

Leading Vendors Validate Power of Clustering Architecture

*Detail of the TPC-C[®] audited benchmark, featuring
IBM Netfinity[®], Pentium[®] III Xeon[™] servers,
Windows[®] 2000 Advanced Server
and
DB2[®] Universal Database[™]*

July 2000

DB2[®]
Universal Database[™]

Microsoft[®]

Netfinity[®]

intel[®]



Introduction

In an accelerating global economy, where the ability to network with suppliers, customers and colleagues is an increasing competitive advantage, the information technology infrastructure is growing in strategic importance.

This has caused some companies to confront a seeming paradox. On the one hand, the strategic nature of their computing systems requires an “enterprise-class” level of performance, including high levels of availability and scalability. On the other hand, as with any strategic investment, careful consideration must be given to return on investment and total cost of ownership. In building a strategic infrastructure, these two factors have sometimes seemed to be in conflict.

This paper documents a recent record-setting benchmark (code named Firestorm) that demonstrated how this conflict can be resolved. By designing and implementing a technology infrastructure based on industry-standard components a very real competitive advantage is achieved.

The Firestorm record was set by a 32 node cluster of IBM Netfinity 8500R servers, each with four Intel® Pentium III Xeon™ processors, interconnected with Gigaset® cLAN™ host adapters and cluster switches, running Microsoft® Windows® 2000 Advanced Server and IBM DB2® Universal Database™ Enterprise-Extended Edition Version 7.1 (EEE).

The results were achieved on the TPC Benchmark C® (TPC-C®)¹, which measures the ability of a system to process new order transactions while also executing payment, order-status, delivery and stock-level transactions, all the while meeting strictly defined response times.

In other words, the TPC-C is designed to simulate the type of workload that would typically be encountered in a real-world on-line transaction processing (OLTP) application, e.g. enterprise resource planning, supply chain management or customer relationship management system. And just as in those real-world systems, superior results depend not just on processor speed, as is true in some other benchmarks, but on the integration of the various components — such as communications equipment, operating system and database software — in order to provide overall system performance and the ability to scale-out to unprecedented lengths. Results are reported not only in terms of performance, but also in terms of price/performance.

The results achieved demonstrate how the combination of DB2 with the Netfinity 8500R servers using Intel’s latest processor technology and the scale-out capabilities of Windows 2000 Advanced Server provides the performance and scalability required for strategic implementation. The superior price/performance delivers a compelling competitive advantage for deployment today.

¹ TPC Benchmark, TPC-C, and tpmC are trademarks of the Transaction Processing Performance Council. TPC data is current as of July 6, 2000.

Benchmark Results

IBM, Intel and Microsoft announced on July 3, 2000 the world's fastest server cluster for commercial use, recording performance levels that tripled the performance of Oracle running on a Sun Microsystems cluster, at one-third the price.² The "winning" combination of IBM, Microsoft and Intel, achieved a new TPC-C record of 440,879.95 transactions per minute C at \$32.28 per tpmC. These numbers shattered the previous TPC-C performance record for any hardware, operating system or database. Total solution availability date is December 7, 2000.

A very important aspect of this benchmark result is the validation of a distributed approach to scaling workloads often called "scale-out". Scale-out enables the throughput of a collection of industry-standard servers to be aggregated to achieve the highest possible levels of scalability with unmatched price/performance. Scale-out is frequently compared and contrasted with the "scale-up" approach that relies on big monolithic single systems to deliver the scalability needed for very large workloads.

Historically, scale-up has been the method employed by many computer companies and customers to achieve high scalability. In the last few years, the scale-out approach has been rapidly gaining momentum and is now widely accepted in Data Warehousing and Technical Computing environments. This benchmark provides compelling evidence that scale-out is now a realistic alternative to scale-up for the most demanding On-Line Transaction Processing environments.

