

White Paper

FUJITSU Server PRIMERGY & PRIMEQUEST

Benchmark Overview OLTP-2

This document presents the internal Fujitsu database benchmark OLTP-2. Particular emphasis is placed on the common ground with and differences to the TPC-E benchmark and the resulting implications with regard to the usability of OLTP-2 benchmark results.

A comparison of OLTP-2 with TPC-E results is neither permitted nor practical. Both benchmarks are in fact based on a similar operator scenario. However, whereas TPC-E is suited for a comparison of the system configurations of various manufacturers, the focus of the OLTP-2 benchmark is placed on statements as to the scaling and relative performance comparisons within the PRIMERGY and PRIMEQUEST server family.

Version

1.2
2015-02-04



Contents

Document history.....	2
The OLTP-2 Benchmark – An Overview	3
The benchmark model.....	3
Performance metrics.....	3
The OLTP-2 database and its transactions	4
Benchmark environment.....	5
Differences between OLTP, OLTP-2 and TPC-E	6
Concluding remark.....	7
Literature.....	8
Contact	8

Document history

Version 1.0 (August 2007)

- First report version

Version 1.1 (June 2009)

- New layout

Version 1.2 (February 2015)

- Minor corrections
- New layout

The OLTP-2 Benchmark – An Overview

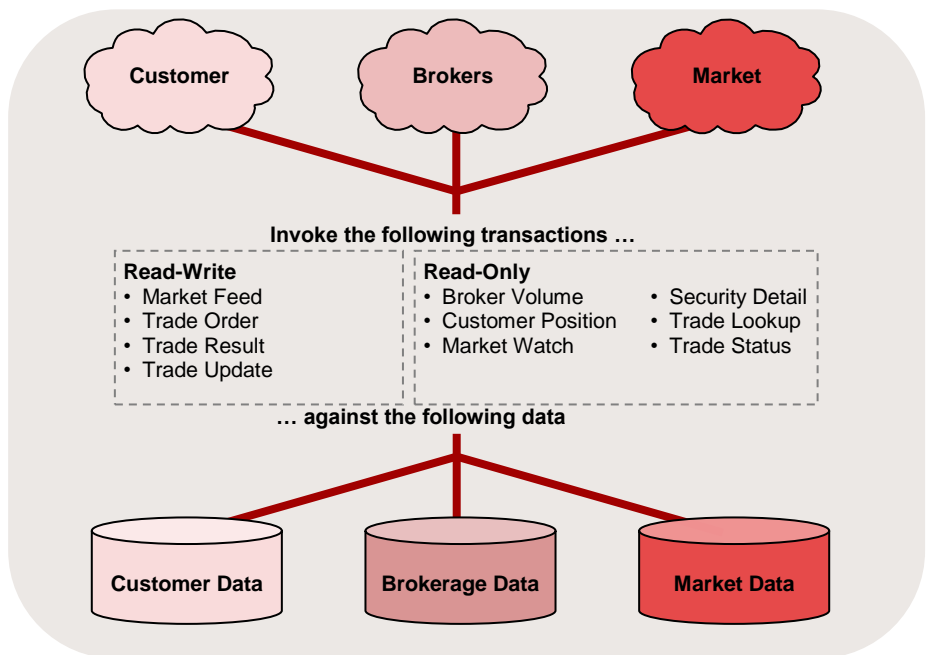
Similar to TPC-E, the OLTP-2 benchmark measures the performance of online transaction processing systems or OLTP. Since hardware and time involved with TPC-E benchmarks merely enable the measurement of specially selected system configurations, Fujitsu developed a measuring method with which it is possible to perform a considerably larger number of measurements. Thus the OLTP-2 benchmark is based on the same complex database and the same transaction types as TPC-E. As with TPC-E, keys to performance are the database server with disk I/O and network communication. In contrast to TPC-E, however, the OLTP-2 benchmark is limited to pure performance measurements, which are implemented according to a less stringent code of procedure, less runtime and without any auditing by an independent body. Although it is as a result not possible to compare the results of OLTP-2 and TPC-E, the OLTP-2 benchmark helps to gain insights into the scaling characteristics of the PRIMERGY models in a complex database environment.

The OLTP-2 benchmark replaces the older OLTP benchmark, which is no longer usable with current system configurations.

The benchmark model

OLTP-2 models the activity of a brokerage firm that must manage customer accounts, execute customer trade orders and be responsible for the interactions of customers with financial markets. The customers generate transactions related to trades, account inquiries, and market research. The brokerage firm in turns interacts with financial markets to execute orders on behalf of the customers and updates relevant account information. The number of customers defined for the brokerage firm can be varied to represent the workloads of different size businesses.

