What Customers Need to Know About Benchmarks

An IDC White Paper

Analyst: Vernon Turner

Introduction

Today's data centers present a range of challenges for technology decision makers. Data center complexity is being driven by demands on business processes being integrated to manage and create more value-added services. As new hardware platforms are added to improve service levels, end users will decide which hardware vendors can deliver the required performance. To judge the performance of competitive hardware platforms, companies can either pretest all of their critical workloads on all components of the technology infrastructure or, if that opportunity doesn't exist, use industry-standard benchmarks created by independent software vendors or standards bodies.

Benchmarking has long been a trademark for vendors to differentiate their products when customers have had to make decisions based on head-to-head comparisons. However, as the Internet infrastructure matures and business processes attempt to develop seamless systems with a wide mix of server, storage, network, and application portfolios, single benchmark results need to be viewed with caution.

Systems technology has undergone significant changes during the past five years. Large multiprocessing servers are no longer the domain of the high-end server market. Clustered systems are being proposed to create an IT environment that allows groups of smaller servers to challenge larger servers for their workloads. Finally, ebusiness solutions are being developed to create end-to-end processing across both types of server landscapes. As a result of these changes, the growing transaction-processing complexity delivered by a mix of system types masks the end-user simplicity.

Management Summary

Benchmarks are useful for comparing relatively simple, static, and tightly managed server workloads. Today's online transaction processing systems are more complicated than the early "mainframe-class" platforms. Often the transactions used in older benchmarks do not reflect the typical or modern IT operational environment. In order to judge the performance of systems, a more holistic approach is needed. For example, examining a variety of benchmarks guards against the flaws of over optimization of a specific benchmark. In addition, too much

IDC
Analyze the Future

5 Speen Street • Framingham, MA 01701 • Phone (508)872-8200 • Fax (508)935-4015

emphasis has been placed on absolute performance per dollar (or currency) to compute a handful of transaction types, which do not reflect the typical IT operational environment.

Additionally, challenges for the customers arise when vendors no longer disclose system list prices of the technology used for the benchmark. Instead, the vendors opt to follow an inconsistent discount pricing scheme where technology pricing is the key denominator in the metric. Yet customers should know the discounting level used to recreate the original price.

Customers should also know how the transactions tested relate to their own online environment and the make-up of the technology base price. Additional disclosures should also be made, including:

- Transaction complexity
- Server partitioning rules
- Highly specific database functions
- Usage of clustering for high availability compared to performance
- Benchmark performance of a range of typical systems management functions
- Partitioning rules used on the online transaction processing database to establish if clustering techniques can be applicable to the customer's environment

Finally, customers should use benchmark results as only one of several key differentiating points and not in isolation. In addition, the weight that the results carry will vary greatly between IT environments. Customers need to press their technology provider for as much disclosure as possible to validate the fit between the controlled benchmark environment and the dynamic environment of their real data center. For true benchmark validity, customers should use a variety of independent software vendor (ISV) benchmarks, such as Oracle, Baan, or PeopleSoft, in conjunction with industry standards from organizations such as TPC and SPEC.

Current Use of Benchmarks

Evaluating Systems Performance

Benchmarking online transaction processing (OLTP) databases has been conducted for over 20 years. During that time many attempts have been made to create the most meaningful benchmark, and

perhaps the most famous and frequently used has been the TPC-C. The Transaction Processing Performance Council has published the TPC-C benchmark and has established a long tradition of evaluating a few important, but simplistic, OLTP transactions. Indeed, the TPC-C benchmark is narrowly defined and performed in a controlled environment to produce comparable results.

However, to effectively measure an OLTP environment in today's Internet infrastructure, disclosure about many baseline settings needs to be made to ensure that neutrality is being maintained. These settings include relating the transaction types that were tested to the relevance of the end-user environment and running databases on clustered platforms rather than a server with a single operating system.

