TPC-W*: Benchmarking An Ecommerce Solution

By Wayne D. Smith, Intel Corporation

Revision 1.2

1 INTRODUCTION

How does one determine server performance and price/performance for an Internet commerce, Ecommerce, environment? The most notable benchmarks, (SPECWeb and TPC-C), have been referenced by many vendors promoting the performance of their Ecommerce servers. However these benchmarks do not come close to representing the complex environment of an Ecommerce workload. In Feb 2000, the Transaction Performance Council, TPC, introduced the TPC-W benchmark targeted at the Ecommerce environment. TPC-W specifies an Ecommerce workload that simulates customers browsing and buying products from a website. The solution for this workload is a number of servers, (Web Servers, Web Caches, Image Servers & Database Server), working in concert to provide an Ecommerce solution. It is very similar to how an actual Ecommerce site would operate. On first glance at a TPC-W publication, there is the immediate reaction to ask what is it benchmarking? Whereas previous benchmarks target one specific server, TPC-W benchmark results contain a host of servers performing different functions. However, unlike other TPC benchmarks, TPC-W requires the test sponsor to report performance data for the various servers involved. Thus by peeling back the covers of TPC-W, one can obtain information of how each of the servers perform in a complex Ecommerce environment. TPC-W's strength is that one can determine the performance characteristics of specific Ecommerce servers in a simulated Ecommerce workload.

2 WHAT IS TPC-W?

TPC-W specifies an Ecommerce workload that simulates the activities of a retail store website. Emulated users can browse and order products from the website. In the case of TPC-W the products are books. A user is emulated via a Remote Browser Emulator, RBE, that simulates the same HTTP network traffic as would be seen by a real customer using a browser. In fact one can connect a real browser and walk through the TPC-W website browsing and ordering books. The TPC-W specification describes in detail the 14 different web pages of the website. There also exists several documents at the TPC website, (www.tpc.org), that describe the TPC-W benchmark. The easiest way to think of the benchmark is to imagine you are browsing an Ecommerce website. The expected introduction or "Home" page is the first page you will see. It includes the company logo, promotional items and navigation options to the top best selling books, a list of new books, search pages, your shopping cart, and order status pages. You can browse pages containing a list of new or best selling books grouped by subject, or perform searches against all books based upon a title, author or subject. A product page will give you detailed information for the book along with a picture of the book's front cover. You may order books by entering the secure, SSL, order pages. If you're a new customer you will have to fill out a customer registration page while for returning customers your information will be retrieved from the database and filled in automatically for you. You can change the quantity of your order or delete a book from your shopping cart. When you wish to buy, you enter your credit card information and submit the order. The system will obtain

* TPC-W is a trademark of the Transaction Processing Performance Council

credit card authorization from a Payment Gateway Emulator, PGE, and present you with an order confirmation page. At a later date you can view the status of your last order. Two additional web pages are provided for the system administrator to change a books front cover picture and price. This change is reflected in the new product book list. Essentially the basic functionality of an Internet Ecommerce website is required by the TPC-W benchmark.

The TPC-W primary metrics are the WIPS rating and system cost per WIPS. A WIPS is the number of Web Interactions Per Second that can be sustained by the System Under Test, SUT, (the collection of servers that provide the TPC-W Ecommerce solution). The cost per WIPS is essentially the cost of the SUT divided by the WIPS rate.

The TPC-W benchmark applies the workload via emulated browsers. An RBE must emulate users with separate connections to the SUT. The RBE simulates user think time by sleeping for a number of seconds based upon a distribution with an average think time of 7 seconds and a maximum of 70 seconds. Thus 7 emulated users are required to generate one WIPS. The number of users applied to the SUT can be found at the top of the Executive Summary in the "Number of Users" bucket. If the SUT was performing perfectly you would see a WIPS rate of Number of Users / 7. Each emulated user simulates think time and then randomly selects a TPC-W navigation option to another TPC-W page. This is not a uniform random selection over the 14 different web pages, but a random selection within the legal navigation options of the current page, (i.e. the emulated users follow the same navigation rules as would a real user). The result is that the accesses to the 14 TPC-W web pages fall into a distribution referred to as the "Web Interaction Mix". Thus 80% of the web page accesses are to the Home, New Products, Best Sellers and Search pages while the remaining 20% of the accesses are to the Shopping Cart, Order, Buy and Admin web pages. Of the 20% ordering web pages, 5% of the accesses are to secure web pages requiring SSL encryption.

