

SPOTLIGHT ON MICROSOFT SQL SERVER

*Results from the 2003
Winter Corp. TopTen™ Program*

SQL SERVER, long considered a departmental database not scalable enough for enterprise application deployment, was featured prominently in the 2003 results of the Winter Corp. TopTen Program survey of the world's largest and most heavily used databases. Microsoft's efforts to overcome the obstacle of a "not ready for prime time" reputation are apparently succeeding.

Background

The Winter Corp. TopTen Program identifies the world's largest databases and recognizes the database practitioners who are defining the frontiers of database scalability. The 2003 campaign was the fifth such program executed by Winter Corp., a leading independent authority on large database technology and implementation. Winter Corp. TopTen Program findings and research have been featured in *The New York Times*, *Information Week*, *ComputerWorld*, *eWeek*, *Intelligent Enterprise*, *DM Review*, *CIO*, and other leading industry publications.

WinterCorp. TopTen Program participants, who are owners and managers of large databases, complete a brief survey on database size, performance, and operating environment. Participants are required to validate their submissions by running queries and executing other procedures developed jointly by the database vendor and Winter Corp. The survey results are categorized by usage—online transaction processing (OLTP) or decision support system (DSS)—and operating system platform. Awards are presented for the leaders in data volume (compressed and uncompressed), number of rows/records, and workload. In the 2003 campaign, there were 21 TopTen lists—the largest and most heavily used databases within the different usage categories.

Program Results

SQL SERVER REPRESENTATION IN THE 2003 TOPTEN PROGRAM

Data from the TopTen Program confirmed the emergence of the Microsoft Windows operating environment, specifically the SQL Server database, as a legitimate platform for large-scale databases. The 2003 TopTen program was open to Windows systems in production that contain 500 or more Gigabytes in the database. The number of SQL Server participants grew substantially since 2001 for both OLTP and DSS databases. The number of OLTP SQL respondents in 2003 increased 140%, representing 32% of the total pool of validated OLTP respondents—second only to the Oracle database platform, which occurs predominantly on Unix OLTP systems (*Figure 1*). The number of DSS SQL respondents increased 43% since 2001, representing 10% of the 2003 total of validated DSS respondents. The survey identified one Windows site with an OLTP database larger than 5 Terabytes—at Verizon Communications

— and this is a SQL Server implementation. These results are indicative of the substantial progress Microsoft has made to date in the production OLTP space.

LARGER SQL SERVER DATABASES— AND PROJECTIONS OF FURTHER GROWTH

The average size of the SQL Server databases for OLTP respondents increased by 675%—from 236 GB to 1.8 TB—whereas the average size for SQL Server DSS respondents increased by 205%—from 293 GB to 893 GB (Table 1). These percentages are inherently overstated since the minimum size for qualification increased from 100 GB in 2001 to 500 GB in 2003. But as noted earlier, the number of respondents increased significantly in 2003 despite the more stringent qualification criteria. The largest SQL Server database was the OLTP database of 5.3 TB at Verizon Communications, whereas the largest production SQL Server DSS database identified in the survey was 1.6 TB of data at Arlight Systems LLC.


**Table 1: Growth in Database Size
for SQL Server Participants**

	2001	2003	— PROJECTED —	
			2005	2006
OLTP				
Size (average)	236 GB	1.8 TB	2.5 TB	2.9 TB
Growth rate	—	675%	37%	56%
DSS				
Size (average)	293 GB	893 GB	2.1 TB	3.0 TB
Growth rate	—	205%	132%	236%

Source: Winter Corporation, 2004

SQL Server OLTP respondents projected continued growth in their already terabyte-plus (on average) databases, estimating a 37% increase by 2005 (to 2.5 TB) and 56% by 2006 (to 2.9 TB). But, the projections of growth for SQL Server DSS systems were particularly impressive—132% by 2005 (to 2.1 TB) and 236% by 2006 (to 3 TB). These projections indicate a level of confidence that users have in the scalability of the platform. The average growth projections for DSS from 2003 to 2006 on the leading database platforms are shown in Figure 2.

Figure 1: TopTen OLTP Participants by Database Platform




Source: Winter Corporation, 2004

ADDITIONAL DIMENSIONS OF GROWTH THAT AFFECT SCALABILITY REQUIREMENTS

The size of the database is the first of three factors affecting scalability requirements that result from the increased integration of information in the enterprise. The second dimension, the magnitude of the workload measured by transaction rate for OLTP and by concurrent, inflight queries for DSS, reflects the growing population of database users. The majority of SQL Server users in the survey projected significant increases in OLTP (transactions per second) and DSS (concurrent queries) workload in 2005 and 2006. For OLTP, users projected the transaction rate to increase 114% by 2005 and 138% by 2006. The DSS SQL Server users projected query workloads to increase 136% by 2005 and 147% by 2006. In each case, respondents expected the workload volume to more than double between 2003 and 2005 and to continue to grow in 2006.