Clarity of Clustering Versus Nonclustering Results

During the past decade, industry-standard volume server technologies have improved to allow "scale-out" deployments in the data center. Here, small servers are connected to create a network of processing power and are managed by clustering software. This is in contrast to traditional "scale-up" deployments, which tend to use functions developed for symmetrical multiprocessing (SMP) in larger servers housing all of the compute power. Both techniques have their advantages, which can only be exploited when the IT developer and operations staff fully understand the characteristics of the workload being run on each platform type.

However, benchmarking OLTP databases on a scale-out platform calls into question the validity of most transaction-based benchmarks. Typical OLTP databases are written for SMP-type servers and often on an ad hoc basis. This requires tightly integrated systems to take advantage of interprocessing connections within the server and not clustered server configurations.

To its credit, the TPC-C benchmark has attempted to show application or workload scalability by allowing the benchmark to be partitioned across a large cluster of small nodes connected with a network switch. This configuration may meet testing requirements to show that some transactions can be run across many servers. However, in reality, few mission-critical databases can be set up in this fashion.

Complete Small Transactions, but Run Them Hard and Often

The TPC-C transaction mix is made up of valid OLTP record types, such as posting an order or payment, checking the status of an order and processing the delivery of an order. However, these record types tend to be simple in nature, and, unlike similar live database functions, they consume little system resources. Unfortunately, this unrealistic, simple transaction constructs databases in a "federated" style, whereby each part of a federated database is a separate database. Once in production,

federated databases require careful, ongoing management for support, load balancing, maintenance, and recovery. Comparing clustering results that use this unrealistic loophole to SMP results is not a reasonable comparison. The results of these large clusters should be regarded very cautiously.

These Are Not Your Fathers' OLTP Databases

As vendors strive to higher performance on the TPC-C benchmark, they have also added database techniques that are not too common in most data centers. One such technique involves "hash indexes," which combine both data and index in a manner that benefits the benchmark.

The outcome of these optimizations and simple transactions is that the TPC-C has many more transactions per minute then a realistic data center. Finally, as system users demand ad hoc processing to meet Internet-driven requests, the real-world database transactions are becoming much more complicated than the comparable TPC-C benchmark transaction types.

Today's database administrators make sure that access, security, backup, and recovery are never compromised. The OLTP databases are the lifeline of all ebusiness environments, and, as such, their availability and response times are critical. In comparison to live OLTP databases, the performance of the TPC-C benchmark is challenged in several ways. For example, the TPC-C test makes a sole pass of a database. This forces the customer application to run a large number of active users to enable the database performance to reach a stressed level. This test would be contrary to a database administrator's goal.

In addition, disk storage management for OLTP databases has advanced significantly as a result of better storage management and cheaper disks. As a result, every database file is protected from hardware failure by incorporating some level of redundant array of independent/inexpensive disk (RAID) processing ranging from RAID-1 (data mirroring) through RAID-5 (data striping with parity).

Despite major improvements in reducing or eliminating data write penalties, storage applications are safer but do behave differently in a fully RAID environment. Currently, the TPC-C benchmark attempts to protect database log files that are important for journal processing and database recovery. It does not attempt to protect the critical database files in live OLTP databases that need to be protected. The TPC-C benchmark setup may be viewed as being behind today's minimum database management requirements.

In conjunction with other database administrators' storage management needs, it's important for OLTP databases to have backup and restore capabilities. Testing this function can determine the ability of both the storage platform and the server to multitask database functions. For example, the database could be writing to a file while attempting to recover or restore another file. As with the level of RAID processing

offered, the TPC-C benchmark is lacking in this important data write/recovery function.

Role of Benchmarks in the Buying Process

Benchmarking server products is a necessary process in today's cluttered technology market. IT management typically doesn't have the financial and technical resources to perform tests on every piece of server technology that it considers. Rather, organizations are using benchmark results as a way of screening out the outlying platforms to allow them to concentrate on just a handful of players.