This benchmark is composed of a set of transactions that are executed against three sets of database tables that represent market data, customer data, and broker data. An other set of tables contains generic dimension data such as addresses and zip codes. The diagram illustrates the key components of the environment:



The benchmark has been reduced to simplified form of this application environment. To measure the OLTP-2 performance of the system, a simple driver generates transactions and their inputs, submits them to the system under test (SUT) and measures the rate of

completed transactions being returned. To simplify the benchmark and focus on the core transactional performance, all application functions related to user-interface and display-functions have been excluded from the benchmark. The system under test is focused on portraying the components found on the sever side of a transaction monitor or application server.

Performance metrics

The performance metric is tps, which is the average number of Trade Result transactions executed within one second.

OLTP-2-metric
Throughput
tps

The OLTP-2 database and its transactions

The database is divided into the 4 groups of tables Customer, Broker, Market and Dimension. The database contains 33 variously structured tables and thus also different types of data records. The size and number of the data records vary depending on the table. A mix of 11 concurrent transactions of varying type and complexity is executed on the database. Due to their competing for the limited system resources many system components are stressed and data changes are executed in a variety of ways.

Structure of the OLTP-2 database	
Tables	Contents
9 Customer tables	Information about customers and brokerage firm
	ACCOUNT_PERMISSION CUSTOMER CUSTOMER_ACCOUNT CUSTOMER_TAXRATE HOLDING HOLDING_HISTORY HOLDING_SUMMARY WATCH_ITEM WATCH_LIST
9 Broker tables	Information about brokerage firm and broker related data
	BROKER CASH_TRANSACTION CHARGE COMMISSION_RATE SETTLEMENT TRADE TRADE_HISTORY TRADE_REQUEST TRADE_TYPE
11 Market tables	Information about companies, market exchange and industry sectors
	COMPANY COMPANY_COMPETITOR DAILY_MARKET EXCHANGE FINANCIAL INDUSTRY LAST_TRADE NEWS_ITEM NEWS_XREF SECTOR SECURITY
4 Dimension tables	Common information such as addresses and zip codes
	ADDRESS STATUS_TYPE TAXRATE ZIP_CODE

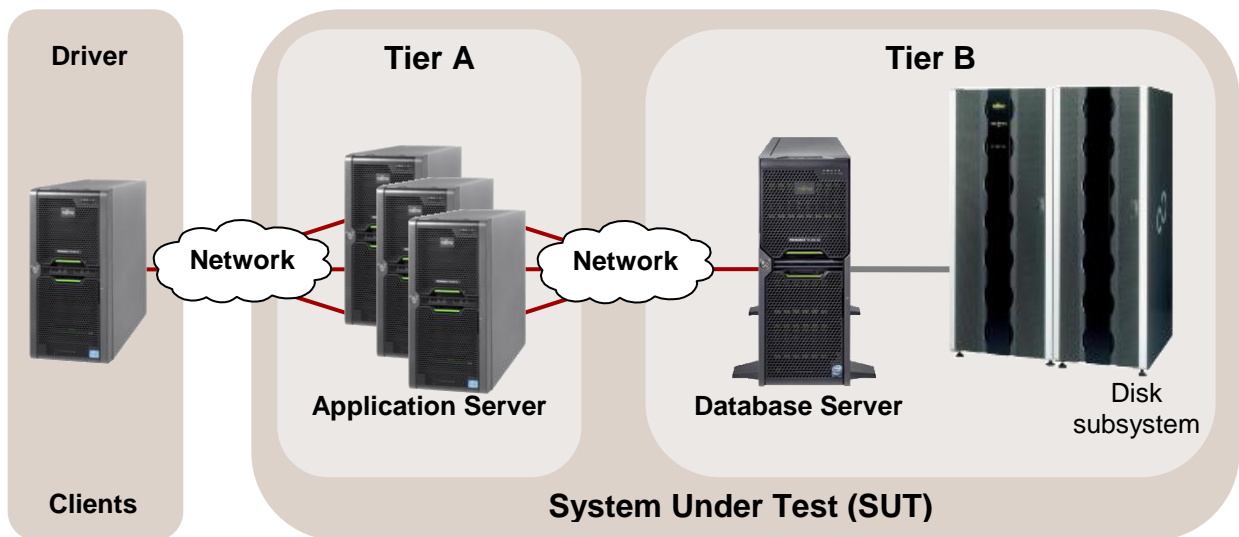
The OLTP-2 benchmark consists of 11 Transactions and one cleanup function. To generate a reasonably balanced workload that resembles real production environments, the Transactions have to cover a wide variety of system functions. Ten of the Transactions follow a specific mix to generate the desired workload while keeping the benchmark environment simple, repeatable and easy to execute. The eleventh Transaction is not part of the Transaction Mix, but is executed at fixed intervals. This Transaction, called "Data-Maintenance", simulates administrative updates to tables that are not otherwise modified by the Transactions in the mix. A cleanup Transaction, called "Trade-Cleanup", is provided to clean up pending and submitted trades that may exist from an earlier run.