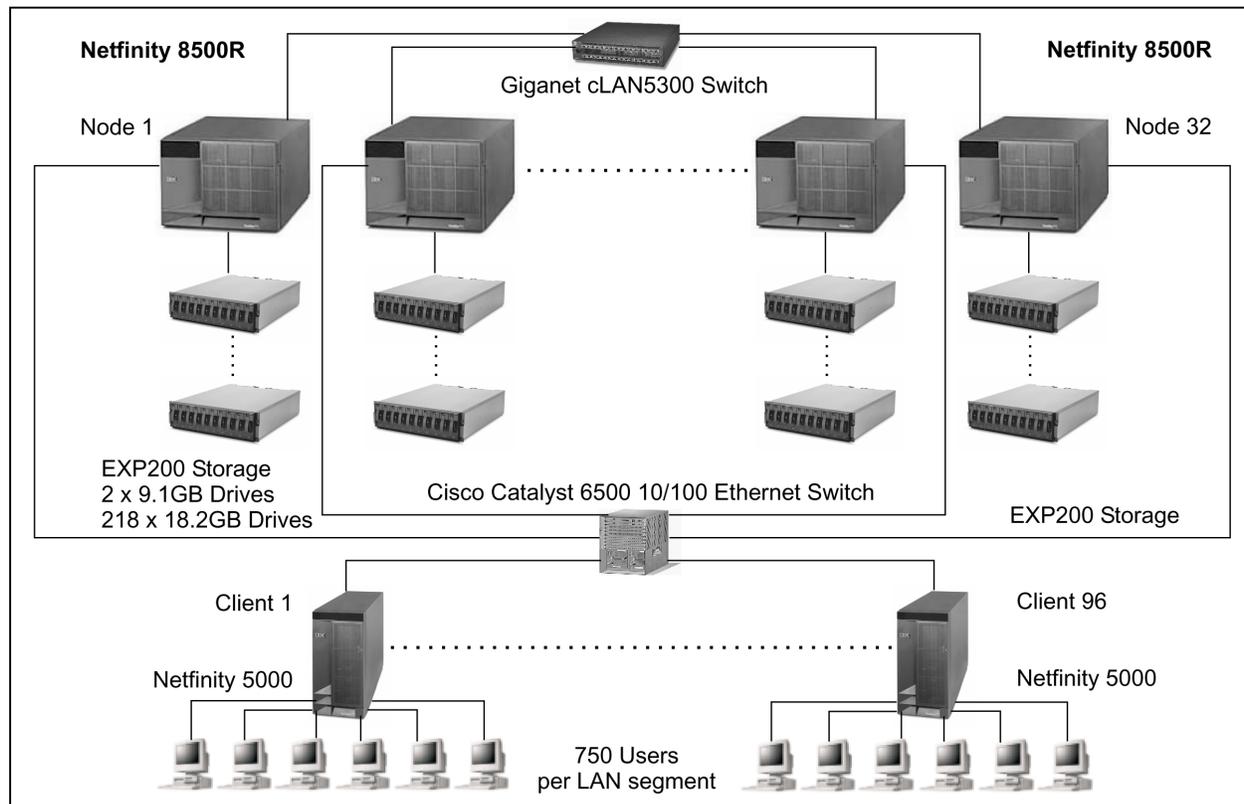
Benchmark Configuration

The benchmark configuration consisted of 32 IBM Netfinity 8500R servers, running Microsoft Windows 2000 Advanced Server and DB2 Enterprise-Extended Edition Version 7.1. Each server was identically configured using:

- Four 700MHz/2MB L2 cache Pentium® III Xeon™ processors
- 4GB ECC SDRAM memory
- Memory coherency filters between the processor buses and the memory and I/O buses
- Eight Netfinity ServeRAID™ -3HB Ultra2 SCSI Adapters
- Twenty-two Netfinity EXP200 Rack Storage Enclosures housed in two Netfinity Racks
- Two 9.1GB 10K Ultra160 SCSI drives and 218 18.2GB 10K Ultra160 SCSI drives
- One EtherJet 10/100 PCI Management Adapter connected to Cisco® Catalyst 6500 Ethernet Switch
- One Gigaset cLAN-1000 Host Adapter.

² Sun's Enterprise 6500 cluster achieved 135,461 tpmC at a price performance of \$97.10 tpmC with availability on January 31, 2000. Data is current as of the date of this document. For the latest information see www.tpc.org

Hardware Schematic



Facts at a Glance:

Firestorm configuration occupied 1,792 square feet (over 160 square meters)

64 Netfinity Racks

Power: 16 – 220 volt/15 A electrical circuits
182 – 110 Volt/15 A electrical circuits
32 – APC Smart-UPS 3000RMB

4+ miles of ethernet cable were used

Details:

Amount of available disk space: More than 116TB of physical disk space was spread over 6,976 18.2GB drives and 64 9.1GB drives. As it is configured, with RAID (a mix of RAID1 for the data and RAID5 for the log), the entire system is protected against disk failures. This made nearly 60TB of space available on the system.

