Microsoft®

Windows NT® Server

Server Operating System

Scalability in Microsoft® Windows NT® Server 4.0, Enterprise Edition

White Paper

Abstract

Microsoft® Windows NT® Server version 4.0, Enterprise Edition meets the needs of the overwhelming majority of the enterprise computing marketplace, demonstrating scalability through SMP and/or clustering in a wide variety of database, application, Web and messaging workloads, as measured by industry standard and/or independently audited benchmarks.

Microsoft Corporation • One Microsoft Way • Redmond, WA 98052-6399 • USA

0998

CONTENTS

INTRODUCTION

Microsoft® Windows NT® Server 4.0, Enterprise Edition meets the needs of the majority of the enterprise computing marketplace today, providing near-linear scalability for very large database, application, and messaging systems, as measured by industry standard and independently audited benchmarks.

Specifically, Windows NT Server-based systems deliver computational, transaction, and I/O capacity for the majority of application workload profiles. The operating system allows applications to scale to support thousands of online transaction processing users, or mail and messaging client systems on a single, large symmetric multiprocessing (SMP) server. In replicated, distributed and/or partitioned systems, Windows NT Server, Enterprise Edition can support even larger user populations, which, in some cases, exceed more than 100,000 users across multiple servers.

Although a more complete analysis is treated later in this paper, Table 1 summarizes scalability gains reported for all major workloads on Windows NT Server- and Window NT Server, Enterprise Edition-based systems since Scalability Day. Results were achieved using industry standard and independently audited benchmarks. Database software used for online transaction processing (TPC-C) and enterprise resource planning (SAP Sales and Distribution Users) benchmarks were Microsoft SQL Server™ 6.5, Enterprise Edition and Microsoft SQL Server 7.0 Beta 3, Enterprise Edition, respectively. ZD Benchmark Operation developed ServerBench, which measured application server performance. Microsoft LoadSim is a tool that measures the performance of Microsoft Exchange Server. Microsoft Windows NT Server's Internet Information Server (IIS) version 3.0 and IIS version 4.0 were the Web servers benchmarked using the SPECweb96 file-mix. Numbers are for Intel-based servers only.

Table 1 – Windows NT Server 4.0, Enterprise Edition Scalability Performance

Workload	Benchmark	Scalability Day	Today (9/21/98)	Gain
Online Transaction Processing	TPC-C	8,028 transactions per minute	18,332 transactions per minute	128%
Enterprise Resource Planning	SAP Sales & Distribution (SD) Users	1,011 SD users	2,288 SD users	126%
Application Server Performance	ServerBench	905 transactions per second	1,655 transactions per second	83%
Exchange Server Performance[1]	LoadSim	3,000 concurrent users per server	15,000 concurrent users per server	400%
Web Server Performance	SPECweb96	1,422 operations per second	3,435 operations per second	142%

This paper proceeds to outline issues surrounding overall system scalability, presenting a detailed analysis of the available data, and a technical discussion of scalability in Windows NT Server Enterprise Edition at the hardware and operating system level.

[1] LoadSim results for are Microsoft Exchange 5.0 and Microsoft Exchange Server 5.5, Enterprise Edition, respectively. Exchange 5.0 did not take advantage of memory beyond 256MB; Microsoft Exchange 5.5 and Exchange 5.5, Enterprise Edition did.

WHAT IS SCALABILITY ABOUT?

Scalability refers to how well the performance of a computer system responds to changes in configuration variables, such as memory size or numbers of processors in an SMP-based system. Most often, system architects are interested in the response of a critical *application* to an increase or decrease in resources available to end users. Typical questions that IT managers ask in this regard include:

- As the size of my data warehouse grows, is it possible to keep response time constant merely by adding CPU and memory resources?
- How well can our messaging server respond to a burgeoning user population? How well will the system respond to expanding use of the messaging system beyond mail, as the platform for workflow and work group collaboration services?
- How will empowering stakeholders, such as suppliers and key customers, to connect into systems through the Internet (with appropriate security controls) affect system performance?
- How will increasing complexity in database transactions to accommodate new user requirements affect response time, batch throughput or the number of simultaneous users who can be supported by a fixed server configuration?

There is no single answer to these questions, because answers depend on the complex interaction of multiple variables, such as system components of the underlying hardware, characteristics of the network, application design, and the architecture and capabilities of the operating system. This applies to Windows NT Server-based systems as well as other systems that customers may consider deploying in the data center.