One of the key performance characteristics of database systems is the ratio of reads and writes generated by the workload. To emulate such a ratio, OLTP-2 has defined Transactions with read-only profiles as well as Transactions with read-write profiles. In addition, the Transactions apply varying loads on the processor.

OLTP-2 transactions and required distribution		
Name of transaction	distribution of transaction	Access
Broker Volume	4.9 %	Read
Customer Position	13.0 %	Read
Market-Watch	18.0 %	Read
Security-Detail	14.0 %	Read
Trade-Status	19.0 %	Read
Trade-Lookup	8.0 %	Read
Market Feed	1.0%	Read/Write
Trade Order	10.1%	Read/Write
Trade Result	10.0%	Read/Write
Trade Update	2.0%	Read/Write
Data Maintenance	Once every 60 seconds	Read/Write
Trade Cleanup	Once per test run	

Benchmark environment

The following figure shows the physical components that could be assembled to implement a hypothetical test configuration. The Driver is defined to be all hardware and software needed to implement the driving & reporting and up-stream connector functional components. Tier A is defined to be all hardware and software needed to implement the down-stream connector, transaction implementation and database interface functional components. Tier A can be one or more separate servers as well as part of the database server. Tier B is defined to be all hardware and software needed to implement the database server functional component. This includes data storage media sufficient to satisfy the initial database population requirements and the business day growth requirements.



Differences between OLTP, OLTP-2 and TPC-E

The database benchmark OLTP is a Fujitsu internal benchmark, which is based on the database and transaction model of the TPC-C benchmark. Since 1992 TPC-C has been used to measure and publish performance of database server. Meanwhile system performance increased in a dramatic way and requires large investigations to provide the test environment. In addition the transaction model is not up-to-date. TPC-E is a new development that takes care of the current situation by using a more complex structure while hardware efforts are reduced (see [Benchmark Overview TPC-E](#)). This is the base of the Fujitsu internal OLTP-2 benchmark.

TPC-C used 5 transaction types and 9 tables. TPC-E uses 11 transaction types and 33 tables. The transactions are more CPU intensive and reduce the disk environment to 10-30% depending on server performance. The application functions related to User-Interface and Display-Functions have been excluded from the benchmark and replaced by a TPC provided driver code.

OLTP and OLTP-2 numbers cannot be compared or converted. The values are completely different because of the different databases, access profiles and performance metrics. Both benchmarks use different system resources, and scaling values of various configurations are not the same. A comparison can only be done within OLTP resp. OLTP-2 itself.

In contrast to TPC-E, the OLTP-2 benchmark is not a standard benchmark. Adherence to the code of procedures is not monitored by an independent body.

A TPC-E measurement can take up to 3 months from preparation through to the official acceptance by the TPC committee. Thus TPC-E is not suited to measure all the PRIMERGY and PRIMEQUEST systems appropriate for a database scenario in all their possible CPU variations and memory configurations. However, in order to have a yardstick similar to the TPC-E benchmark the OLTP-2 benchmark was developed, which with acceptable outlay and in a reasonable amount of time enables the measurement of a great many system configurations and thus the classification of their capacity.

One aim of OLTP-2 measurements is to gain insights into the scaling of the PRIMERGY and PRIMEQUEST systems within a system family. In this way, it is possible to quantify the increase in performance achieved in a system component by measuring various processor and memory configurations. OLTP-2 measurements also make it possible to compare the performance of successive system generations. Of course, a comparison can only refer to entire systems and not to individual components because both the hardware used (chip set, controller, etc.) and the software (OS, database, etc.) are continuously being developed. A system configuration measured with the OLTP-2 benchmark need not inevitably be orderable. Single-processor and flat rack systems in particular are not necessarily suitable as servers for large databases. When new PRIMERGY and PRIMEQUEST systems are developed, OLTP-2 measurements help indicate any faults in the system configuration. Findings as to whether certain components are not yet fully developed or cause bottlenecks help the development department make a new system ready for the market. Moreover, during the preparatory phase for the complex TPC-E measurements, OLTP-2 measurements are also a suitable method of optimizing a system and its environment.

