RESEARCH REPORT

WinterCorp

www.wintercorp.com

WINTER CORPORATION

PRACTICES AT THE FRONTIERS OF DATABASE SCALABILITY 2006

TopTen Program Members' Report

The Survey of the World's Largest Databases

Practices at the Frontiers of Database Scalability 2006

TopTen Program Members' Report

Kathy Auerbach

September 2006

411 WAVERLEY OAKS ROAD, SUITE 327
WALTHAM, MA 02452
781-642-0300

1 Highlights

This report presents key findings from the 2005 WinterCorp TopTen™ Program, the world's only independent survey of the largest and most heavily used databases in operation today. The report identifies the products, platforms and architectures used by 175 terabyte-scale database practitioners who span the globe and represent all major industries. Unfiltered by vendors and validated by WinterCorp's staff, the survey results provide visibility into today's leading database solutions and best practices—enabling executive owners and sponsors of large-scale data management systems to make informed, fact-based technology decisions.

- The largest data warehouse in the 2005 program contained 100.4 TB of data, a growth of 243% in two years. The average size of data warehouses grew 82% to 9.5 TB.
- The largest transaction processing system reached 23.1 TB, a 26% gain since 2003. Average size of operational systems increased 4% to 4.1 TB.
- Data warehousing and transaction processing are in different stages of their life cycles. DW is in a robust growth phase while transaction processing has matured and is expanding gradually.
- Over the past two years, the number of UNIX data warehouses with more than 10 TB of data grew 175%. Windows systems tripled their presence in data warehousing and for the first time, appeared in the TopTen Program list for DW size.
- Use of Windows for transaction processing rose 129% since 2003. The number of UNIX OLTP systems nearly doubled, with major gains in the number of the largest systems. The percent of z/OS systems used for transaction processing declined.
- The 2005 program reported the initial appearance of Linux as a large database environment for both data warehousing, 9%, and online transaction processing, 5%.
- 90% of implementations surveyed employed a centralized database architecture.
- Addressing the growing demands for and benefits of a multi-platform offering, Oracle was the only vendor with participating databases on UNIX, Windows and Linux.
- Oracle RAC, the cluster option of Oracle Database, was established as a large DBMS choice for data warehouses as well as operational systems, 10% each.
- Nearly half of all data warehouse sites, 47%, ran Oracle Database. SQL Server, the second most widely used DW DBMS, bolstered its data warehouse presence significantly. Oracle and Daytona controlled the largest data warehouses.
- Oracle Database/Oracle RAC owned the highest percent of transaction processing implementations, 59%. DB2 and Oracle managed the largest OLTP systems.
- The data warehouse leader, 100.4 TB of data, was managed by Oracle Database and supported by UNIX on a Fujitsu Siemens PRIMEPOWER 1500 server. The data was stored in an EMC DMX disk system.
- The largest transaction processing database, 23.1 TB, was controlled by DB2 for z/OS and hosted on an IBM zSeries 990-308. Storage was provided by an IBM DS8300 2107-9A2 system and Hitachi 9980V disk array.
- In row count, the OLTP leader, a DB2 for z/OS site, doubled the 2003 mark to 90 billion. Among data warehouses, an HP NonStop SQL database on NonStop OS contained almost three trillion rows, followed by Daytona, a UNIX system, with nearly two trillion.
- HP manufactured nearly half of all servers, 46%, in the program; IBM, 19%, and Sun, 18%, comprised the second tier. Forty-three percent, 43%, of the largest databases ran on HP.

- SMP remained the prevailing architectural choice for servers hosting large databases, 71%.
- The average storage capacity per system was 23.2 TB, an increase of 10% in the past two years.
- Top workload for transaction processing exceeded one billion SQL statements in an hour, 1.13 billion. The mark was achieved by DB2 for z/OS on an IBM zSeries 990 and IBM storage systems.
- The highest performing data warehouse completed 28.8 million SQL statements per hour. The DBMS was DB2 for UNIX, supported by an IBM pSeries 670 and EMC storage.
- Respondents reported that the benefits delivered by large database projects significantly outweighed the problems encountered. On a scale of 1 (low) to 5 (high), the average of benefits gained was 4.1 compared to problems experienced, 3.1.
- The frontiers of transaction processing in 2008 will be 25 TB of data, a growth of 74%. Maximum workload will grow 174% and approach 200 million SQL statements per hour.
- By 2008, the frontiers for data warehouse systems will jump by 300% in size to reach 200 TB of data, and double in workload to exceed six million hourly SQL statements.

TopTen Programs show that over the past decade, databases have grown significantly in both size and power. In fact, participating databases have become *increasingly* larger and faster. Based on projections from the 2005 respondents and related input, WinterCorp expects that growth rates will continue to escalate. In the next three to four years, operational systems could double in workload and data warehouses could triple in size. This growth will push the frontiers of database scalability far beyond the current boundaries, into territory that seemed virtually unimaginable just a few short years ago.

2 Table of Contents

1 Highlights 5

2 Table of Contents 7

3 Overview and Methodology 9

4 Demographics 10

4.1 By industry and geographic area, where are the respondents located? 10

4.2 Is there a correlation between industry and database usage? 11

5 Operating System 12

5.1 By usage, what operating systems support large databases? 12

5.2 By database size, what operating systems support large transaction processing systems? 13

5.3 By database size, what operating systems support large data warehouse systems? 14

6 Database Architecture and Size 15

6.1 What architectures are large databases using, 2001 to 2005? 16

6.2 By usage, how much have large databases grown since 2001? 17

6.3 What size are the largest data warehouses, 1998 to 2005? 18

6.4 By usage and operating system, what are the ranges in size of the largest databases in 2003 and 2005? 19

6.5 By usage and operating system, what are the ranges in number of rows of the largest databases in 2003 and 2005? 20

6.6 By usage, what size is the largest table in each database? 21

7 Database Management System (DBMS) Overview 22

7.1 What database management systems support large transaction processing databases in 2003 and 2005? 22

7.2 What database management systems support large data warehouse systems in 2003 and 2005? 23

7.3 What are the largest OLTP databases supported by the leading database management systems? 24

7.4 What are the largest DW databases supported by the leading database management systems? 25

7.5 By usage and operating system, which DBMS vendors support large databases? 26

7.6 By database size, what DBMS are transaction processing systems using? 27

7.7 By database size, what DBMS are data warehouse systems using? 28

8 Hardware 29
8.1 What is the architecture of the server (or servers) that support the database? 29
8.2 By operating system, which server vendors support large databases? 30
8.3 What are the largest transaction processing systems supported by the leading server platforms? 31
8.4 What are the largest data warehouse systems supported by the leading server platforms? 32
8.5 What is the capacity of the systems that store large databases? 33
8.6 What is the total storage capacity of the 2005 TopTen Program participants? 34
9 Workload 35
9.1 By usage, what are the workloads of the most heavily used databases? 35
10 Project Assessment 36
10.1 By usage, what business benefits have you achieved with this large database project? . . . 36
10.2 By industry, what business benefits have you achieved with this large database project? 37
10.3 By usage, what problems have you experienced with this large database project? 38
10.4 By industry, what problems have you experienced with this large database project? 39
11 Frontier 40
11.1 What are the frontiers for transaction processing systems, 2005 through 2008? 40
11.2 What are the frontiers for data warehouse systems, 2005 through 2008? 41
12 Conclusion 42
13 TopTen Lists 43
14 Questionnaire 64
15 Appendix 68

3 Overview and Methodology

WinterCorp is an independent consulting firm that specializes in the performance and scalability of terabyte- and petabyte-scale data management systems throughout their lifecycle. For the past twelve years, we have conducted quantitative surveys of the world's leading databases. Program findings are based on respondent input and disclose the products, platforms and practices of the leading implementations. TopTen Programs recognize the organizations that operate the leading databases as well as the vendors who support them.

The 2005 TopTen Program, the sixth run by WinterCorp, tracked four database metrics:

- Database size: amount of actual data (includes user data, aggregates, summaries and indices; excludes freespace and redundancy)
- Normalized data volume: the amount of data managed by the database system, prior to transformations that may reduce data size
- Number of rows, records or objects managed by the database
- Workload: peak number of SQL statements or database operations completed in an hour

WinterCorp publicized the 2005 TopTen Program on its web site, and via mailings and outbound promotional initiatives. We worked closely with our industry partners in their efforts to solicit potential participants.

One of the distinguishing features of WinterCorp TopTen Programs is that respondents are required to validate their database metrics. This process ensures that program findings are based on actual achievements, not best estimates. Validation requires that participants run scripts developed by WinterCorp and their DBMS vendor. The scripts produce statistics about the size and performance of the participant's database under a specific DBMS. Alternatively, some participants submit system-generated documentation that supports their database metrics.

WinterCorp launched the 2005 TopTen Program on February 15th, 2005 and accepted surveys until August 15th, 2005. To qualify for the 2005 program, a database had to be in production and contain a minimum of one terabyte of data. Validation was finalized in early September and the award winners were announced on September 8th. We received 249 completed surveys, of which 175 met the stringent validation requirements, a 24% increase since the last program. The surveys originated from 22 countries and represented the major DBMS, server and storage vendors.

WinterCorp published 42 lists of the world's largest and best performing databases (up to 11 award winners per list). The lists were defined by the four metrics above, plus:

Primary Usage:

- Business Intelligence, Data Warehousing, Decision Support, or
- Operational System, Transaction Processing, e-Commerce, or
- Scientific, Content Stores, Other

Platform:

- All Environments
- UNIX Only
- Windows Only
- Linux Only

Program findings revealed the opinions of participants and information about their database systems and not the industry at large.

The 2005 TopTen Program was sponsored by HP, IBM, Microsoft, Oracle, Sun and Sybase. Sponsorship conferred promotional benefits but did not influence the findings in any way.

Note: in the graphs that present data from past TopTen Programs, the findings include only those sites that met the 2005 criterion of a minimum of one terabyte of data.

4 Demographics

4.1 BY INDUSTRY AND GEOGRAPHIC AREA, WHERE ARE THE RESPONDENTS LOCATED?

Large Databases Were Most Prevalent in TelCom, Finance and IT

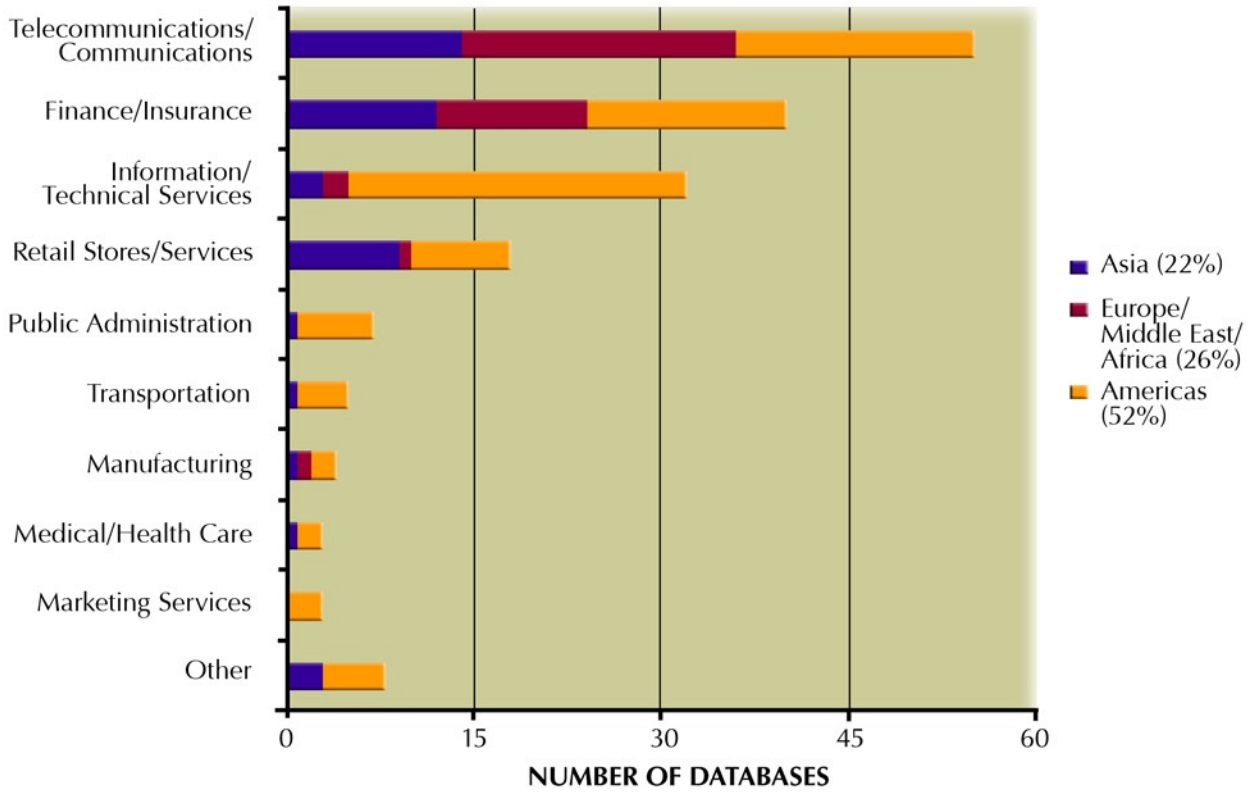

N = 175

Other = Hotel, Education, Media, Gas Industry, Real Estate and Science

- Just over one-half of respondents were in the Americas. About one-quarter hailed from Europe/Middle East/Africa and Asia, each.
- Nearly a third of databases in the program, 55, supported Telecommunications/ Communications. Of these, 22 were situated in Europe/Middle East/Africa.
- About a quarter of databases, 40, ran Finance/Insurance applications, distributed relatively evenly among the three geographic areas.
- Information/Technical Services databases, 32, accounted for nearly one-fifth of sites in the program. The overwhelming majority of them, 27, were located in the Americas.

4.2 IS THERE A CORRELATION BETWEEN INDUSTRY AND DATABASE USAGE?

Database Usage was Not Industry-Specific

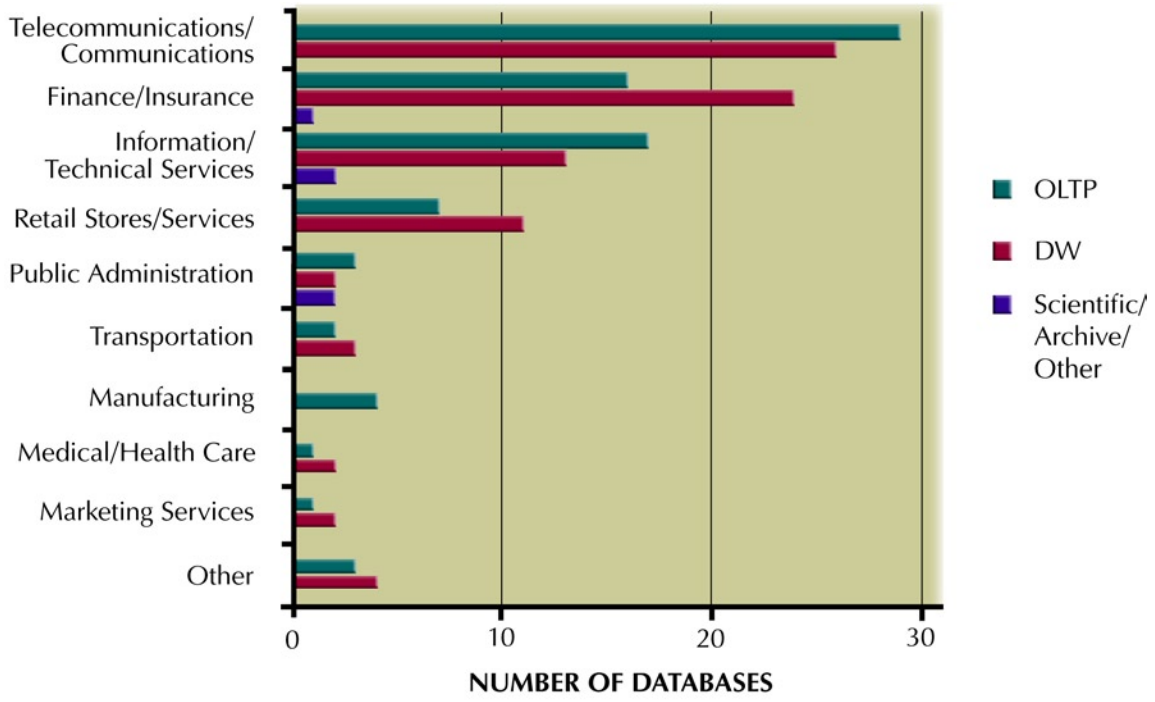

N = 175

- Databases in Telecommunications/Communications and Information/Technical Services were nearly evenly divided between OLTP and DW.
- More than half of Finance/Insurance and Retails Stores/Services databases were data warehouses.
- All Manufacturing databases supported online transaction processing.
- The few Scientific/Archive/Other databases in the program were mostly found in Information/Technical Services and Public Administration.

5 Operating System

z/OS systems increased in number but declined in percent from the 2003 to the 2005 program. The number of UNIX environments grew significantly for both data warehousing and online transaction processing. Windows strengthened its presence in almost all size categories, appearing as a platform for the largest data warehouses for the first time. Linux was established as a large database environment, in particular for emerging large databases.

5.1 BY USAGE, WHAT OPERATING SYSTEMS SUPPORT LARGE DATABASES?

Use of Windows and UNIX Rose; Linux Established Presence

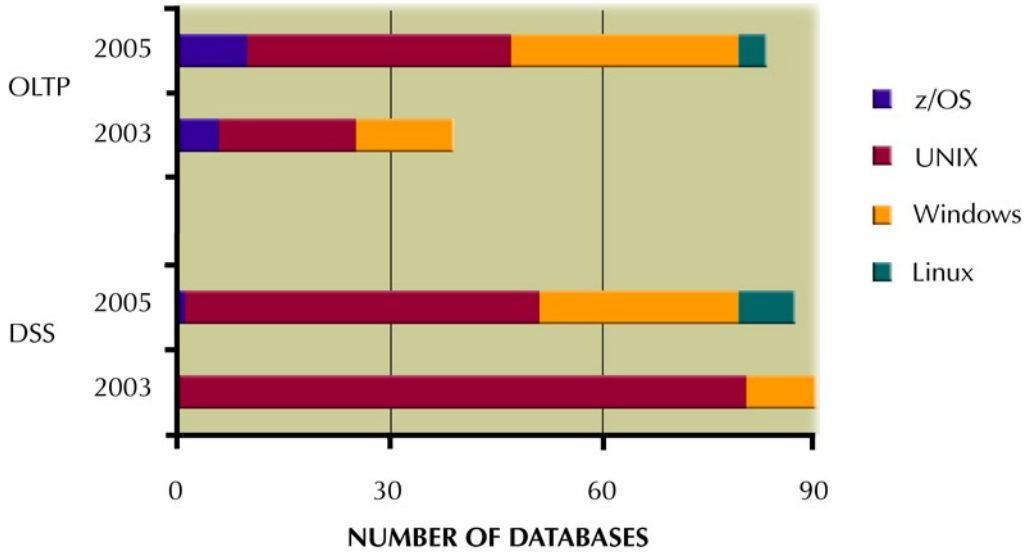

N = 170

- The percent of mainframes that hosted large databases fell from the 2003 to 2005 program.
- The number of UNIX OLTP systems nearly doubled from 2003 to 2005.
- The incidence of Windows-based large databases grew moderately in transaction processing and substantially in data warehousing.
- The 2005 program confirmed Linux as a large database platform.
- UNIX data warehouses were underrepresented because Teradata did not participate in the 2005 TopTen Program.

5.2 BY DATABASE SIZE, WHAT OPERATING SYSTEMS SUPPORT LARGE TRANSACTION PROCESSING SYSTEMS?

UNIX and Windows OLTP Systems Increased in Size and Number

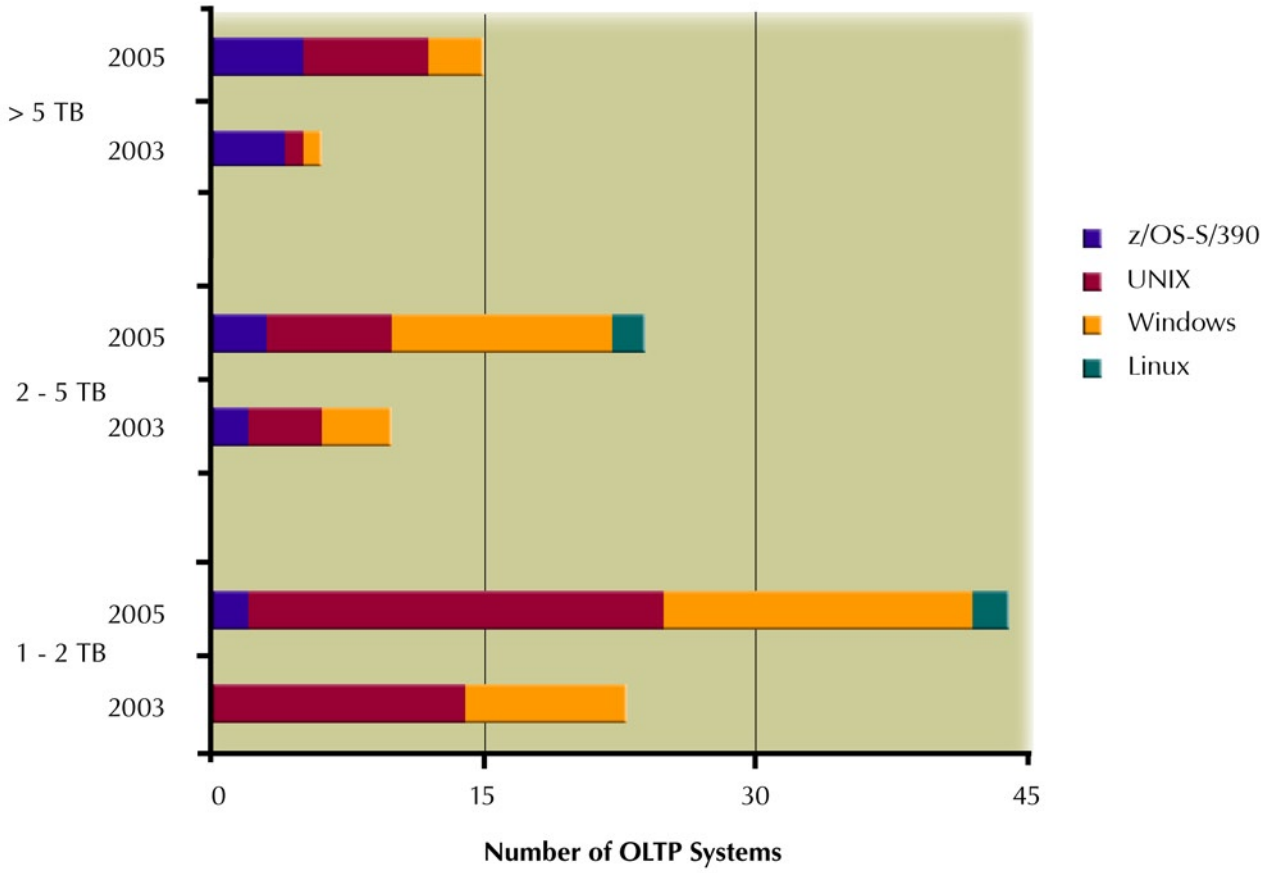

2005 N = 83
2003 N = 39

- Although the number of z/OS OLTP systems in the program nearly doubled from 2003 to 2005, the percent of z/OS systems surveyed declined from 15% to 11%.
- Emerging and mid-sized UNIX OLTP systems increased in number but dropped in percent. The largest OLTP UNIX systems jumped in both number and percent.
- The presence of Windows OLTP systems grew in all size categories, particularly in the 2 - 5 TB range.
- Linux made minor inroads among OLTP systems with up to 5 TB of data.