Microsoft's design goal has been to architect Windows NT Server so that customers have the flexibility to scale Windows NT Server-based systems *up*, *out*, or *down* without compromising the multipurpose and price performance advantages of the Windows NT Server platform.

"Scaling up" is achieved by adding more resources, such as memory, processors, and disk drives to a system. This technique enables Windows NT Server, Enterprise Edition to support very large decision support, enterprise resource planning (ERP), and data mining applications in clustered, highly available and manageable systems on industry-standard hardware. But scalability is more than brute force. Application design, database tuning, network configuration, and well-developed data center procedures matter materially.

"Scaling out" delivers high performance when the throughput requirements of an application exceed the capabilities of an individual system. By distributing resources across multiple systems, contention for these resources can be reduced, and availability improved. Clustering and system services, such as reliable transaction message queuing, allow Windows NT Server-based applications to scale out in this manner.

“Scaling down” can also deliver tangible business benefits. For example, a company may divest a division or move a formerly centralized function from the data center to a division. In another instance, the IT department may decide to transition, or possible repartition workloads to improve overall performance and enhance life-cycle management of server systems. In either case, Windows NT Server-based systems make this possible without having to completely overhaul the system, or reprogram applications.

SCALABLE HARDWARE PLATFORMS FOR WINDOWS NT SERVER

Scalable, well-designed hardware is a huge advantage for the customers of any operating system. It provides the benefits of lower cost and broad choice, and a constant stream of innovation. These benefits are very much a part of the Windows NT story, and a significant part of its growing popularity. Today, industry leaders such as Compaq, Data General, Dell, Fujitsu, Hewlett-Packard, Hitachi, IBM, Intergraph, NCR, NEC, Sequent, SNI, and Unisys build systems that support Windows NT Server, Enterprise Edition.

Specifically, this section explores the fundamental hardware elements required to support scalable computing, and briefly highlights selected platforms to demonstrate how Windows NT Server takes advantage of the advanced capabilities that these systems provide. However, this section is not a substitute for detailed information on hardware scalability from Microsoft and its systems vendors.

Processor Characteristics

Multiprocessing means allowing multiple instruction streams to execute simultaneously. Symmetric multiprocessing (SMP) is a technology used to support multiprocessing by applying multiple instances of identical processors to a problem set.2 Therefore, the processor itself is the basic building block of all SMP configurations.

Multiprocessor systems for Windows NT Server-based systems rely on the Intel Architecture or Compaq Alpha processor families. Both processor families use sophisticated on-chip cache and bus-controller logic designed specifically for SMP configurations.

For example, earlier this year Hewlett-Packard posted 16,257.20 tpmC throughput on an 8-way SMP NetServer LXr Pro8 with Pentium Pro 200-MHz processors using Windows NT Server 4.0, Enterprise Edition and SQL Server 6.5, Enterprise Edition. Three months later, Compaq posted tpmC results of 18,127.40 with a 4-way SMP ProLiant 7000 system configured with Pentium II Xeon 400-MHz microprocessors. Windows NT Server 4.0, Enterprise Edition was the operating system and SQL Server 7.0 Enterprise Edition was the database server.3

Today, Windows NT Server and Windows NT Server, Enterprise Edition support the 64-bit RISC-based Alpha microprocessor from Compaq. Microsoft is also developing a 64-bit version of Windows NT Server, which will support Intel Architecture 64 processors, starting with the microprocessor code-named "Merced" and the 64-bit Alpha microprocessor from Compaq.

---

[2] This is different from *asymmetric processing*. Asymmetric processing requires that one or more processors be set aside for exclusive use by the operating system or for specific functions such as I/O processing, while the remainder of the processors run user applications.

[3] For more information on the hardware configuration used in each of these benchmarks, visit the Transaction Processing Council.

Interconnection Strategies

Designing hardware for scalable computing depends on the interconnect strategy. Specifically, as the number of processors increases, and individual processor speeds increase, the bottleneck preventing further growth in throughput moves to the bus architecture, which ferries information between processors and memory. Today, there are essentially two interconnect strategies:

- All processors use a single shared address space; or
- Each processor has a private memory—processes communicate with one another using a message-passing mechanism.

The overwhelming majority of today's multiprocessing configurations use the first strategy. Scalable clustered configurations use the second strategy. For example, Compaq's ServerNet System Area Network (SAN) is a high speed, scalable, and reliable interconnection technology that allows for interprocessor communication and device I/O transfer rates up to 125 MB per second. In April, Compaq posted benchmark results of a ProLiant 6500 6-Node Cluster using ServerNet technology. This system, which delivered 27,383.40 tpmC at a cost of $71.50 per transaction, used Windows NT Server 4.0, Enterprise Edition.

