New TPC Benchmarks for Decision Support and Web Commerce

Meikel Poess
Oracle Corporation, Redwood Shores, CA 94065
meikel.poess@oracle.com

Chris Floyd
IBM Corporation, Research Triangle Park, NC 27709
csfloyd@us.ibm.com

Guest Column Introduction

For as long as there have been DBMS's and applications that use them, there has been interest in the performance characteristics that these systems exhibit. This month's column describes some of the recent work that has taken place in TPC, the Transaction Processing Performance Council.

TPC-A and TPC-B are obsolete benchmarks that you might have heard about in the past. TPC-C V3.5 is the current benchmark for OLTP systems. Introduced in 1992, it has been run on many hardware platforms and DBMS's. Indeed, the TPC web site currently lists 202 TPC-C benchmark results. Due to its maturity, TPC-C will not be discussed in this article.

We've asked two very knowledgeable individuals to write this article. Meikel Poess is the chair of the TPC-H and TPC-R Subcommittees and Chris Floyd is the chair of the TPC-W Subcommittee. We greatly appreciate their efforts.

A wealth of information can be found at the TPC web site [1]. This information includes the benchmark specifications themselves, TPC membership information, and benchmark results.

Andrew Eisenberg and Jim Melton

The TPC-H and TPC-R Decision Support Benchmarks

In April 1999, two new benchmarks, TPC-R and TPC-H, replaced TPC-D as the industry's standard benchmarks for decision support applications. Unable to integrate two workloads typical for decision support systems, the TPC constituted two separate benchmarks, TPC-R for a *reporting workload* and TPC-H for an *ad-hoc querying workload*.

An ad-hoc querying workload simulates an environment in which users connected to the database system send individual queries that are not known in advance. Hence, the system's database administrator (DBA) cannot optimize the database system specifically for this set of queries. Consequently, execution time for those queries can be very long.

On the other hand, queries in a reporting workload are very well known in advance. As a result, the DBA can optimize the database system specifically for these queries to execute them very rapidly.

This article describes the two new benchmarks, TPC-H and TPC-R. It focuses on the differences and commonalties between these benchmarks. Differences between these benchmarks and TPC-D are mentioned where necessary. The next three sections describe the business environment, the schema and the workload of the new benchmarks. The two sections that follow them describe the implementation rules and the benchmark metric.

Business Environment

Instead of representing the activity of any particular business segment, TPC-H and TPC-R model any industry which manages, sells, or distributes products worldwide (e.g., car rental, food distribution, parts, suppliers, etc.). Figure 1 describes the business environment of the modeled business. It is divided into two large areas, a *business operation* area and a *decision support* area.

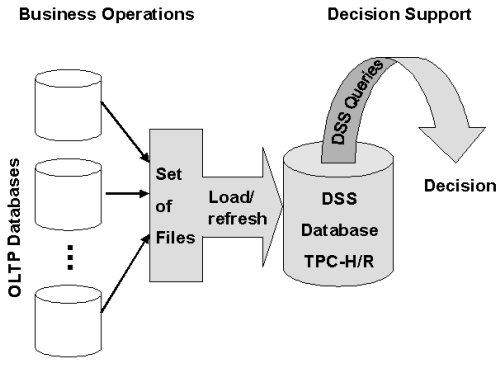

Figure 1: The TPC-H and TPC-R Business Environment

The *business operation* models the operational end of the business environment where transactions are executed on a real time basis. Performance of systems managing this area is represented in other TPC benchmarks such as TPC-C and TPC-W.

TPC-H and TPC-R model the *decision support* area where business trends are computed and refined data are produced to support the making of sound business decisions. The scope of TPC-H and TPC-R is indicated with the gray shaded parts of Figure 1.

The purpose of these benchmarks is to reduce the diversity of operations found in a typical decision support application, while retaining the application's essential performance characteristics, namely: the level of system utilization and the complexity of operations. The core of the benchmarks is comprised of a set of business queries designed to exercise system functionalities in a manner representative of complex decision support applications. These queries have been given a realistic context, portraying the activity of a wholesale supplier to help the audience relate intuitively to the components of the benchmarks.

The selected queries assist decisions makers in the following domains of business analysis:

- pricing and promotions;
- supply and demand management;
- profit and revenue management;
- customer satisfaction study;
- market share study;
- shipping management.

