White Paper
FUJITSU Server PRIMERGY & PRIMEQUEST Benchmark Overview TPC-E

This document presents the database benchmark TPC-E which was published by the Transaction Processing Performance Council (TPC). This is a new development compared with the older TPC-C benchmark. TPC-E has an actual transaction model, in which complete OLTP (Online Transaction Processing) system configurations are measured.

In 2009 TPC augmented the existing TPC benchmarks with “TPC-Energy”, a specification of energy metrics. The reporting of the energy metrics is optional.

As a result of the benchmark concept realized by TPC, 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.

Version
2.1
2015-01-29

Contents

Document history

Version 1.0

- First report version

Version 1.1

- Minor changes

Version 2.0

- TPC-Energy integrated and document structure adapted
- Minor corrections

Version 2.1

- New layout

TPC-E Overview

The database benchmark TPC-E was developed and published in 2007 by the Transaction Processing Performance Council (TPC). The TPC-E 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-E 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-E does not just measure an individual server, but a rather extensive system configuration. Keys to performance in this respect are the database server, disk I/O and network communication.

In 2009 therefore TPC developed a specification “TPC-Energy” that contains rules and methodology for measuring and reporting energy metrics in TPC benchmarks, see chapter “TPC-Energy”.

Benchmark model

TPC-E 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 turn 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. Another set of tables contains generic dimension data such as addresses and zip codes. The diagram opposite illustrates the key components of the environment.

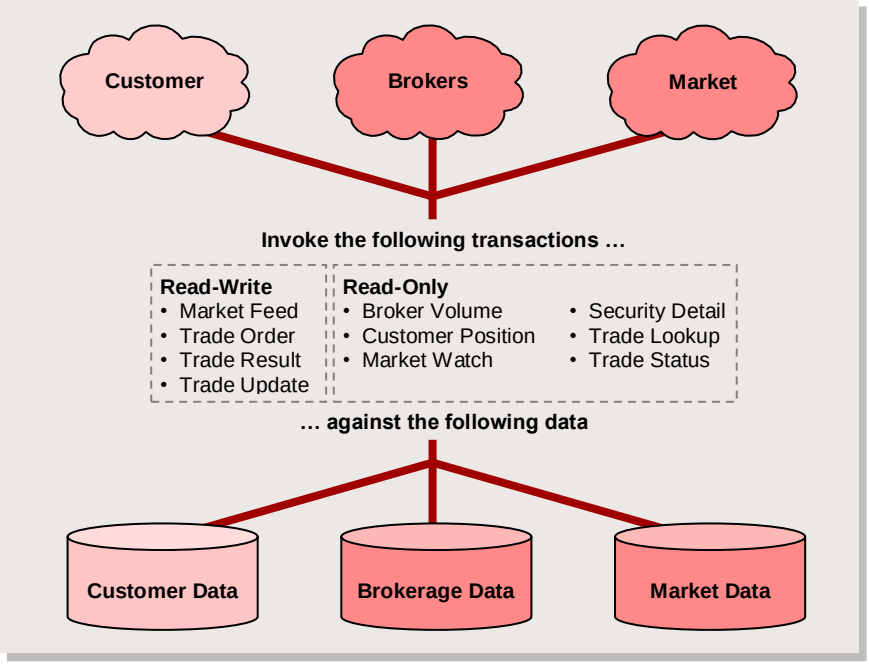

The benchmark has been reduced to a simplified form of this application environment. To measure the performance of the OLTP system, a simple driver generates transactions and their inputs, submits them to the System under Test (SUT) and measures the rate of completed transactions being returned. To simplify 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 SUT is focused on portraying the components found on the server side.

TPC-E database and transactions

The database is divided into the four 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 eleven 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.

Structure of the TPC-E 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 TPC-E benchmark consists of eleven 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 and reproducible mix to generate the desired workload. 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, TPC-E 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.

TPC-E 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	

Performance metrics

The performance metrics is tpsE. tps means transactions per second. tpsE is the average number of Trade Result transactions executed within one second. To be compliant with the TPC-E standard, all references to tpsE results must include the tpsE rate, the associated price-per-tpsE, and the availability date of the priced configuration.

TPC-E-metrics		
Throughput	Cost of Ownership	Availability Date
tpsE	$/tpsE	date

TPC-Energy

Energy efficiency has become besides performance and price another significant factor in evaluating computing hardware. In 2009 therefore TPC developed a specification “TPC-Energy” that contains rules and methodology for measuring and reporting energy metrics in TPC benchmarks. The TPC-Energy specification applies to all published TPC benchmarks, at present TPC-C, TPC-E and TPC-H. TPC-Energy is not a benchmark of its own but can only be performed in connection with other TPC benchmarks. The measuring and publishing of the TPC-Energy metrics in the TPC Benchmarks are optional and not required to publish a TPC result. If published, TPC-Energy measurements and results must be audited like other TPC benchmarks.