Today, chip designers are embedding greater "intelligence" into microprocessors. For example, the Pentium Pro and Pentium II from Intel, and the Digital Alpha from Compaq prefetch data and instructions into the caches to minimize the time that processors must spend waiting for data to cross the bus. Windows NT Server 4.0 and Windows NT Server 4.0, Enterprise Edition have been designed to take advantage of these features in order to maximize the effectiveness of these processor caches.

In the short term, Unisys and others are developing a next generation of data center hardware for Windows NT-based computing, incorporating a variety of advanced technologies. The Unisys systems use a technology that they call Cellular Multiprocessing; the technology combines a third level of caching, directory-based cache coherency, and a fully connected point-to-point crossbar, eliminating some of the bottlenecks associated with traditional bus architectures. Unisys asserts that this method of interconnection is likely to enable breakthroughs in both application and I/O scalability, and as much as, 6.4 GB per second overall throughput on up to 24 parallel PCI I/O buses. Figure 1 gives a conceptual view of this architecture.[4]

[4] Up to four Intel processors share two Intel memory buses connected to a single Level 3 memory cache. Two such processor "pods" and a pair of I/O pipes, each connected to three PCI I/O buses, are fully connected through a crossbar to each of the four memory storage units. Up to four crossbars can be fully interconnected with up to 32 GB of main memory, 96 primary PCI I/O slots, and 32 IA-32 or IA-64 processors.

Figure 1 –Cellular Multiprocessing Architecture

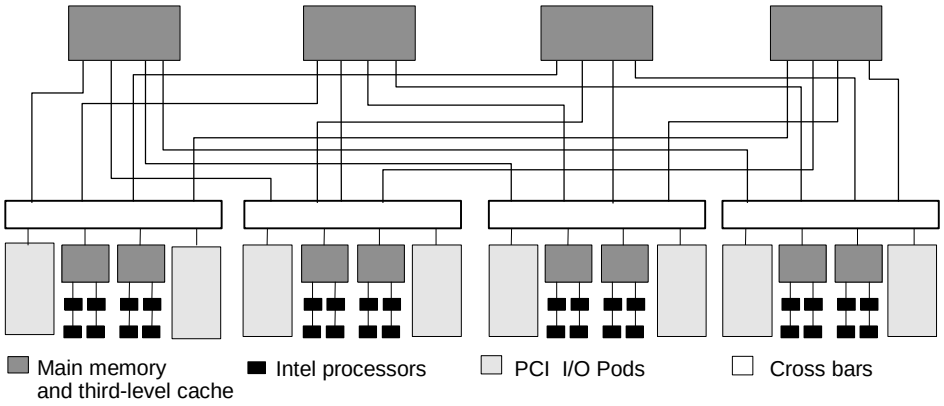

Source: Unisys

Unisys states that this system can be configured as a single 32-way SMP system, or divided into as many as eight subsystems, each with its own copy of the operating system. In the second case, each instance of the operating system will have the ability to operate in complete isolation from the other parts of the system. This allows a single system to reflect easily the heterogeneous nature of the software needs of a data center. For example, such a system can host one instance of Windows NT Server 4.0 with SQL Server 6.5, another with Exchange Server, and another with Windows NT Server 5.0.

Scalable Memory and Storage Subsystems

Increasing physical memory capacity can significantly benefit database and Web server applications because I/O to memory is much faster than I/O to disk. The ability to address these memories is, of course, required; most 32-bit operating systems can only address 4 GB of physical memory. Following is a look at how the standard edition of Windows NT Server differs from the Enterprise Edition with respect to this issue, and further, how Windows NT Server 5.0 differs from Windows NT 4.0 on this issue.

Windows NT Server 4.0 provides a virtual 2-GB address space to every application. An additional 2 GB (for a 4-GB total between system and applications) is reserved by the system itself. Windows NT Server 4.0, Enterprise Edition extends this capability by allowing large memory-aware applications to use up to 3 GB, reserving only 1 GB for the operating system.

Microsoft recently announced its Enterprise Memory Architecture (EMA). Available in Windows NT Server 5.0 Enterprise Edition, EMA makes it possible for Microsoft to support systems that use the Compaq Alpha VLM and Intel Pentium II Xeon PSE-36 (Page Size Extension 36)-capable microprocessors. It will support up to

64 GB of physical memory, independent of the hardware platform, through special Win32® application programming interfaces (APIs).