Size of database: The DB2 database was approximately 6.5TB, consisting of 1.8TB of log space and 4.7TB allocated for data and indexes. For the 36,864 warehouses modelled by the TPC-C configuration used in the benchmark, the database consists of over 18Billion rows.

Users: The system was configured to have the RTE (Remote Terminal Emulator) model 368,640 users. These users were spread over 96 client machines (3 client machines per DB2 server node) and the COM+ transaction manager multiplexed the users, however, there were 4,608 concurrent database connections to DB2.

Transactions: The measured throughput was 440,879.95 new-order transactions per minute. New-order is only one of the 5 defined TPC-C transactions, and in Firestorm the new-order transaction represented only 44.6% of the total

transactions. Therefore, the system was executing a *total* of 988,518 transactions per minute, representing a total of over 1.4Billion transactions per day.

System configuration: The database system consisted of 32 server nodes, each with 4x 700MHz processors, 220 disks and 4GB RAM for a total of 128 processors, 7,040 disks and 128GB of real memory.

Lab space: Each of the server nodes occupied 2 full Netfinity racks, additional racks were required for the client and RTE machines. In all, approximately 1,792 square feet of raised floor lab space was used to house the configuration.

Cabling: Each node required over 90 meters of SCSI cable for a total of nearly 3KM (1.83 miles) of SCSI cable. To connect all 225 systems together required over 4 miles of ethernet cable.

Client Support

Ninety-six IBM Netfinity 5000 servers were used as TPC-C clients, which executed the Web-serving, terminal I/O and transaction monitor functions. The operating system installed on each client was Microsoft Windows 2000 Server.

Storage Subsystem

The storage subsystem consisted of 6,976 18.2GB 10K Ultra160 SCSI disk drives, 64 9.1GB 10K Ultra160 SCSI disk drives, 704 Netfinity EXP200 Rack Storage Enclosures and 64 Netfinity Racks. RAID-1 and RAID-5 technologies were used to protect the data from disk drive failures. The operating system and database tables were installed on RAID-1 protected disks. The database log was installed on RAID-5 protected disks. For each of the 32 nodes, the operating system was installed on a two-disk RAID-1 disk array, the database log was installed on an eight-disk RAID-5 disk array, and the database tables were installed on a set of 14 RAID-1 arrays each of which contained 15 disks.

Write-caching was enabled on the ServeRAID-3HB Ultra2 SCSI Adapters. The ServeRAID controllers come standard with a removable, redundant battery-backed cache that protects cached data if the primary cache or the controller fails. The redundant cache offers improved durability over other vendor's single-cache designs.

Database System

At the time of the benchmark announcement, the database is among the worlds' largest OLTP databases both in terms of size and transactions. The Firestorm database was 6.5TB. Load of the over 3TB of base-table data took approximately 6 hours. Backup of the database was also approximately 6 hours and it took only 3.5 hours to restore the database.

Total Disk Space:

Purpose	# of Physical Disks	Formatted Space	Available Space	Used by DB2
OS	64 (32 x 2 9.1GB)	553GB (32 x 17.3GB)	275GB (32 x 8.6GB w/RAID1)	
Log	256 (32 x 8 18.2GB)	4.3 TB (32 x 8 x17.3GB)	3.7 TB (32 x121GB w/RAID5)	1.8 TB (58GB x32)@ 128GB / hour (4GBx32) (more than 14 hours)
Data	6,720 (32 x 210 18.2GB)	112.8 TB (32 x 210 x 17.3GB)	56.8 TB (32x14x130G w/RAID1E)	4.7 TB (3,326GB data, 570GB index, 998GB free)
Total	7,040 Disks	116 TB	59.85 TB	6.5TB

Size of database:

- 41,760 warehouses loaded (32 x 1305 wh), 36,864 warehouses used (32 x 1152 wh)
- Approximately 21 billion records
- 4.7 Terabytes total (32 x 152GB per node)
- 3.2 Terabytes data (32 x 103GB)
- 570GB index (32 x 18GB)
- 998GB free space (32 x ~31GB).