For instance, one of the queries in TPC-H and TPC-R quantifies the amount of revenue increase that would have resulted from eliminating certain company-wide discounts in a given percentage range in a given year. Asking this type of "what if" query can be used to look for ways to increase revenues.

TPC-H and TPC-R require the database to support queries and refresh functions against all tables on a 7 day by 24 hour (7 x 24) basis. Although both benchmarks model a business environment in which refresh functions are an integral part of data maintenance, the refresh functions actually required in the benchmark do not attempt to fully model this aspect of the business environment. Their purpose is rather to demonstrate the update functionality of the DBMS.

Schema

Both the TPC-H and TPC-R benchmarks use the same 3rd normal form schema as their predecessor, TPC-D. It consists of eight base tables. Figure 2 illustrates the entity relationship (ER) diagram of this schema. TPC supplies a data generator (dbgen) that generates data for all base tables depending on a scale factor (SF). The scale factor determines the size of the raw data outside the database, e.g. SF=100 means that the sum of all base tables equals 100 gigabytes. It must be chosen from the following set of fixed scale factors: 1, 10, 30, 100, 300, 1000, 3000, 10000 – all scale factors are in gigabytes. The two largest tables, *lineitem* and *orders* contain about 83 percent of the data. The size of all tables, except for nation and region, scales proportionally with the scale factor.

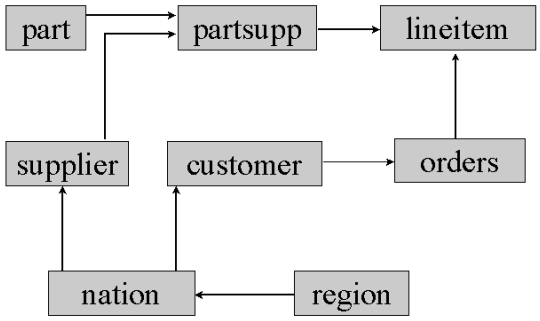

Figure 2: TPC-H and TPC-R database schema

Workload

The workloads in both benchmarks, TPC-H and TPC-R, consist of a database load, the execution of 22 read-only queries in both single and multi-user mode and two refresh functions.

The process of building the test database is known as "database load". The database load time includes all of the elapsed time to create the tables, load data, create indices, define and validate constraints, gather statistics, configure the system, and ensure that the test database meets the ACID requirements including syncing loaded data on RAID devices and the taking of a backup of the database, when necessary. Both benchmarks define the same database load.

Two refresh functions (RF1, RF2) model the loading of new sales information (RF1) and the purging of stale or obsolete sales information (RF2) from the database. RF1 inserts new rows into the tables *lineitem* and *orders* while RF2 removes the same number of rows from those tables.

Four components, a business question, a functional query definition, substitution parameters and a query validation, characterize the 22 queries used in both benchmarks.

The business question illustrates the business context in which the query is used. The functional query definition defines the function to be performed by the query. It is defined in SQL-92 language (ISO/IEC 9075:1992). Each query is defined as a query template. Before a query can be executed against a database, substitution parameters need to be completed. The substituted values are generated by the supplied program *qgen* in such way that the performance for a query with different substitution parameters is comparable. The query validation describes how to validate each query against a 1 gigabyte database (qualification database).

TPC-H and TPC-R define 22 queries, 16 of which were carried over from TPC-D. They are run stand-alone to demonstrate the ability of the test system to use all of the resources for a single user and in multiple concurrent sessions to demonstrate the ability of the systems to use all of the resources to satisfy concurrent users. The six new queries were defined to add more functionality and query complexity to the workload. The resulting 22 queries are intended to test most query capabilities in respect to a typical decision support system. They:

- give answers to real-world business questions;
- simulate generated ad-hoc queries (TPC-H) and reporting queries (TPC-R) (e.g., via a point and click GUI interface);
- are far more complex than most OLTP transactions;
- include a rich breadth of operators and selectivity constraints;
- generate intensive disk and CPU activity on the part of the database server;
- are implemented with constraints derived from staying closely synchronized with an on-line production database.

Implementation Rules

In order to implement the two quite different workloads, the *reporting workload* of TPC-R and the *ad-hoc querying workload* for TPC-H, the two benchmarks differ in their implementation rules.

Both benchmarks require a commercially available database management system (DBMS) and the same physical implementation of the test database. However, TPC-R and TPC-H allow very different partitioning schemes and auxiliary structures.