Energy metrics

The primary metric reported with TPC-Energy is in the form of “Watts per performance”, where the performance units are particular to each TPC benchmark. For TPC-E the primary metric is “Watts/tpsE”.

More detailed: the metric is calculated as the ratio of the energy consumed by all components of the benchmark system (typically measured in watts-seconds) to the total work completed (typically measured as a number of transactions). The benchmark system (system under test) includes servers, storage systems, and also network components like switches.

Besides the primary metric, the power consumption in “watts” of the configuration in an idle, but ready to work state must be measured and stated. This is useful to customers who have systems that have idle periods, but require the system to respond to a request for work at any time.

Furthermore, the ambient temperature has to be recorded during energy measurement, because air-cooled components have lower power consumption in an intensively cooled environment. To achieve a valid result the temperature must not drop below the limit of 20°C.

To provide more detailed comparisons and analyses of the results optional information about single subsystems can be listed. These partial results are denoted in TPC-Energy as secondary metrics. The benchmark system is partitioned in four subsystems: Database Server, Application Server, Storage and Miscellaneous, e.g. monitors or switches. The energy consumption for each subsystem is divided by the total work completed and showed as “Watts/tpsE”. The sum of all secondary metrics is identical to the reported primary metric. In addition the idle power consumption of each subsystem is reported.

Implementation

To facilitate the implementation of TPC-Energy measurements TPC provides the software package EMS (Energy Measurement System), which includes the Power Temperature Daemon (PTDaemon) licensed from the Standard Performance Evaluation Corporation (SPEC). The main features of EMS are interface to power instrumentation as power analyzers and temperature probes, logging of power and temperature readings and report generation. For a TPC-Energy measurement the use of EMS package is mandatory.

The energy data may be gathered from power devices. The power meters used in the measurements need to satisfy the accuracy defined in the specification. All power meters must have been calibrated by an authorized specialist within the past year.

The performance and energy consumption is measured during the measurement interval. For TPC-E it must begin after the system reaches steady state and last at least for 120 minutes. The idle power measurement must be started within 30 minutes of completion of the benchmark run for a duration of at least 10 minutes.

Benchmark environment

The following figure shows the physical components of a possible configuration.

The driver is defined as all the hardware and software needed to implement the driving & reporting and upstream connector functional components.

Tier A is defined as all the hardware and software needed to implement the downstream connector, transaction implementation and database interface functional components. Tier A can be one or more separate servers.

Tier B is defined as all the hardware and software needed to implement the database server functional components. This includes data storage media sufficient to satisfy the initial database population requirements and the growth requirements.

Together, tier A and tier B constitute the system under test (SUT), which is viewed and measured as a whole in the TPC-E benchmark.