Sizes of tables

Name	# of Records	Total # of Pages with Rows (npage in 4K)	Total # of Page (fpages in 4K)	Record Length
warehouse	41,760	1,152	1,536	97
customer	1,252,800,000	250,636,256	325,827,168	673
district	417,600	11,936	15,552	103
history	1,252,800,000	19,584,800	25,460,256	52
item	100,000	2,206	2,206	86
new_order	375,840,000	1,870,784	2,432,032	10
order_line	12,529,684,864	212,475,584	276,218,272	58
orders	1,252,800,000	11,824,896	15,372,384	28
stock	4,176,000,000	335,203,968	435,765,184	308
=====				
Total	20,840,484,224 ~21 billion records	831,611,582 of 4K pages used 3.2 Terabytes data	1,081,094,590 of 4K pages total fpage - npage = free pages = 998GB	

Selecting the Ingredients for Success

In collaboration, the following partnership of industry technology leaders was instrumental in providing the unsurpassed performance and scalability of industry-standard based servers reflected by the TPC-C benchmark.

Hardware and Software Components

IBM Netfinity servers

The results of the TPC-C benchmark demonstrate the extreme performance, expansive scalability and reliability of Netfinity clustered solutions with proven solution vendors.

IBM has dramatically expanded the role of industry-standard computing with highly available and scalable clustering solutions. Through IBM's X-architecture strategy of bringing proven, enterprise class capabilities to industry-standard computing, Netfinity servers are the choice for Windows 2000 users seeking a powerful, reliable and scalable platform that is optimized for their mission critical solutions.

IBM's Netfinity 8500R server was selected because of its advanced design. Based on the X-architecture blueprint, the Netfinity 8500R optimizes the processor and I/O-intensive operations required to maximize performance, whether setting a new performance record or conducting your e-business in today's competitive marketplace. A server designed with the proper balance of processing power and I/O is critical to great performance. Scalability is also very important in a benchmark of this type as well as in today's dynamic market. Planning for rapid growth is only one of the many challenges that face an IT department. For this performance record, only 4 CPUs were used which demonstrates the server has headroom for today's challenges as well as growth for tomorrow.

The storage subsystem consisted of Netfinity ServeRAID-3H Ultra2 SCSI Adapters – eight per server. These hardware RAID controllers provided the I/O performance necessary to maintain a transaction rate to maximize the processing capabilities of the host server and to connect to the external SCSI storage housed in twenty-two Netfinity EXP200 Rack Storage Enclosures. Two Netfinity 9.1GB 10K Ultra160 SCSI drives and 218 18.2GB 10K Ultra160 SCSI drives per node completed each node's storage configuration.

IBM Netfinity 5000 servers were used to drive the workloads required to stress the servers. Each server simulated approximately 3,840 TPC-C users.

Find out more about IBM Netfinity at: www.ibm.com/pc/us/netfinity

APC Smart-UPS 3000RMB provided the rack-mountable Uninterruptible Power Source and protected the IBM systems from outages due to power related problems.

Intel® Pentium® III Xeon™ 700MHz Processor: The Computing Brain behind each Server Node in the Cluster

Each server node in this benchmark is an eight-way symmetrical multiprocessing (SMP) capable Netfinity 8500R server built on **Intel's Profusion® chipset** architecture. Each node was populated with four **Intel Pentium III Xeon 700MHz** processors.

The efficient design of the .18u Pentium III Xeon processor offered at 700MHz, or higher core frequency, enables Intel to build a fast, 2MB L2 cache directly into the processor itself. Unlike earlier Intel Pentium III Xeon and Pentium II Xeon processors, which accessed external L2 cache, the new Pentium III Xeon processor features an internal cache that runs at full speed. Measurements obtained during this benchmark indicate that this feature of the 700MHz Pentium III Xeon processor, with the increased core frequency, added more than 20 percent in transaction processing performance over earlier 550MHz Pentium III Xeon processors.

The integrated design reduces memory latency by four times and improves performance across the board. The L2 cache is called the Advanced Transfer Cache (ATC), and was used because of its high level of performance via the internal full speed cache, higher core frequency and integrated design. This L2 cache employs a 256-bit link to the processor that is four times as wide as the connection used on earlier Pentium III Xeon and Pentium II Xeon processors. The cache also employs higher set associativity than earlier L2 cache stores, speeding access to cached data. Bigger fill buffers and bus queue entries enable prompt and smooth data flow. In summary, the Advanced Transfer Cache consists of micro-architectural improvements to provide a higher data bandwidth interface between the level 2 cache and the processor core that is completely scalable with the processor core frequency resulting in improved transaction processing performance.