5.3 BY DATABASE SIZE, WHAT OPERATING SYSTEMS SUPPORT LARGE DATA WAREHOUSE SYSTEMS?

Windows Data Warehouses Gained Ground on Leader, UNIX

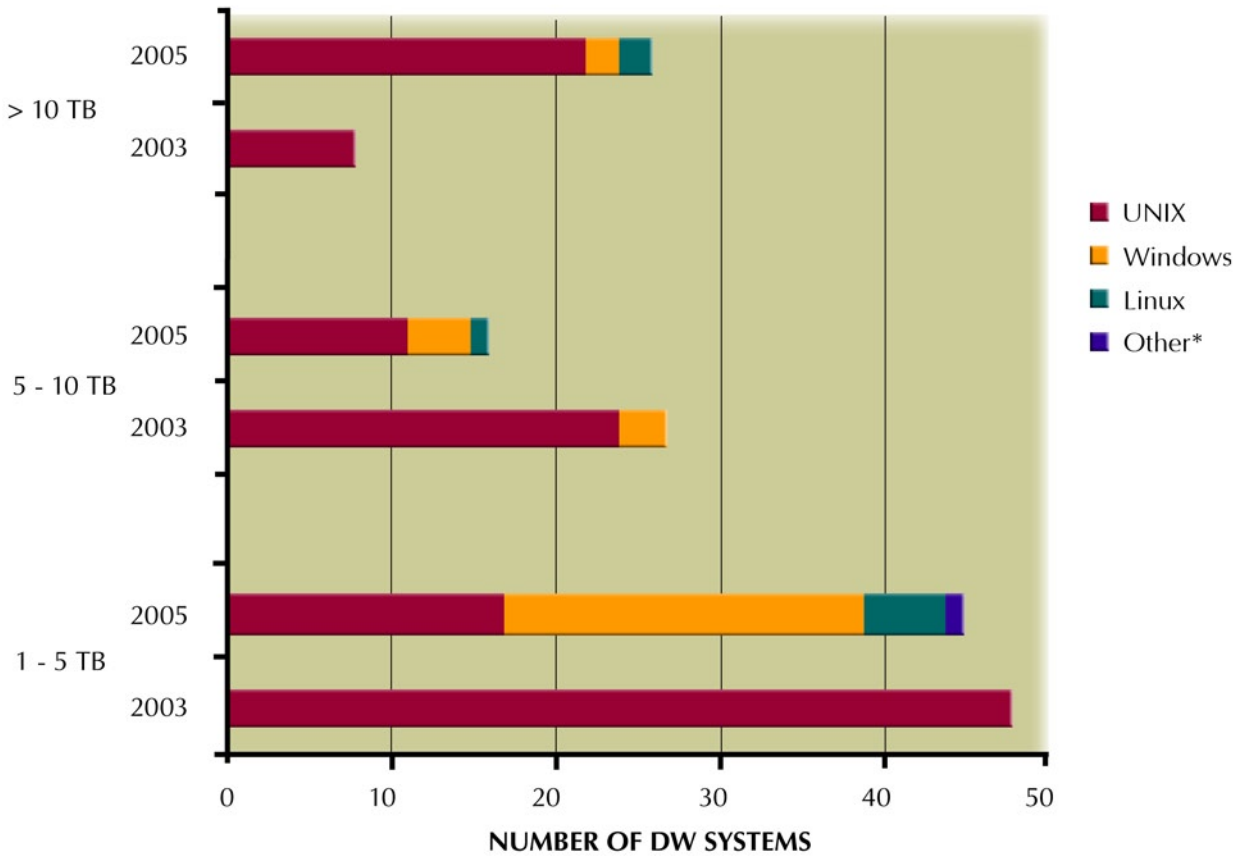

*Other = Nonstop OS

- Windows jumped in number and percent of emerging data warehouses, 1 – 5 TB data. For the first time, Windows appeared as a platform for the largest data warehouses.
- The number of UNIX data warehouses with more than 10 TB of data rose significantly, offsetting declines in the two other size ranges.
- Linux established a presence as a data warehouse platform in 2005, most notably for systems with up to 5 TB of data.
- There was one data warehouse on NonStop OS in the 2005 program.

6 Database Architecture and Size

Centralized further strengthened its position as the prevailing database architecture, found at 90% of surveyed databases.

The 2005 program showed that both operational and data warehouse systems continue to grow larger, in response to factors such as demands for near real-time data infrastructures, better interoperability and more widespread globalization.

However, the growth rates for online transaction processing and data warehouse systems differed sharply. While size of the largest operational system doubled from 2001 to 2005, the size of the average system was essentially the same. In contrast, the leading data warehouse ballooned nearly seven-fold over the past four years, while average size tripled. In fact, figure 6.3 shows that over the past eight years, the sizes of the largest data warehouses in TopTen Programs have grown exponentially.

The largest transaction processing in the program, 23 TB, was hosted on z/OS, as in the last program. The biggest UNIX OLTP system was 202% larger than its 2003 predecessor. The top Windows operational system managed 50% more data than the 2003 frontrunner.

The data warehouse leader, a 100 TB UNIX implementation, tripled the 2003 high watermark. The topmost Windows data warehouse doubled in size in two years. There were no z/OS data warehouses in the 2005 program.

Row count underwent enormous growth from 2003 to 2005. The OLTP leader, a z/OS site, doubled to 90 billion rows. Close behind was an 83 billion row database, representing more than a five-fold increase for the leading UNIX system.

Among data warehouses, a NonStop OS system contained *almost three trillion rows* and a UNIX database comprised nearly two trillion. These figures outdistanced the Linux and Windows row leaders 30 and 40 times, respectively.

Table size expanded noticeably in the past two years. Average size grew about 16% among databases surveyed. The five largest DW tables averaged nearly 20 TB, a 42% gain in two years. The five largest OLTP tables measured about 2 TB, a 28% increase.

6.1 WHAT ARCHITECTURES ARE LARGE DATABASES USING, 2001 TO 2005?

Centralized Remained Dominant Database Architecture

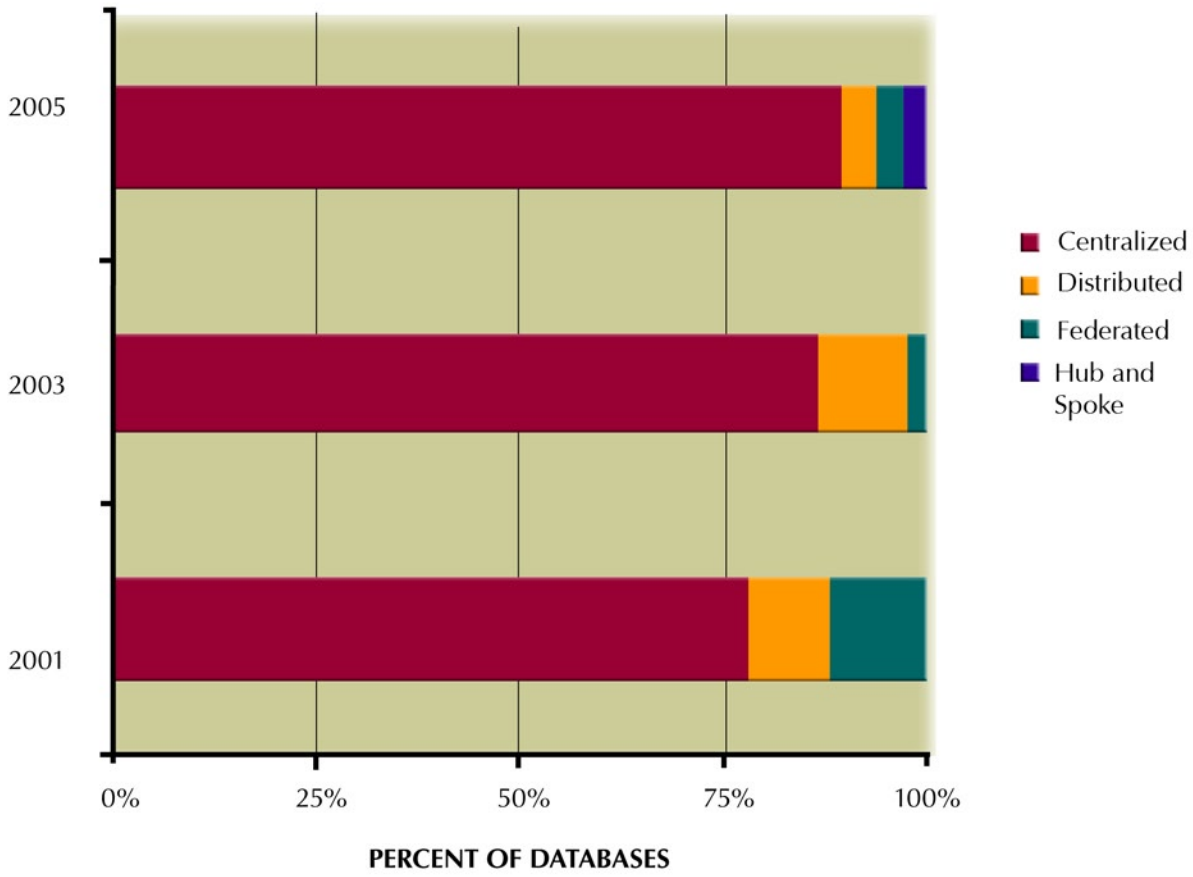

2005 N = 175
2003 N = 129
2001 N = 59

- Centralized database architecture became even more widespread in 2005. Its presence has increased steadily, from 78% in 2001 to 87% in 2003 to 90% in 2005.
- Distributed and Federated architectures, each found at about 11% of databases surveyed in 2001, declined to only 3-4%, each, of 2005 implementations.
- The 2005 program was the first survey to track Hub and Spoke, found at 3% of participating databases.

6.2 BY USAGE, HOW MUCH HAVE LARGE DATABASES GROWN SINCE 2001?

OLTP Growth has Slowed

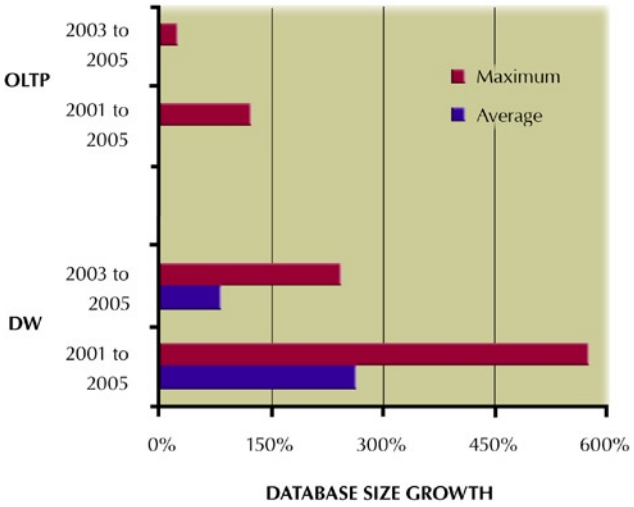

2005 N = 170
2003 N = 129
2001 N = 59

- From 2003 to 2005, the largest OLTP database increased 26% while the average OLTP database grew only marginally, 4%, from 3.0 to 3.1 TB. A similar profile occurred from 2001 to 2005: the largest transaction processing system doubled in size whereas average size remained constant, 1% increase.

Data Warehouses Soared to 100 TB

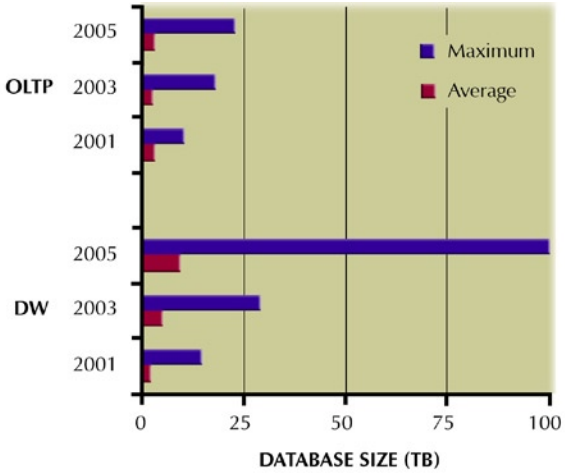

- Growth of data warehouses has been dramatic. From 2003 to 2005, the largest went from 29.2 to 100.4 TB and the average went from 5.2 TB to 9.5 TB. Over the past four years, average size rose 243% while maximum size grew an astonishing 578%.

6.3 WHAT SIZE ARE THE LARGEST DATA WAREHOUSES, 1998 TO 2005?

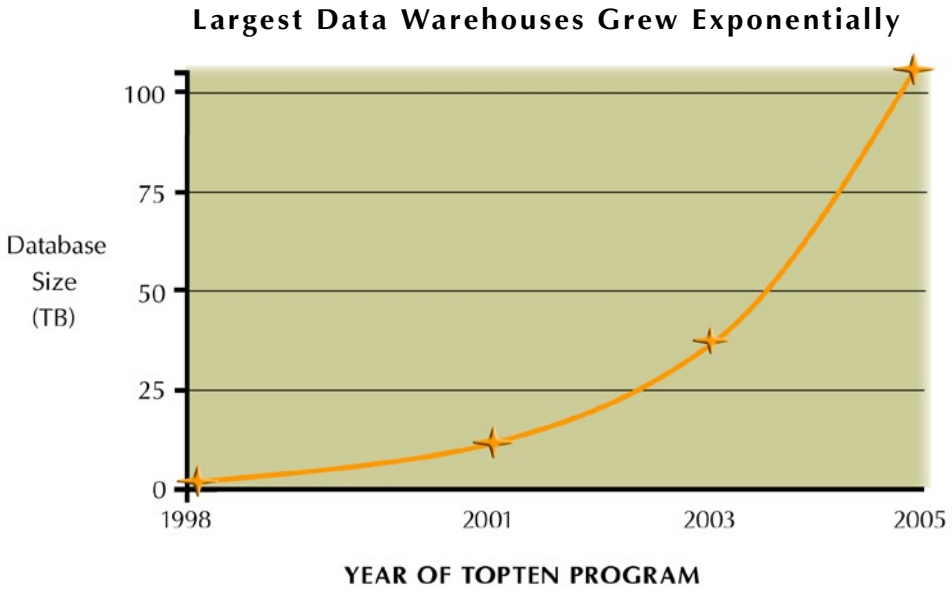

- In the five years between the 1998 and 2003 programs, the largest data warehouse grew from 5 TB to 30 TB, an exponential increase.
- Between 2001 and 2005, a period of just four years, the largest data warehouse also grew exponentially, from 10 TB to 100 TB.

6.4 BY USAGE AND OPERATING SYSTEM, WHAT ARE THE RANGES IN SIZE OF THE LARGEST DATABASES IN 2003 AND 2005?

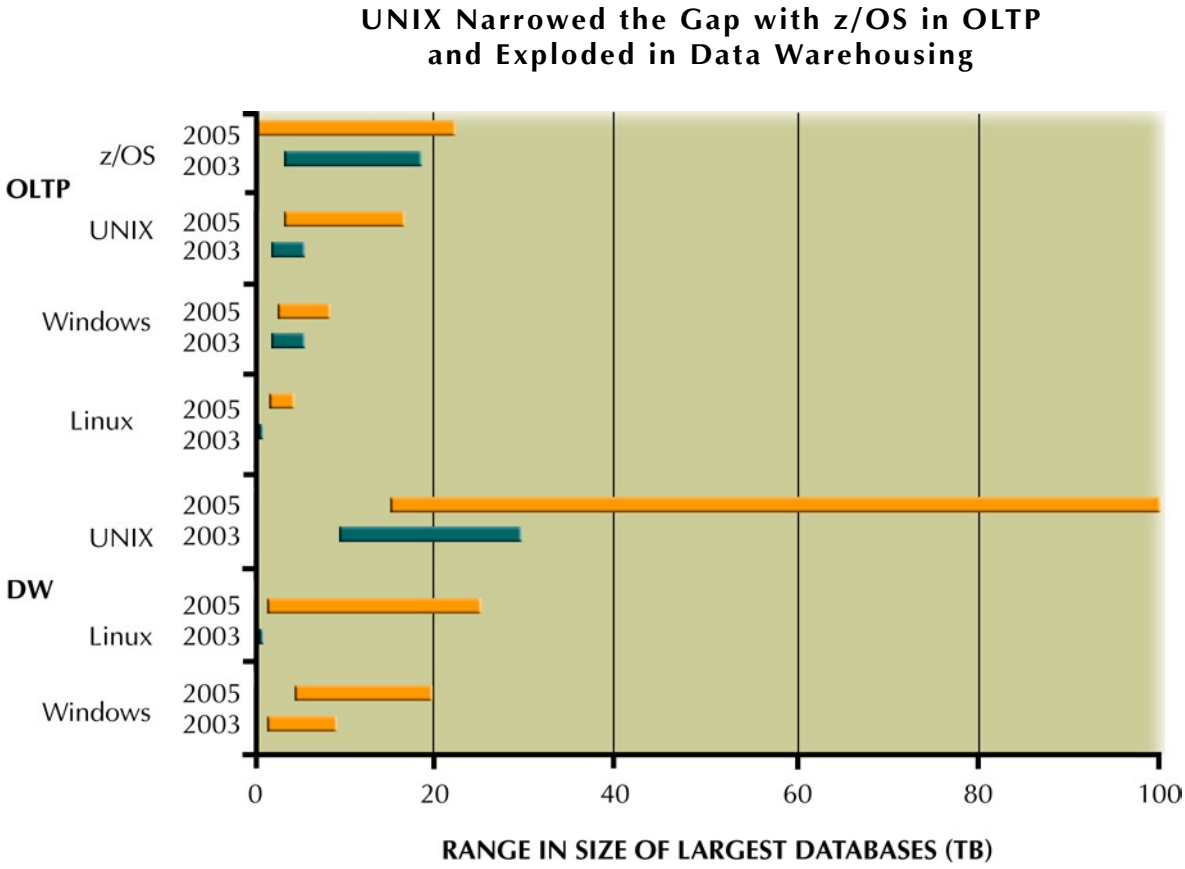

- The largest operational system in 2005, 23.1 TB, was implemented on z/OS. The biggest increase since 2003 went to the UNIX leader, which tripled in size. There were no mainframe data warehouses in the 2005 program and only one in 2003.
- Among data warehouses, the size leader contained 100.4 TB of data. It was implemented on UNIX, the platform that supported the largest growth increase since 2003, 243%.
- Sizes of Windows systems grew steadily, doubling in data warehousing and growing by more than half in OLTP.
- Linux leaders supported almost 25 TB of warehouse data and 4 TB of operational data.