The SUT cannot just arbitrarily respond in any fashion, the TPC-W benchmark specifies strict response time criteria for each of the 14 web pages. If the SUT cannot meet the response time criteria, the workload must be decreased until the SUT can respond in the specified time.

The RBEs emulate user sessions. A user can be an existing customer (80% of the time), or a new customer unknown to the SUT (20% of the time). The user connection time is a distribution with an average of 15 minutes truncated at a maximum of 60 minutes. When an emulated user ends its session, the RBE will disconnect all connections and re-establish a new user session with the SUT to maintain a constant number of users.

2.1 WIPS b, WIPS o

The TPC-W benchmark specifies two secondary web interaction mixes that are measured to produce the corresponding secondary metrics of WIPSb and WIPSo. The WIPSb mix is intended to simulate a website where there are few buy orders and the majority of the customer requests are browsing the website. This is accomplished by having 95% of the web pages accessed be the browsing pages, (Home, New Products, Best Sellers, Product Detail and Search pages) while only 5% of the Web accesses are to the order web pages. This mix tends to place more pressure on the front-end Web Servers, Image Servers and Web Caches. The WIPSo mix is intended to simulate a website with a significant percentage of order requests. This is accomplished by having 50% of the web page accesses be the browsing pages and 50% of the accesses be to the order web pages. The WIPSo mix stresses the Database Server.

2.2 TPC-W SCALE FACTOR

TPC-W scales in both fixed scale sizes as in the TPC-H/TPC-R benchmarks and scales in the number of the database rows relative to throughput as in the TPC-C benchmark. The fixed scaling changes the number of books in the website inventory. Valid scales are 1,000, 10,000, 100,000, 1,000,000 and 10,000,000 books or product items. TPC-W results can only be compared at the same scale factor, (i.e. you can not compare a 1,000 item result with a 10,000 item result). The fixed scale sizes effect the memory requirements for the Web Servers, Image Servers and Web Caches, as more memory cache is required to hold the larger book inventory. The increase in the database items tables are insignificant compared to the customer and order database tables which must be scaled with the reported throughput. For each of the emulated browsers, the database must maintain 2880 customer records and associated order information. Thus as the number of emulated browsers are added to increase the workload, the database size increases accordingly. The intent of requiring the customer and order database tables to scale with the workload is to prevent outrageous results reported on a small database. The reported WIPS throughput must satisfy the requirement: (Number of Browsers / 14) < WIPS < (Number of Browsers / 7). Remember that the emulated browser think time is an average of 7 seconds, thus a perfect run would see a WIPS rate of Number of Browsers / 7. The lower Number of Browsers / 14) is required to prevent bogus results reported on a very large database. The usual practice is to run as close to the (Number of Browsers/7) limit as possible in order to achieve the highest performance result with the minimum number of resources.

2.3 TPC-W CAVEATS

TPC-W is by no means the perfect Ecommerce benchmark. However it comes closer to simulating an Ecommerce workload than any other benchmark publicly available. TPC-W does not require that an Ecommerce package be used to implement the Web Server application. Most TPC-W solutions are custom coded applications written as close to the Web Server Internet layer as possible. The benchmark provides a great test of Web Server hardware, operating system and basic Internet software capability but does not provide any information on the higher-level Ecommerce packages. In defense of TPC-W, at the time the benchmark was being developed, Ecommerce packages were still evolving and there did not exist a good definition of an "Ecommerce package". The requirement for an Ecommerce package is being addressed in subsequent versions of the TPC-W benchmark.

The intent behind the WIPSb and WIPSo secondary metrics was to provide an indication of how the SUT would perform under just a browse environment or a heavy Ecommerce environment. However in practice, the time and expense involved to publish a TPC-W result influences the test sponsor to concentrating on the primary WIPS metric and expend the minimum effort on the secondary WIPSb and WIPSo metrics. If the SUT is balanced correctly where the front-end servers and Database Server are running near 100%, the WIPSb result will be roughly equivalent to the WIPS result since the WIPSb mix is constrained by the front-end Web, Image and Cache servers. The WIPSo mix generates a significant load on the Database Server. You will find that the WIPSo results are one third to one half of the WIPS results. In most solutions the WIPSo results violates the TPC-W scaling rules such that the test sponsor would have to rebuild and retune the system. For ease of benchmarking, the TPC-W specification relaxes the TPC-W scaling rules for the secondary metrics. This usually results in a drop in the number of browsers in order to drive the WIPSo mix.