Features of the ATC include:

- Non-Blocking, full speed, on-die level 2 cache 0
- 8-way set associativity
- 256-bit data bus to the level 2 cache
- Reduced latency interface to cache data (as compared to discrete caches).

The heart of the memory and I/O at each node in the cluster was the Intel supplied Profusion chipset and the base-board. It blends the power of Pentium III Xeon processor based, up to eight-way SMP, with the largest memory subsystem offered today for an Intel based server platform. The chipset also delivers a powerful I/O subsystem that balances the processor capability with powerful I/O bandwidth.

This architecture includes three 100MHz system buses connected through a crossbar switch to the CPU unit, I/O and memory. Of these buses, two support the CPUs (a set of four CPUs each), and the third is dedicated to the I/O subsystem in order to reduce processor disruption, with cache coherency between buses. The Profusion chip supports the latest Pentium III Xeon 700MHz processors. Of special importance in this architecture are the memory and I/O subsystems. The memory subsystem supports up to 32GB at 100MHz, two high-bandwidth SDRAM subsystems and a maximum bandwidth of 1.6GB per second.

Additionally, the I/O sub-system is modular and expandable, based on the customer requirements for supporting either legacy PCI cards or the latest 64-bit PCI cards. This subsystem includes specifications for two 64-bit PCI expansion slots at 33MHz (expandable to six), two 64-bit PCI slots at a speed of 66MHz (expandable to four) and up to four peer 66MHz, 64-bit PCI buses.

With a maximum bandwidth of 800MB per second, the PCI buses interface to a 100MHz dedicated PCI bus. A number of other architectural features enhance the enterprise caliber of the Profusion platform. Profusion supports Extended Server Memory Architecture (ESMA), which allows full 36-bit memory addressing for enterprise applications allowing 32GB of memory. This feature breaks the legacy 32-bit (4GB) memory barrier, giving applications memory headroom and improving the performance of OLTP applications significantly.

Find out more about Intel at: <http://www.intel.com>

Giganet

Giganet develops and markets products and technology that enable the creation of flexible, high performance Virtual Interface (VI) based server farm networks. Server farm networks are comprised of multiple rack-mounted systems linked via high-performance communications, allowing businesses to scale computing resources. The Giganet products provide businesses with the ability to cost effectively deploy server farm networks that deliver on-demand scalability and unmatched availability in support of mission critical enterprise and Internet applications.

Giganet's cLAN technology was selected for the cluster interconnects. The Giganet cLAN-1000 Host Bus Adapters were used in each node connected to a Giganet cLAN5300 Switch. Giganet is an IBM ServerProven™ partner and offers a fast interconnect that provides a streamlined protocol for internodal communication.

The Gigaset cLAN products include a full line of ultra high-speed host adapters and cluster switches, available in 8- and 30-port configurations, and can be used as flexible, high performance building blocks for deploying large plug-and-play server networks that meet today's demanding environments.

Gigaset's products are based on the industry-standard Virtual Interface Architecture as well as TCP/IP, built on top of their VI foundation. By implementing VI in hardware, Gigaset increases the speed and efficiency of server-to-server communications. Gigaset server farm networks deliver high performance by allowing applications to bypass the operating system, resulting in very high throughput rates (up to 100MB/s) with very low latency (delay) and ultra low CPU utilization.

Gigaset also incorporates *RapidFlow*[™] architecture, which is a unique switching functionality with fast flow control, non-blocking and queuing mechanisms that brings efficiency and performance to a new level for e-commerce environments. Gigaset's level of efficiency frees up valuable database and application system resources and improves overall end user performance.

Gigaset products support Windows NT, Windows 2000 and Linux environments.

Gigaset cLAN High Performance Cluster Switch and cLAN Host Adapters functionality summary:

- Native VI implementation
- Ultra high performance
- 8 30-port switch configurations
- 1.25GB link speed
- Switch port-to-port latency of 500 nanoseconds
- Ultra low CPU utilization, freeing up the system for application processing
- Fast flow control and queuing functionality allow for faster network speeds
- Hot pluggable design
- Auto detecting/configuring for easy setup and administration.