Partitioning Schemes

In TPC-H, horizontal partitioning is allowed with some restrictions. If data is assigned to partitions based on the values of a partitioning field, this field must be one and only one of the following:

- a primary key column as defined in the benchmark specification;
- a foreign key as defined in the benchmark specification;
- a single date column.

If the partitioning schemes require the use of directives that specify explicit values for the partitioning field, the use of such directives is restricted. They may not rely on any knowledge of the data stored in the table except the minimum and maximum values of columns used for the partitioning field. Within the limitations of integer division, they must define each partition to accept an equal portion of the range between the minimum and maximum values of the partitioning column(s). Furthermore, the directives must allow the insertion of values of the partitioning column(s) outside the range covered by the minimum and maximum values.

In TPC-R, horizontal partitioning is allowed, too. Although there are fewer restrictions than in the TPC-H partitioning schemes, a partitioning schema in TPC-R cannot rely on knowledge of the data stored in the partitioned columns either. A table can be partitioned on any table column or combination of table columns. As with TPC-H, any partition directives must allow the insertion of values of the partitioning column(s) outside the range covered by the minimum and maximum values.

Auxiliary Structures

The permissible types of auxiliary data structures managed by the DBMS differ substantially between TPC-H and TPC-R. Both benchmarks prohibit the logical replication of database objects (i.e., tables, rows, or columns). The physical implementation of auxiliary structures to tables, which may involve data replication of selected data from the tables, is permitted if:

- all replicated data are managed by the DBMS;
- all replications are transparent to all data manipulation operation;
- by the time the transaction is committed, data modifications are reflected in all auxiliary structures;
- all copies of replicated data maintain full ACID properties.

The only type of auxiliary structure allowed in TPC-H is an index. Indexes can refer to only one base table. They can refer to a primary or foreign key as defined in the TPC-H specification, a

compound primary or foreign key as defined in the TPC-H specification or any single date datatype column. In TPC-R none of the above restrictions on the creation of indexes exist, e.g. indexes can be built on any table column or any combination of table columns.

While TPC-H explicitly prohibits the creation of materialized views, TPC-R allows a wide variety of auxiliary structures including indexes, join-indexes and materialized views, that pre-compute the answers to the given queries. However, all of the above characteristics of auxiliary structures need to be fulfilled.

Primary Performance Metric

The primary performance metric in TPC-R and TPC-H is the *composite* performance metric (QphH/QphR). It equally weights the contribution of the old TPC-D single user power metric and the old TPC-D multi-user throughput metric. It expresses the number of queries the system can perform per hour. Both benchmarks compute the composite metric the same way. Hence, the following paragraphs apply to both benchmarks if not otherwise mentioned.

In order to compute the composite metric for a test system at a given scale factor, one needs to run a power test followed by a throughput test. Figure 3 demonstrates the two tests. The results of the power test and the throughput test are then combined to calculate the composite metric.

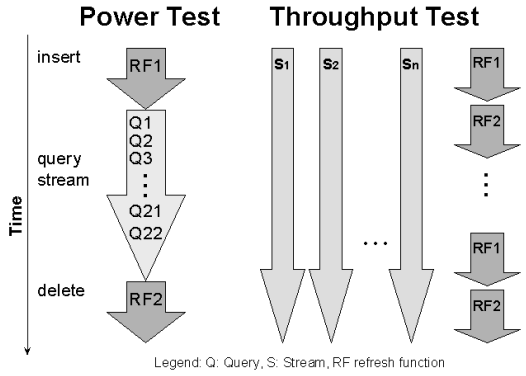

Figure 3: Benchmark Execution Plan of TPC-H and TPC-R

A power test consists of the execution of the refresh function RF1, followed by the execution of all 22 queries in single user mode and concluded by the execution of the refresh function RF2. During a throughput test multiple concurrent sessions execute all 22 queries in a pre-defined order. One sequence of queries in a session is also called a 'query-stream' (S_1, …, S_n in Figure 3). One pair of update functions (RF1, RF2) is executed for each query-stream. All update-pairs are executed sequentially in a separate session. The sequence of update functions is also called the 'update-stream'. Unlike TPC-D, where a multi-stream run was optional, TPC-H and TPC-R require a multi-stream run with a minimum number of streams. The minimum number of streams depends on the scale factor (see Table 1).