The TPC-W specification encourages the use of Full Text Search products for the Author and Title searches. However the results of these searches may be cached for the life of the benchmark measurement interval. In practice, all of the search pages are computed during the benchmark warm-up phase and cached in the Web Cache servers during the measurement interval. A slower database query can be used for the initial Author and Title searches. For some publications, if a Full Text Search product is bundled with the database or OS, it is used to reduce the benchmark warm-up phase.

The TPC-W benchmark does not specify use of buy confirm Emails which are currently in use at most Ecommerce websites. The standard practice is to respond to the user immediately indicating that the order has been recorded and then to follow up with a subsequent Email to confirm the buy order. The TPC-W benchmark requires that the order be confirmed immediately and that 90% of the Buy Confirm requests must return in 5 seconds. It is one of the goals of the TPC to push technology and having an immediate confirmation of an order increases customer satisfaction vs. waiting for several hours or days for an Email confirmation.

A real world Ecommerce website would have many more images and/or pages than what is simulated by TPC-W. In order to compensate for this fact, the TPC-W benchmark disallows any caching by the RBE that would normally be performed by a user's browser. This was done by design to inflate the TPC-W network traffic to simulate the network traffic from a more complex real world Ecommerce website.

3 A TYPICAL TPC-W SOLUTION

Similar to other TPC benchmarks, TPC-W is a specification, which a test sponsor must provide a solution. No specific network topology or server configuration is mandated by the TPC-W specification. The most common configuration found in current TPC-W publications is depicted in Figure 1.