In other words, the vendors see the benchmark results as a ticket or pass through the data center's glass doors. They also see the benchmark results as a way to create marketing and public relations announcements in either constructive or aggressive plans. After all, the benchmarks were very expensive to perform, and the vendor needs to derive some form of payback.

But do customers really pay any attention to the tables published by the testing agency? If the answer is yes, then what weight do the results carry during their decision-making processes? Customers who take the time to follow a managed request for proposal (RFP) are usually more serious about the results from a testing agency.

However, these customers are also savvy to the fact that results like the TPC-C benchmarks are exactly that — they are benchmarks, and the decision is not based solely on this lone reference point. Often, these proposals will ask for a myriad of benchmarks to cover the platform's true capabilities. This is in response to the ultimate use of these servers as they are deployed to run multiple applications (often multiple databases) or workloads. The goal of these questions is to give the customer a feel for the server's ability to allow complete workload integration and not an isolated function. In conclusion, customers are going to combine the benchmark results with other topics, such as server reliability, ISV portfolios, financing arrangements, research and development plans, service, support, and other customer testimonials.

Can Customers Trust the Benchmark Results?

Customers can sign onto a testing agency's Web site to see the benchmark results. However, should they feel like they are buying a car from the dealer and being shown the price the dealer paid for the car? The two events are similar except that the car dealer often looks for the printout to show either the bigger sales amount to demonstrate the discount you are getting or the lowest price to show you how good a job you have done squeezing the margins out of the deal.

When it comes to the benchmarking world, the testing agency has no control over the price validity of the server submitted by the vendor. Both the testing agency and the reader have no indication if the cost of the servers to perform the tests was discounted or not. However, smart

customers can go further and calculate server prices used to generate the number of transactions per dollar metric. Indeed, it would appear that some vendors are using discounted prices instead of list prices to improve their overall standings.

Conclusion

Customers need holistic and unbiased information about a server platform's technical ability to complete their procurement and deployment tasks. Customers can obtain this information by looking at a variety of benchmarks beyond the TPC-C type, such as those from the ISV community. In doing so, discussions with the end users will take on different meanings for different levels of management. For example, financially orientated managers will want to understand the true list price for components as a starting point, and they will not use prices that are affected by benchmark discounting games. Likewise, chief systems architects may be more concerned about operational preparation as a more influencing factor. Benchmarks can help each of them reach their objectives.

However, the Internet infrastructure has changed how customers evaluate the results from benchmarks. When considering a server platform to perform and support OLTP database functions, customers are aware of the increased application complexity, including:

- Expanding the transaction mix
- Changing database structures to support nonstop and highly available requirements
- Supporting multiple access and response times of the application

When customers look to metrics, such as the TPC-C, they may be concerned that it may not provide them with enough assurance that the vendors had the client's best interest at heart. Logical or operational improvements to the benchmark could include:

- More robust storage management testing
- Broader transaction access methods (such as adding transaction types that don't make heavy use of hash indices)
- Server list prices and not discounted prices
- Balance of I/O with processor configuration testing
- Establishment of accurate OLTP testing for the scale-up versus the scale-out environments

When customers look at benchmark results, they should consider them and their subsequent rankings in the context they are worth. If the vendor performance on a single benchmark is the sole basis for making a purchasing decision, they are not considering all of their options and, therefore, are being too influenced by both the testing agency and the vendor.

Finally, customers buy servers because they meet their actual business requirements. These requirements are made up of a very large ecosystem, where benchmark results are very much on the periphery.

What Customers Need to Know About Benchmarks

Before you commit to any system based on the benchmark results, you should consider the following issues:

1) Is there any difference between the price of the hardware used in the benchmark and the one being proposed?
2) Identify the relevant benchmark tests to your environment. Where do their synergies exist?
3) Use a variety of benchmark results; develop your own total cost of ownership model, including implementation and staff costs, for use in the evaluation process.
4) Rank the importance of benchmarking results in the proposed solution.