The third dimension affecting scalability requirements is the nature of the workload. More subject areas and the exploratory nature of data analysis have led to increased complexity of individual queries and transactions. There is no explicit measure of query complexity included in this round of the TopTen Survey program. It is interesting that, whereas the majority (71%) of DSS respondents using SQL Server in 2001 characterized the application as an Independent Data Mart, the vast majority (80%) of SQL Server DSS respondents in 2003 characterized the application as an Enterprise Data

Figure 2: Projected Growth by 2006 in DSS Database Size by Platform



Source: Winter Corporation, 2004

Warehouse (EDW). This change reflects the increase in mission criticality of the DSS systems deployed on Windows and SQL Server.

SQL SERVER PRIZE WINNERS

In the 2003 Winter Corp. TopTen Program, 17 unique users of Windows/SQL Server platforms received "TopTen winner" recognition out of 69 total unique winners. Four of the 21 Grand Prize winners represented SQL Server implementations. In total, of 192 awards in the program, 51 were conferred upon SQL Server users.

SQL SERVER TOPTEN AWARD HIGHLIGHTS INCLUDE:

- For OLTP, SQL Server captured 87% (26 of 30) in the Windows categories: database size, number of rows, and peak workload.
- For OLTP, SQL Server captured 33% (10 of 30) in the All Environments categories, the only Windows customers in any "All" lists. This is the first time that a Windows-based system was an All Environment TopTen winner.
- SQL Server captured 60% (36 of 60) in OLTP categories in which it participated.
- For DSS, SQL Server captured 47% (14 of 30) in the Windows categories.
- For DSS, a SQL Server user, Stratapult Enterprises, was the first Windows-based TopTen winner in an All Environment category (peak workload).
- SQL Server captured 25% (15 of 60) in DSS categories in which it participated.

As an example of SQL Server's leadership as a Windows database platform for OLTP, the TopTen lists for database size and peak workload, Windows only, OLTP are included as *Tables 2 and 3*.

The Verizon Communications 5.3 TB database (first in *Table 2*) ranked sixth in the Database Size (All Environments) list for OLTP systems. Three SQL Server customers appear in the TopTen list for Number of Rows (All Environments) for OLTP (*Table 4*).

Table 2: TopTen Database Size, Windows Platform, OLTP

Organization	Size in TB	DBMS	DBMS Vendor	System Vendor	System Architecture	Storage Vendor
Verizon Communications	5.3	SQL Server	Microsoft	HP	SMP	EMC
HP	3.2	Oracle	Oracle	HP	SMP	HP
Lucent Technologies	2.9	SQL Server	Microsoft	Dell	Cluster	Dell
Commander Communications	2.9	SQL Server	Microsoft	Dell	Cluster	Dell
Verizon Communications	2.2	SQL Server	Microsoft	HP	SMP	EMC
Verizon Communications	1.8	SQL Server	Microsoft	HP	SMP	EMC
HP	1.7	Oracle	Oracle	HP	SMP	HP
Internet Auction	1.7	SQL Server	Microsoft	Unisys	SMP	EMC
HP	1.5	Oracle	Oracle	HP	SMP	HP
Verizon Data Services	1.5	SQL Server	Microsoft	HP	SMP	EMC

Source: Winter Corporation, 2004

They Keep Growing and Growing...

Since the completion of the 2003 Winter Corp. TopTen Program, many SQL Server customers have reported that their production databases are continuing to increase in size. For example, the largest DSS database at Edgars Consolidated Stores (EdCon) of South Africa is up to 4.5 TB—from just over 1 TB in the 2003 survey. Additionally, several other new DSS deployments are expected to be included in the 2005 TopTen Program. Those include Sage Telecom's Call Detail Record data warehouse at 3 TB, Shoprite of South Africa's Retail data warehouse at 5.3 TB, and the United States Department of Agriculture's 11 TB Geospatial data warehouse. Microsoft and Winter Corp. encourage SQL Server customers with more than a Terabyte of data to nominate their databases for the 2005 TopTen Program.