6.5 BY USAGE AND OPERATING SYSTEM, WHAT ARE THE RANGES IN NUMBER OF ROWS OF THE LARGEST DATABASES IN 2003 AND 2005?

NonStop OS Data Warehouse Topped Rows

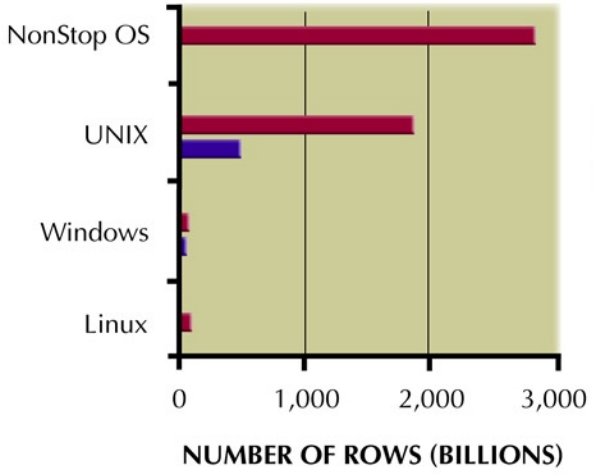

2005 N = 52
2003 N = 88

z/OS Maintained Slight Edge in OLTP Row Leadership

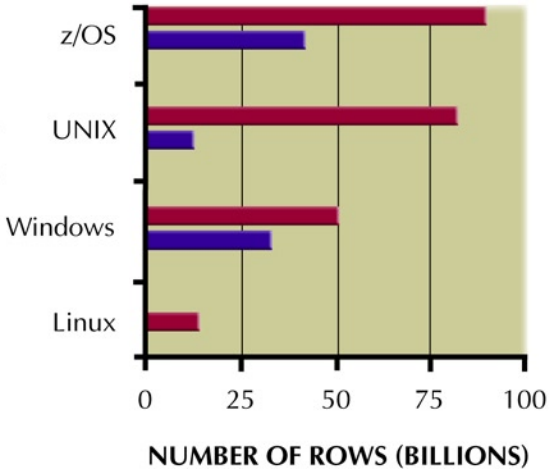

2005 N = 62
2003 N = 37

- Among data warehouse, a NonStop OS system hosted a database with almost *three trillion rows*, 2,848 billion.
- The leading UNIX warehouse nearly quadrupled the 2003 figure to 1,883 rows. The largest Linux and Windows data warehouses contained 97 billion and 73 billion rows, respectively.
- As in 2003, z/OS supported the 2005 OLTP row frontrunner. The system contained 90 billion rows, twice as many as two years ago.
- The top UNIX operational system multiplied the 2003 row mark 6½ times to 83 million. The biggest Windows OLTP implementation comprised 51 billion rows, a 50% gain since 2003.

6.6 BY USAGE, WHAT SIZE IS THE LARGEST TABLE IN EACH DATABASE?

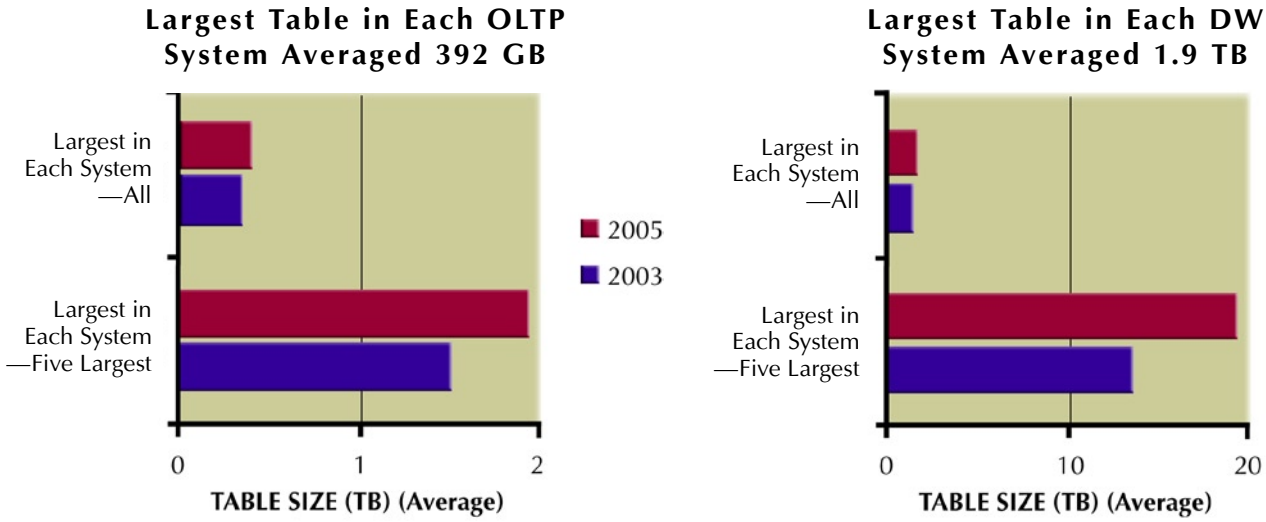

- The average size of the largest table in each OLTP database was 392 GB, an increase of 15% since 2003.
- Among the largest tables in each OLTP database, the five largest ranged from 872 GB to 5 TB. Their average size was 1.9 TB, a growth of 28% in two years.
- Data warehouse tables significantly outweighed their OLTP counterparts. The average size of the largest table in each DW database was 1.9 TB, an increase of 17% since 2003.
- Among the largest tables in each DW database, the five largest ranged from 5.8 to 25.1 TB. Their average size was 19.5 TB, a jump of 42% in two years.

7 *Database Management System (DBMS) Overview*

Oracle database products had the strongest presence of any DBMS in the 2005 TopTen Program. Oracle was the only vendor with participating UNIX-, Windows- and Linux-based systems. Oracle Database also ran the greatest number of OLTP implementations in the program. The largest OLTP systems favored Oracle and DB2. SQL Server led among mid-sized databases while Oracle managed the majority of emerging large systems.

Oracle Database also led in number of data warehouses - nearly half of those surveyed - and was the most widely used DBMS for mid- and large-size systems. SQL Server achieved the largest increase in penetration since the 2003 program. Usage of SQL Server for data warehousing rose from 4% to 24%, making it the second most widely used DW solution in the 2005 program.

Oracle RAC secured a position as a large database DBMS in the 2005 program, implemented on approximately one in ten of participating OLTP and DW systems, each.

DB2 managed the largest OLTP system, Land Registry, 23.1 TB of data. Oracle Database ran the second largest operational system, 16.4 TB. Oracle controlled the largest data warehouse, Yahoo!, with 100.4 TB of data. Daytona helmed the second largest DW, 98.9 TB of data.

7.1 WHAT DATABASE MANAGEMENT SYSTEMS SUPPORT LARGE TRANSACTION PROCESSING DATABASES IN 2003 AND 2005?

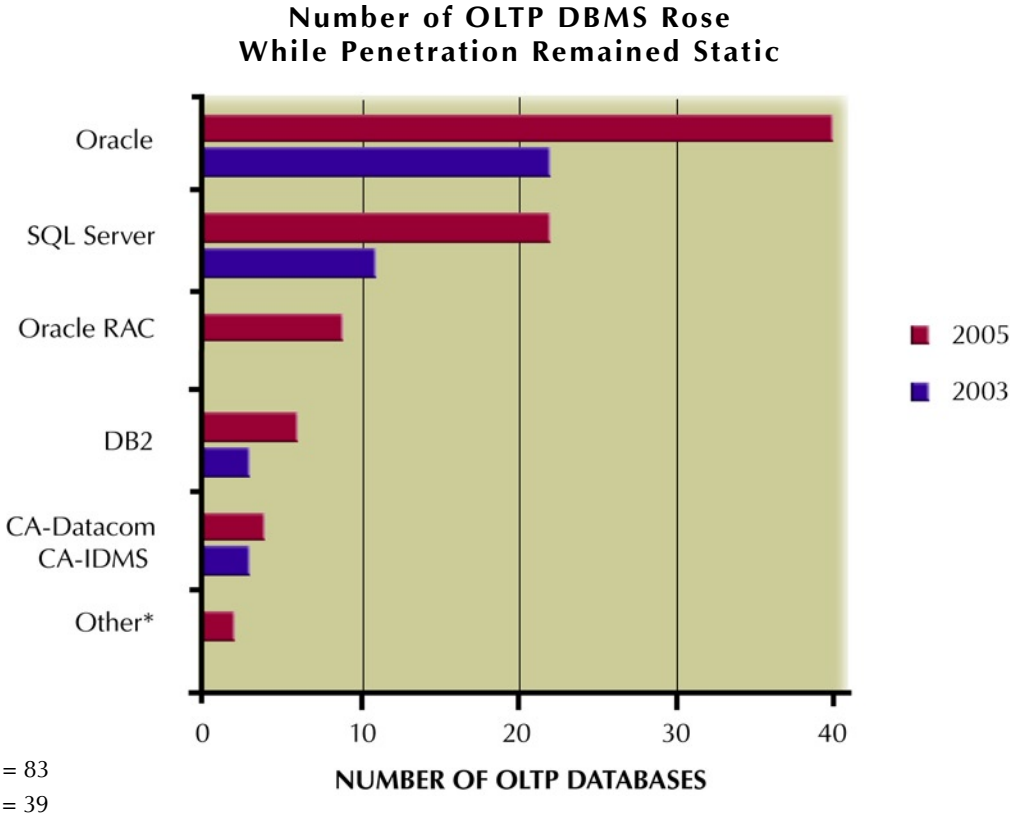

2005 N = 83
2003 N = 39

- Most transaction processing DBMS maintained about the same presence in the 2005 survey pool as in 2003. The exceptions were Oracle Database and CA-IDMS/Datacom, which ceded some penetration to Oracle RAC.

7.2 WHAT DATABASE MANAGEMENT SYSTEMS SUPPORT LARGE DATA WAREHOUSE SYSTEMS IN 2003 AND 2005?

SQL Server Increased Data Warehousing Presence

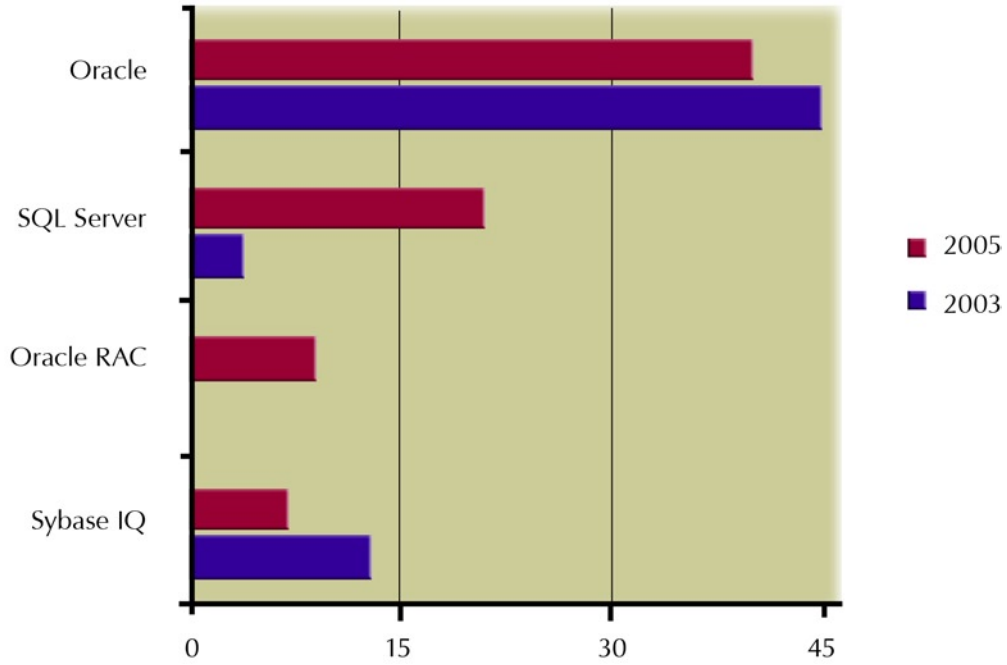

2005 N = 77
2003 N = 62

- About three-quarters of participating data warehouses in 2003 were Oracle Database and one-quarter were Sybase IQ.
- By 2005, the database breakdown had broadened to include Oracle, one-half, SQL Server, one-quarter, and Oracle RAC and Sybase, about one-tenth, each.
- Teradata did not participate in the 2005 TopTen Program.

7.3 WHAT ARE THE LARGEST OLTP DATABASES SUPPORTED BY THE LEADING DATABASE MANAGEMENT SYSTEMS?

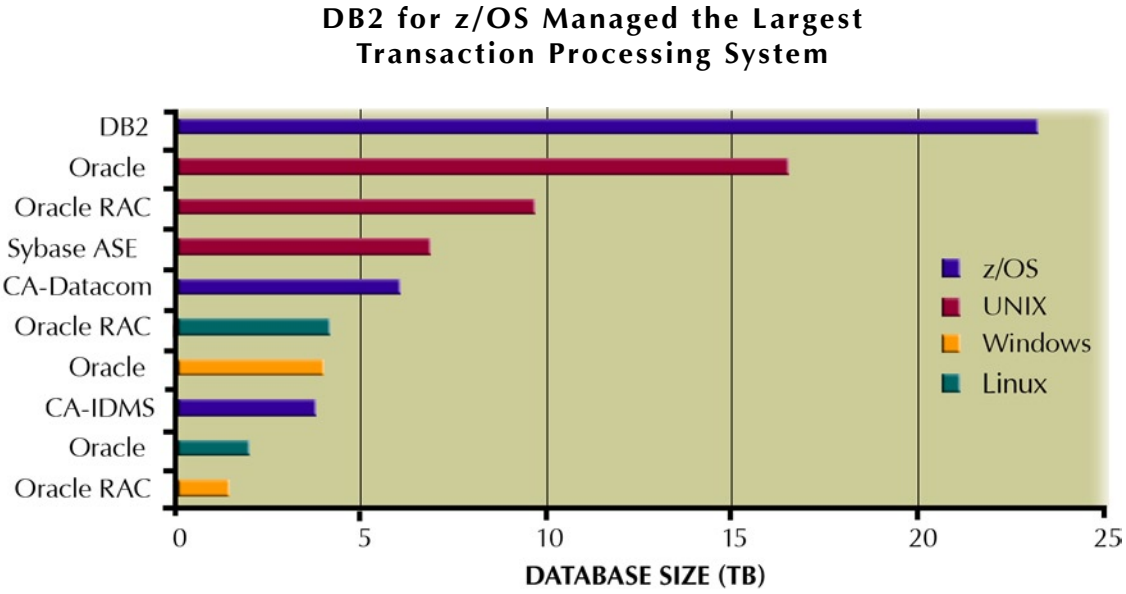

- DB2 for z/OS helmed the largest OLTP system, 23.1 TB, at the Land Registry. The government-owned system maintains property and land ownership records for England and Wales. The centralized database supported more than 10 applications and comprised at least 1,000 tables. An IBM eServer zSeries 990-308 hosted the database, most of which was stored in an IBM DS8300 2107-9A2 system.
- Oracle Database managed the second largest OLTP system, 16.4 TB of data, at the U.S. Patent and Trademark Office. The UNIX-based system supported about five applications and has been in production about 2 ½ years. There was 5 TB of data in the largest table, making it a large database onto itself. At peak, the database supported nearly 30,000 concurrent users. An IBM eServer pSeries 690 housed the database, which utilized an EMC DX2000 storage solution.
- Oracle RAC controlled the third largest OLTP implementation at Elsevier. The 9.67 TB database produces and manages more than 20,000 scientific publications and services available worldwide. The UNIX-based system contained nearly 8 TB of uncompressed user data. A Sun Fire V1280 server handled the system, which was stored in an IBM ES 800 disk array.

7.4 WHAT ARE THE LARGEST DW DATABASES SUPPORTED BY THE LEADING DATABASE MANAGEMENT SYSTEMS?

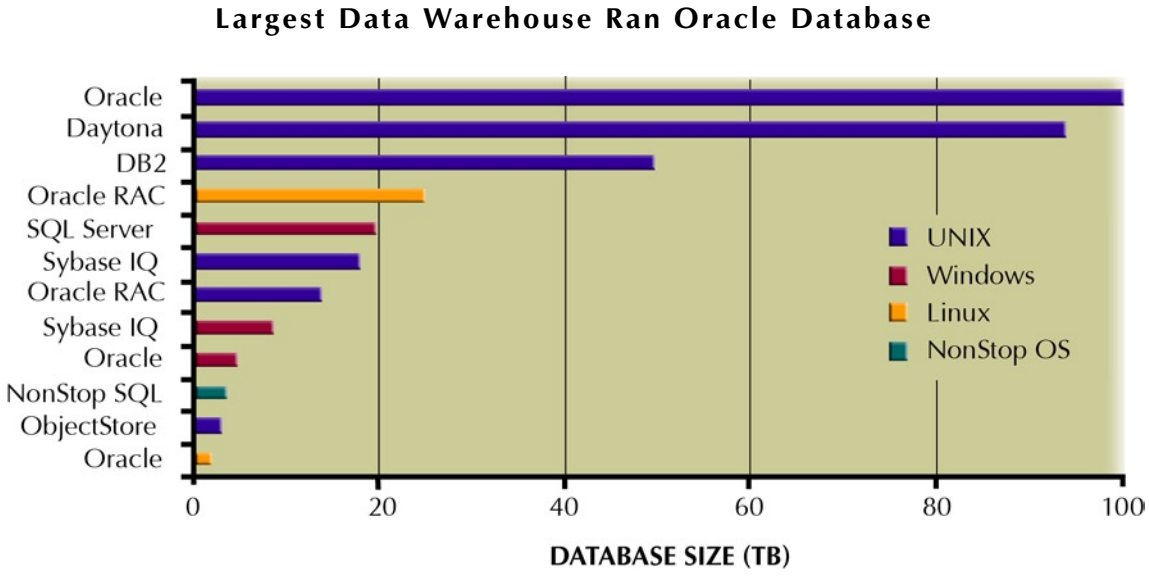

- Oracle Database controlled the largest data warehouse, 100.4 TB, at Yahoo! The UNIX-based system supports one of the most powerful web search engines, which allows users to tailor results to their particular interests. The data warehouse has been in production more than five years and is updated continuously. It was supported by a Fujitsu Siemens PRIMEPOWER 1500 and stored in an EMC DMX disk array.
- Daytona hosted the second largest data warehouse, 98.9 TB, at AT&T. The federated database contains call detail records and manages business processes. At peak usage, the system processed 24 million SQL statements in one hour. There were 1.9 trillion rows in the system, including over 725 billion in one 50+ TB table. HP XP12K and XP1024 disks stored the data, overseen by Veritas Volume Manager software.
- DB2 for AIX managed the third largest data warehouse, 49.4 TB of data, at KT-IT Group in Korea. There were 136 billion rows and almost 16,000 tables in the database. The UNIX-based system, which supports more than 10 applications and feeds a number of business-critical data marts, is updated daily. Hitachi stored the data in a 9980V disk array, managed by Hitachi storage software.

7.5 BY USAGE AND OPERATING SYSTEM, WHICH DBMS VENDORS SUPPORT LARGE DATABASES?

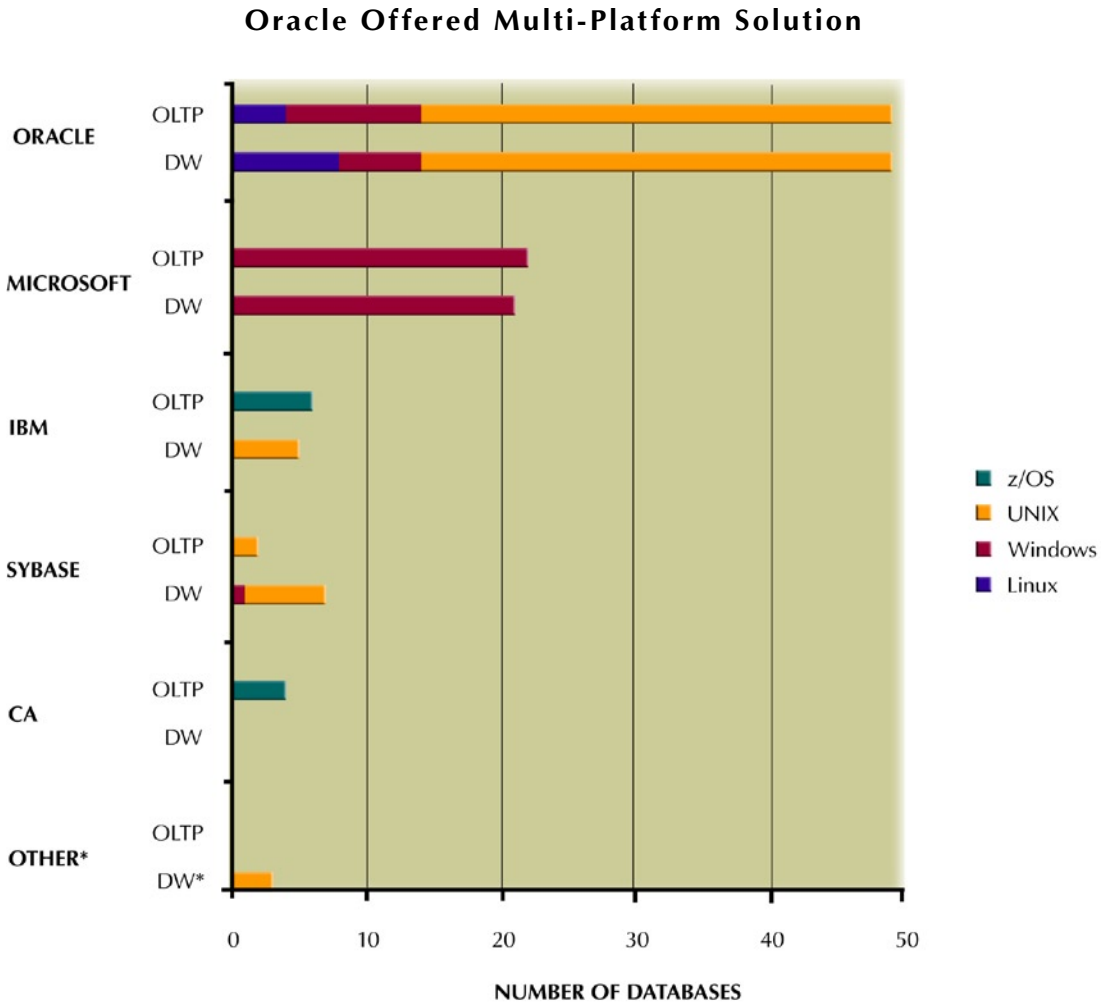

N = 169

Other* = ObjectStore and AT&T on UNIX, and HP on NonStop OS

- Oracle was the only vendor represented by participating databases on three different platforms. Oracle UNIX included 35 OLTP and DW systems, each. There was a slight preference for using Oracle on Windows for transaction processing and Oracle on Linux for data warehousing.
- Microsoft databases, which were about evenly divided between OLTP and DW, were all Windows-based.
- IBM systems were categorized by usage: operational systems ran only on z/OS while data warehouses were implemented on UNIX.
- Sybase databases were hosted on UNIX except for one data warehouse on Windows.
- All CA implementations were used for transaction processing and supported by z/OS.