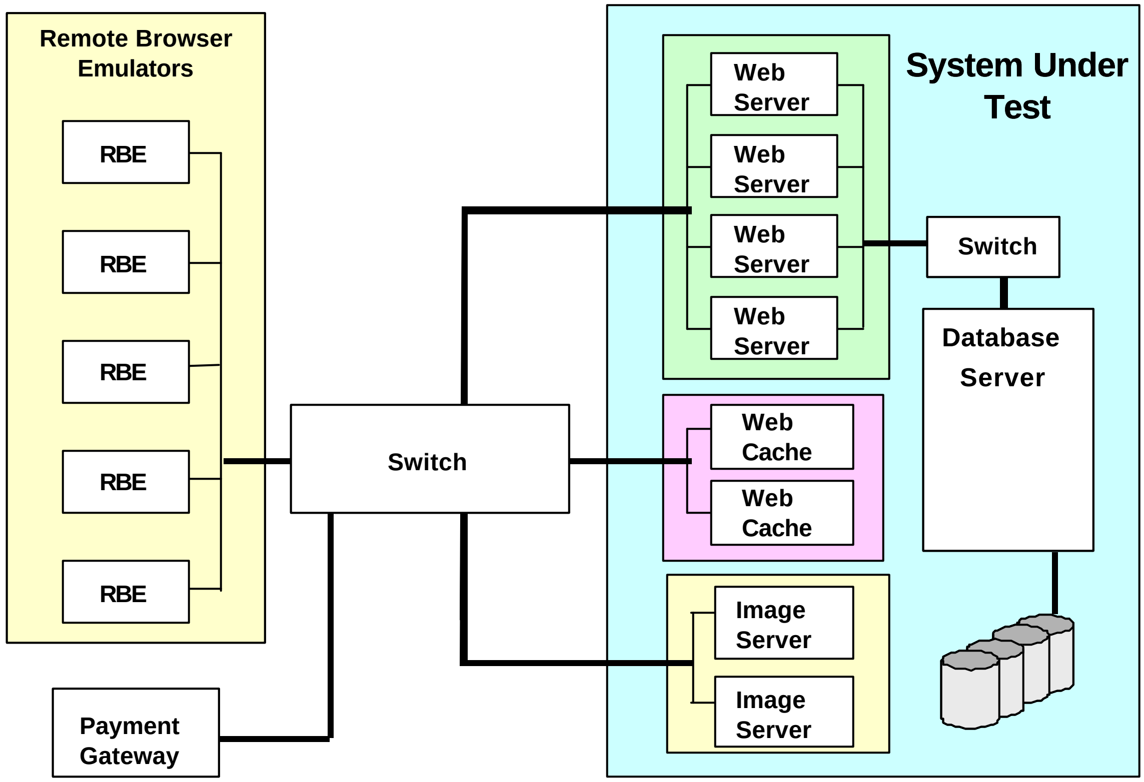

Figure 1

The RBEs are at the left along with the Payment Gateway Emulator, PGE, which simulates the credit card authorization for a Buy Confirm web interaction. These systems are not included in the price of the SUT. The SUT itself is depicted at the right and includes several servers performing different functions. The Web Servers are the point of contact for the RBEs. All web interactions start with a request to a Web Server. It is the Web Server that sends requests to the Database Server and PGE and instructs the RBE to obtain html frames and images from the Web Cache and Image Servers. In some solutions the Web Server can be broken up into two separate servers where a Web Server is fielding RBE requests and forwards only order requests to an Application Server. The Image Servers respond to RBE requests for the Promotional and Product Images. The Web Caches respond to requests from the RBEs for html frames for Best Sellers, New Products and Search data that can be cached. The Best Sellers, New Products and Subject Search results can only be cached for 30 seconds before the data must be refreshed by the SUT. The Title and Author search results can be cached with an infinite timeout. The Database Server responds to requests from the Web Servers. The database contains the customer information and order information. The SUT must also include some type of load balancing software or hardware to load balance the RBE connections to the various servers. This can be a stand-alone server, run on one of the Web Caches or Image Servers or be in the switch itself.

4 TPC-W PERFORMANCE INFORMATION

The TPC-W specification specifies the various performance statistics to be reported in the Full Disclosure Report, FDR. In most FDRs the performance data can be found after the WIPS, WIPSb and WIPSo throughput graphs. The performance data consists of tables that present the performance metrics for the various servers for the WIPS, WIPSb and WIPSo measurement intervals. If multiple systems are performing the same function using the same hardware and software configuration, the performance metrics can be the average of all the systems in that group or can be captured from the system in that group with the highest average CPU utilization. The FDR must state whether the performance data is an average or the data from the server with the highest average CPU utilization. This performance data can be used to determine the performance characteristics of the various servers. The following sections of this paper describe how simple back of the envelope calculations can be computed to provide a wealth of performance information. For the back of the envelope analysis, the difference between average data or data from the system with the highest average CPU utilization is one of degree. For average data, assume there is some variation while for the data from the server with the highest average CPU utilization, assume that the performance data represents the worst-case scenario. Following the performance data tables in the FDR one can find the graph of the CPU utilization for the server with the highest average CPU utilization over the measurement interval.

For a first pass over the data, the performance metrics of interest are the CPU utilization, HTTP bytes/sec and Disk I/Os per second. From the CPU utilization one can determine if the server is balanced in terms of I/O and CPU. Is the server sitting idle most of the time waiting for I/O or has the server been tuned such that the CPU usage is overlapped with I/O, i.e. high CPU utilization. The HTTP bytes/sec data can be converted to Mbytes/sec throughput. With a little knowledge of Network Interface Cards, (NICs), one can quickly determine if the server is network bound or if the NIC is underutilized. The Disk I/O is important for database performance. Given the WIPS rate and some knowledge of the benchmark one can compute connection performance. The following sections discuss in detail the various performance aspects of the TPC-W servers.