Find out more about Gigaset at: <http://www.gigaset.com>

Microsoft Windows 2000 Advanced Server

Windows 2000 Advanced Server is Microsoft's operating system for line-of-business and dot-com solutions. It supports today's most demanding transaction workloads on the latest industry-standard hardware. With Windows 2000 Advanced Server, your business can take full advantage of the Internet economy. Based on the solid Windows NT foundation, it incorporates the Internet technologies you need to build Web sites and Web-enabled applications. It's reliable, so your business stays up and running. Smart management features help you cut operating costs and you can take advantage in the latest hardware—from tiny smart cards to midrange 8-way servers.

Windows 2000 Advanced Server was selected because industry has recognized it as one of the most widely used operating system platforms and because of its scale-out capability. The Advanced Server is optimized for large amounts of memory (up to 8GB of system memory) and SMP support (up to 8-way SMP processors).

Windows 2000 Advanced Server is a member of the Windows 2000 Server Family, which also includes Windows 2000 Server and Windows 2000 Datacenter Server. The Windows 2000 Server Family lets you do this:

- **Internet-enable your business** with essential technologies woven throughout the operating system
- **Cut costs** with improved management systems for networks, servers and Windows desktops
- **Sustain up-time** with extensive reliability and availability improvements

- **Take advantage of new hardware** with broad support for existing and emerging hardware and communication products.

The Windows 2000 Server Family builds on the strengths of Windows NT technology. It integrates standards-based directory, Web, application, communications, and file and print services with high reliability, efficient management and support for the latest advances in networking hardware to provide the best foundation for integrating your business with the Internet. Windows 2000 Advanced Server brings the specific following benefits to the family:

- Best platform for line-of-business and dot-com backend usage. Provides enhanced performance and scalability through SMP and extended memory support.
- Scales from 1 to 8 processors and up to 8 gigabytes (GB) of memory.
- Enhanced reliability and availability — two-node clustering, 32-node network load balancing, and components that are tested through WHQL.

Windows 2000 Advanced Server supports up to 8GB of memory with processors supporting Intel's Physical Address Extensions (PAE). Combined with support for 8-way SMP, enhanced large-memory support ensures that demanding memory- and processor-intensive applications can be run on Windows 2000 Advanced Server. Symmetric multiprocessing (SMP) is a technology that lets software use multiple processors on a single server in order to improve performance. Improvements in the implementation of the SMP code in Windows 2000 allow for improved scaling linearity, making Windows 2000 Advanced Server an even more powerful platform for business-critical applications, databases and Web services.

The Windows 2000 platform now holds the top five slots for the absolute performance category of the TPC-C benchmark. In addition, Windows NT and Windows 2000 continue to hold more than 50 of the top slots for best price/performance results of the TPC-C benchmark, highlighting how Windows NT and the Windows 2000 platform consistently provide the industry's best value for the money. Combined with relatively inexpensive PC hardware, including the industry-standard Intel-based IBM Netfinity servers used in this benchmark, Windows 2000 Advanced Server provides organizations a powerful and scalable alternative to more expensive proprietary solutions.

Find out more about the Windows 2000 platform at: <http://www.microsoft.com>

IBM DB2 Universal Database

DB2 Universal Database was selected for its scalability and performance. It treats all nodes involved as parts of a single database, rather than a federation of separate databases. The shared nothing architecture improves scalability by having less server-to-server network messages due to the use of function shipping instead of I/O shipping. It can further minimize message passing and bottlenecks by passing partial results between remote nodes for further processing instead of always returning to the coordinator node for central processing. Across the 32 servers, there is only one DB2 instance and one DB2 database. Each server is configured to manage one partition, representing 1/32 of the database, except for the first server, which has two logical partitions; one is a repository for the catalog (meta data), and the other is a regular partition.

DB2 is optimized for the clustering capabilities of the Microsoft Windows 2000 operating system and was the first database certified for the Windows 2000 platform. DB2's shared-nothing clustering database architecture spreads a single database image across multiple servers with parallel implementation of SQL and the utilities for maximum performance and scalability. The hash based partitioning of DB2 minimizes administration and requires no code changes for spreading the database to multiple nodes. The shared-nothing architecture allows each node to own a portion, or partition, of the database without the need to share disk or memory access. This architecture bypasses the scalability limitations imposed by a shared-disk architecture, allowing any number of additional nodes to be added to the cluster as scalability needs increase. The data is distributed among the nodes by DB2 using a hash based partitioning algorithm. This technique has the benefits of being simple to administer and minimizing "hot spots" in the data that are common to range based partitioning. This architecture also means that referential and integrity checking constraints, indexes and authorizations are only done once and apply to all nodes in the cluster.