NORTH AMERICA

Corporate Headquarters
5 Speen Street
Framingham, MA 01701
508-872-8200

IDC Canada
36 Toronto Street, Suite 950
Toronto, Ontario
Canada M5C2C5
416-369-0033

IDC Irvine
18831 Von Karman Ave, Ste 200
Irvine, CA 92612
949-250-1960

IDC Mountain View
2131 Landings Drive
Mountain View, CA 94043
650-691-0500

IDC New Jersey
120 Wood Ave South, Suite 509
Iselin, NJ 08830
732-632-9222

IDC New York
2 Park Avenue
Suite 1505
New York, NY 10016
212-726-0900

IDC Texas
100 Congress Ave, Suite 2000
Austin, TX 78701
512-469-6333

IDC Washington
8304 Professional Hill Drive
Fairfax, VA 22031
703-280-5161

ASIA/PACIFIC

IDC Asia/Pacific (Hong Kong)
12/Floor, St. John's Building, 33 Garden Road
Central, Hong Kong
852-2530-3831

IDC Asia/Pacific (Singapore)
71 Bencoolen Street, #02-01
Singapore 189643
65-226-0330

IDC Australia
Level 4, 76 Berry Street
North Sydney
NSW 2060, Australia
61-2-9922-5300

IDC China
Room 611, Beijing Times Square,
88 West Chang'an Avenue, Beijing,
P.R. China, 100031
86-10-8391-3456

IDC (India) Limited
Cyber House
B-35, Sector 32 - Institutional
Gurgaon - 122002
Haryana, India
91-124-6381673 to 80

IDC Japan
10F The Itoyama Tower
3-7-18, Mita Minato-ku
Tokyo 108-0073, Japan
81-3-5440-3400

IDC Korea Ltd
Suite 704, Korea Trade Center
159-1, Samsung-Dong, Kangnam-Ku
Seoul, Korea 135-729
82-2-55-14380

IDC Malaysia
Suite 13-03, Level 13, Wisma KiaPeng
No. 3, Jalan Kia Peng
50450 Kuala Lumpur, Malaysia
6-03-2163 3715

IDC New Zealand
Level 7, 246 Queen Street
Auckland, New Zealand
64-9-309-8252

IDC Philippines
7F, SEDCCO 1Bldg
Rada Street Corner
Legaspi Street
Legaspi Village
Makati City, Philippines
632-894-4808

IDC Taiwan Ltd.
10F, 31
Jen-Ai Rd, Sec 4,
Taipei 106, Taiwan, R.O.C.
886-2-2731-7288

IDC Thailand
27 Soi Charoen Nakorn 14
Charoen Nakorn Road, Klongtonsai
Klongsan Bangkok 10600, Thailand
66-2-439-4591-2

IDC Vietnam
37 Ton Duc Thang Street
Unit 1606
District-1 Hochiminh City Vietnam
84-8-910-1235

LATIN AMERICA

IDC Miami
Latin America Headquarters
8200 NW 41 Street
Suite 300
Miami, FL 33126
305-267-2616

IDC Argentina
Trends Consulting
Rivadavia 413, 4th Floor, Suite 6
C1002AAC, Buenos Aires, Argentina
54-11-4343-8899

IDC Brasil
Alameda Ribeirão Preto, 130 cj 41
01331-000 São Paulo
SP Brazil
55-11-253-7869

International Data Corp. Chile
Luis Thayer Ojeda 166 Piso 12
Providencia, Santiago 9, Chile
56-2-231-0111

IDC Colombia
Carrera 40 # 103-78
Bogota, Colombia
571-533-2326

IDC Mexico
Select - IDC
Av. Nuevo Leon No. 54 Desp. 501
Col. Hipodromo, Condesa
C.P. 06100 Mexico, D.F.
52-5-256-1426

IDC Venezuela
Calle Guaicapuro
Edif. Torre Seguros Alianza
Piso 6, Ofc. 6-D, El Rosal
Caracas 1060, Venezuela
58-2-951-3270

