THE POSSIBILITIES ARE INFINITE FUJITSU

WHITE PAPER

Benchmark Overview TPC-C

Version 1.2
October 2003

Pages 5

Abstract

This document presents the database benchmark TPC-C which was developed and published by the Transaction Processing Performance Council (TPC).

This benchmark, in which complete OLTP system configurations are measured, is a benchmark that is generally accepted by the industry. As a result of the benchmark concept realized in TPC-C it is possible to make across-the-board manufacturer comparisons. The auditing of the results by the TPC as an independent body ensures both objectivity and reproducibility.

Contents

The TPC-C benchmark – An overview

The TPC-C benchmark measures the performance of **o**n**l**ine **t**ransaction **p**rocessing systems or OLTP and is based on a complex database and a number of different transaction types that are executed on it. TPC-C is not only a hardware-independent but also a software-independent benchmark and can thus be run on every test platform, i.e. proprietary or open. In addition to the results of the measurement, all the details of the systems measured and the measuring method must also be explained in a measurement report (**F**ull **D**isclosure **R**eport or FDR). Consequently, this ensures that the measurement meets all benchmark requirements and is reproducible. TPC-C does not just measure an individual server, but a rather extensive system configuration. Keys to performance in this respect are the user interface, network communication and disk I/O as well as the critical factors of backup and recovery.

The benchmark model

TPC-C 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. TPC-C 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.

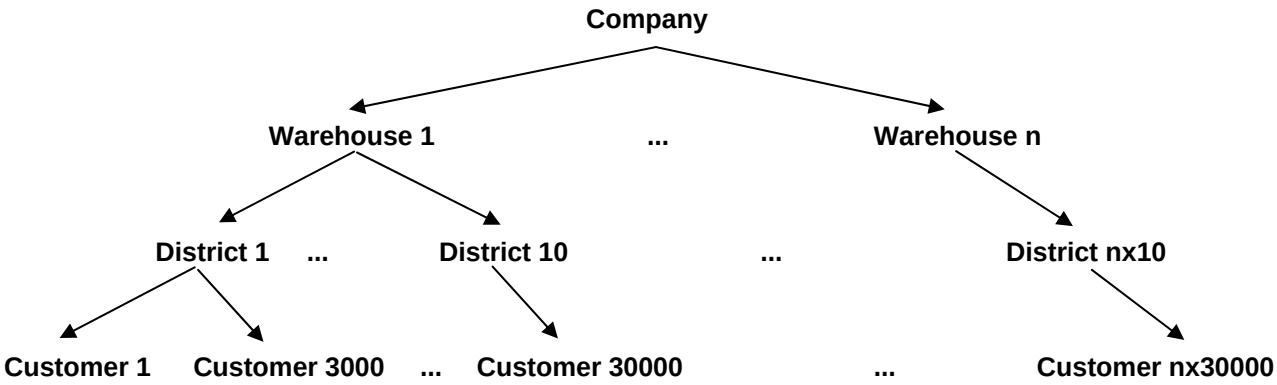

The most frequent transaction is the new order, which on average comprises 10 different items. Each warehouse attempts, if possible, to deal with the delivery from its own stock. Since this is hardly realistic with such a large number of items, delivery in virtually 10% of all cases is effected via the 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 TPC-C's applicability to systems of differing capacity, TPC-C implementations must scale both the number of terminals and the size of the database proportionally to the computing power of the system to be measured. In addition to performance, the definitive data security aspects of a fully production-ready system are also tested. The system features required in this regard are denoted as "ACID" properties (**A**tomicity, **C**onsistency, **I**solation and **D**urability). In order to furnish proof of the accuracy of the benchmark results it is necessary for the measurement configuration and procedure to be documented in its entirety in a full disclosure report (FDR) so that reproducibility by third parties is ensured.

Performance metrics

The unit used in TPC-C to measure performance specifies the number of processed new-order transactions per minute and is expressed in tpmC. The measured performance must always be reported together with the cost of ownership and the report date. The total system configuration costs are the actual retail prices that are definitive for the customer. They include all hardware and software components, maintenance expenditure over a period of three years and the storage capacities required maintaining the documented

throughput for a period of sixty 8-hour days. The cost of ownership, stated in currency / tpmC, follows from the total costs in relation to throughput.

TPC-C metrics		
Throughput	Cost of ownership	Report date
tpmC	$/tpmC	-

The TPC-C database and its transactions

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 TPC-C database	
Table	Number of entries
Warehouse	n (specified in a measurement)
Item	100,000
Stock	n x 100,000
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 TPC-C 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. A remote terminal emulator (RTE) is used to simulate the terminal activities and to maintain the required mix of transactions over the performance measurement period. 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.

TPC-C 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 is also used to measure transaction response time and to simulate keying and think times of the operator. Keying time is the time required to enter the data at the terminal. Think time is defined as the time an operator needs to read and evaluate the results of a transaction. Each transaction has a minimum keying time and a minimum think time. In addition, the response times of each transaction must be below a defined

threshold. For 90% of transactions this threshold is less than 5 seconds and less than 20 seconds for the stock-level transaction.

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

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 units of measurement used in TPC-C are 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 TPC-C throughput values chiefly depends on how similar a customer's database and application in fact are in comparison with those of TPC-C. Although TPC-C results provide an indication of the throughput values that can be achieved in customer environments, the simple extrapolation of these is not recommended. System performance and thus also benchmark results very much depend on system load, application-specific requirements, system design and implementation. TPC-C therefore cannot replace the benchmarking of a customer application.

Users of benchmark information and results, whether they are members of the press, market researchers, or commercial users, have to be able to rely on the validity of the benchmark results. To meet this demand, TPC-C was designed in such a way that system configurations with all the necessary production-oriented features, including backup and recovery features are measured and fully documented in the FDR (full disclosure report). These reports are reviewed by TPC itself as an independent body and are made available to the public. All these requirements ensure valid, objective measures of performance.

Literature

PRIMERGY Systems

http://ts.fujitsu.com/primergy

PRIMERGY Performance

http://ts.fujitsu.com/products/standard_servers/primergy_bov.html

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

OLTP-2

Benchmark Overview OLTP-2
http://docs.ts.fujitsu.com/dl.aspx?id=e6f7a4c9-aff6-4598-b199-836053214d3f

Contact

PRIMERGY Hardware

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

PRIMERGY Performance and Benchmarks

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

Published by department:

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

Internet:
http://ts.fujitsu.com/primergy

Extranet:
http://partners.ts.fujitsu.com/com/products/servers/primergy