There are differences between how the Compaq Alpha and Intel Pentium II Xeon microprocessors implement this technology. Special Win32 APIs support for EMA on Compaq Alpha-based systems. These APIs are extensions to the existing APIs that deal with virtual memory. For example, VirtualAllocVlm and VirtualFreeVlm act like their 32-bit equivalents, but use 64-bit addresses and size parameters. In the case of Pentium II Xeon-based systems, an Intel-developed memory driver provides access to the 33 to 36-bit address.

Table 2 shows the increase in available memory to applications in various versions of Windows NT Server on different processor architectures.

Table 2 – Application Memory Capacity

Processor Architecture	Platform	Operating System	Application Memory Capacity
Intel	All	Windows NT Server 4.0, Windows NT Server 5.0	2 GB
Alpha	All	Windows NT Server 4.0, Windows NT Server 5.0	2 GB
Intel	All	Windows NT Server, Enterprise Edition 4.0	3 GB[5]
Intel	Xeon	Windows NT Server, Enterprise Edition 4.0	63 GB[6]
Alpha	All	Windows NT Server, Enterprise Edition 4.0	2 GB[7]
Alpha	All	Windows NT Server, Enterprise Edition 4.0	6 GB[8]
Intel	Xeon	Windows NT Server, Enterprise Edition 5.0	63 GB[9]
Alpha	All	Windows NT Server, Enterprise Edition 5.0	30 GB[10]

Source: Microsoft

Customers seeking highly scalable solutions will use software and hardware solution in combination. For example, the Windows NT File System (NTFS) uses 64-bit addresses and file offsets. This allows for theoretically immense file and volume sizes. Today, there are external limitations on volume and file sizes imposed by the logical disk manager's disk partitioning system and by the underlying hardware. NTFS will continue to scale as these limitations are broken down. Enhancements to the NTFS file system in Windows NT Server 5.0 improve performance, and incorporate new features like per-user disk quotas and distributed link tracking.

5 CISC removes some “slack” from the memory addresses in kernel.
6 Using IntelPSE36 driver, which Microsoft does not specifically support; also, not all chipsets support addresses above 4 GB.
7 This is due to RISC code size and impact on “compressing” the kernel.
8 Using Oracle VLM, which Microsoft does not specifically support; also, not all AXPs have the physical ability to support this much memory.
9 Only data, not .exe, .com, .sys or .dll code, can be in the addresses above 4GB.
10 *Ibid.* Also, both platforms have chipset, BIOS, and physical capacity issues concerning total RAM capacity which are unrelated to Windows NT Server 4.0, Enterprise Edition.

Disk striping, and disk mirroring, increase the scalability of storage facilities. Striping is a technique used to scale file systems by allowing data from a single logical disk volume to be physically stored across multiple physical disks. Large data objects, for example x-ray images in a health care application, are an example of the kind of data that can be read much faster on striped disks, in which transfers can occur from all members of the strip set in parallel. Mirroring is a technique used to enhance both file system performance and reliability by maintaining consistent copies of data across multiple disks. On read operations, either copy of a mirror set can be read, allowing for greater performance. A fault-tolerant store maintains error-correction information on one disk for data stored on the remaining disks in a set. Windows NT Server supports software RAID, as well as many third-party hardware RAID solutions

Caching controllers and Intelligent Input/Output (I20) can improve scalability by offloading I/O operations from the main processor, freeing it to do real user work instead of merely moving bits. This results in greater throughput and lower CPU utilization. I2O features for Windows NT Server 5.0 include base support, specialized board support, network adapters, and Redundant Array of Inexpensive Disks (RAID) cards. Fast interconnects are also important. Support for Fibre Channel in Windows NT is an example. Specifically, Fiber Channel is a technology for 1-gigabit-per-second data transfer that maps common transport protocols such as SCSI and IP, thus merging networking and high-speed I/O in a single connectivity technology.

Two emerging technologies that show great promise are Storage Area Networks (SANs) and Network Attached Storage (NAS). Both SAN and NAS allow system administrators to collect large amounts of disk storage in highly localized and manageable units, such as large servers full of RAID arrays. Application servers and users can then store their data on these large bit buckets. First, NAS devices are analogous to a dedicated file server. A device includes large amounts of storage managed by a local file system. Multiple application servers can share NAS devices using standard network-level protocols. Second, SAN devices are essentially large storage buckets. The file system runs on the application server and uses block level I/O to access SAN devices. SANs typically use fibre channel and fibre switches to connect to application servers. NAS is more appropriate for traditional LANs with heterogeneous hardware and software. SAN is more appropriate for homogeneous hardware and software used in a computing cluster.