Scale Factor	S (Streams)
1	2
10	3
100	4
300	6
1000	7
3000	8
10000	9

Table 1: Minimum Required Query-Stream Count

The results of the power and throughput tests are used to compute the processing and throughput power. The processing power (Power@Size) is computed as the geometric mean of the elapsed times for all queries and both refresh functions obtained from the power run. The unit is queries per hour. The geometric mean reduces the effects of queries with long and short elapsed times[1]. For a given scale-factor (SF) it is computed as:

- **i** = 1,2, … ,22; **j** = 1,2
- **QI(i,0):** elapsed time of query Q_i obtained in the power run in seconds.
- **RI(j,0):** elapsed time of the refresh function RF_j obtained in the power run in seconds.
- **3600**: seconds (= 1 hour)

$$\text{Power@Size} = \frac{3600 * SF}{\sqrt[24]{\prod_{i=1}^{i=22} QI(i,0) * \prod_{j=1}^{j=2} RI(j,0)}}$$

The throughput power (Throughput@Size) is computed as the ratio of the total number of queries executed over the length of the measurement interval of the multi-stream run. The unit is queries per hour. For a given scale factor (SF), it is computed as:

[1] Assuming we had only three queries, Q_1, Q_2 and Q_3 with elapsed times of 10, 12 and 500 seconds. The arithmetic mean of the three query times is 174 seconds while the geometric mean is 39.15 seconds.

T_s: elapsed time of the multi stream run
S: stream count
3600: seconds (= 1 hour)
22: number of queries per stream

$$\text{Throughput@Size} = \frac{S * 22 * 3600}{T_s} * SF$$

The Composite Query-Per-Hour Performance Metric (QphH for TPC-H and QphR for TPC-R) is calculated as:

QphH@Size, QphR@Size =

$$\sqrt{Power@Size * Throughput@Size}$$

Lastly, TPC-H and TPC-R define a price/performance metric as the ratio of the total system price divided by the composite metric.

In comparison, TPC-D defined four metrics, three performance metrics and one price performance metric. The power metric (QppD) as the first primary metric represented the response time performance of the test system for a single user. The throughput metric (QthD) as the second primary metric represented the throughput capacity of the test system with multiple parallel users. The power metric and the throughput metric were combined to calculate the composite metric as the secondary metric (QphD). The price performance metric measured the price per QppD.

While the three performance metrics truly reflected different performance evaluation perspectives, ultimately, the metrics generated confusion in the marketplace. Both vendors and users alike want one primary performance number, and the TPC responded by creating one *composite* performance metric for TPC-R and TPC-H.

Published Results

Since TPC introduced TPC-H and TPC-R in April 1999, there have been 16 TPC-H and 2 TPC-R results published. Table 2 shows the number of TPC-H results at different scale factors between April 1999 and September 2000. The two TPC-R results were published at a scale factor of 1000. In the first year after TPC-D was introduced there were approximately 11 TPC-D results, while there were about 8 TPC-H and TPC-R results published in the first year after TPC-H and TPC-R were introduced.

Scale Factor	Number of Results
100	6
300	7
1000	3
Total	16

Table 2: TPC-H Results since April 1999

Conclusion

In April 1999 TPC introduced two new benchmarks, TPC-R and TPC-H, to replace TPC-D as the industry's standard benchmarks for decision support applications. TPC-H and TPC-R implement two workloads, typical in decision support applications, a *reporting workload* and an *ad-hoc querying workload*.

The two benchmarks are very similar. They model the same business application, they are based on the same database model and they use the same queries, load and refresh model. However, they differ significantly in their implementation rules. While TPC-H limits the use of indexes and partitioning schemes, TPC-R does not restrict the use of those database structures. The use of materialized views is prohibited in TPC-H. TPC-R allows materialized views if the DBMS' implementation supports the rules for replicated data.

The two workloads are implemented through the above restrictions on allowed auxiliary structures. A common technique to optimize a decision support DBMS for a given set of queries is the building of partitions, indexes and materialized views. TPC-H restricts the use of these database structures to simulate an *ad-hoc workload* while TPC-R allows them to implement a *reporting workload*.

As beseeched by both vendors and users alike, TPC defines only one primary metric for TPC-H and TPC-R, the Composite Query-Per-Hour Performance Metric - QphH for TPC-H and QphR for TPC-R. The metric equally weights the contribution of the single-user and the multi-user tests. Although the metrics have identical meanings in both benchmarks, they cannot be compared to each other, because they are obtained by different workloads.