4.1 WEB SERVER PERFORMANCE

The Web Server performance metrics of interest are its CPU utilization and the HTTP bytes/sec data. Both can be found in the Web Server performance metric tables usually just after the WIPS, WIPSb and WIPSo throughput graphs in the FDR. Given the HTTP bytes/sec data, one can compute the Mbytes/sec or Mbits/sec network traffic per server and per CPU. For example given a 2 processor Web Server running at 87% CPU utilization with a HTTP bytes/sec value of 4,160,450, one can compute, (4,160,450 / (1024*1024)), a network throughput rate of 3.96 Mbytes/sec or 31.7 Mbits/sec sustained by the 2P Web Server at 87% utilization. The effective throughput rate of a 100 Mbit Network Interface Card, NIC, is roughly 60 Mbits/sec before the NIC latency perturbs performance and for a Gigabit NIC the effective throughput rate is effectively 480 Mbits/sec. You can determine what type and how many NICs have been used in the Web Server by examining the price data on the second page of the Executive Summary. Given the Web Server CPU utilization and network throughput rate, one can determine if the Web Server has any extra headroom or is it configured near its maximum capabilities. In addition, by multiplying the value by the number of Web Servers and computing similar data on the Image Servers and Web Caches, you can obtain an idea of the overall network traffic supported by the front-end servers and the switch.

As mentioned previously, the number of emulated browsers generating the TPC-W workload is specified on the first page of the Executive Summary. The Number of Users / Number of Web Servers gives the number of emulated users supported by each Web Server. The TPC-W benchmark allows keep-alive connections for the duration of the user session. Keep-alive connections are used to reduce the CPU overhead required to process a connection. This is confirmed by the low TCP/IP Connections/sec rate in the performance data. Each user can have one non-secure and one secure connection, thus the total number of connections supported by a Web Server is the 2 * (Number of Browsers / Web Servers). For example given a result of 4,800 WIPS with 35,000 emulated browsers in a configuration of 10 Web Servers, each Web Server is supporting 2 * (35,000 / 10) = 7,000 internet connections.

Dividing the Web Server network traffic and connection data by the number of processors in the Web Server will give you the network traffic per processor and the number of supported connections per processor. This is useful when comparing different Web Server processors or comparing Web Servers with different number of processors. By examining the Web Server performance characteristics and applying some back of the envelope calculations, one can estimate the relative performance of a Web Server under a simulated TPC-W Ecommerce workload. The Web Server is not isolated by itself, repeating the same function over and over, but it is working in concert with other Web Servers to provide an Ecommerce solution for the TPC-W workload.

4.2 IMAGE SERVER PERFORMANCE

The Image Server requirements are dependent upon the TPC-W scale factor as well as the WIPs throughput rate. The Image Server responds to RBE requests for the Promotional and Product Images. For each book there is a Promotional Image and a Product Image. All Promotional Images are 5K while the Product Images vary from 5K to 256K. The cache memory requirements are roughly 2.5 Mbytes per 100 Items, thus a 100,000-item scale factor would require the Image Server to cache 2,500 Mbytes of images. However the working set is roughly 70% of the total number of images, which is brought into the cache during the benchmark warm up phase. During the measurement interval the remainder of the images are accessed at a very slow rate. In addition to the book images, the Image Server may contain the web page navigational button images.

The performance metrics of interest are similar to the Web Server performance metrics. The CPU utilization and HTTP Bytes/sec data can be converted to Mbytes/sec (or Mbits/sec) per server or per CPU. The number of connections supported by the Image Server is computed by the Number of Browsers / Number of Image Servers since the Image Server usually only handles the non-secure web page images. If however the TPC-W solution has the Image Server also handling the navigational button images for the secure pages, an additional connection per user is required along with extra CPU cycles to handle the SSL encryption. Unfortunately it is difficult to determine if the Image Server is handling the button images from the secure pages without a detailed investigation of the source code presented in the FDR. For our back of the envelope calculations we ignore this fact and assume that the Image Server is at least handling one connection per emulated browser.

By examining the Image Server performance characteristics and computing some simple metrics one can estimate the relative performance of the Image Server in a particular TPC-W benchmark result. How does it compare with the network requirements of a real Ecommerce site? The key is that these computed performance metrics are derived while the Image Server is operating in a simulated Ecommerce workload.

4.3 DATABASER SERVER PERFORMANCE

The Database Server maintains all of the tables required by the TPC-W benchmark. The Database Server only responds to Web Server requests. The TPC-W workload does not require a database with a significant amount of memory nor does it require high performance network or disk I/O subsystems. The performance metrics of interest are the CPU utilization, HTTP bytes/sec and Disk I/O rate. Similar to the Web Server and Image Server, the HTTP bytes/sec can be converted to

Mbytes/sec or Mbits/sec. The overall network traffic is low compared to the front-end server network traffic. The Database Server is usually connected to the Web Servers via a 100 Mbit switch or hub. If the test sponsor has configured the SUT correctly the database should be running at or near 100% CPU utilization. Current network load balancing technology makes it extremely easy to add front-end servers vs. implementing a distributed database. Not that a distributed database will be implemented, just that whatever database configuration is used, it is easy to add front-end servers to push the Database Server(s) to 100% CPU utilization.