7.6 BY DATABASE SIZE, WHAT DBMS ARE TRANSACTION PROCESSING SYSTEMS USING?

**Oracle Used by Half of OLTP Systems;
DB2, Oracle and SQL Server Managed Size Leaders**

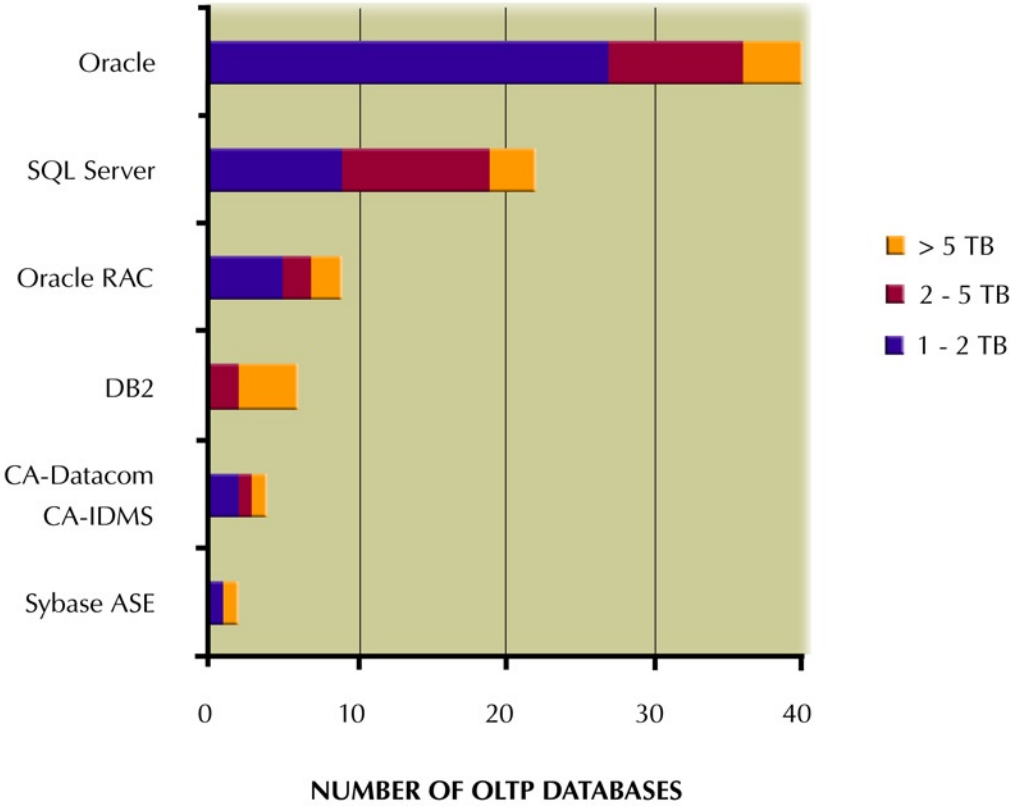

N = 83

- Oracle Database was implemented by almost half of operational systems surveyed, 48%. It managed 27% of the largest databases, 38% of mid-size ones and 61% of emerging databases.
- SQL Server was the leading DBMS at mid-size operational systems, 42%. It was also found at 20% of the largest and the smallest OLTP databases.
- Oracle RAC managed about 10% of OLTP databases in all size categories.
- DB2 was the DBMS at more than a quarter, 27%, of the largest OLTP databases.

7.7 BY DATABASE SIZE, WHAT DBMS ARE DATA WAREHOUSE SYSTEMS USING?

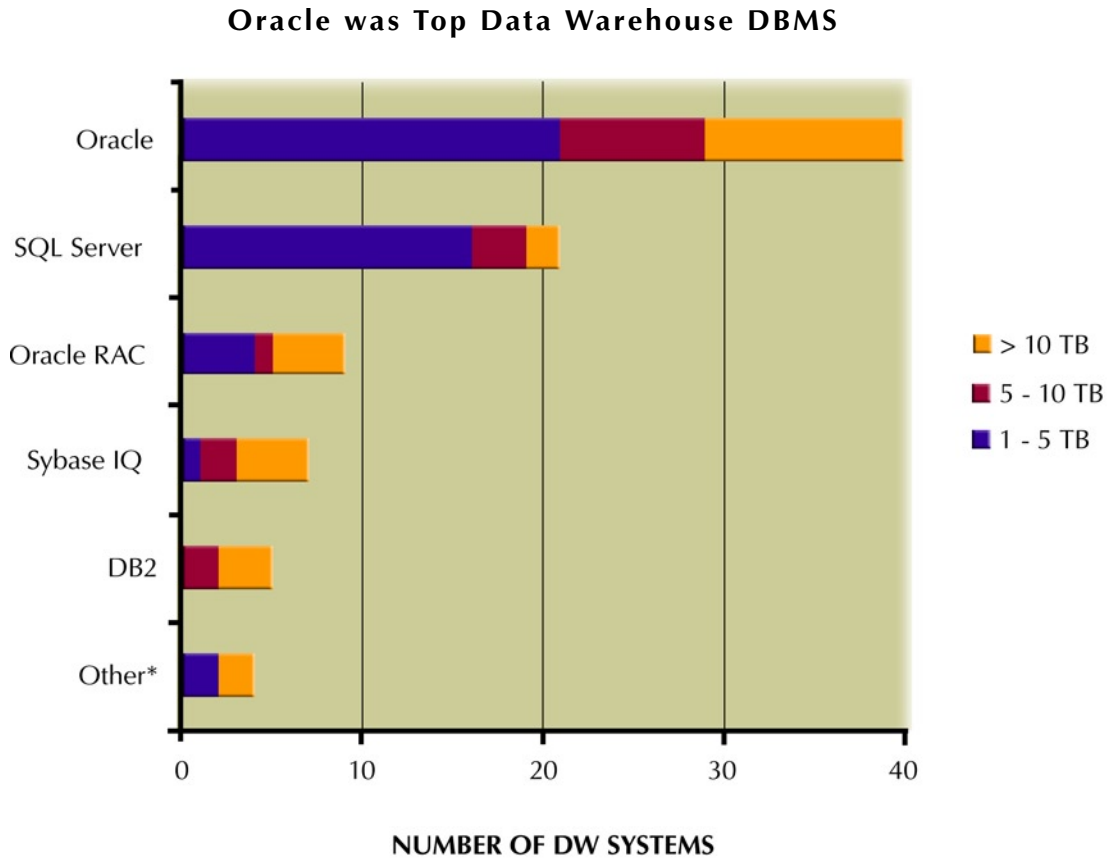

N = 86

Other* includes Daytona, ObjectStore and NonStop SQL

- Oracle Database managed nearly half, 47%, of data warehouses in the 2005 program. It was the most widely used DBMS in all size categories.
- SQL Server achieved the greatest jump in usage as a data warehouse solution since the 2003 program, from 4% to 24%. It was one of the leading choices for small and mid-sized warehouses.
- Oracle RAC was used by 10% of participating data warehouses and had a noticeable presence among the largest DW systems, 15%.
- Sybase IQ and DB2 were implemented at about 13% of both middle and large size data warehouses.
- Teradata did not participate in the 2005 TopTen Program.

8 *Hardware*

HP led many of the server-related findings in the 2005 survey. Overall, HP supplied nearly half of all servers in the program, and supported between one-third and one-half of databases in each size category.

On UNIX, HP and Sun each supplied over one-third of participating servers and IBM provided another 20%.

HP was the dominant server vendor for Windows databases, hosting 60% on the platform. Linux servers were either HP or Dell machines, plus one Fujitsu Siemens box.

The two largest operational systems were implemented on IBM servers. Their data warehouse counterparts ran on Fujitsu Siemens and HP servers, respectively.

Storage capacity, another measure of database scale, grew noticeably since 2003. The average storage capacity of databases surveyed was 23.2 TB, a 10% increase in two years.

8.1 WHAT IS THE ARCHITECTURE OF THE SERVER (OR SERVERS) THAT SUPPORT THE DATABASE?

SMP was Most Widely-Used Server Architecture

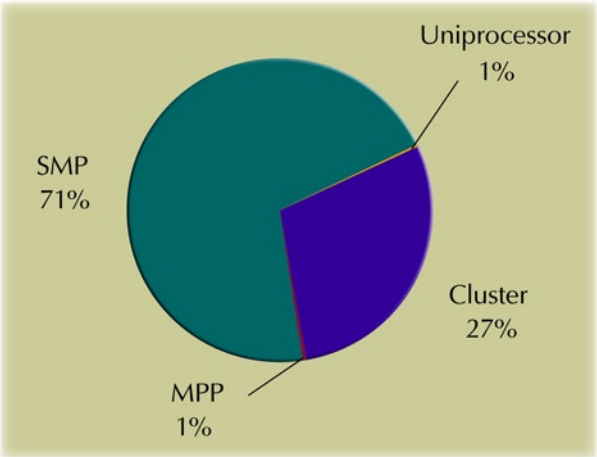

N = 168

- Almost three-quarters of servers in the program, 71%, had SMP architecture. One quarter, 26%, were Cluster systems.

8.2 BY OPERATING SYSTEM, WHICH SERVER VENDORS SUPPORT LARGE DATABASES?

Top Server Provider was HP;
Sun and IBM Comprised Second Tier

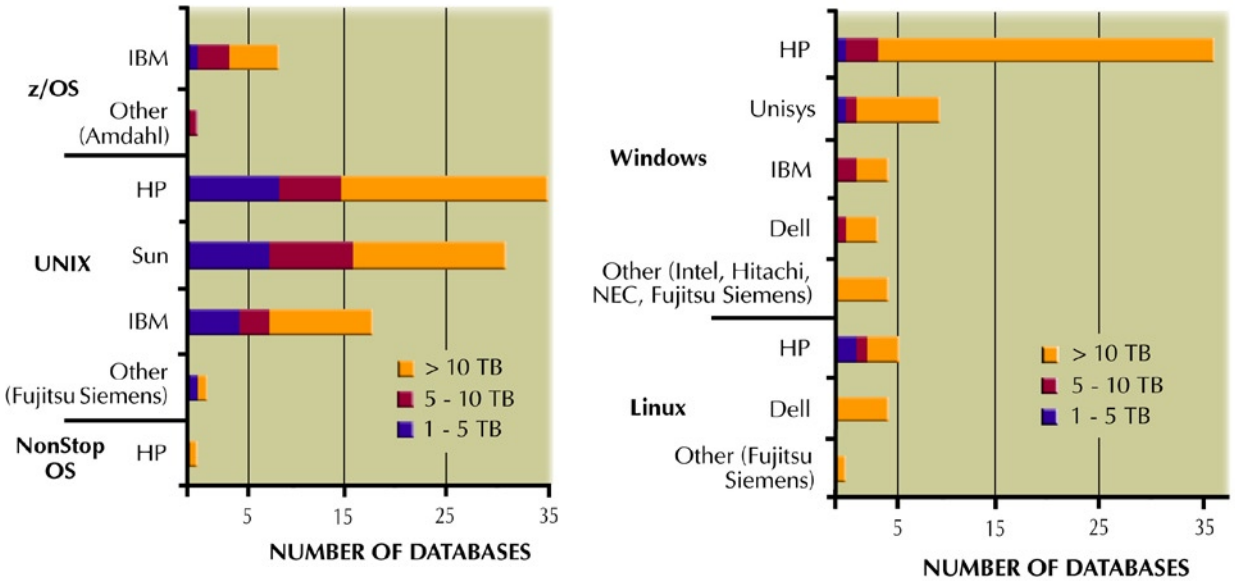

N = 168

- Overall, nearly half of all servers in the program were from HP. IBM and Sun provided almost 20% of machines, each.
- IBM was the leading manufacturer of z/OS servers. IBM systems supported databases split about evenly between those with more than and less than 5 TB of data.
- HP and Sun provided the majority of UNIX servers in the program, 35 and 31, respectively. Along with IBM, about one-quarter of the servers supported the largest and mid-size databases, and the remaining half hosted emerging large databases.
- HP was the leading supplier of Windows-based servers in the program, providing 60% of those used. Unisys systems supported close to 20% of participating Windows databases.

8.3 WHAT ARE THE LARGEST TRANSACTION PROCESSING SYSTEMS SUPPORTED BY THE LEADING SERVER PLATFORMS?

IBM eServer zSeries Hosted the Largest Online Transaction Processing System

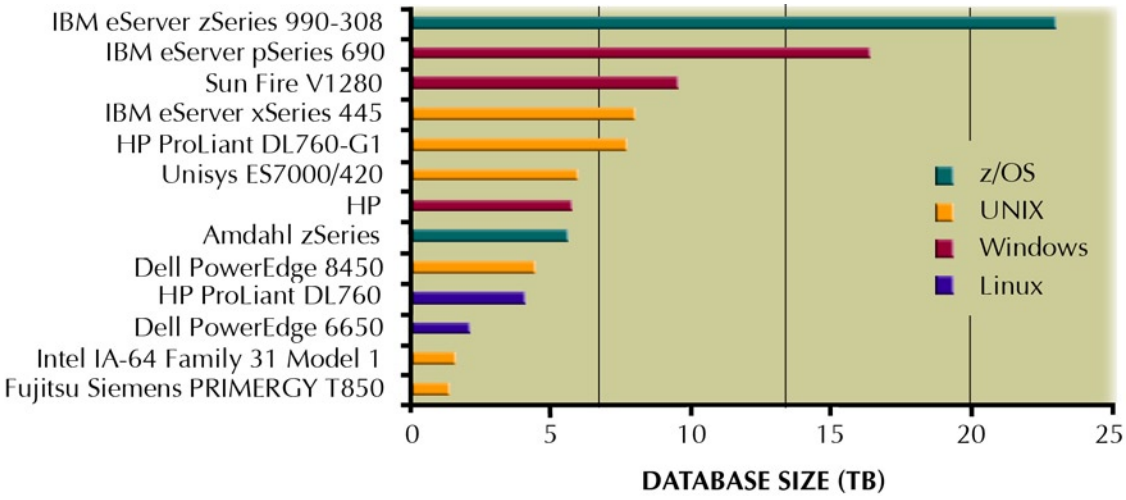

- The largest operational system, 23.1 TB was hosted on a 15-processor IBM eServer zSeries 990-308. The database, found at Land Registry in the UK, is a government system that maintains property and land ownership records. IBM's DB2 for z/OS managed the data, which was stored in an IBM DS8300 2107-9A2 system and a Hitachi 9980V array.
- An IBM eServer pSeries 690 supported the OLTP system with second most data, 16.4 TB, at the United States Trade and Patent Office. Oracle Database managed the data, which was stored in an EMC DMX2000 disk array and managed by Veritas.
- Sun hosted the third largest operational system on a 3-node cluster of Sun Fire V1280 systems. The 9.6 TB database, owned by Elsevier, holds all scientific publications that Elsevier has published and is the center of all production activities. Oracle RAC managed the data, which was stored in an IBM ESS 800 system managed by Veritas.

8.4 WHAT ARE THE LARGEST DATA WAREHOUSE SYSTEMS SUPPORTED BY THE LEADING SERVER PLATFORMS?

Fujitsu Siemens PRIMEPOWER
Supported Largest Data Warehouse

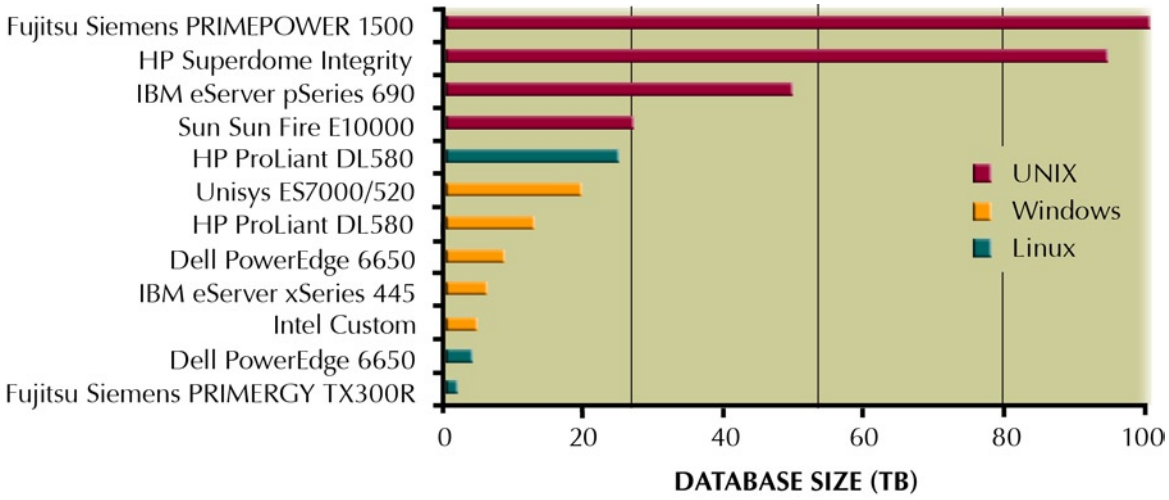

- A Fujitsu Siemens PRIMEPOWER 1500 supported the largest data warehouse, 100.4 TB of data, at Yahoo! The 385-billion row database supports a powerful web search engine and customizable search results page. The data was managed by Oracle Database and stored in an EMC DMX storage system.
- An HP Superdome Integrity hosted the 98.9 TB database at AT&T. The data warehouse stores call detail records and is also used to manage business processes. The data was managed by the Daytona DBMS. Storage was provided by HP XP12K and XP1024 disk systems and managed by Veritas Volume Manager.
- An IBM eServer pSeries 690 cluster was the platform of the third largest data warehouse. The system was implemented at KT IT-Group. An Hitachi 9980V provided storage capabilities, which were also managed by Hitachi.

8.5 WHAT IS THE CAPACITY OF THE SYSTEMS THAT STORE LARGE DATABASES?

Storage to Size Ratio Decreased as Database Size Increased

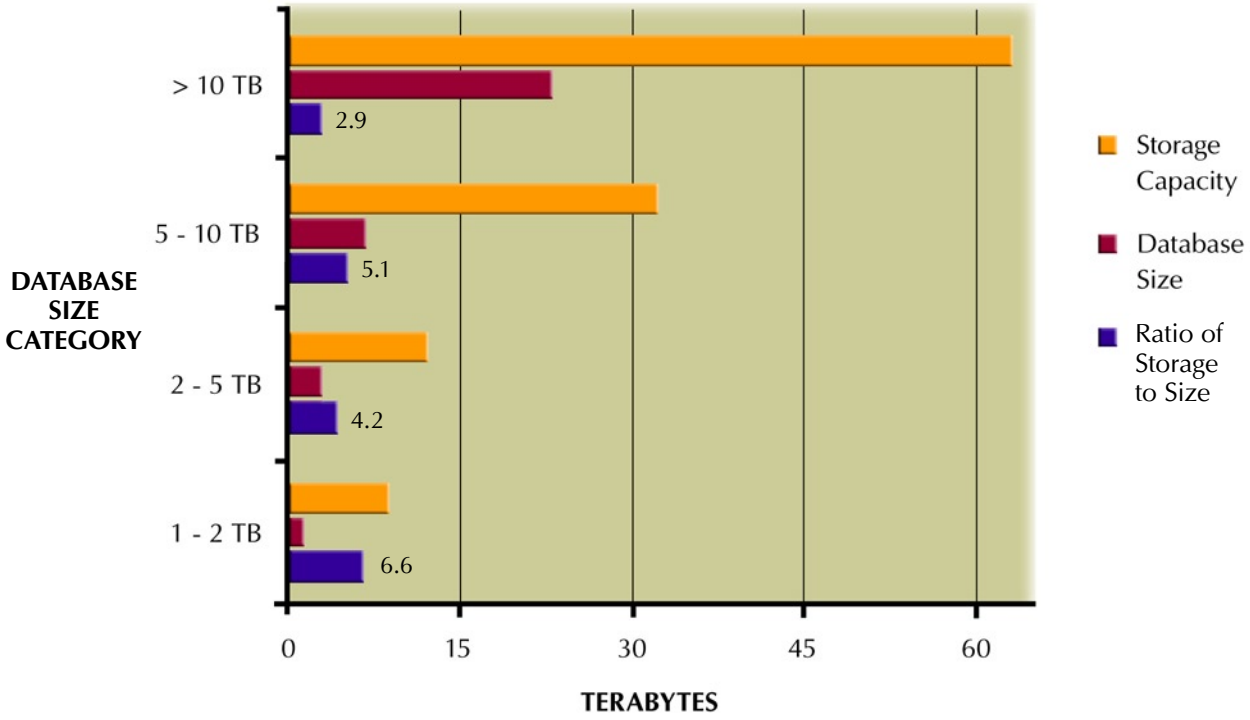

- In the 2005 program, the average storage capacity per database was 23.2 TB, an 11% increase from 2003.
- In general, the larger the database, the lower the ratio of storage capacity to database size. WinterCorp observed the same trend in the 2003 program.