At the time TPC-D was introduced in 1999 it spurred the development of both hardware and software products. Being the industry's benchmark of data warehouse applications, it continued to encourage database and hardware development through 1998. At the end of 1998, with the introduction of new database technology such as materialized views, TPC-D was conquered. Hoping to create a similarly challenging benchmark, TPC

responded by splitting TPC-D into two separate benchmarks and adding more complex queries. At this early stage in their existence TPC-H and TPC-R still have to prove that they can drive hardware and database development as TPC-D did for almost 6 years.

The TPC-W Web Commerce Benchmark

The TPC developed TPC Benchmark W (TPC-W) in response to the increasing need for an industry-standard benchmark that could measure the performance of both hardware and software in an e-commerce environment. The benchmark was approved in February 2000.

TPC-W is based on a business model that employs a shopping scenario typical of an online bookstore. The benchmark includes a mix of static content (predominantly images) and dynamic content, some of which can be cached for a specified period of time. It is important to note that including the dynamic content generated by database transactions distinguishes TPC-W from other e-commerce benchmarks. TPC-W's content mix provides the opportunity to illustrate the significant benefits of using emerging Web cache technologies. TPC-W also requires the use of secure transactions (SSL V3) for a subset of the interactions.

The benchmark allows the use of varying scale factors to represent larger product-set sizes as well as larger customer bases (i.e., the number of books and customers); thus, benchmark sponsors are responsible for deciding which scale factor will result in a database population that represents their customer needs.

The primary metric is reported as Web Interactions per Second, or WIPS; two other mixes used include an e-commerce site with heavy browse traffic (reported as WIPSb) and a B2B model with heavy ordering (reported as WIPSo). The transaction percentages used by these models were calculated by analyzing statistics from actual production Web sites. Application code can be implemented using either custom code, or commercially available e-commerce software packages.

The main components of the system under test (SUT) are the Web server(s), the application server(s), communication interface(s), and database server(s), including their respective hardware and software. Additional components, such as load balancing, index searching, or Web cache products are included in the benchmark configuration. All defined components must be commercially available. A series of tests that have been defined to illustrate that ACID (atomicity, consistency, isolation and durability) properties are met. The entire system must display full ACID properties. Clustering in any or all components is allowed as long as ACID properties can be met.

Benchmark Issues

When you consider the wide variety of e-commerce solutions employed today, you can understand that the options for implementing a TPC-W benchmark are equally varied. In developing the benchmark, the TPC had to address many issues to ensure that TPC-W results would be relevant to a wide range of audiences. The benchmark results are intended to demonstrate a given system's ability to handle thousands of simultaneous requests for dynamic data. Although the system is comprised of various hardware and software components, most of the Web interactions result in some form of database activity. This means that the database itself is one of the components most stressed in the running of the benchmark. Among the many challenges of such an environment is finding a way to minimize resource contention (e.g., database lock contention, CPU utilization). Optimization of database layout and queries is essential; however, it does not address all the issues in an e-commerce environment, or in TPC-W benchmark. For example, Web server optimization, which includes load balancing, Web server acceleration, and Web caches, must all be taken into consideration. User load balancing based on Web interaction function is a viable solution, but may not be as optimal as the more typical "Web-farm" approach (in which all Web servers perform the same functions).

To develop a realistic workload for the entire system (i.e., to reduce database utilization), TPC-W allows the use of various commercially available components. For example, commercial index servers can be used to reduce database traffic for defined product search queries. The results from certain Web interactions that generate a result list (e.g., New Products List and Best Sellers List) are allowed to remain "stale" for up to 30 seconds, which enables the showcasing of common Web cache methods as well as commercial Web cache products that can dramatically reduce database utilization and increase overall system performance.

Use of secure transactions (SSL V3) raises issues regarding how these kinds of transactions are processed. For example, SSL encryption that is accomplished by the Web server itself must be weighed against the cost of performing these calculations via SSL hardware devices, which are becoming more commonplace.

The primary metrics of TPC-W are performance (WIPS), price/performance ($/WIPS), and the date of total system availability. The price/performance metric presents an additional variable that must be considered before a given implementation is attempted. Benchmark sponsor companies must decide whether to build the largest possible configuration to demonstrate the highest throughput, or a relatively small configuration to demonstrate a more affordable e-commerce solution. In either case, the resources and effort required to produce a result for this relatively new benchmark are considerable.

The e-commerce environment is complex, and the TPC-W benchmark reflects this complexity to some degree. To date, only IBM Corporation has published a TPC-W result (1,262.79 WIPS @10,000 and $276.28 per WIPS @ 10,000).[2] IBM published this result approximately five months after the TPC-W specification was approved. The expectation is that more results will eventually be published by members of the TPC; however, no new results had been released at the time this article was written.