Conclusions

The price-performance advantages of SQL Server- and Windows-based applications running on Intel-based servers have been acknowledged for years in comparison with Unix/RISC and mainframe-based systems. But, until recently, many practitioners questioned whether Windows and SQL Server were capable of satisfying high-end database requirements. Some of those doubts were mitigated by Windows benchmarks, both TPC and application results, such as SAP benchmarks. But, standard benchmarks address theoretical scalability; in practice, real-world implementations are the most fundamental barometer of ability to scale.

The results of the 2003 Winter Corp. TopTen Program, which surveys the world's largest and most heavily used databases, provide validation that Windows and SQL Server have gained greater acceptance by users for challenging database implementations. SQL Server appears prominently in the Winter Corp. TopTen (All Environments) lists for OLTP systems, and user projections foretell an increasing presence for Microsoft SQL Server in the TopTen lists (All Environments) for DSS systems moving forward. OLTP and DSS are often mission-critical applications, implying that, for many of the largest database implementations, the reliability and availability characteristics of Windows and SQL Server have progressed as well. Future Winter Corp. research will assess whether Microsoft is able to continue the significant momentum demonstrated by the 2003 TopTen Program results.

Table 3: TopTen Peak Workload, Windows Platform, OLTP

Organization	Tps	DBMS	DBMS Vendor	System Vendor	System Architecture	Storage Vendor
Anonymous	4,010	SQL Server	Microsoft	HP	SMP	HP
Internet Auction	3,634	SQL Server	Microsoft	Unisys	SMP	EMC
Commander Communications	2,604	SQL Server	Microsoft	Dell	Cluster	Dell
Scottish and Southern Energy	2,152	SQL Server	Microsoft	HP	SMP	HP
2001OUTLET Co.	1,204	SQL Server	Microsoft	HP	SMP	EMC
Anonymous	1,200	SQL Server	Microsoft	HP	Cluster	EMC
Verizon Communications	206	SQL Server	Microsoft	HP	SMP	EMC
Verizon Communications	131	SQL Server	Microsoft	HP	SMP	EMC
Verizon Communications	128	SQL Server	Microsoft	HP	SMP	EMC
Verizon Communications	113	SQL Server	Microsoft	HP	SMP	EMC

Source: Winter Corporation, 2004

Table 4: TopTen Number of Rows, All Environments, OLTP

Organization	Rows (mil)	DBMS	DBMS Vendor	System Vendor	System Architecture	Storage Vendor
UPS	41,985	DB2 for z/OS	IBM	IBM	SMP	EMC
Caixa Econômica Federal	34,152	CA-IDMS	Computer Associates	IBM	SMP	EMC
Verizon Communications	33,341	SQL Server	Microsoft	HP	SMP	EMC
Commander Communications	17,033	SQL Server	Microsoft	Dell	Cluster	Dell
Bureau of Customs & Border Protection	16,663	CA-Datcom	Computer Associates	IBM	SMP	Hitachi
GevityHR, Inc.	12,703	Oracle	Oracle	HP	SMP	HP
BT plc	11,898	CA-IDMS	Computer Associates	IBM	SMP	EMC
American Electric Power	10,316	DB2 for z/OS	IBM	IBM	SMP	EMC
Verizon Communications	9,157	SQL Server	Microsoft	HP	SMP	EMC
Turkcell Iletisim Hizmetleri A.S.	8,147	Oracle	Oracle	Sun	SMP	EMC

Source: Winter Corporation, 2004

ACKNOWLEDGEMENTS

Microsoft would like to thank its customers who participated in the 2003 Winter Corp. TopTen Program. The participants included:

- 2001OUTLET Co. Ltd.
- Arclight Systems, LLC
- Clalit
- Commander Communications Ltd.
- DataQuick
- Edgars Consolidated Stores
- Fair Isaac (NAREX, Inc.)
- First Premier Bankcard
- Gomez
- Internet Auction Co., Ltd.
- Lucent Technologies
- Ordina
- Scottish and Southern Energy plc
- Stratapult Inc.
- Verizon Communications.

*A leading center of expertise in very large databases,
Winter Corporation provides services in
consulting, research, architecture and engineering.*

*We help users and vendors understand their opportunities;
select their database and data warehouse platforms;
define and measure the value of their strategies, architectures and products;
plan, architect and design their implementations;
and manage their scalability, performance and availability issues.*

Our focus is databases near, at and beyond the frontier of database scalability.



WINTER CORPORATION

411 WAVERLEY OAKS ROAD, SUITE 327
WALTHAM, MA 02452
617-695-1800

visit us at www.wintercorp.com