EUROPE, MIDDLE EAST, AND AFRICA

IDC Austria
c/o Loisel, Spiel, Zach Consulting
Mayerhofgasse 6
A-1040 Vienna, Austria
43-1-50-50-900

IDC Benelux (Belgium)
29 Avenue Louis Gribaumont
B-1150 Brussels, Belgium
32-2-779-46-04

IDC Benelux (The Netherlands)
A. Fokkerweg 1
1059 CM Amsterdam
The Netherlands
31-20-669-2721

IDC Central Europe (ECE)
Male Namesti 13
Praha 1 110 00, Czech Republic
420-2-2142-3140

IDC Central Europe (Germany)
Nibelungenplatz 3, 11th Floor
60318 Frankfurt, Germany
49-69-90502-0

IDC Central Europe (Switzerland)
Niederlassung Züerich
WTC, Leutschenbachstrasse 95
CH - 8050 Züerich
Switzerland
41-1-307-1000

IDC Egypt
39 Iraq Street
Mohandesseen, Cairo, Egypt
20-2-336-7355

IDC France
Immeuble La Fayette
2, Place des Vosges, Cedex 65
92051 Paris la Defense 5, France
33-14-904-8000

IDC Hungary
Nador utca 23, 5th Floor
H-1051 Budapest, Hungary
36-1-473-2370

IDC Israel
4 Gershon Street
Tel Aviv 67017, Israel
972-3-5611660

IDC Italy
Viale Monza, 14
20127 Milano, Italy
390-2-284-571

IDC Nigeria
House 2, 'C' Close, 403 Road, 4th Avenue
New Extension, Festac Town
Lagos, Nigeria
234-1-883585

IDC Nordic (Denmark)
Jagtvej 169B
DK-2100 Copenhagen, Denmark
45-39-162222

IDC Nordic (Finland)
Jarrumiehenkatu 2
FIN-00520
Helsinki, Finland
358-9-8770-466

IDC Nordic (Sweden)
Box 1096 Kistagången 21
S-164 25 Kista, Sweden
46-8-751-0415

IDC Poland/ProMarket
Wrobla 43
02-736 Warsaw, Poland
48-22-754-0518

IDC Portugal
Av. Antonio Serpa, 36 Piso 9
1050-027 Lisbon
Portugal
351-21-796-5487

IDC Russia
c/o PX Post, RDS 186
Ulitsa Zorge 10
Moscow 125525
Russian Federation
7-501-929-9959

IDC South Africa
c/o BMI-TechKnowledge
3rd Floor, 356 Rivonia Blvd.
PO Box 4603, Rivonia, 2128
South Africa
27-11-803-6412

IDC Spain
Ochandiano, 6
Centro Empresarial El Plantio
28023 Madrid
34-91-7080007

IDC Turkey
Tevfik Erdonmez Sok. 2/1 Gul Apt.
Kat 9D; 46 Esentepe
Istanbul, Turkey
90-212-275-0995

IDC U.K.
British Standards House
389 Chiswick High Road
London W4 4AE
United Kingdom
44-20-8987-7100

IDC is the foremost global market intelligence and advisory firm helping clients gain insight into technology and ebusiness trends to develop sound business strategies. Using a combination of rigorous primary research, in-depth analysis, and client interaction, IDC forecasts worldwide markets and trends to deliver dependable service and client advice. More than 700 analysts in 43 countries provide global research with local content. IDC's customers comprise the world's leading IT suppliers, IT organizations, ebusiness companies and the financial community. Additional information can be found at www.idc.com.

IDC is a division of IDG, the world's leading IT media, research and exposition company.

Sponsored by Sun Microsystems Inc.
01C3036SYSTEM3036
August 2001

Part Number FE1628-0

IDC
5 Speen Street • Framingham, MA 01701
(508) 872-8200 • Fax (508) 935-4015 •
www.idc.com