insight 2

White Paper • Benchmark Overview • Issue 2.1 October 2003

Benchmark Overview
OLTP

Abstract

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

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

White Paper

www.fujitsu-siemens.com

FUJITSU COMPUTERS SIEMENS

Contents

The OLTP benchmark – An overview

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

The benchmark model

The OLTP benchmark simulates an environment in which the operator performs various transactions against a database. The central elements of the benchmark are the typical transactions of a wholesale company concerning order entries (order acceptance, delivery, recording payments, checking the status of orders and monitoring stock levels). The simulated company operates out of a number of warehouses and their allocated districts. The OLTP benchmark is designed in such a way that the size of the company (i.e. the number of its warehouses) may vary. Set parameters on the other hand are the 100,000 items as well as ten sales districts per warehouse and 3,000 customers per district. Every operator can at any time implement one of five transactions on the company's goods ordering system. Both the transactions and their frequency are based on a realistic scenario.

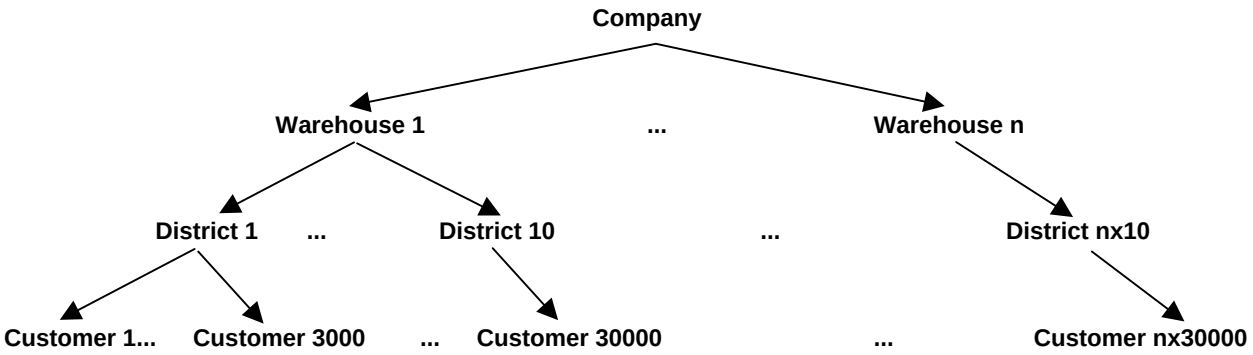

The most frequent transaction is the new order, which on average comprises 10 different items. Each warehouse attempts, if possible, to deal with the delivery from its own stock. Since this is hardly realistic with such a large number of items, delivery in virtually 10% of all cases is effected via the company's other warehouses. Another frequent transaction is the recording of a payment. Order status queries, the processing of delivery orders and checking of local stock levels for possible bottlenecks are less frequent. The entire business activity is modeled on the basis of these five transactions. In order to maintain the applicability of the OLTP benchmark to systems of differing capacity, TPC-C implementations must scale the size of the database proportionally to the computing power of the system to be measured.

The parameter used in the OLTP benchmark to measure performance specifies the number of new-order transactions processed per minute and it is expressed in tpm.

OLTP metrics
Throughput
tpm

The OLTP database and its transactions

The definitions of database and transactions are identical to those of TPC-C. There are differences in the generation of the load for the system to be measured.

The database consists of nine variously structured tables and thus also nine types of data records. The size and number of the data records vary depending on the table. A mix of five concurrent transactions of varying type and complexity is executed on the database - largely online or in part queued for deferred batch processing. 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 database	
Table	Number of entries
Warehouse	n (specified in a measurement)
Item	100,00
Stock	n x 100,00
District	n x 10
Customer	3,000 per district, 30,000 per warehouse
Order	number of customers (initial value)
New order	30% of the orders (initial value)
Order line	approx. 10 per order
History	number of customers (initial value)

The way in which data are entered by operators in the TLC-C benchmark is based on the most basic characteristics of real-life data-input situations. For example, it is possible for invalid item numbers to be entered, which then results in the cancellation of the transaction. In order to model as realistic a scenario as possible, the artificial simplifications used in many other benchmarks were largely omitted. Thus, for example, it must always be possible for all terminal input to also be entered by real-life operators. To this end, all entry screens must include specified field definitions as well as labeled input and output fields and also have the common cursor motion and field correction mechanisms.