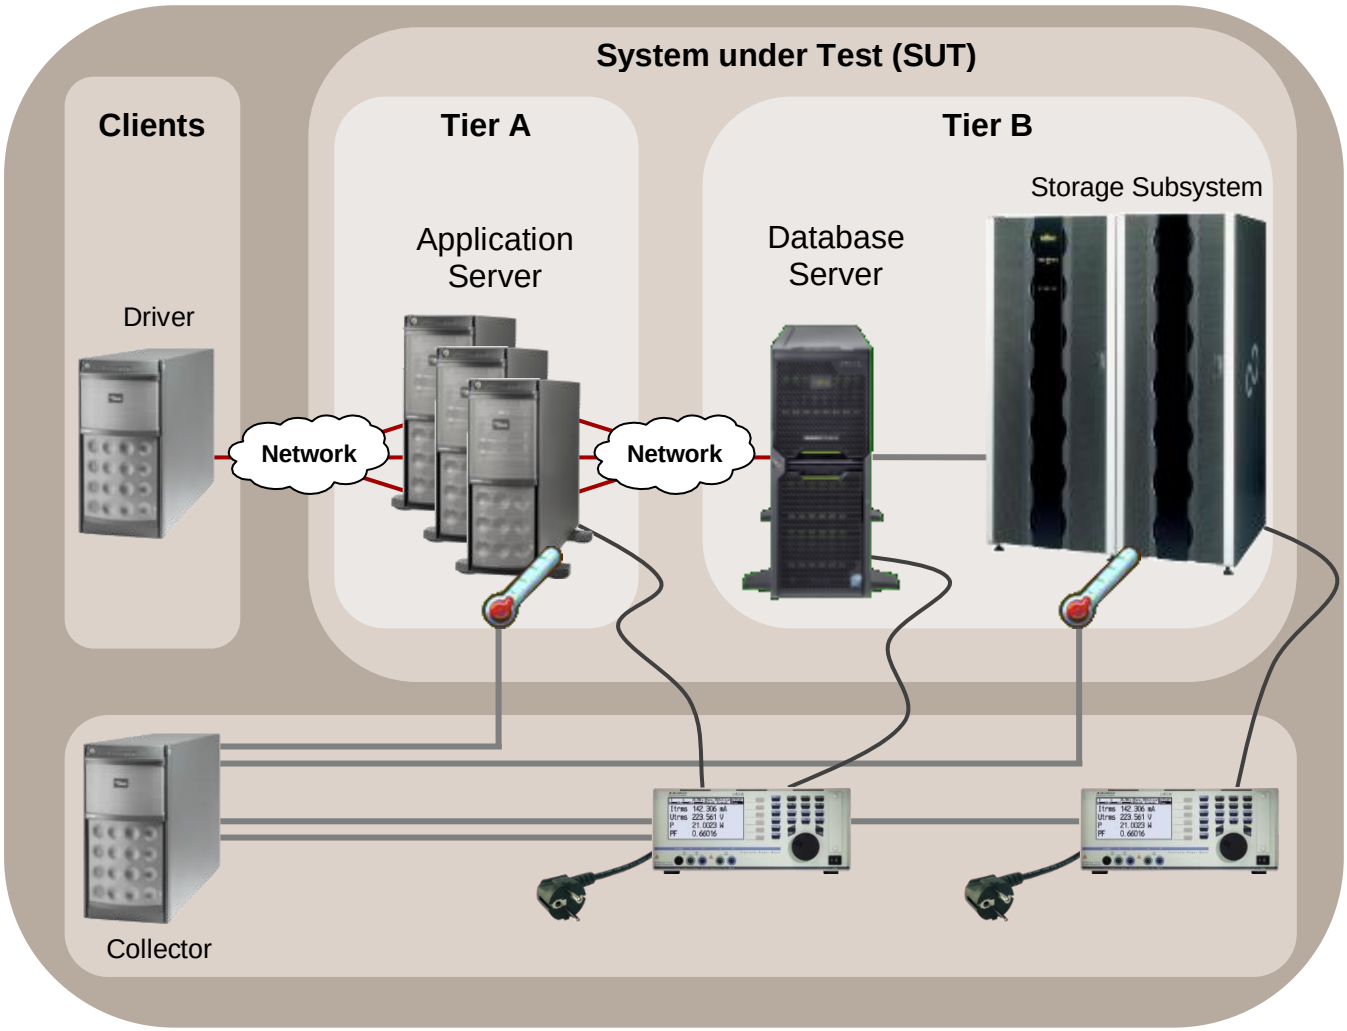

The TPC-Energy component is optional. If the power consumption of the system under test for TPC-Energy is measured in the TPC-E benchmark, the measuring environment is extended to include a collector system for the collection of the measurement data as well as one or more power measuring devices and temperature sensors, which are connected to the collector system.

Differences to TPC-C

In addition to the TPC-E benchmark, there is also the TPC-C benchmark, which dates back to 1992 and which had established itself as the accepted database benchmark over the years. Meanwhile system performance has increased in a dramatic way and requires a huge effort to provide the test environment for TPC-C. 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.

TPC-C used five transaction types and nine tables. TPC-E uses eleven 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 driver code provided by TPC.

TPC-E is not a new version of TPC-C and does not replace it. Furthermore, it is possible to run TPC-C measurements and publish results. Current TPC-C values remain valid.

Concluding remark

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-E 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.

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 measurement results of TPC-E 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-E throughput values chiefly depends on how similar a customer's database and application in fact are in comparison with those of TPC-E. Although TPC-E 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-E therefore cannot replace the performance analysis of a customer application.

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=da0ce7b7-3d80-48cd-9b3a-d12e0b40ed6d

http://docs.ts.fujitsu.com/dl.aspx?id=08c95eef-5f18-4453-bed6-cbf9363f4e2f

http://docs.ts.fujitsu.com/dl.aspx?id=4d3c087a-6978-4923-849f-e4288af8615a

TPC-E

http://www.tpc.org/tpce

TPC-Energy

TPC-Energy Specification
http://www.tpc.org/tpc_energy

Energy Benchmarks: A Detailed Analysis
http://www.tpc.org/tpc_energy/Presentations/EnergyBenchmarks.pdf

TPC-Energy Benchmark Development
http://www.tpc.org/tpc_energy/Presentations/TPC-Energy-2.pdf

TPC-C

http://www.tpc.org/tpcc

Benchmark Overview TPC-C
http://docs.ts.fujitsu.com/dl.aspx?id=8d633e1e-8b44-4c5a-800f-7a1319649081

Standard Performance Evaluation Corporation (SPEC)

http://www.spec.org/

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