TPC-W emulates database configurations with a large amount of on-line storage (roughly 200 to 250 gigabytes per 1,000 WIPS) but with a low I/O rate (100 to 150 I/Os/sec per 1,000 WIPs). Current publications have used database configurations of 2 to 4 gigabytes of memory. If the Disk I/O rates are high given the database memory configuration, then this might be indicative of a sub optimal database server configuration. Similar to the Web Server and Image Server, the Database performance characteristics can be examined to obtain some idea of its performance characteristics while running in an Ecommerce environment.

4.4 WEB CACHE SERVER PERFORMANCE

The Web Cache's primary function is to reduce the load on the Database Server. Without some form of caching, the Database Server will be the bottleneck and reduce performance by a factor of 10 or more. As some Web Caches may be plug in appliances, the TPC-W specification does not require performance data for Web Caches that do not run application code or contain application disk storage. This is unfortunate since most TPC-W publications do not include performance statistics for the Web Cache whether the data exists or not.

With a bit of knowledge about the TPC-W benchmark along with some back of the envelope calculations, one can estimate the minimum network traffic supported by the Web Cache. The Web Cache is primarily used to cache the result frames for the Searches, the Best Sellers lists and the New Products lists. The bulk of the cached frames are from the Title and Author searches, (1/3 the scale factor), which are cached for an infinite amount of time. The Best Sellers, New Products and Subject Search result frames, (72 total), are cached for 30 seconds. If the Web Cache finds these frames older than 30 seconds, then it forwards a request to a Web Server to refresh the frame. The Web Server will forward a request to the Database Server, which executes a database query to re-compute the result set. The size of each frame varies a bit with the size of the scale factor, however an estimate for the average cached frame size is 10 to 12 Kbytes. The number of accesses to the Web cache can be computed from the web interaction mix and the throughput rate. For WIPS, 27% of the web page accesses are to the Searches, Best Sellers and New Product pages. Thus the WIPs rate times the 0.27 will result in the number of Web Cache accesses per second. At roughly 10 to 12 Kbytes per access the equation (WIPs *.27 * 10K)/1024 will give you an estimate for the minimum Mbytes/sec network traffic supported by the Web Cache. For the WIPSb mix the estimated network traffic is (WIPSb * .33 * 10K)/1024 while for WIPSo the estimated network traffic is (WIPSo * .14 * 10K)/1024.

The Web Cache is increasingly becoming more important as subsequent TPC-W publications push the envelope on what is cached. For example a small frame containing a list of promotional items for the Search, Best Sellers & New Products web pages can be cached. This saves a database lookup when generating the Promotional images for these web pages. A push cache technology can be used to cache the product detail page in order to save a database lookup for the Product Detail web interaction. Since the product information changes due to the infrequent Admin web interactions, the database must push updates made to the product tables to the cached pages residing in the Web Cache. Unfortunately, determining whether the Ecommerce solution is caching the promotional frames or the product pages can only be determined by a detailed analysis of the source code in the FDR. Without detailed knowledge of the Ecommerce solution, one can only compute the minimum value for the Web Cache network traffic.

The Web Cache does not cache any encrypted data, thus the number of connections supported by Web Cache is equal to the Number of Browsers.

.

5 IN SUMMARY

As shown, by having a little knowledge about the TPC-W benchmark and computing a few metrics, one can determine what is being benchmarked in a TPC-W publication. TPC-W gives one the performance characteristics of various servers in the context of a simulate Ecommerce workload. More importantly is that the companies who do run TPC-W have gained knowledge for configuring and tuning high performance Ecommerce environments. TPC-W forces one to load balance the Web Servers, move images off to a separate Image Server and offload database information to a Web Cache. TPC-W requires a network topology that supports several 100 Mbytes/sec of data. TPC-W benchmarking/experience provides a wealth of performance knowledge regarding Ecommerce solutions.