DB2 addresses OLTP workload scalability in a number of ways. First, the engine supports multiple concurrent requests on any given node at any time allowing for concurrency within a node. Second, any OLTP request that cannot be satisfied by the receiving node (coordinator) is forwarded to the node(s) that can satisfy the request thus enabling the concurrency of multiple cluster nodes. Third, any node in a DB2 cluster can receive requests allowing clients to be balanced across the nodes in a cluster to distribute the coordinator workload. Specifically, in the Firestorm benchmark, there were 96 machines sending transactions to the cluster. Each of the 32 nodes had 3 of these machines connecting to it, so that the workload was balanced evenly across all 32 nodes.

DB2's optimizer dispatches and manages multiple requests in parallel via either a multithread or multiprocessing model. This allows DB2 to efficiently service large volume OLTP environments. For complicated SQL requests, DB2 also has the ability to break a single SQL statement into multiple sub-tasks that are then executed in parallel to deliver excellent performance for a complex workload.

DB2 allows multiple storage models that allow for database space to either be dynamically allocated or pre-allocated using either the native file system or raw I/O. DB2's optimizer automatically considers the number of I/O devices, the latency and throughput of the I/O devices when determining how much I/O parallelism a particular system can support. The optimizer includes a number of techniques for identifying opportunities for asynchronous disk access. DB2 supports efficient memory utilization by support of large or multiple data buffers, or buffer pools. This capability allows hot tables or indexes to be given a dedicated area of memory so they will not be paged out.

To ensure database availability, DB2 always exploits the capabilities native to the platform. On Windows 2000, that means the Microsoft Cluster Service (MSCS) available in either Windows 2000 Advanced Server or Windows 2000 Datacenter Server. DB2 integrates with MSCS to deliver high availability of a single DB2 node or multiple nodes of a partitioned database. DB2 can span multiple of the 2 or 4 node availability clusters provided by MSCS to assure high availability for even the largest DB2 EEE clusters.

Find out more about DB2 at: www.ibm.com/software/db2

TPC-C Definition

The TPC Benchmark C was developed by the Transaction Processing Performance Council (TPC), which was founded to define transaction processing benchmarks and to disseminate objective, verifiable performance data to the industry. The TPC describes this benchmark in Clause 0.1 of the specification as follows: TPC Benchmark C is an On-Line Transaction Processing (OLTP) workload. It is a mixture of read-only and update-intensive transactions that simulate the activities found in complex OLTP application environments. It does so by exercising a breadth of system components associated with environments, which are characterized by:

- The simultaneous execution of multiple transaction types that span a breadth of complexity
- On-line and deferred transaction execution modes
- Multiple on-line terminal sessions
- Moderate system and application execution time
- Significant disk input/output
- Transaction integrity (ACID properties)
- Non-uniform distribution of data access through primary and secondary keys
- Databases consisting of many tables with a wide variety of sizes, attributes and relationships
- Contention on data access and update.

The performance metric reported by TPC-C is a "business throughput" measuring the number of orders processed per minute. Multiple transactions are used to simulate the business activity of processing an order, and each transaction is subject to a response time constraint. The performance metric for this benchmark is expressed in

transactions-per-minute-C (tpmC). To be compliant with the TPC-C standard, all references to tpmC results must include the tpmC rate, the associated price-per-tpmC and the availability date of the priced configuration.

Despite the fact that this benchmark offers a rich environment that emulates many OLTP applications, this benchmark does not reflect the entire range of OLTP requirements. In addition, the extent to which a customer can achieve the results reported by a vendor is highly dependent on how closely TPC-C approximates the customer application. The relative performance of systems derived from this benchmark does not necessarily hold for other workloads or environments. Extrapolations to any other environment are not recommended.

Benchmark results are highly dependent upon workload, specific application requirements, systems design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, TPC-C should not be used as a substitute for a specific customer application benchmarking when critical capacity planning and/or product evaluation decisions are contemplated.

For more information: <http://www.tpc.org>

Summary

The Firestorm project brought together the industry leaders IBM, Intel and Microsoft. The goal was to prove that industry-standard hardware and software could scale to meet the most demanding application environments. Where as, many customers will not need the performance delivered in a 32 node cluster, our testing has proven the scalability of the hardware and software platform for smaller clusters of 4 – 8 or 16 nodes. Combining the enterprise-class operating environment of Windows 2000 with IBM DB2 on IBM Netfinity 8500R servers demonstrates the commitment of industry leaders to deliver the best in power, performance and scalability.