8.6 WHAT IS THE TOTAL STORAGE CAPACITY OF THE 2005 TOPTEN PROGRAM PARTICIPANTS?

2005 TopTen Program Participants Reported Almost 4 Petabytes of Storage

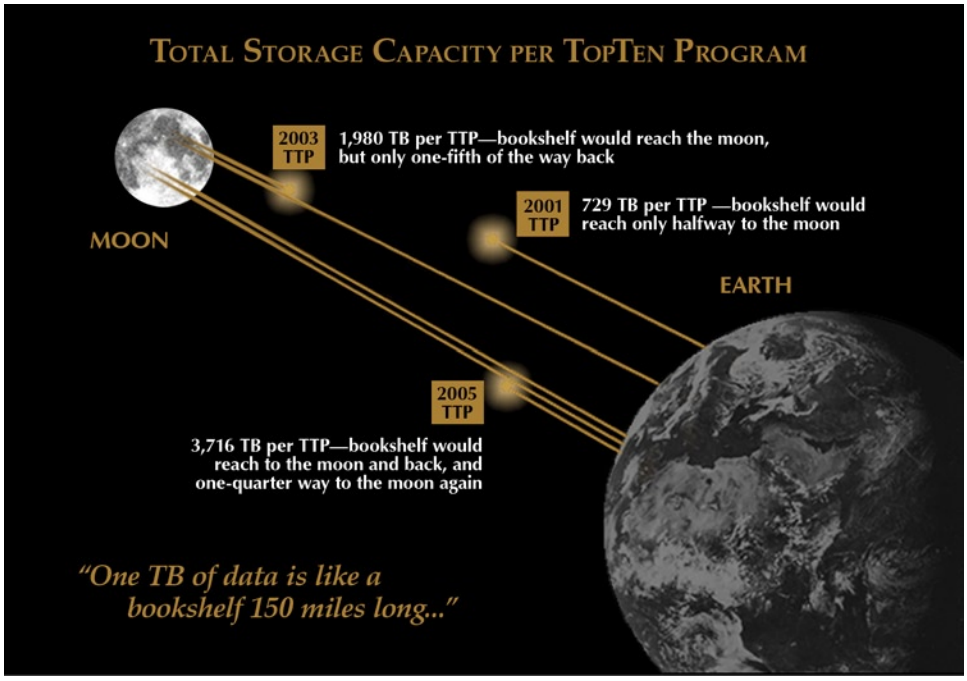

- In the 2001 program, the average storage capacity per database was 13.5 TB. By 2003, the average had increased to 21.2 TB per database.
- In the 2005 program, participating databases had an average of 23.3 TB of storage capacity. This represents an 11% increase since 2003 and a 72% increase since the 2001 survey.

9 *Workload*

The highest performing transaction processing system in the 2005 program, a z/OS implementation, executed over one billion SQL statements in an hour. The average for operational systems was 35 million SQL statements or database operations per hour.

The most heavily used data warehouse in the survey, a UNIX system, processed almost 30 million SQL statements per hour. The hourly average for surveyed data warehouses was about 2 million statements.

9.1 BY USAGE, WHAT ARE THE WORKLOADS OF THE MOST HEAVILY USED DATABASES?

Leading OLTP System, DB2 for z/OS, Processed Over One *Billion* SQL Statements per Hour

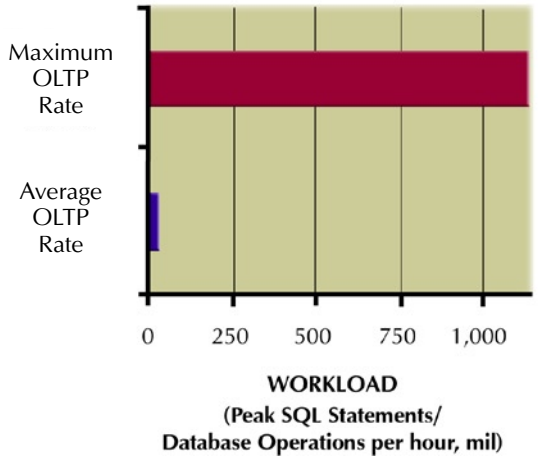

Highest Performing Data Warehouse, DB2 for UNIX, Neared 30 Million Hourly Mark

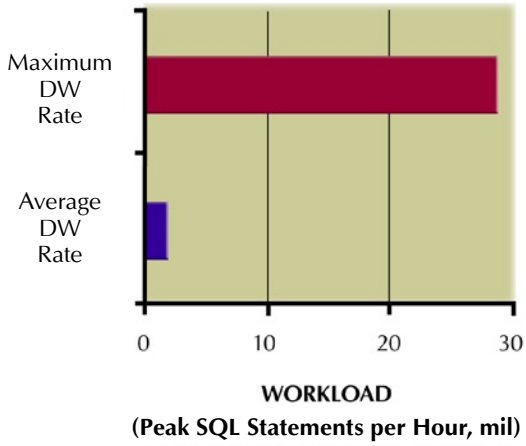

OLTP N = 62
DW N = 63

- A DB2 system for z/OS reached the peak OLTP workload, 1,134,034,718 SQL statements per hour. Among participating OLTP databases, the average was 35,240,621 SQL statements or database operations per hour.
- The hardest working data warehouse, DB2 for UNIX, executed 28,797,833 SQL statements per hour. The average workload for data warehouses was 1,973,814 SQL statements per hour.

10 Project Assessment

Overall, respondents felt that large database projects had delivered noticeable business benefits. On a 1 (low) to 5 (high) scale, they awarded 4.2 to large OLTP and 3.9 to large DW projects. Customer-oriented results (Improved Product/Service Quality, Increased Customer Satisfaction) were noted more often than corporate financial measures were. Among the industries examined, Information/Technical Services benefited the most.

Problems with large databases were deemed slightly higher than the midpoint for OLTP systems, 3.3, and at the midpoint, 3.0, for data warehouses. Owners/managers of transaction processing databases identified Cost and Scalability as above average problems, while data warehouse practitioners named Response Time/Query Throughput and Data Quality as the most troublesome. Telecommunications/Communications organizations experienced the least difficulty with the database project.

10.1 BY USAGE, WHAT BUSINESS BENEFITS HAVE YOU ACHIEVED WITH THIS LARGE DATABASE PROJECT?

Benefits Reported Slightly Higher with OLTP Systems

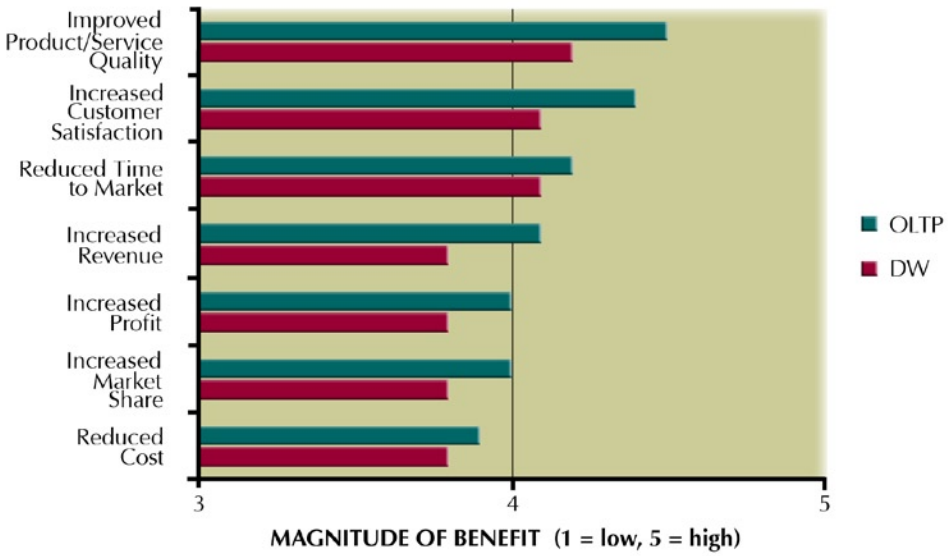

- On average, transaction processing systems delivered slightly more business benefits than data warehouse systems did, 4.2 and 3.9, respectively.
- Regardless of usage, the leading benefit reported was Improvements in Product and Service Quality.
- Large database projects led to more benefits for customers (Better Product/Service Quality, Increased Customer Satisfaction) than to improvements of the organization's financial performance.
- The degree of business benefits derived from large database projects averaged 4.1, significantly higher than problems experienced, 3.1 (see graph 10.3).

10.2 BY INDUSTRY, WHAT BUSINESS BENEFITS HAVE YOU ACHIEVED WITH THIS LARGE DATABASE PROJECT?

Information/Technical Services Reported the Greatest Gains from Large Database Projects

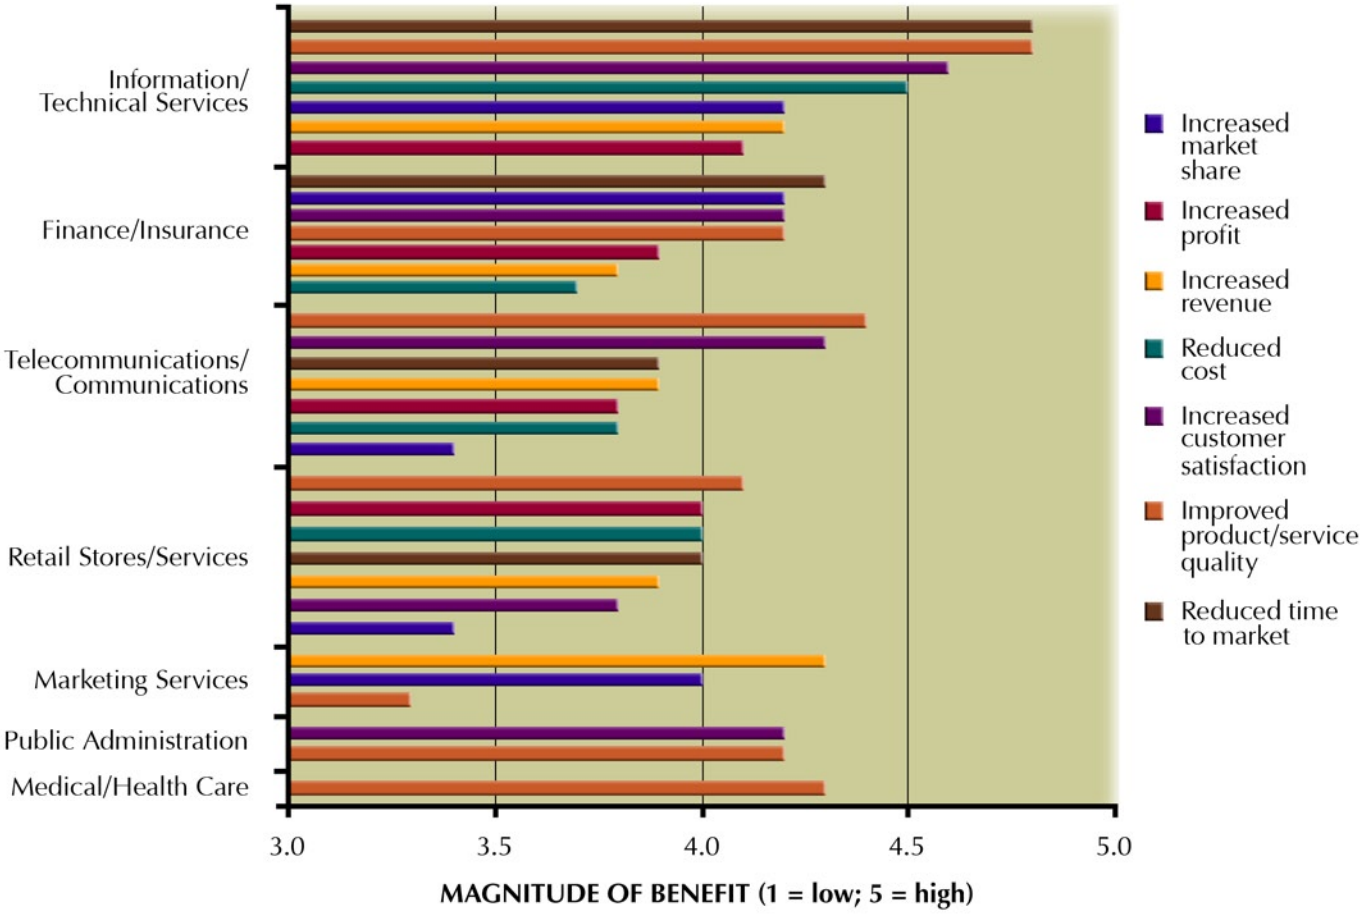

- Among the industries examined, Information/Technical Services firms experienced the greatest business benefits from their large databases. In particular, there were improvements to the quality of products and services, increased market share, and reduced costs.
- Respondents in the Finance/Insurance industry reported that large database projects were beneficial to product-related aspects of the business (Product/Service Quality, Time to Market) as well as to financial measures.
- Customer-oriented benefits (improved products and services, increased satisfaction) were noted by organizations in Telecommunications/Communications, Public Administration and Medical/Health Care.
- Large database projects noticeably reduced time to market and increased the revenues of Marketing Services firms.

10.3 BY USAGE, WHAT PROBLEMS HAVE YOU EXPERIENCED WITH THIS LARGE DATABASE PROJECT?

Problems with Large Databases Deemed Moderate, Especially for Data Warehouses

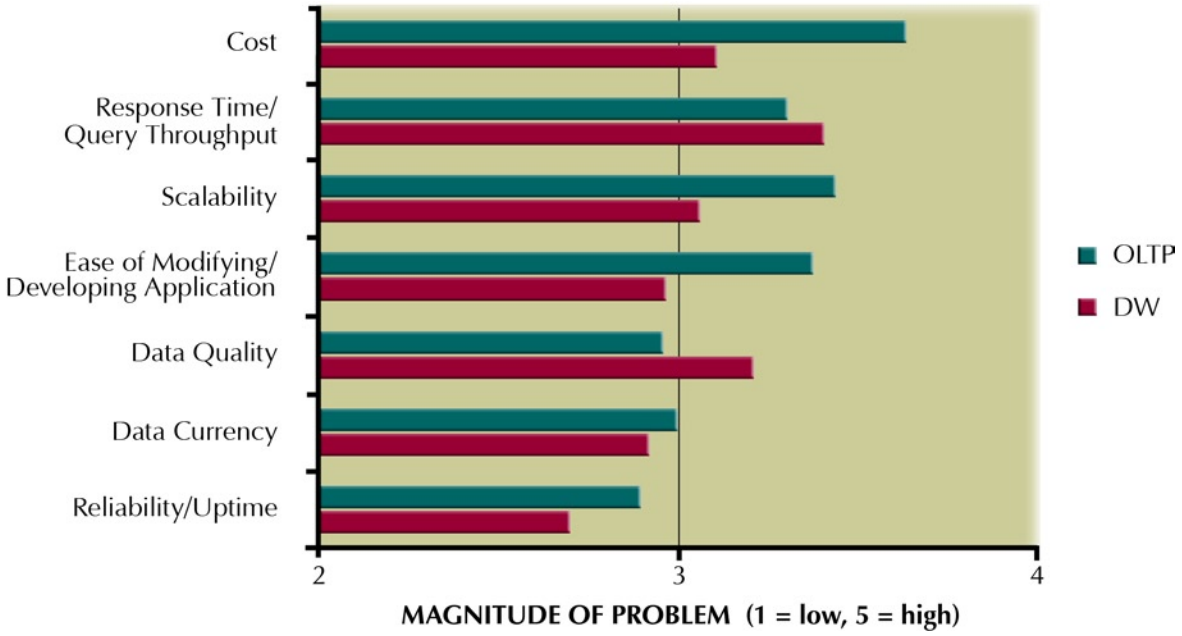

- Overall, the degree of problems experienced with large database projects, 3.1 average, was outweighed by the benefits, 4.1 average.
- Owners/managers of transaction processing databases experienced somewhat more intense problems, 3.3, than did their data warehouse systems counterparts, 3.0. WinterCorp noted the same trend with business benefits, graph 10.1.
- The most challenging issue for transaction processing systems was Cost, 3.4. For data warehouses, it was Response Time/Query Throughput, 3.4. The least problematic, regardless of usage, was Reliability/Uptime, 2.8 average.

10.4 BY INDUSTRY, WHAT PROBLEMS HAVE YOU EXPERIENCED WITH THIS LARGE DATABASE PROJECT?

Industries Grouped Around Midpoint of Problem Assessment

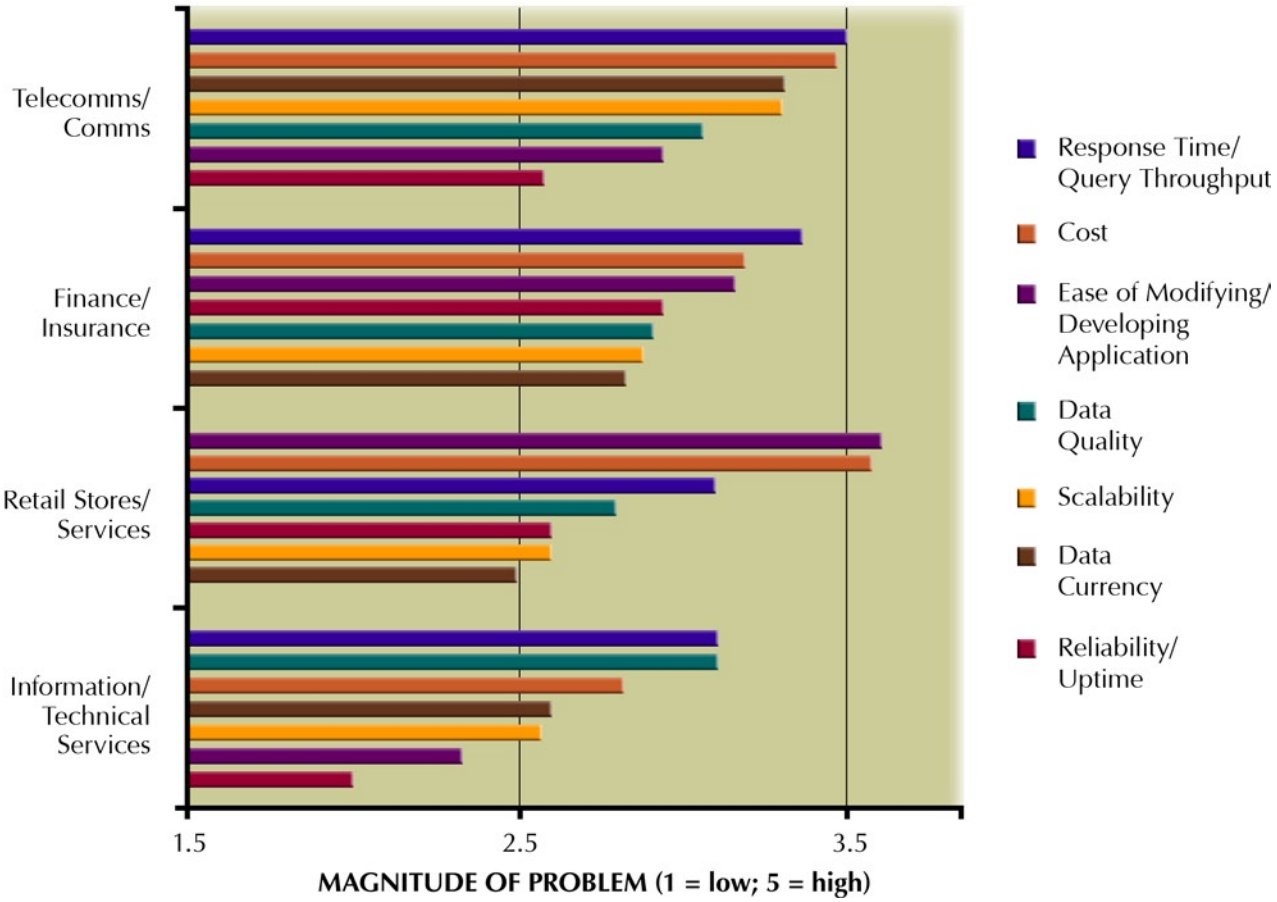

- Telecommunications/Communications and Finance/Insurance reported the most difficulties with Response Time/Query Throughput and Cost.
- Cost and Ease of Modifying/Developing Applications/Database Design were difficult for Retail Stores/Services.
- Information/Technical Services were most challenged by Response Time/Query Throughput and Data Quality.
- Overall, the degree of problems experienced with large database projects, 3.1 average, was outweighed by the benefits, 4.1 average.

11 *Frontier*

In keeping with observations noted in recent TopTen Programs, WinterCorp expects database frontiers to continue expanding at escalating rates. Demands for systems that can mange vast amounts of data, perform near real-time analytics or process billions of daily transactions are constantly redefining what we think of as large-scale data management.

By 2008, operational systems will have seen the biggest upsurge in performance, inflating the 2005 workload boundary nearly three-fold. In data warehousing, size will undergo the greatest growth, pushing the 2008 frontier to a once unthinkable 200 TB.

