

An Overview of the TPC-H Benchmark on HP ProLiant Servers and Blades



Table of Contents

TPC-H Overview	2
How TPC-H Performance Is Measured	2
The TPC-H Metrics	3
Comparing TPC-H Results	3
About the TPC-H Workload	4
TPC-H on HP ProLiant Servers & Blades	4
Conclusion	4
For more information	4

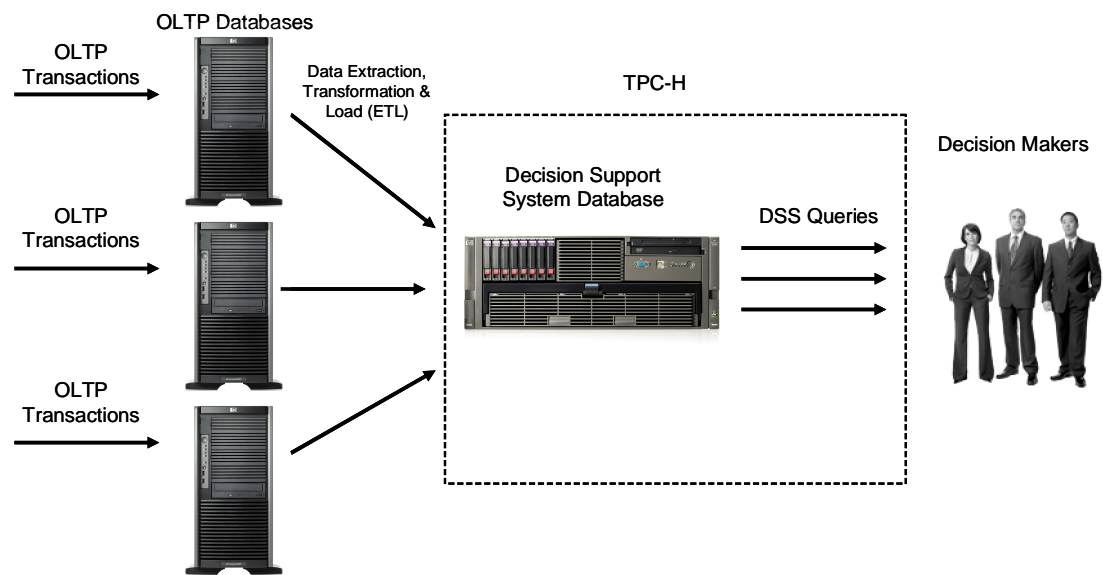
TPC-H Overview

TPC-H is a benchmark that simulates a Decision Support System or Business Intelligence database environment. The performance of a system is measured when the system is tasked with providing answers for business analyses on a dataset. These analyses include

- Pricing & Promotions Analysis
- Supply & Demand Management Analysis
- Profit & Revenue Management Analysis
- Customer Satisfaction Studies
- Market Share Studies
- Shipping Management Analysis

The server system runs a read-intensive Decision Support System (DSS) style database to provide the results for the business analyses. The DSS database is designed to mimic a repository of commercial order-processing Online Transaction Processing Databases. The analyses are performed on 100GB, 300GB, 1000GB, 10,000GB, 30,000GB or even 100,000GB scale factor datasets.

Figure 1. The TPC-H Simulated Workflow



How TPC-H Performance Is Measured

The TPC-H benchmark runs the business analyses in two different ways – it performs a Power Test and a Throughput Test.

The TPC-H Power Test measures the query execution power of the system when connected with a single user. It runs the analyses in a serial manner – the queries and update functions run one at a time and the elapsed time is measured.

The TPC-H Throughput Test measures the ability of the system to process the most queries (and update functions) in the least amount of time in a multi-user environment. Each simulated user runs its own version of the TPC-H Power Test simultaneously. The number of simulated users for this test is selectable with a minimum number of users or “streams” detailed by the TPC-H specification for each scale factor.

The TPC-H Metrics

The primary performance metric for the TPC-H benchmark is the Composite Queries-Per-Hour (QphH). Composite Queries Per Hour (QphH) is calculated as a geometric mean from the results from the TPC-H Power Test and the TPC-H Throughput Test. The more QphH, the better the benchmark result. As per TPC rules, QphH is always represented as QphH@Scale Factor (ex: 5028QphH@100GB).

TPC-H also has a pricing component, which is the total cost of the benchmark system (hardware and software) including 3 years worth of 24x7 4-hour response time support, divided by the Composite Queries Per Hour (ie price/performance). This is represented as \$(currency)/QphH@Scale Factor (example: \$15USD/QphH@100GB). The lower the price/performance, the more cost effective the system.

Comparing TPC-H Results

When comparing TPC-H results, be sure to compare results of the same scale factor. For example, comparing a system that ran 19323QphH@100GB with a system that ran 12,225QphH@300GB is not a fair comparison, whereas comparing a system that ran 19323QphH@100GB with a system that ran 17,686QphH@100GB is.

- The overall TPC-H performance of a group of systems may be compared using the TPC-H Composite Queries-Per-Hour (QphH) metric.
- The price/performance of a group of systems may be compared using the \$/QphH metric.
- The system power of a single user may be compared using the TPC-H Power (QppH) metric.
- The throughput and ability to process the most queries in the shortest amount of time may be compared using the TPC-H Throughput (QptH) metric.

All published TPC-H results are stored on the TPC website at <http://www.tpc.org>. When making comparisons, it is useful to download the 3-4 page benchmark Executive Summary rather than try to wade through the more detailed Full Disclosure Report. The Executive Summary will include the price/performance, QphH, QppH, and QptH metrics, along with a full price sheet including all of the hardware and software used in the benchmark.

About the TPC-H Workload

The business analysis workload of the TPC-H consists of “ad-hoc” database queries with a high degree of complexity. The intention is that the exact queries are not known beforehand, which is often the case for DSS databases.

The TPC-H Power Test workload performs significant sequential disk I/O as the single user’s queries perform index and table scans over the substantial datasets. The DSS database’s query optimizer typically splits the query into a number of parts and runs them in parallel, distributing the data retrieval across the disk arrays attached to the server system. The workload generally consists of read activity, though there are bursts of database inserts and updates performed during queries called “refresh functions”, which simulate OLTP databases periodically feeding data into the DSS system. There is typically a considerable amount of tempdb activity as the results of complex joins, aggregations and interim results are stored for the business analyses.

The TPC-H Throughput Test is essentially several instances of the Power Test being run simultaneously using different query parameters. This results in database buffer cache contention in system memory as the various streams compete for buffer space. The sequential read activity tends to become more random as the multiple streams access the disk drives simultaneously. The processors on the system also become busier from running more processes and switching tasks.

TPC-H on HP ProLiant Servers & Blades

HP ProLiant Servers and Blades have a long history of high performance and record-breaking results on 100GB, 300GB, 1,000GB and 3,000GB TPC-H benchmarks. Between the ProLiant Server, BladeSystem and Integrity Server products, HP consistently publishes more than twice as many TPC-H benchmarks as any other vendor.

Conclusion

TPC-H benchmarks are useful for comparing systems for a Decision Support System or Business Intelligence database environment by presenting system performance metrics including overall system performance and system price/performance.

For more information

www.hp.com/products/servers/benchmarks

www.tpc.org

© 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

July 2007

