM
	SQL random write	360,793	dbps	STO-SCATTER

peakmarks® Key Performance Metrics for critical database background processes

Category	Key Performance Metric	Latency		Throughput		peakmarks® Workload
Log Writer (LGWR)	Small-sized transactions - commit throughput and latency	0.323 - 1.443	ms	209,735	tps	LGWR-LAT1
	Throughput of REDO data			1,443	MBps	LGWR-THR
Database Writer (DBWR)	Buffer cache writes			328,586	dbps	DBWR-THR

peakmarks® Key Performance Metrics for representative database operations

Category	Key Performance Metric	Throughput	peakmarks® Workload
Data Load	Transactional data load - using the buffer cache (5 rpt)	136 MBps 476,194 rps	DL-BUFFER
	Data Warehouse data load - bypassing the buffer cache	1,981 MBps 6,923,907 rps	DL-DIRECT

Category	Key Performance Metric	Throughput	peakmarks® Workload
Data Analytics	Data scan - using storage system	12,519 MBps 36,773,204 rps	DA-STORAGE
	Data scan - using smart scan	68,539 MBps 201,183,310 rps	DA-OFFLOAD
	Data scan - using row store	255,266 MBps 845,157,690 rps	DA-ROWSTORE
	Data scan - using column store	39,916,054 MBps 131,717,969,993 rps	DA-COLSTORE

Category	Key Performance Metric	Response time	Throughput	peakmarks® Workload
Transaction Processing	Read-intensive transaction mix	0.058 - 0.091 ms	2,335,079 tps	TP-MIXED1
	Write-intensive transaction mix	0.762 - 2.478 ms	51,392 tps	TP-MIXED2

peakmarks® Key Performance Metrics for PL/SQL application code

Category	Key Performance Metric	Throughput	Execution Time	peakmarks® Workload
PL/SQL	PL/SQL operations throughput	6,937 Mops	-	PLS-MIXED
	PL/SQL algorithm execution time Fibonacci NUMBER for n = 42		69 s	PLS-FIBO