11.1 WHAT ARE THE FRONTIERS FOR TRANSACTION PROCESSING SYSTEMS, 2005 THROUGH 2008?

OLTP Workload Frontier Forecast to Triple by 2008

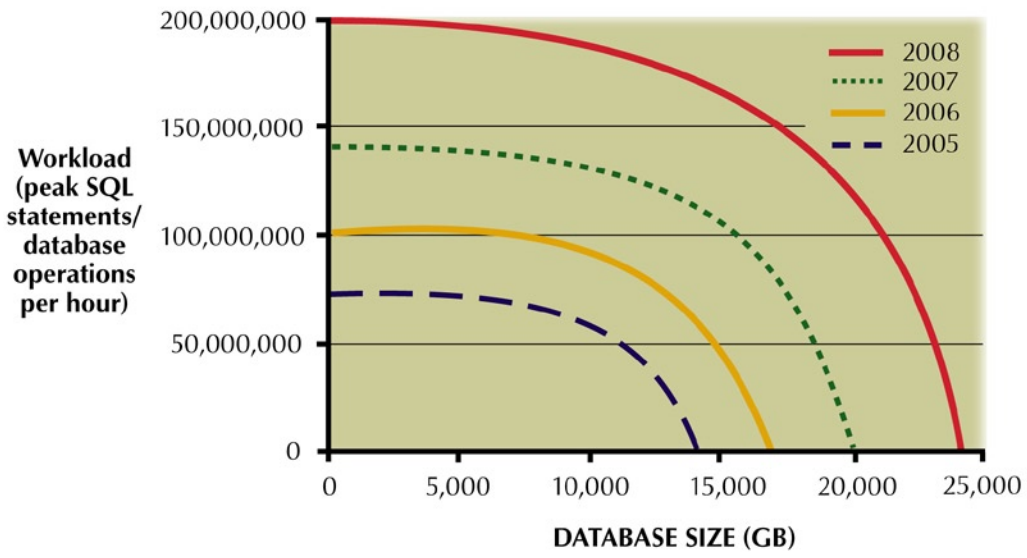

- WinterCorp places the current frontier for OLTP systems at 14 TB of data and 70 million SQL statements or database operations per hour.
- By 2008, the OLTP size frontier will almost double in size and approach 25 TB. Workload is projected to mushroom, 174% growth rate, and approach the 200 million mark.
- WinterCorp frontiers are based on responses from the 2005 survey, historical program data and independently obtained information. They exclude some TopTen databases considered to lie outside the boundaries of standard practice.

11.2 WHAT ARE THE FRONTIERS FOR DATA WAREHOUSE SYSTEMS, 2005 THROUGH 2008?

Largest Data Warehouses to Surpass 200 TB by 2008

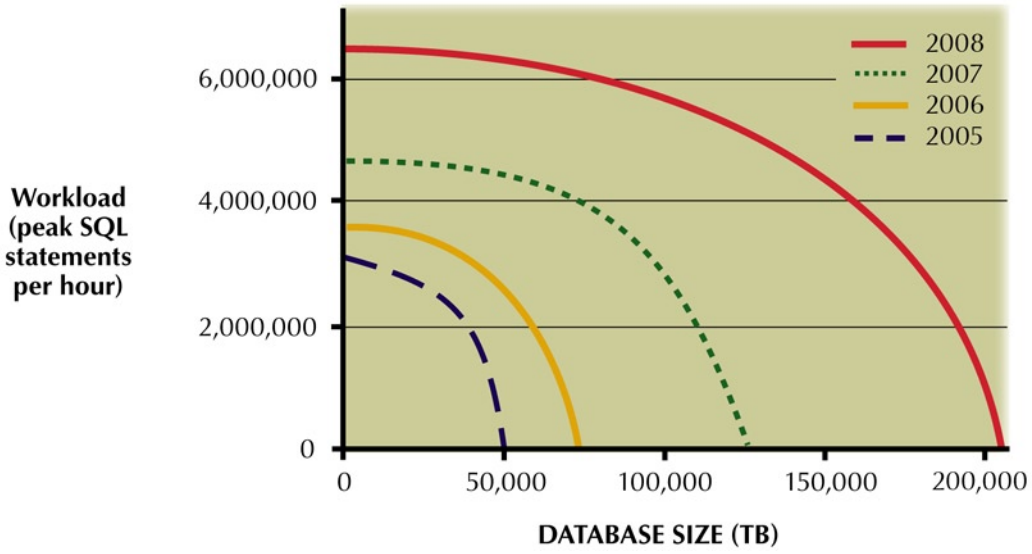

- WinterCorp sets the current frontier for data warehouse systems at 50 TB of data and 3 million SQL statements per hour.
- By 2008, the biggest data warehouses will quadruple in size and hold 200 TB of data. Workload will more than double and reach well over six million SQL statements per hour.
- Frontier dimensions take into account respondents' projections, historical program data and independently obtained information. They exclude some TopTen databases whose metrics exceed the boundaries of conventional data management.

12 Conclusion

Findings from the 2005 TopTen Program confirm that large databases are growing briskly in number, size and power. Our daily lives are filled with the benefits we reap from this advancing technology. Large databases are enabling medical professionals to create and share a worldwide knowledge base. They are facilitating the proliferation and refinement of wide-ranging telecommunications services. Large databases have transformed shopping into a home-based pastime in which goods are brought—virtually—to us.

But as these data-laden systems grow, so do the challenges and risks facing the practitioners who manage them. Issues such as availability, price-performance, system utilization and disaster recovery have been intensified by more data, additional users and higher load levels.

The last four TopTen Programs have revealed *exponential* growth of data warehouses. In the next few years, we could see a data warehouse system with more than 500 TB of data or an operational database that can process 2 billion SQL statements per hour. The combination of an unending thirst for data and our enduring drive to enhance our lives with technology is likely to sustain this course in the coming decade.

13

TopTen Lists

NOTES

Grand Prize Winners are noted in **red**.

If the Grand Prize winner is a federated database, WinterCorp will also identify a non-federated Grand Prize winner.

Anonymous sites are not eligible to be Grand Prize winners.

The majority of Workloads reported in the tables which follow are *peak SQL statements* per hour. However, where indicated by ‡, the reported Workloads are *peak database operations* per hour.

13.1 DATABASE SIZE, ALL ENVIRONMENTS, DATA WAREHOUSING

Company/ Organization	Database Size (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Rows (millions)	Workload (peak SQL statements per hour)
Yahoo!	100,386	UNIX	Oracle	Oracle	Fujitsu Siemens	PRIMEPOWER 1500	SMP	EMC	110	385,318	–
AT&T	93,876	UNIX	AT&T	Daytona	HP	Superdome Integrity	SMP	HP	299	1,882,638	24,000,000
KT IT-Group	49,397	UNIX	IBM	DB2	IBM	pSeries 690	Cluster	Hitachi	210	136,641	315,703
AT&T	26,713	UNIX	AT&T	Daytona	Sun	Sun Fire E10000	SMP	Sun	99	533,723	24,000,000
LGR - Cingular Wireless	25,203	UNIX	Oracle	Oracle	HP	Superdome 9000	SMP	HP	28	50,512	446,448
Amazon.com	24,773	Linux	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	42	81,534	–
Anonymous	19,654	UNIX	IBM	DB2	IBM	pSeries 670	MPP	EMC	82	78,654	28,797,833
UPSS	19,467	Windows	Microsoft	SQL Server	Unisys	ES7000/520	SMP	EMC	40	67,831	658,967
Amazon.com	18,558	Linux	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	67	97,152	–
Nielsen Media Research	17,685	UNIX	Sybase	Sybase IQ	Sun	Sun Fire 4800	SMP	EMC	21	502,407	–

13.2 DATABASE SIZE, UNIX, DATA WAREHOUSING

Company/ Organization	Database Size (GB)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Rows (millions)	Workload (peak SQL statements per hour)
Yahoo!	100,386	Oracle	Oracle	Fujitsu Siemens	PRIMEPOWER 1500	SMP	EMC	110	385,318	–
AT&T	93,876	AT&T	Daytona	HP	Superdome Integrity	SMP	HP	299	1,882,638	24,000,000
KT IT—Group	49,397	IBM	DB2	IBM	pSeries 690	Cluster	Hitachi	210	136,641	315,703
AT&T	26,713	AT&T	Daytona	Sun	Sun Fire E10000	SMP	Sun	99	533,723	24,000,000
LGR — Cingular Wireless	25,203	Oracle	Oracle	HP	Superdome 9000	SMP	HP	28	50,512	446,448
Anonymous	19,654	IBM	DB2	IBM	pSeries 670	MPP	EMC	82	78,654	28,797,833
Nielsen Media Research	17,685	Sybase	Sybase IQ	Sun	Sun Fire 4800	SMP	EMC	21	502,407	–
Health Insurance Review Agency	16,979	Sybase	Sybase IQ	HP	Superdome 9000	SMP	HDS	84	94,644	–
Anonymous	15,197	Sybase	Sybase IQ	HP	Superdome 9000	SMP	HP	21	134,880	–
UBS AG	14,806	Oracle	Oracle	Sun	Sun Fire E6900	SMP	EMC	37	120,200	80,820

13.3 DATABASE SIZE, WINDOWS, DATA WAREHOUSING

Company/ Organization	Database Size (GB)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Rows (millions)	Workload (peak SQL statements per hour)
UPSS	19,467	Microsoft	SQL Server	Unisys	ES7000/520	SMP	EMC	40	67,831	658,967
USDA	12,670	Microsoft	SQL Server	HP	ProLiant DL580	SMP	EMC	44	566	1,433
ComScore Networks Inc.	8,576	Sybase	Sybase IQ	Dell	PowerEdge 6650	SMP	EMC	10	73,214	973
Edcon	6,199	Microsoft	SQL Server	IBM	xSeries 445	SMP	Hitachi	20	31,588	—
HP	6,087	Microsoft	SQL Server	HP	Superdome Integrity	SMP	HP	11	11,013	—
Sage Telecom	5,762	Microsoft	SQL Server	HP	Superdome Integrity	SMP	EMC	10	15,979	49,895
DPF Data Services Group	4,737	Microsoft	SQL Server	Intel	custom Intel	SMP	EMC	20	341	—
OTP Bank	4,490	Oracle	Oracle	Unisys	ES7000	SMP	EMC	5	23,089	22,585,648
Anonymous	4,382	Oracle	Oracle	IBM	xSeries 445	SMP	EMC	10	11,082	30,451
Premier Bankcard Inc.	4,220	Microsoft	SQL Server	HP	Superdome Integrity 8620	SMP	HP	23	14,499	55,896

13.4 DATABASE SIZE, LINUX, DATA WAREHOUSING

Company/ Organization	Database Size (GB)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Rows (millions)	Workload (peak SQL statements per hour)
Amazon.com	24,773	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	42	81,534	—
Amazon.com	18,558	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	67	97,152	—
Amazon.com	7,857	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	42	35,874	—
Telstra Corporation	3,769	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	Network Appliance	48	7,643	773,428
Dell	3,160	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	4	10,819	20,263
Globus SB — Warenhaus Holding GmbH & Co. KG	1,887	Oracle	Oracle	Fujitsu Siemens	PRIMERGY TX300R	SMP	Hitachi		3,195	13,387
Dell	1,623	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	3	1,733	52,286
Gas Natural Informatica s.a.	1,117	Oracle	Oracle RAC	HP	ProLiant DL740	Cluster	HP	6	5,503	176,607

13.5 DATABASE SIZE, ALL ENVIRONMENTS, ONLINE TRANSACTION PROCESSING

Company/ Organization	Database Size (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Rows (millions)	Workload (peak SQL statements per hour)
Land Registry	23,101	z/OS	IBM	DB2	IBM	zSeries 990—308	Cluster	IBM	72	6,479	6,464,623
United States Patent and Trademark Office	16,424	UNIX	Oracle	Oracle	IBM	pSeries 690	SMP	EMC	28	63	21
Elsevier	9,616	UNIX	Oracle	Oracle RAC	Sun	Sun Fire V1280	Cluster	IBM	20	458	3,102,248
UPS	9,284	z/OS	IBM	DB2	IBM	zSeries 990	SMP	IBM	103	89,621	1,134,034,718
KTF	8,706	UNIX	Oracle	Oracle	Sun	Sun Fire 15000	SMP	EMC	33	—	7,160,276
AIM Healthcare Services	8,026	Windows	Microsoft	SQL Server	IBM	xSeries 445	SMP	EMC	200	14,286	—
Verizon Communications	7,781	Windows	Microsoft	SQL Server	HP	ProLiant DL760—G1	SMP	EMC	12	50,747	—
Anonymous	6,800	UNIX	Sybase	Sybase ASE	IBM	pSeries 690	SMP	Hitachi	9	82,639	—
US Bureau of Customs & Border Protection	5,986	z/OS	CA	CA—Datacom	IBM	zSeries 990	SMP	Hitachi	67	24,364	340,838,403‡
Anonymous	5,973	Windows	Microsoft	SQL Server	Unisys	ES7000/420	SMP	Hitachi	10	10,337	—

13.6 DATABASE SIZE, UNIX, ONLINE TRANSACTION PROCESSING

Company/ Organization	Database Size (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Rows (millions)	Workload (peak SQL statements per hour)
United States Patent and Trademark Office	16,424	UNIX	Oracle	Oracle	IBM	pSeries 690	SMP	EMC	28	63	21
Elsevier	9,616	UNIX	Oracle	Oracle RAC	Sun	Sun Fire V1280	Cluster	IBM	20	458	3,102,248
KTF	8,706	UNIX	Oracle	Oracle	Sun	Sun Fire E15000	SMP	EMC	33	—	7,160,276
Anonymous	6,800	UNIX	Sybase	Sybase ASE	IBM	pSeries 690	SMP	Hitachi	9	82,639	—
Hutchison Max Telecom Ltd.	5,824	UNIX	Oracle	Oracle	HP	Superdome 9000	SMP	EMC	8	294	2,865
BPU Banca	5,329	UNIX	Oracle	Oracle RAC	Sun	Sun Fire V880	Cluster	EMC	80	1,980	6,274,536
Anonymous	5,093	UNIX	Oracle	Oracle	Sun	Sun Fire E15000	SMP	EMC	61	—	1,503,352
Coop	3,677	UNIX	Oracle	Oracle	Sun	Sun Fire E15000	SMP	Hitachi	30	5,248	952,373
Turkcell	3,361	UNIX	Oracle	Oracle	HP	AlphaServer	SMP	EMC	3	16,855	3,194,200
ChevronTexaco Information Technology	2,914	UNIX	Oracle	Oracle	HP	HP 9000 rp7410	SMP	Hitachi	80	5,849	31,913

13.7 DATABASE SIZE, WINDOWS, ONLINE TRANSACTION PROCESSING

Company/ Organization	Database Size (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Rows (millions)	Workload (peak SQL statements per hour)
AIM Healthcare Services	8,026	Windows	Microsoft	SQL Server	IBM	xSeries 445	SMP	EMC	200	14,286	–
Verizon Communications	7,781	Windows	Microsoft	SQL Server	HP	ProLiant DL760–G1	SMP	EMC	12	50,747	–
Anonymous	5,973	Windows	Microsoft	SQL Server	Unisys	ES7000/420	SMP	Hitachi	10	10,337	–
Commander Communications Ltd.	4,511	Windows	Microsoft	SQL Server	Dell	PowerEdge 8450	SMP	EMC	18	31,849	–
HP	3,934	Windows	Oracle	Oracle	HP	ProLiant DL760	SMP	HP	10	1,237	22,352
Verizon Communications	3,858	Windows	Microsoft	SQL Server	HP	ProLiant DL760–G1	SMP	EMC	4	5,044	–
Verizon Communications	2,922	Windows	Microsoft	SQL Server	HP	ProLiant DL760	SMP	EMC	4	4,070	–
Verizon Communications	2,761	Windows	Microsoft	SQL Server	HP	ProLiant DL760	SMP	EMC	3	10,450	–
Verizon Communications	2,601	Windows	Microsoft	SQL Server	HP	ProLiant DL760–G1	SMP	EMC	3	8,091	–
HP	2,363	Windows	Oracle	Oracle	HP	ProLiant DL760–G	SMP	HP	6	1,509	35,216

13.8 DATABASE SIZE, LINUX, ONLINE TRANSACTION PROCESSING

Company/ Organization	Database Size (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Rows (millions)	Workload (peak SQL statements per hour)
Amazon.com	4,082	Linux	Oracle	Oracle RAC	HP	ProLiant DL760	Cluster	HP	6	14,279	119,615
Dell	2,149	Linux	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	3	10,467	89,266
Amazon.com	1,938	Linux	Oracle	Oracle	HP	ProLiant DL760	SMP	HP	4	2,273	4,028,971
Dell	1,229	Linux	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	4	292	1,450,423

13.9 DATABASE SIZE, ALL ENVIRONMENTS, SCIENTIFIC AND OTHER

Company/ Organization	Database Size (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Rows (millions)	Workload (peak SQL statements per hour)
Max Planck Institute for Meteorology*	**222,835**	**Linux**	**Oracle**	**Oracle**	**NEC**	**TX7**	**SMP**	**NEC**	–	–	262,220
USGS/EROS	**17,197**	**UNIX**	**Oracle**	**Oracle**	**Sun**	**Sun Fire 4800**	**SMP**	**StorageTek**	**25**	**1,115**	**30,136**
USGS/EROS	17,033	UNIX	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	35	1,617	1,498,720
HP	1,108	NonStop OS	HP	NonStop SQL	HP	NonStop S86000	MPP	HP	7	2,501	3,569,955
T–Systems DDM GmbH	1,003	UNIX	Oracle	Oracle RAC	Sun	Sun Fire V480	Cluster	Hitachi	8	4,455	40,076

*Federated System

13.10 DATABASE SIZE, UNIX, SCIENTIFIC AND OTHER

Company/ Organization	Database Size (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Rows (millions)	Workload (peak SQL statements per hour)
USGS/EROS	17,197	UNIX	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	25	1,115	30,136
USGS/EROS	17,033	UNIX	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	35	1,617	1,498,720
T–Systems DDM GmbH	1,003	UNIX	Oracle	Oracle RAC	Sun	Sun Fire V480	Cluster	Hitachi	8	4,455	40,076

13.11 DATABASE SIZE, LINUX, SCIENTIFIC AND OTHER

Company/ Organization	Database Size (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Rows (millions)	Workload (peak SQL statements per hour)
Max Planck Institute for Meteorology	222,835	Linux	Oracle	Oracle	NEC	TX7	SMP	NEC	–	–	262,220

13.12 DATA VOLUME, ALL ENVIRONMENTS, DATA WAREHOUSING

Company/ Organization	Data Volume (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)	Workload (peak SQL statements per hour)
AT&T	**330,644**	**UNIX**	**AT&T**	**Daytona**	**HP**	**Superdome Integrity**	**SMP**	**HP**	**299**	**93,876**	**1,882,638**	**24,000,000**
AT&T	93,468	UNIX	AT&T	Daytona	Sun	Sun Fire E10000	SMP	Sun	99	26,713	533,723	24,000,000
Amazon.com	**28,184**	**Linux**	**Oracle**	**Oracle RAC**	**HP**	**ProLiant DL580**	**Cluster**	**HP**	**42**	**24,773**	**81,534**	**–**
Nielsen Media Research	17,969	UNIX	Sybase	Sybase IQ	Sun	Sun Fire 4800	SMP	EMC	21	17,685	502,407	–
Yahoo!	17,014	UNIX	Oracle	Oracle	Fujitsu Siemens	PRIMEPOWER 1500	SMP	EMC	110	100,386	385,318	–
Amazon.com	14,849	Linux	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	67	18,558	97,152	–
UBS AG	14,177	UNIX	Oracle	Oracle	Sun	Sun Fire E6900	SMP	EMC	37	14,806	120,200	80,820
China Telecom Corp Ltd., Guangzhou Research Institute	13,241	UNIX	Sybase	Sybase IQ	Sun	Sun Fire E6900	SMP	Sun	20	9,267	133,924	–
USDA	11,876	Windows	Microsoft	SQL Server	HP	ProLiant DL580	SMP	EMC	44	12,670	566	1,433
Reliance Infocomm Ltd.	11,500	UNIX	Oracle	Oracle	Sun	Sun Fire F6800	SMP	EMC	15	13,339	45,452	9,197
Cellcom	10,345	UNIX	Oracle	Oracle RAC	HP	Superdome 9000	Cluster	EMC	80	13,627	40,212	100,292

13.13 DATA VOLUME, UNIX, DATA WAREHOUSING

Company/ Organization	Data Volume (GB)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)	Workload (peak SQL statements per hour)
AT&T	330,644	AT&T	Daytona	HP	Superdome Integrity	SMP	HP	299	93,876	1,882,638	24,000,000
AT&T	93,468	AT&T	Daytona	Sun	Sun Fire E10000	SMP	Sun	99	26,713	533,723	24,000,000
Nielsen Media Research	17,969	Sybase	Sybase IQ	Sun	Sun Fire 4800	SMP	EMC	21	17,685	502,407	—
Yahoo!	17,014	Oracle	Oracle	Fujitsu Siemens	PRIMEPOWER 1500	SMP	EMC	110	100,386	385,318	—
UBS AG	14,177	Oracle	Oracle	Sun	Sun Fire E6900	SMP	EMC	37	14,806	120,200	80,820
China Telecom Corporation Ltd., Guangzhou Research Institute	13,241	Sybase	Sybase IQ	Sun	Sun Fire E6900	SMP	Sun	20	9,267	133,924	—
Reliance Infocomm Ltd.	11,500	Oracle	Oracle	Sun	Sun Fire F6800	SMP	EMC	15	13,339	45,452	9,197
Cellcom	10,345	Oracle	Oracle RAC	HP	Superdome 9000	Cluster	EMC	80	13,627	40,212	100,292
Turkcell	9,504	Oracle	Oracle	Sun	Sun Fire E15000	SMP	Hitachi	20	10,948	181,083	146,090
JPMorganChase	8,875	IBM	DB2	IBM	pSeries 690	MPP	IBM	20	10,560	61,743	303,789