The throughput of TPC-C is a direct result of the level of activity at the terminals. Each warehouse has ten terminals and all five transactions can be executed at each terminal. Remote terminal emulators (RTE) are used to maintain the required mix of transactions over the performance measurement period for TPC-C measurements. With the OLTP benchmark the load is on the other hand generated in batch mode on the clients. The transaction mix represents the complete business processing of an order from its entry through to delivery. More specifically, the required mix is defined to produce an equal number of new-order and payment transactions and to produce one delivery transaction, one order-status transaction, and one stock-level transaction for every ten new-order transactions.

OLTP transactions and required distribution	
Name of transaction	Share of all transactions
New order	≤ 45%
Payment	≥ 43%
Order status	≥ 4%
Delivery	≥ 4% (batch transaction)
Stock level	≥ 4%

The RTE or the analogue load generation program with the OLTP benchmark is also used to measure the transaction response times. However, with the OLTP benchmark there is no simulation of the keying and think times of the operator.

All possible performance-relevant database design techniques, e.g. partitioning or replication, are permitted. Unrealistically large performance advantages are not to be expected because of the way in which the use of data records by the transactions is defined.

Differences between OLTP and TPC-C

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

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

One aim of OLTP measurements is to gain insights into the scaling of the PRIMERGY 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 measurements also make it possible to compare the performance of successive PRIMERGY 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 benchmark need not inevitably be orderable. Dual-processor and flat rack systems in particular are not necessarily suitable as servers for large databases. When new PRIMERGY systems are developed, OLTP 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-C measurements, OLTP measurements are also a suitable method of optimizing a system and its environment.

The OLTP benchmark does not provide any information about the costs of the test configuration. In contrast to TPC-C, a cost/benefit ratio (price per tpmC) is also not possible. In addition, aspects of fail-safety are - other than with TPC-C - not an integral part of the OLTP 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 PRIMERGY system, similar fail-safety as to that of TPC-C is therefore in principle to be expected.

The measuring environment for TPC-C is a 3-tier environment consisting of server, clients and RTEs for the simulation of the operator load and network topology. With OLTP measurements it is a 2-tier environment in which load generation is taken over from the clients by the RTEs in batch mode. Thus there is no terminal emulation here.

According to the code of procedure of the TPC consortium, a measurement that has not been audited in keeping with TPC-C rules may not be put in relation to a TPC-C result. Although both benchmarks provide a number of transactions per minute and the same application software is typically used for the measurements, both benchmarks use different methods to simulate the operator load. Hence the similar results, which as experience has shown only slightly differ from each other. TPC-C values must always be used in conjunction with price performance data (price per tpmC). The results of the OLTP benchmarks do not include any price details. They are merely used to evaluate scaling within the PRIMERGY 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 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 OLTP environments. To what extent customers can achieve typical OLTP benchmark throughput values chiefly depends on how similar a customer's database and application in fact are in comparison with those of the OLTP benchmark. Although OLTP 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 benchmark therefore cannot replace the benchmarking of a customer application.

TPC-C is an ideal benchmark for the comparison of OLTP system configurations of various manufacturers. However, due to the measuring complexity with TPC-C there is a only a relatively small number of TPC-C results. The system configurations measured in this regard are generally high-end configurations which frequently exceed customer requirements by far. TPC-C therefore does not provide a good basis for valid statements as to the scaling of systems. This intrinsic shortcoming is overcome by the OLTP benchmark with its very simplified measuring methodology. Thus OLTP 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 server family.

Literature

TPC-C	http://www.tpc.org/tpcc
	http://vilpublic.fujitsu-siemens.com/vil/pc/vil/primergy/performance/Benchmark_Overview_TPC-C.pdf
PRIMERGY Systems	http://www.fujitsu-siemens.com/primergy
PRIMERGY Performance	http://www.fujitsu-siemens.com/rl/products/primergy/benchmarks.html

Contact

PRIMERGY Hardware
PRIMERGY Product Marketing
mailto:Primergy-PM@fujitsu-siemens.com

PRIMERGY Performance and Benchmarks
PRIMERGY Performance and Benchmarks
mailto:primergy.benchmark@fujitsu-siemens.com

Published by department:

Enterprise Products
PRIMERGY Server
PRIMERGY Performance Lab
mailto:primergy.benchmark@fujitsu-siemens.com

Extranet:
http://extranet.fujitsu-siemens.com/primergy
Internet:
http://www.fujitsu-siemens.com/primergy