Considerations for Benchmark Implementation

IBM had to overcome a number of obstacles on the way to publishing a TPC-W result. First, clarification was needed as to whether the specification allowed certain implementations. The TPC has an established a TPC-W maintenance subcommittee whose charter is to deal with specification improvements and clarifications. Several changes were made to the specification in June 2000. At that time, a result could be published using either the old version (v1.0.1) or the new version (v1.1) of the specification. After the clarification and implementation modifications to the specification were made, IBM chose to publish with version 1.1 of the specification. This is an example of the kinds of issues that an interpretation can present during the process of conducting the benchmark.

The TPC has two other mechanisms for addressing questions related to benchmark implementation. The Technical Advisory Board (TAB) may be asked to interpret the wording of the specification. A test sponsor can obtain an opinion from an independent group of seven member companies as to how they would interpret the wording of the specification. Or, if a member company feels a test sponsor has not complied with the specification, they also have the option of challenging an implementation through the TAB.

The cost in resources to produce a TPC-W benchmark can vary greatly depending on the hardware vendor, the scale of the database, internal pricing, and other factors, such as the cost of developing application code, which is not reflected the published total system cost. The total cost to produce a TPC-W result, with all resources factored in, is substantial, making it a difficult, if not impossible, undertaking for a small company. All TPC member companies are among the largest hardware and software companies in the industry. Since TPC-W utilizes far less physical I/O per transaction than does TPC-C, disk subsystems will likely not be as large (and expensive) as those required by TPC-C. Although, with the additional e-commerce system components included, it is difficult to determine whether the overall cost of producing a TPC-W benchmark will be less than TPC-C. The 3-Year Cost of Ownership published for the IBM Netfinity TPC-W result was $348,879. However, this cost does not represent the benchmark sponsor's cost to produce the result. The TPC-W executive summary for this result is available online [2].

Current Activities

The TPC-W subcommittee is currently working on resolving a few remaining version 1 issues. Most of these issues are related to the use of caching methods or commercial Web cache products. E-commerce configurations integrate a wide variety of solutions to maximize the benefits obtained by utilizing Web caches. A balance must be found between caching data and providing the most current data to the end user. The TPC-W subcommittee is charged with the task of finding this balance in a way that does not overburden or under-utilize the database management system (DBMS), but still represents a realistic workload that demonstrates the benefit of cache technologies.

The subcommittee also looks to the future in its ongoing development of TPC-W version 2. For example, in developing version 1, the subcommittee considered it highly desirable to require e-commerce software for the middle-tier application code. At the time of development, however, no easily definable set of functionality seemed to be included in all e-commerce packages. One of the goals for version 2 is to require the use of e-commerce software suites in order to showcase more concrete and obtainable results. The thinking is that customers will be more inclined to embrace results that have been produced

[2] The availability date for the IBM Netfinity TPC-W result is August 1, 2000.

with commercially available e-commerce software only.

Another recommendation for version 2 is the addition of more B2B processing. Concerns have been voiced that the current version is too complicated to implement with the secondary metrics, WIPSo and WIPSb. Although the secondary metrics were included to provide more information to the customer in the form of additional workload models, it may be more beneficial to either change the primary model, or create a second benchmark dedicated to B2B. The integration of additional Web technologies and methodologies (e.g., XML, EDI, WAP, multimedia objects, click stream analysis) is being considered for version 2. The subcommittee will begin reviewing findings in these areas and will define accepted goals in the near future.

Conclusion

The TPC community produced TPC-W in response to a pressing need for an industry-standard e-commerce benchmark. Gaining consensus from several of the largest companies in the industry on a benchmark standard for a dynamic environment such as the Internet was no small task. The ongoing work to revise the specification to incorporate emerging technologies, and measure the performance capabilities of e-commerce systems, promises to be extremely interesting as well as challenging. The integration of new technologies and benchmarks is not only a goal of the TPC, but a cornerstone of all TPC benchmarks. In the coming months, we anticipate seeing how member companies will implement TPC-W and market the results — and how TPC-W will change to meet the dynamic e-commerce world.

Web References

[1] Welcome to the TPC Main Page!: http://www.tpc.org

[2] TPC-W Summary Listing: http://www.tpc.org/New_Result/tpcw_summary_results.asp