13.14 DATA VOLUME, WINDOWS, DATA WAREHOUSING

Company/ Organization	Data Volume (GB)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)	Workload (peak SQL statements per hour)
USDA	11,876	Microsoft	SQL Server	HP	ProLiant DL580	SMP	EMC	44	12,670	566	1,433
UPSS	8,924	Microsoft	SQL Server	Unisys	ES7000/520	SMP	EMC	40	19,467	67,831	658,967
OTP Bank	3,547	Oracle	Oracle	Unisys	ES7000	SMP	EMC	5	4,490	23,089	22,585,648
DPF Data Services Group	3,443	Microsoft	SQL Server	Intel	custom Intel	SMP	EMC	20	4,737	341	—
Sage Telecom	3,043	Microsoft	SQL Server	HP	Superdome Integrity	SMP	EMC	10	5,762	15,979	49,895
Omnicom Engineering Ltd.	2,892	Oracle	Oracle	HP	ProLiant DL380	SMP	HP	6	3,275	287	281,516
USDA Natural Resource Conservation Service	2,659	Microsoft	SQL Server	Dell	PowerEdge 2650	SMP	EMC	42	3,703	821	994
Microsoft Corporation	2,290	Microsoft	SQL Server	HP	ProLiant DL580—G2	Cluster	HP	6	2,190	264	—
Anonymous	2,279	Oracle	Oracle	IBM	xSeries 445	SMP	EMC	10	4,382	11,082	30,451
KKH — Kaufmaennische Krankenkasse	1,216	Oracle	Oracle	Unisys	ES7000	SMP	EMC	8	1,401	8,160	18,445

13.15 DATA VOLUME, LINUX, DATA WAREHOUSING

Company/ Organization	Data Volume (GB)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)	Workload (peak SQL statements per hour)
Amazon.com	28,184	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	42	24,773	81,534	–
Amazon.com	14,849	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	67	18,558	97,152	–
Amazon.com	5,326	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	42	7,857	35,874	–
Dell	1,756	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	4	3,160	10,819	20,263
Gas Natural Informatica s.a.	863	Oracle	Oracle RAC	HP	ProLiant DL740	Cluster	HP	6	1,117	5,503	176,607
Dell	721	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	3	1,623	1,733	52,286
Globus SB – Warenhaus Holding GmbH & Co. KG	704	Oracle	Oracle	Fujitsu Siemens	PRIMERGY TX300R	SMP	Hitachi		1,887	3,195	13,387
Telstra Corporation	529	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	Network Appliance	48	3,769	7,643	773,428

13.16 DATA VOLUME, ALL ENVIRONMENTS, ONLINE TRANSACTION PROCESSING

Company/ Organization	Data Volume (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)	Workload (peak SQL statements per hour)
UPS	**29,329**	**z/OS**	**IBM**	**DB2**	**IBM**	**zSeries 990**	**SMP**	**IBM**	**103**	**9,284**	**89,621**	**1,134,034,718**
US Bureau of Customs & Border Protection	**10,642**	**z/OS**	**CA**	**CA–Datacom**	**IBM**	**zSeries 990**	**SMP**	**Hitachi**	**67**	**5,986**	**24,364**	**340,838,403‡**
Elsevier	7,873	UNIX	Oracle	Oracle RAC	Sun	Sun Fire V1280	Cluster	IBM	20	9,616	458	3,102,248
Anonymous	5,761	UNIX	Sybase	Sybase ASE	IBM	pSeries 690	SMP	Hitachi	9	6,800	82,639	–
Turkcell	4,654	UNIX	Oracle	Oracle	HP	AlphaServer	SMP	EMC	3	3,361	16,855	3,194,200
AIM Healthcare Services	4,080	Windows	Microsoft	SQL Server	IBM	xSeries 445	SMP	EMC	200	8,026	14,286	–
Caixa Econômica Federal	3,733	z/OS	CA	CA–IDMS	IBM	zSeries 900	SMP	EMC	40	3,733	20,582	131,847,300‡
Amazon.com	2,176	Linux	Oracle	Oracle RAC	HP	ProLiant DL760	Cluster	HP	6	4,082	14,279	119,615
Stadtwerke Munich	1,300	Windows	Oracle	Oracle RAC	Fujitsu Siemens	PRIMERGY T850	Cluster	EMC	50	1,420	3,974	80,422
Cellcom	1,279	UNIX	Oracle	Oracle	HP	Superdome 9000	SMP	EMC	20	1,556	5,593	501,722
Coop	1,231	UNIX	Oracle	Oracle	Sun	Sun Fire 25000	SMP	Hitachi	10	1,608	7,841	220,273

13.17 DATA VOLUME, UNIX, ONLINE TRANSACTION PROCESSING

Company/ Organization	Data Volume (GB)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)	Workload (peak SQL statements per hour)
Elsevier	7,873	Oracle	Oracle RAC	Sun	Sun Fire V1280	Cluster	IBM	20	9,616	458	3,102,248
Anonymous	5,761	Sybase	Sybase ASE	IBM	pSeries 690	SMP	Hitachi	9	6,800	82,639	–
Turkcell	4,654	Oracle	Oracle	HP	AlphaServer	SMP	EMC	3	3,361	16,855	3,194,200
Cellcom	1,279	Oracle	Oracle	HP	Superdome 9000	SMP	EMC	20	1,556	5,593	501,722
Coop	1,231	Oracle	Oracle	Sun	Sun Fire 25000	SMP	Hitachi	10	1,608	7,841	220,273
Hutchison Telecom India	1,169	Oracle	Oracle	HP	9000 Superdome	SMP	EMC	4	1,070	1,419	15,784
ChevronTexaco Information Technology	1,039	Oracle	Oracle	HP	HP 9000 rp7410	SMP	Hitachi	80	2,914	5,849	31,913
Nonghyup	996	Oracle	Oracle RAC	IBM	pSeries 690 Model 681	Cluster	EMC	20	1,814	5,445	72,513
Coop	989	Oracle	Oracle	Sun	Sun Fire E15000	SMP	Hitachi	30	3,677	5,248	952,373
Anonymous	902	Sybase	Sybase ASE	IBM	pSeries 690	SMP	HP	7	1,267	9,826	–

13.18 DATA VOLUME, WINDOWS, ONLINE TRANSACTION PROCESSING

Company/ Organization	Data Volume (GB)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)	Workload (peak SQL statements per hour)
AIM Healthcare Services	4,080	Microsoft	SQL Server	IBM	eServer xSeries 445	SMP	EMC	200	8,026	14,286	
Stadtwerke Munich	1,300	Oracle	Oracle RAC	Fujitsu Siemens	PRIMERGY T850	Cluster	EMC	50	1,420	3,974	80,422
ETCC	1,028	Oracle	Oracle	Dell	PowerEdge 6650	SMP	Network Appliance	8	1,824	9,002	1,902,254
Cellcom	552	Oracle	Oracle	Intel	IA–64 Family 31	SMP	EMC	60	1,597	3,978	279,186
2001OUTLET Co. Ltd.	531	Microsoft	SQL Server	Unisys	ES7000	SMP	EMC	6	1,342	2,709	414,430
123 Multimedia	387	Microsoft	SQL Server	HP	ProLiant DL585	Cluster	HP	17	1,517	1,892	
DataQuick	383	Microsoft	SQL Server	HP	ProLiant	SMP	EMC	10	1,330	4,249	4,410
HP	255	Oracle	Oracle	HP	ProLiant 8500	SMP	HP	4	1,306	3,241	201,443
HP	252	Oracle	Oracle	HP	ProLiant DL760	SMP	HP	2	1,238	3,162	3,017,684
HP	214	Oracle	Oracle	HP	ProLiant 8500	SMP	HP	4	1,228	2,595	44,663

13.19 DATA VOLUME, LINUX, ONLINE TRANSACTION PROCESSING

Company/ Organization	Data Volume (GB)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)	Workload (peak SQL statements per hour)
AIM Healthcare Services	4,080	Microsoft	SQL Server	IBM	xSeries 445	SMP	EMC	200	8,026	14,286	—
Stadtwerke Munich	1,300	Oracle	Oracle RAC	Fujitsu Siemens	PRIMERGY T850	Cluster	EMC	50	1,420	3,974	80,422
ETCC	1,028	Oracle	Oracle	Dell	PowerEdge 6650	SMP	Network Appliance	8	1,824	9,002	1,902,254
Cellcom	552	Oracle	Oracle	Intel	IA—64 Family 31	SMP	EMC	60	1,597	3,978	279,186
2001OUTLET Co. Ltd.	531	Microsoft	SQL Server	Unisys	ES7000	SMP	EMC	6	1,342	2,709	414,430
123 Multimedia	387	Microsoft	SQL Server	HP	ProLiant DL585	Cluster	HP	17	1,517	1,892	
DataQuick	383	Microsoft	SQL Server	HP	ProLiant	SMP	EMC	10	1,330	4,249	4,410
HP	255	Oracle	Oracle	HP	ProLiant 8500	SMP	HP	4	1,306	3,241	201,443
HP	252	Oracle	Oracle	HP	ProLiant DL760	SMP	HP	2	1,238	3,162	3,017,684
HP	214	Oracle	Oracle	HP	ProLiant 8500	SMP	HP	4	1,228	2,595	44,663

13.20 DATA VOLUME, ALL ENVIRONMENTS, SCIENTIFIC AND OTHER

Company/ Organization	Data Volume (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)	Workload (peak SQL statements per hour)
HP	**1,128**	**NonStop OS**	**HP**	**NonStop SQL**	**HP**	**NonStop S86000**	**MPP**	**HP**	**7**	**1,108**	**2,501**	**3,569,955**
T—Systems DDM GmbH	364	UNIX	Oracle	Oracle RAC	Sun	Sun Fire V480	Cluster	Hitachi	8	1,003	4,455	40,076
USGS/EROS	60	UNIX	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	35	17,033	1,617	1,498,720
USGS/EROS	16	UNIX	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	25	17,197	1,115	30,136

13.21 DATA VOLUME, UNIX, SCIENTIFIC AND OTHER

Company/ Organization	Data Volume (GB)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)	Workload (peak SQL statements per hour)
T—Systems DDM GmbH	364	UNIX	Oracle	Oracle RAC	Sun	Sun Fire V480	Cluster	Hitachi	8	1,003	4,455	40,076
USGS/EROS	60	UNIX	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	35	17,033	1,617	1,498,720
USGS/EROS	16	UNIX	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	25	17,197	1,115	30,136

13.22 ROWS, ALL ENVIRONMENTS, DATA WAREHOUSING

Company/ Organization	Rows (millions)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Workload (peak SQL statements per hour)
Sprint	**2,847,553**	**NonStop OS**	**HP**	**NonStop SQL**	**HP**	**NonStop S7400**	**Cluster**	**HP**	**8**	**3,494**	**12,663,305**
AT&T	1,882,638	UNIX	AT&T	Daytona	HP	Superdome Integrity	SMP	HP	299	93,876	24,000,000
AT&T	533,723	UNIX	AT&T	Daytona	Sun	Sun Fire E10000	SMP	Sun	99	26,713	24,000,000
Nielsen Media Research	502,407	UNIX	Sybase	Sybase IQ	Sun	Sun Fire 4800	SMP	EMC	21	17,685	—
Yahoo!	385,318	UNIX	Oracle	Oracle	Fujitsu Siemens	PRIMEPOWER 1500	SMP	EMC	110	100,386	—
Turkcell	181,083	UNIX	Oracle	Oracle	Sun	Sun Fire E15000	SMP	Hitachi	20	10,948	146,090
Anonymous	167,173	UNIX	Sybase	Sybase IQ	Sun	Sun Fire E15000	SMP	EMC	20	14,424	—
KT IT—Group	136,641	UNIX	IBM	DB2	IBM	pSeries 690	Cluster	Hitachi	210	49,397	315,703
Anonymous	134,880	UNIX	Sybase	Sybase IQ	HP	Superdome 9000	SMP	HP	21	15,197	—
China Telecom Corporation Ltd., Guangzhou Research Institute	133,924	UNIX	Sybase	Sybase IQ	Sun	Sun Fire E6900	SMP	Sun	20	9,267	—

13.23 ROWS, UNIX, DATA WAREHOUSING

Company/ Organization	Rows (millions)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Workload (peak SQL statements per hour)
AT&T	1,882,638	AT&T	Daytona	HP	Superdome Integrity	SMP	HP	299	93,876	24,000,000
AT&T	533,723	AT&T	Daytona	Sun	Sun Fire E10000	SMP	Sun	99	26,713	24,000,000
Nielsen Media Research	502,407	Sybase	Sybase IQ	Sun	Sun Fire 4800	SMP	EMC	21	17,685	–
Yahoo!	385,318	Oracle	Oracle	Fujitsu Siemens	PRIMEPOWER 1500	SMP	EMC	110	100,386	–
Turkcell	181,083	Oracle	Oracle	Sun	Sun Fire E15000	SMP	Hitachi	20	10,948	146,090
Anonymous	167,173	Sybase	Sybase IQ	Sun	Sun Fire E15000	SMP	EMC	20	14,424	–
KT IT–Group	136,641	IBM	DB2	IBM	pSeries 690	Cluster	Hitachi	210	49,397	315,703
Anonymous	134,880	Sybase	Sybase IQ	HP	Superdome 9000	SMP	HP	21	15,197	–
China Telecom Corporation Ltd., Guangzhou Research Institute	133,924	Sybase	Sybase IQ	Sun	Sun Fire E6900	SMP	Sun	20	9,267	–
UBS AG	120,200	Oracle	Oracle	Sun	Sun Fire E6900	SMP	EMC	37	14,806	80,820

13.24 ROWS, WINDOWS, DATA WAREHOUSING

Company/Organization	Rows (millions)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Workload (peak SQL statements per hour)
ComScore Networks Inc.	73,214	Sybase	Sybase IQ	Dell	PowerEdge 6650	SMP	EMC	10	8,576	973
UPSS	67,831	Microsoft	SQL Server	Unisys	ES7000/520	SMP	EMC	40	19,467	658,967
Shoprite	41,658	Microsoft	SQL Server	IBM	xSeries 370	SMP	IBM	12	3,874	–
Edcon	31,588	Microsoft	SQL Server	IBM	xSeries 445	SMP	Hitachi	20	6,199	–
Microsoft Corporation	26,088	Microsoft	SQL Server	HP	Superdome 9000–32	SMP	EMC	20	1,989	–
OTP Bank	23,089	Oracle	Oracle	Unisys	ES7000	SMP	EMC	5	4,490	22,585,648
Sage Telecom	15,979	Microsoft	SQL Server	HP	Superdome Integrity	SMP	EMC	10	5,762	49,895
Microsoft Corporation	15,932	Microsoft	SQL Server	Unisys	ES7000/500	Cluster	HP	9	2,749	–
Premier Bankcard Inc.	14,499	Microsoft	SQL Server	HP	Superdome Integrity 8620	SMP	HP	23	4,220	55,896
Anonymous	11,082	Oracle	Oracle	IBM	xSeries 445	SMP	EMC	10	4,382	30,451

13.25 ROWS, LINUX, DATA WAREHOUSING

Company/ Organization	Rows (millions)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Workload (peak SQL statements per hour)
Amazon.com	97,152	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	67	18,558	–
Amazon.com	81,534	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	42	24,773	–
Amazon.com	35,874	Oracle	Oracle RAC	HP	ProLiant DL580	Cluster	HP	42	7,857	–
Dell	10,819	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	4	3,160	20,263
Telstra Corporation	7,643	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	Network Appliance	48	3,769	773,428
Gas Natural Informatica s.a.	5,503	Oracle	Oracle RAC	HP	ProLiant DL740	Cluster	HP	6	1,117	176,607
Globus SB – Warenhaus Holding GmbH & Co. KG	3,195	Oracle	Oracle	Fujitsu Siemens	PRIMERGY TX300R	SMP	Hitachi		1,887	13,387
Dell	1,733	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	3	1,623	52,286

13.26 ROWS, ALL ENVIRONMENTS, ONLINE TRANSACTION PROCESSING

Company/ Organization	Rows (millions)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Workload (peak SQL statements per hour)
UPS*	**89,621**	**z/OS**	**IBM**	**DB2**	**IBM**	**zSeries 990**	**SMP**	**IBM**	**103**	**9,284**	**1,134,034,718**
Anonymous	82,639	UNIX	Sybase	Sybase ASE	IBM	pSeries 690	SMP	Hitachi	9	6,800	–
Verizon Communications	**50,747**	**Windows**	**Microsoft**	**SQL Server**	**HP**	**ProLiant DL760–G1**	**SMP**	**EMC**	**12**	**7,781**	**–**
Anonymous	43,549	z/OS	IBM	DB2	Amdahl	zSeries	SMP	IBM	17	5,652	–
Experian	35,398	z/OS	IBM	DB2	IBM	zSeries 990	SMP	EMC	31	5,264	202,214,000
Commander Communications Ltd.	31,849	Windows	Microsoft	SQL Server	Dell	PowerEdge 8450	SMP	EMC	18	4,511	–
LG Credit Card	30,546	z/OS	IBM	DB2	IBM	zSeries 990	Cluster	EMC	90	4,136	36,639,038
US Bureau of Customs & Border Protection	24,364	z/OS	CA	CA–Datacom	IBM	zSeries 990	SMP	Hitachi	67	5,986	340,838,403‡
Caixa Econômica Federal	20,582	z/OS	CA	CA–IDMS	IBM	zSeries 900	SMP	EMC	40	3,733	131,847,300‡
Turkcell	16,855	UNIX	Oracle	Oracle	HP	AlphaServer	SMP	EMC	3	3,361	3,194,200
AIM Healthcare Services	14,286	Windows	Microsoft	SQL Server	IBM	xSeries 445	SMP	EMC	200	8,026	–

*Federated System

13.27 ROWS, UNIX, ONLINE TRANSACTION PROCESSING

Company/ Organization	Rows (millions)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Workload (peak SQL statements per hour)
Anonymous	82,639	Sybase	Sybase ASE	IBM	pSeries 690	SMP	Hitachi	9	6,800	–
Turkcell	16,855	Oracle	Oracle	HP	AlphaServer	SMP	EMC	3	3,361	3,194,200
SBC	13,189	Oracle	Oracle	HP	9000 rp7420	SMP	HP	8	1,392	7,015
Anonymous	9,826	Sybase	Sybase ASE	IBM	pSeries 690	SMP	HP	7	1,267	–
Starwood Hotels and Resorts Worldwide	8,815	Oracle	Oracle RAC	HP	Integrity rx8620	Cluster	EMC	28	1,310	1,211,318
Coop	7,841	Oracle	Oracle	Sun	Sun Fire 25000	SMP	Hitachi	10	1,608	220,273
ICICI Bank Ltd.	7,053	Oracle	Oracle	Sun	Sun Fire E15000	SMP	HP	12	2,615	5,477,020
Bharti Televentures India	6,193	Oracle	Oracle	HP	AlphaServer	SMP	EMC	10	2,835	458,098
ChevronTexaco Information Technology	5,849	Oracle	Oracle	HP	9000 rp7410	SMP	Hitachi	80	2,914	31,913
Cellcom	5,593	Oracle	Oracle	HP	Superdome 9000	SMP	EMC	20	1,556	501,722

13.28 ROWS, WINDOWS, ONLINE TRANSACTION PROCESSING

Company/ Organization	Rows (millions)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Workload (peak SQL statements per hour)
Verizon Communications	50,747	Microsoft	SQL Server	HP	ProLiant DL760–G1	SMP	EMC	12	7,781	–
Commander Communications Ltd.	31,849	Microsoft	SQL Server	Dell	PowerEdge 8450	SMP	EMC	18	4,511	–
AIM Healthcare Services	14,286	Microsoft	SQL Server	IBM	xSeries 445	SMP	EMC	200	8,026	–
Verizon Communications	10,450	Microsoft	SQL Server	HP	ProLiant DL760	SMP	EMC	3	2,761	–
Verizon Communications	10,352	Microsoft	SQL Server	HP	ProLiant DL760–G1	SMP	EMC	3	2,150	–
Anonymous	10,337	Microsoft	SQL Server	Unisys	ES7000/420	SMP	Hitachi	10	5,973	–
ETCC	9,002	Oracle	Oracle	Dell	PowerEdge 6650	SMP	Network Appliance	8	1,824	1,902,254
Verizon Communications	8,091	Microsoft	SQL Server	HP	ProLiant DL760–G1	SMP	EMC	3	2,601	–
Verizon Communications	7,001	Microsoft	SQL Server	HP	ProLiant DL760	SMP	EMC	3	2,203	–
Verizon Communications	6,554	Microsoft	SQL Server	HP	ProLiant DL760	SMP	EMC	2	1,930	–

13.29 ROWS, LINUX, ONLINE TRANSACTION PROCESSING

Company/ Organization	Rows (millions)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Workload (peak SQL statements per hour)
Amazon.com	14,279	Oracle	Oracle RAC	HP	ProLiant DL760	Cluster	HP	6	4,082	119,615
Dell	10,467	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	3	2,149	89,266
Amazon.com	2,273	Oracle	Oracle	HP	ProLiant DL760	SMP	HP	4	1,938	4,028,971
Dell	292	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	4	1,229	1,450,423

13.30 ROWS, ALL ENVIRONMENTS, SCIENTIFIC AND OTHER

Company/ Organization	Rows (millions)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Workload (peak SQL statements per hour)
T—Systems DDM GmbH	4,455	UNIX	Oracle	Oracle RAC	Sun	Sun Fire V480	Cluster	Hitachi	8	1,003	40,076
HP	2,501	NonStop OS	HP	NonStop SQL	HP	NonStop S86000	MPP	HP	7	1,108	3,569,955
USGS/EROS	1,617	UNIX	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	35	17,033	1,498,720
USGS/EROS	1,115	UNIX	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	25	17,197	30,136

13.31 ROWS, UNIX, SCIENTIFIC AND OTHER

Company/ Organization	Rows (millions)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Workload (peak SQL statements per hour)
T—Systems DDM GmbH	4,455	Oracle	Oracle RAC	Sun	Sun Fire V480	Cluster	Hitachi	8	1,003	40,076
USGS/EROS	1,617	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	35	17,033	1,498,720
USGS/EROS	1,115	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	25	17,197	30,136

13.32 WORKLOAD, MAINFRAME AND OTHER, DATA WAREHOUSING

Company/ Organization	Workload (peak SQL statements per hour)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)
Sprint	12,663,305	Nonstop OS	HP	NonStop SQL	HP	NonStop S7400	Cluster	HP	8	3,494