As a Microsoft Windows 2000 Global Launch Partner, IBM Netfinity has been supporting large enterprise customers for more than a year to prepare for migration to Windows 2000. And, with the IBM ServerProven™ and ClusterProven™ programs, Netfinity can ensure that the right software partners, including the IBM Software Group, have fully tested business solutions on the Netfinity server platform. Eleven Solution Partnership Centers around the world are working with ServerProven software partners to enable them to migrate their applications to Windows 2000.

Today, more than 4,000 IBM and Business Partner specialists have been trained and certified in the TechConnect® program for IBM Netfinity. IBM and Microsoft are team players driving technology development and broadening industry standards across a variety of platforms and market spaces. This TPC-C benchmark achievement is another example of the strong technical collaboration between IBM, Intel and Microsoft to develop industry-leading solutions for today's enterprise customers.

Customer Responses

Our customers express the importance of the Firestorm proof-point in their quotes below:

“Scalability concerns for e-businesses are a worry of the past. With this OLTP Benchmark we receive the cooperative efforts of IBM, Intel and Microsoft yielding a standardized, available and tested solution with double the transaction capabilities of anything else before. These technologies aren’t ideas still being tested, or dreams of engineers -- this solution exists now.”

-Perry Cain, Vice President, **Neoteric Solutions**

“The scalability of the hardware and software in this OLTP Benchmark extends the growth potential that new and existing companies need to stay competitive in the new world of e-business. In this new world we store data that is dynamic rather than static, and in turn require scalable and dynamic data structures. Having given us this very thing and more, is what makes this IBM, Intel and Microsoft collaboration a milestone in e-business.”

-Alan Ivie, Senior Consultant, Research & Development, **Preisaukunft.de**

“This OLTP Benchmark announcement assures us that we have made the right choice. With it we have been able to start small and scale up fast, without painful redevelopment costs. And as the expected growth of Suppleye.com continues, we can enjoy the hard earned profitability we have all made so many sacrifices for. We now see the world’s fastest and largest database running on the very same technology from DB2, Intel, and Windows that we used in developing Suppleye.Com. “Dot Coms” take notice! The low cost of a cluster of this type should have you looking at profits *while* growing.”

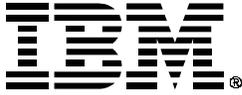
-John Meek, President, **Suppleye.com**

“The OLTP Benchmark constitutes a solution that will entirely bypass the normal glitches and costs of second implementations that accompany exponential transaction growth rates. It also offers scalability for e-businesses affected heavily by the transaction spikes associated with the holiday seasons. This is the type of cooperation between industry leaders that we should expect. With IBM, Intel and Microsoft making a move like this, others are bound to follow.

-Marshall Freiman, CTO, **Web Emporium LLC**

“We had originally chosen IBM because we wanted solutions that would support industry standards. Using the OLTP Benchmark will add Intel and Microsoft two additional standard supporting industry giants. With this solution a company like ours could enjoy a scalability that would achieve ROI much sooner than in the past. This is a very exciting industry development.”

-Michelle Farabaugh, Senior VP of Marketing and Strategic Planning, **West Marine Inc.**



© International Business Machines Corporation 2000

IBM Corporation
Marketing Communications
Software Group
Route 100
Somers, NY 10589

Printed in the United States of America.
07-00
All Rights Reserved

This publication was produced in the United States of America.
IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products or services available in your area.

IBM hardware products are manufactured from new parts or new and serviceable used parts. Regardless, our warranty terms apply.

Actual performance and environmental costs will vary depending on individual customer configurations and conditions.

® ClusterProven, DB2, DB2 Universal Database, IBM, the IBM logo, Netfinity, ServerProven, ServRAID, and WebSphere are trademarks of International Business Machines Corporation in the United States, other countries or both.

™ IMS, RMF, S/390 Parallel Enterprise Server, ClusterProven, and TXSeries are trademarks of the International Business Machines Corporation in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries or both.

Intel, Pentium and Xeon are trademarks or registered trademarks of Intel Corporation.

Other company, product, and service names may be trademarks or service marks of others.