The OLTP-2 benchmark does not provide any information about the costs of the test configuration. In contrast to TPC-E, a cost/benefit ratio (price per tpsE) is also not possible. In addition, aspects of fail-safety are - other than with TPC-E - not an integral part of the OLTP-2 benchmark because of the hardware and time outlay involved. The hardware outlay required here primarily influences the costs of the system. Inasmuch as the necessary fail-safe hardware components can be ordered for the respective system, similar fail-safety as to that of TPC-E is therefore in principle to be expected.

The measurement environment used with TPC-E takes care of the costs, because the price per tpsE is also published and depends on the performance of the server to be measured. The OLTP-2 measurement environment will not be change for different system configurations, like CPU scaling and is oversized in some cases. OLTP-2 is used to measure the server performance, which is not IO or network bound. This is the reason why often more disks are used than needed.

According to the code of procedure of the TPC consortium, a measurement that has not been audited in keeping with TPC-E rules may not be put in relation to a TPC-E result. Although the OLTP-2 benchmark uses the same application software and load profile measurement runtime is shorter and there is no audit. Hence the similar results, which, as experience have shown only slightly, differ from each other. TPC-E values must always be used in conjunction with price performance data (price per tpsE) and the availability date. The results of the OLTP-2 benchmarks do not include any price details. They are merely used to evaluate scaling within the PRIMERGY and PRIMEQUEST family.

Concluding remark

If you consider the great complexity of the database and the transaction mix, the measured values achieved thus come close to throughput values from real-life business processes. The unit of measurement used in the OLTP-2 benchmark is therefore not synthetic, but altogether representative of the real business world. However, it must be noted that this applies to many, but not to all database environments. To what extent customers can achieve typical OLTP-2 benchmark throughput values chiefly depends on how similar a customer's database and application in fact are in comparison with those of the OLTP-2 benchmark. Although OLTP-2 results provide an indication of the throughput values that can be achieved in customer environments. The simple extrapolation of these, however, is not recommended. System performance and thus also benchmark results very much depend on system load, application-specific requirements, system design and implementation. The OLTP-2 benchmark therefore cannot replace the benchmarking of a customer application.

TPC-E is an ideal benchmark for the comparison of system configurations of various manufacturers. However, due to the measuring complexity with TPC-E there is only a relatively small number of TPC-E results. The system configurations measured in this regard are generally high-end configurations which frequently exceed customer requirements by far. TPC-E therefore does not provide a good basis for valid statements as to the scaling of systems. This intrinsic shortcoming is overcome by the OLTP-2 benchmark with its very simplified measuring methodology. Thus OLTP-2 results enable both fine tuning at system component level and well-founded statements as to the gain in performance of new models of the PRIMERGY or PRIMEQUEST server family.

Literature

PRIMERGY & PRIMEQUEST Servers

<http://www.fujitsu.com/fts/products/computing/servers/>

PRIMERGY & PRIMEQUEST Performance


<http://www.fujitsu.com/fts/x86-server-benchmarks>


Benchmark descriptions

<http://www.fujitsu.com/fts/products/computing/servers/primergy/benchmarks/benchmark-descriptions.html>

This White Paper:

 <http://docs.ts.fujitsu.com/dl.aspx?id=e6f7a4c9-aff6-4598-b199-836053214d3f>

 <http://docs.ts.fujitsu.com/dl.aspx?id=743d7d46-56e8-41d2-9d50-9ab29ccf4d18>

 <http://docs.ts.fujitsu.com/dl.aspx?id=9775e8b9-d222-49db-98b1-4796fbc6d67a>

TPC-C

<http://www.tpc.org/tpcc>

Benchmark Overview TPC-C

<http://docs.ts.fujitsu.com/dl.aspx?id=8d633e1e-8b44-4c5a-800f-7a1319649081>

OLTP

Benchmark Overview OLTP

<http://docs.ts.fujitsu.com/dl.aspx?id=e58838c3-0b11-460d-8b8c-a68d774a473e>

TPC-E

<http://www.tpc.org/tpce>

Benchmark Overview TPC-E

<http://docs.ts.fujitsu.com/dl.aspx?id=da0ce7b7-3d80-48cd-9b3a-d12e0b40ed6d>

Contact

FUJITSU

Website: <http://www.fujitsu.com/>

PRIMERGY & PRIMEQUEST Product Marketing

<mailto:Primergy-PM@ts.fujitsu.com>

PRIMERGY Performance and Benchmarks

<mailto:primergy.benchmark@ts.fujitsu.com>