13.33 WORKLOAD, UNIX, DATA WAREHOUSING

Company/ Organization	Workload (peak SQL statements per hour)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)
Anonymous	28,797,833	IBM	DB2	IBM	pSeries 670	MPP	EMC	82	19,654	78,654
AT&T	**24,000,000**	**AT&T**	**Daytona**	**HP**	**Superdome Integrity**	**SMP**	**HP**	**299**	**93,876**	**1,882,638**
AT&T	24,000,000	AT&T	Daytona	Sun	Sun Fire E10000	SMP	Sun	99	26,713	533,723
UPS	2,944,113	Oracle	Oracle	HP	Superdome 9000	SMP	IBM	12	1,356	3,488
Cisco Systems	829,528	Oracle	Oracle	HP	Superdome 9000	Cluster	EMC	8	6,836	19,439
Absa Ltd.	567,740	ObjectStore	ObjectStore	IBM	pSeries S80	SMP	IBM	8	2,659	9,835
Postal Technology Centre	525,593	Oracle	Oracle	Sun	Sun Fire E25000	SMP	Hitachi		10,478	64,775
TATA Tele Services Ltd.	501,010	Oracle	Oracle	Sun	Sun Fire E15000	SMP	EMC	24	3,710	31,542
KTF	446,625	Oracle	Oracle RAC	IBM	pSeries 690	Cluster	EMC	50	13,430	24,291
LGR — Cingular Wireless	446,448	Oracle	Oracle	HP	Superdome 9000	SMP	HP	28	25,203	50,512

13.34 WORKLOAD, WINDOWS, DATA WAREHOUSING

Company/ Organization	Workload (peak SQL statements per hour)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)
OTP Bank	**22,585,648**	**Oracle**	**Oracle**	**Unisys**	**ES7000**	**SMP**	**EMC**	**5**	**4,490**	**23,089**
UPSS	658,967	Microsoft	SQL Server	Unisys	ES7000/520	SMP	EMC	40	19,467	67,831
Omnicom Engineering Ltd.	281,516	Oracle	Oracle	HP	ProLiant DL380	SMP	HP	6	3,275	287
HP	166,105	Oracle	Oracle	HP	ProLiant 8500	SMP	HP	7	2,633	1
Premier Bankcard Inc.	55,896	Microsoft	SQL Server	HP	Superdome Integrity 8620	SMP	HP	23	4,220	14,499
Sage Telecom	49,895	Microsoft	SQL Server	HP	Superdome Integrity	SMP	EMC	10	5,762	15,979
Premera Blue Cross	42,100	Microsoft	SQL Server	Unisys	ES7000 Orion 420	SMP	EMC	20	2,155	4,091
Microsoft Corporation	36,049	Microsoft	SQL Server	Unisys	ES7000/540	SMP	SAN Storage	3	1,026	4,943
Anonymous	30,451	Oracle	Oracle	IBM	xSeries 445	SMP	EMC	10	4,382	11,082
KKH — Kaufmaennische Krankenkasse	18,850	Oracle	Oracle	Unisys	ES7000	SMP	EMC	8	1,206	9,549

13.35 WORKLOAD, LINUX, DATA WAREHOUSING

Company/ Organization	Workload (peak SQL statements per hour)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)
Telstra Corporation	773,428	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	Network Appliance	48	3,769	7,643
Gas Natural Informatica s.a.	176,607	Oracle	Oracle RAC	HP	ProLiant DL740	Cluster	HP	6	1,117	5,503
Dell	52,286	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	3	1,623	1,733
Dell	20,263	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	4	3,160	10,819
Globus SB — Warenhaus Holding GmbH & Co. KG	13,387	Oracle	Oracle	Fujitsu Siemens	PRIMERGY TX300R	SMP	Hitachi	—	1,887	3,195

13.36 WORKLOAD, MAINFRAME AND OTHER, ONLINE TRANSACTION PROCESSING

Company/ Organization	Workload (peak SQL statements per hour)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)
UPS*	**1,134,034,718**	**z/OS**	**IBM**	**DB2**	**IBM**	**zSeries 990**	**SMP**	**IBM**	**103**	**9,284**	**89,621**
US Bureau of Customs & Border Protection	**340,838,403‡**	**z/OS**	**CA**	**CA—Datacom**	**IBM**	**zSeries 990**	**SMP**	**Hitachi**	**67**	**5,986**	**24,364**
Experian	202,214,000	z/OS	IBM	DB2	IBM	zSeries 990	SMP	EMC	31	5,264	35,398
State Street Corporation	195,430,140‡	z/OS	CA	CA—Datacom	IBM	zSeries 900	SMP	EMC	80	1,077	2,342
Caixa Econômica Federal	131,847,300‡	z/OS	CA	CA—IDMS	IBM	zSeries 900	SMP	EMC	40	3,733	20,582
CheckFree Corporation	66,046,711	z/OS	IBM	DB2	IBM	zSeries 900	Cluster	EMC	11	3,247	10,167
LG Credit Card	36,639,038	z/OS	IBM	DB2	IBM	zSeries 990	Cluster	EMC	90	4,136	30,546
Land Registry	6,464,623	z/OS	IBM	DB2	IBM	zSeries 990—308	Cluster	IBM	72	23,101	6,479
US Department of Treasury/FMS	331,509‡	z/OS	CA	CA—Datacom	IBM	zSeries 900	SMP	IBM	11	1,330	9,954

*Federated System

13.37 WORKLOAD, UNIX, ONLINE TRANSACTION PROCESSING

Company/ Organization	Workload (peak SQL statements per hour)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)
TATA Teleservices Ltd.	8,580,893	Oracle	Oracle RAC	Sun	Sun Fire E15000	Cluster	EMC	4	1,683	4,862
KTF	7,160,276	Oracle	Oracle	Sun	Sun Fire E15000	SMP	EMC	33	8,706	—
BPU Banca	6,274,536	Oracle	Oracle RAC	Sun	Sun Fire V880	Cluster	EMC	80	5,329	1,980
ICICI Bank Ltd.	5,477,020	Oracle	Oracle	Sun	Sun Fire E15000	SMP	HP	12	2,615	7,053
Anonymous	3,902,413	Oracle	Oracle	Sun	Sun Fire E25000	SMP	Hitachi	3	1,200	1,905
ICICI Bank Ltd.	3,205,327	Oracle	Oracle	Sun	Sun Fire E15000	SMP	HP	14	1,010	4,021
Turkcell	3,194,200	Oracle	Oracle	HP	AlphaServer	SMP	EMC	3	3,361	16,855
Elsevier	3,102,248	Oracle	Oracle RAC	Sun	Sun Fire V1280	Cluster	IBM	20	9,616	458
Hutchison Telecom India	3,053,494	Oracle	Oracle	IBM	pSeries 690	SMP	EMC	5	1,201	3,554
Bharti Televentures India	2,368,938	Oracle	Oracle	HP	Superdome 9000	SMP	EMC	5	1,119	—

13.38 WORKLOAD, WINDOWS, ONLINE TRANSACTION PROCESSING

Company/ Organization	Workload (peak SQL statements per hour)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)
HP	3,017,684	Oracle	Oracle	HP	ProLiant DL760	SMP	HP	2	1,238	3,162
Premier Bankcard	2,199,320	Microsoft	SQL Server	HP	ProLiant 760	SMP	HP	5	1,284	4,273
ETCC	1,902,254	Oracle	Oracle	Dell	PowerEdge 6650	SMP	Network Appliance	8	1,824	9,002
2001OUTLET Co. Ltd.	414,430	Microsoft	SQL Server	Unisys	ES7000	SMP	EMC	6	1,342	2,709
Cellcom	279,186	Oracle	Oracle	Intel	IA–64 Family 31 Model 1	SMP	EMC	60	1,597	3,978
HP	201,443	Oracle	Oracle	HP	ProLiant 8500	SMP	HP	4	1,306	3,241
Stadtwerke Munich	80,422	Oracle	Oracle RAC	Fujitsu Siemens	PRIMERGY T850	Cluster	EMC	50	1,420	3,974
HP	51,814	Oracle	Oracle	HP	ProLiant 8500	SMP	HP	4	1,194	12
HP	44,663	Oracle	Oracle	HP	ProLiant 8500	SMP	HP	4	1,228	2,595
HP	35,216	Oracle	Oracle	HP	ProLiant DL760–G	SMP	HP	6	2,363	1,509

13.39 WORKLOAD, LINUX, ONLINE TRANSACTION PROCESSING

Company/ Organization	Workload (peak SQL statements per hour)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)
Amazon.com	4,028,971	Oracle	Oracle	HP	ProLiant DL760	SMP	HP	4	1,938	2,273
Dell	1,450,423	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	4	1,229	292
Amazon.com	119,615	Oracle	Oracle RAC	HP	ProLiant DL760	Cluster	HP	6	4,082	14,279
Dell	89,266	Oracle	Oracle RAC	Dell	PowerEdge 6650	Cluster	EMC	3	2,149	10,467

13.40 WORKLOAD, ALL ENVIRONMENTS, SCIENTIFIC AND OTHER

Company/ Organization	Workload (peak SQL statements per hour)	Platform	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)
HP	3,569,955	NonStop OS	HP	NonStop SQL	HP	NonStop S86000	MPP	HP	7	1,108	2,501

13.41 WORKLOAD, UNIX, SCIENTIFIC AND OTHER

Company/ Organization	Workload (peak SQL statements per hour)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)
USGS/EROS	1,498,720	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	35	17,033	1,617
T–Systems DDM GmbH	40,076	Oracle	Oracle RAC	Sun	Sun Fire V480	Cluster	Hitachi	8	1,003	4,455
USGS/EROS	30,136	Oracle	Oracle	Sun	Sun Fire 4800	SMP	StorageTek	25	17,197	1,115

13.42 WORKLOAD, LINUX, SCIENTIFIC AND OTHER

Company/ Organization	Workload (peak SQL statements per hour)	DBMS Vendor	DBMS	Server Vendor	System	System Architecture	Storage Vendor	Total Storage (TB)	Database Size (GB)	Rows (millions)
Max Planck Institute for Meteorology	262,220	Oracle	Oracle	NEC	TX7	SMP	NEC		222,835	–

14

Questionnaire

WinterCorp 2005
TopTen Program Survey

Winter TopTen™ Survey 2005

PLEASE ANSWER ALL QUESTIONS

I. RESPONDENT PROFILE

Title: _____ First Name:* _____ Last Name:* _____

Job Title:* _____ Company/Organization:* _____

Street Address:* _____

City:* _____ State/Province:* _____

Zip/Postal Code:* _____ Country:* _____

Industry:* _____

Work Phone: _____ Fax: _____

E-mail: (read only) _____ Company/Organizaton URL: _____

[ ] If you wish your survey response to be posted anonymously, please check here.

II. DATABASE OVERVIEW

1. What is the name of the database?* (read only) _____
2. What is the status of the database project?* (select one)
   [ ] In production [ ] Pilot use [ ] Under development
3. If in production, for how long?*
   [ ] < 1 year [ ] 1-2 years [ ] 2-3 years [ ] 3-5 years [ ] > 5 years
4. What is the database primarily used for?* (select one)
   [ ] Operational system/transaction processing/ecommerce
   [ ] Business intelligence/data warehousing/ decision support
   [ ] Scientific/engineering analysis/reference
   [ ] Content/document store, unstructured data repository (XML, email, etc.)
   [ ] Other (specify) _____
5. How frequently is the database updated?*
   [ ] Read only – Never/rarely updated after initial load
   [ ] Monthly or less frequently
   [ ] Weekly
   [ ] Daily
   [ ] Multiple times per day
   [ ] Continuously
6. How many applications are supported by the database?*
   [ ] 1 [ ] 2-3 [ ] 4-6 [ ] 7-10 [ ] > 10
7. What is the database architecture?*
   [ ] Centralized [ ] Hub and spoke [ ] Distributed [ ] Federated
8. What DBMS product *primarily* supports the database?* (select one) _____
9. What is the version number of the DBMS product?* _____

III. DATABASE CHARACTERISTICS

10. How large is the database?*
    [ ] User data (in tables) _____GB
    [ ] Summaries and Aggregates _____GB
    [ ] Indices _____GB
    [ ] Redundancy _____GB
    [ ] Temporary/work space _____GB
    OR:
    [ ] Free/Other space _____GB
    [ ] Total Space _____GB
11. How much user data (in GB) is managed by this database? _____
12. How many tables does the database contain?* _____
13. How many total rows/records/objects does the database contain (in millions)?* _____000,000

* = required

14. How much data is stored in the largest table? _____GB
15. How many rows/records/objects are stored in the largest table (in millions)? _____000,000
16. How many users connect to this database, regularly or occasionally, as part of their job?*_____
17. How many users are concurrently connected to the database on average?*_____ At peak?*_____
18. How many SQL statements or database operations are executed on the database in an average hour?*_____ in the peak hour?* _____
19. How many database transactions are executed on the database in an average hour?_____ in the peak hour?* _____

IV. HARDWARE AND OPERATING SYSTEM

20. What is the architecture of the server (or servers) that supports the database?*
    [ ] Symmetric Multiprocessing (SMP) [ ] Cluster
    [ ] Massively Parallel Processing (MPP) [ ] Uniprocessor
21. Describe the configuration of each node (separate computer) in the system.*

Number of nodes with configuration	Manufacturer	Model	Number processors	Processor type	Memory size (GB)
				[ ] 32-bit [ ] 64-bit	
				[ ] 32-bit [ ] 64-bit	
				[ ] 32-bit [ ] 64-bit	

22. What operating system runs on the server or servers?* (specify name or distribution, and version)
    [ ] UNIX: _____ [ ] z/OS or OS/390: _____
    [ ] Windows: _____ [ ] OS/400: _____
    [ ] Linux: _____ [ ] Other: _____
23. Describe the disk configuration that stores the database.*

Architecture	Manufacturer	Model	Capacity (TB)	Primary vendor
[ ] Networked [ ] Direct attach				[ ]
[ ] Networked [ ] Direct attach				[ ]
[ ] Networked [ ] Direct attach				[ ]
[ ] Networked [ ] Direct attach				[ ]

24. What storage management software product do you primarily use?
    Vendor _____ Product Name _____

V. DATABASE AND SYSTEM GROWTH

25. Estimate the total database size (excluding redundancy) for each of the next three years.
    2006: _____GB 2007: _____GB 2008: _____GB
26. Estimate the peak number of SQL statements (database operations) per hour that will be executed on the database in each of the next three years.
    2006: _____SQL stmnts/hr 2007: _____SQL stmnts/hr 2008: _____SQL stmnts/hr

27. Estimate the peak number of database transactions per hour that will be executed on the database in each of the next three years?
    2006: _____trans/hr 2007: _____trans/hr 2008: _____trans/hr
28. Estimate the number of processors that will be used by the database in each of the next three years.
    2006: _____processors 2007: _____processors 2008: _____processors
29. Estimate the storage capacity in TB that will be used by the database in each of the next three years.
    2006: _____TB 2007: _____TB 2008: _____TB
30. What vendor will you use to support your future server growth?
    Vendor (select one) _____ Other (specify) _____
    [ ] Don't know
31. What vendor will you use to support your future storage growth?
    Vendor (select one) _____ Other (specify) _____
    [ ] Don't know

VI. DATABASE SATISFACTION

32. The largest business benefits we have achieved by implementation of this database project are: (check all that apply)
    For each item checked, indicate the magnitude of the benefit on a 1-5 scale (1=low, 3=medium, 5=high)

[ ] Increased revenue	1	2	3	4	5
[ ] Increased profit	1	2	3	4	5
[ ] Reduced cost	1	2	3	4	5
[ ] Increased customer satisfaction	1	2	3	4	5
[ ] Improved product/service quality	1	2	3	4	5
[ ] Reduced time to market for product/service	1	2	3	4	5
[ ] Increased market share	1	2	3	4	5
[ ] Other (specify) _____	1	2	3	4	5
[ ] Other (specify) _____	1	2	3	4	5

33. The biggest problems I am currently experiencing with my large database projects are: (check all that apply)
    For each item checked, indicate the magnitude of the problem on a 1-5 scale (1=small, 3=medium, 5=large)

[ ] Scalability	1	2	3	4	5
[ ] Response time or query throughput	1	2	3	4	5
[ ] Reliability or uptime	1	2	3	4	5
[ ] Cost of managing and maintaining the database	1	2	3	4	5
[ ] Data quality, such as data inconsistencies	1	2	3	4	5
[ ] Data currency	1	2	3	4	5
[ ] Ease of developing or modifying applications	1	2	3	4	5
[ ] Ease of modifying the database design	1	2	3	4	5
[ ] Other (specify) _____	1	2	3	4	5
[ ] Other (specify) _____	1	2	3	4	5

34. I believe that my current infrastructure: (choose one)
    [ ] Can handle my requirements over the next 3 years
    [ ] Cannot handle my future requirements and I plan to migrate to another platform
    [ ] Cannot handle my future requirements and I don't know what to do
35. I would be willing to participate in Winter Corp. research on the benefits and ROI that I have derived from these large-scale database applications, in exchange for receiving the study findings.
    [ ] Yes [ ] No
36. How did you hear about the TopTen Program?*
    [ ] Past Program Participant [ ] Winter Corp. Personalized Letter [ ] Industry Publication
    [ ] Winter Corp. Press Release [ ] Winter Corp. Web Site (print or electronic)
    [ ] Winter Corp. Post Card [ ] DBMS/Server/Storage Vendor [ ] Other (specify)_____

* = required

15

Appendix

Frequently Asked Questions
about the TopTen Program

15 *Frequently Asked Questions*

WHAT IS THE WINTER TOPTEN PROGRAM?

The Winter TopTen Program surveys users of large databases to understand the characteristics of the world's largest and most heavily used databases. The program provides vendor-independent information that tracks the changing frontiers of database scalability and describes some of the practices employed by the owners and operators of the leading computing installations. See **About the Program** for more information.

Participants complete an online questionnaire about the size and complexity of their large database, their hardware and software environments and configuration, the benefits derived from and problems encountered with the database project, and the expectations for growth of the database. In addition, respondents are required to validate key database scalability metrics by running scripts that collect internal database statistics. The metrics are published in the TopTen lists on the WinterCorp web site.

HOW DOES THE TOPTEN SURVEY WORK?

TopTen Programs run periodically. Data collection for the 2005 survey began in February 15th, 2005, and ended on August 15th. During that period—typically six months—respondents completed a web-based questionnaire and ran validation scripts that assessed select internal database metrics. To qualify for the 2005 survey, databases had to meet a minimum size requirement of 1 terabyte of data.

WHAT IS THE SIGNIFICANCE OF "TOPTEN?"

"TopTen" refers to the ten winners recognized in each of the categories of the program. The first place winner in each of the 12 'All Environments' categories was recognized as the Grand Prize winner of that category.

WHAT DOES THE DATABASE SIZE METRIC MEASURE?

Database size measures the actual storage in use for user tables, indices and aggregates. It excludes all temp tables log data and database free space, as well as duplicated tables, RAID mirroring and other redundancies that augment database availability. Database size is computed from internal database statistics.

WHAT DOES THE NORMALIZED DATA VOLUME METRIC MEASURE?

Normalized data volume estimates a database-independent measure of the volume of user data which the database manages, independent of database compression, indexing and other data transformations that affect the size of the data. It calculates the amount of data in the database by data type and translates it into a standardized format by data type. The information is gathered from internal database statistics and metadata, and sampling of actual data for variable length data types.

WHAT DOES THE NUMBER OF ROWS METRIC MEASURE?

Number of rows counts the total number of rows in user tables. For relational systems, it omits system tables and views. For non-relational systems, the equivalent number of records is captured. The information is gathered from internal database statistics.

WHAT DOES THE PEAK WORKLOAD METRIC MEASURE?

Peak workload counts the number of SQL statements (or equivalent database operations in non-relational systems) executed in the peak hour of the measurement period. For queries, it measures the number of SELECT statements executed, not the number of rows fetched. The information is gathered from internal database statistics.

HOW ARE THE METRICS VALIDATED?

Respondents run database scripts developed by WinterCorp in cooperation with the leading database vendors. The scripts gather information from internal database metadata and operating statistics, and are designed to be quick, read-only, and non-invasive. Respondents submit the output of the scripts to WinterCorp for review. Respondents must complete both the questionnaire and the validation scripts to be included in the program.

WHY ARE THERE NO TERADATA SITES IN THE 2005 SURVEY RESULTS?

WinterCorp invites large database sites worldwide to take part in the TopTen Program. We work with the leading vendors to encourage participation by their customers. NCR Teradata decided not to participate in the 2005 program. In the absence of vendor recruitment, no Teradata users participated.

ARE THE IDENTITIES OF PROGRAM PARTICIPANTS MADE PUBLIC?

Winter Corp publishes only the names, program metrics and database environment of the winning organizations. Companies may request to be listed as 'Anonymous.'

ARE THE RESULTS OF THE PROGRAM MADE PUBLIC?

Program findings are made public through several initiatives. WinterCorp executes a media campaign that publicizes the winners and research findings. We write a trade article, white paper and research report that feature the TopTen winners and select program results.

WinterCorp is an independent consulting firm that specializes in the performance and scalability of terabyte- and petabyte-scale data management systems throughout their lifecycle.

Since our inception in 1992, we have architected many of the world's largest and most challenging databases in production today. Our consulting services help organizations define business-critical database solutions, select their platforms, engineer their implementations, and manage their growth to optimize business value.

With decades of experience working on large-scale database implementations and in-depth knowledge of database products, we deliver unmatched insight into the issues that impede performance and the technologies that enable success.

411 WAVERLEY OAKS ROAD, SUITE 327
WALTHAM, MA 02452
781-642-0300