MEASURING SCALABILITY

Scalability is measured different ways: the size of the problem, the number of operations per unit time, or the size of the user population. As mentioned earlier, many factors have a role in determining a system's overall scalability, including processor throughput and configuration, memory capacity, disk and network I/O bandwidth, operating system features, and application software design.

Although the most accurate comparisons are possible only when analyzing the performance of specific, customer-relevant applications in a specific installation and usage scenario, several industry standard benchmarks measure both relative and absolute performance and price/performance in useful ways.[11] When used knowledgeably, these benchmarks help customers to understand how hardware and software alternatives can improve performance, and help customers to understand how similarly configured servers compare. Benchmark data represents only one element in the buying criteria. Other important factors to consider are price/performance, development environment, application availability, total cost of ownership, vendor reputation, and future technology directions.

This section describes benchmarks from the Transaction Processing Council (TPC-C), SAP (SD Users), ZD Benchmark Operation (ServerBench), Microsoft (LoadSim), and the Open Systems Group (SpecWeb). The following section, *Scalability Workloads*, charts performance metrics for each of these benchmarks.

Online Transaction Processing (TPC-C Benchmark)

The TPC-C benchmark is the most widely accepted measure of a transaction processing system's cost and performance. TPC-C one of several benchmarks measuring relative scalability defined by the Transaction Processing Council. This particular benchmark emulates the characteristics of many real-world transaction-processing systems, making it an excellent measure for scalability of online systems.

Specifically, TPC-C is relatively complex application that models an order management system for a wholesale supplier. The benchmark specifies the user interface, one short-duration transaction type (New Order), and four longer-duration transaction types (Payment, Order-Status, Delivery, Stock-Level) that access nine database tables.

TPC benchmarks focus on two measures of a system: *total throughput* and *price/performance*. For TPC-C, throughput is the number of short-duration transactions per minute (TPM), while the system also executes longer-duration transactions. For a system to qualify under this benchmark, 90 percent of New Order requests must complete in less than 5 seconds while the system fulfills the workload of other transactions.

There are two ways to increase performance in this benchmark–scaling up and

[11] The Transactions Processing Council is an example.

scaling out. First, a single computer system can scale up by adding memory, more processors, high-speed disks. Tuning the operating system, placing data and log files on disk, and adding a TP monitor determines how these resources are used. Second, a transaction-processing system can also "scale out" by partitioning the data and workload across more than one computer. In this case, clustered computers must maintain a degree of location transparency required by this benchmark, and they do so by the use of databases that can span multiple nodes of a cluster.

Price/performance is the cost of a system over three years divided by the number of transactions it can perform in one minute. Terminals, communications equipment, backup storage, software, the computer (or computers), and three years of maintenance are included in the cost of a system.12 Therefore, if the total system costs $500,000 and the throughput is 10,000 TPM, the price/performance of the entire system is $50 per tpmC.

Enterprise Resource Planning (SAP SD Users Benchmark)

Enterprise Resource Planning (ERP) encompasses integrated application suites spanning multiple functional disciplines of the enterprise. ERP "modules" can include financials, planning, inventory, materials production, quality, sales and ordering, distribution, and transportation. Together, these packages allow companies to integrate and streamline core business functions, which improve operating efficiency on a global basis. Baan, PeopleSoft, and SAP are well-known independent software vendors offering ERP solutions.[13]

Each of the major ERP vendors has its own respective programs that measure the performance of different computer systems. When used carefully and responsibly, these benchmarks can help IT managers and users to compare hardware and operating system software. For example, BAAN Reference Users and PeopleSoft Payroll are frequently cited benchmarks measuring ERP performance for Baan and PeopleSoft, respectively.

SAP, for example, has several R/3 System standard application benchmarks, the most frequently cited being the Sales and Distribution (SD) benchmark. SAP and its hardware partners developed this tool to test hardware and database performance. This benchmark simulates the maximum load of concurrent users that a system can tolerate and still provide "reasonable" response time.[14] It works by placing substantial load on a system during the testing of hardware, system software components, and relational database management systems.

---

12 Costs do not include system architecture and ongoing maintenance and service costs.

13 To better understand Microsoft's initiatives in this space, visit the Enterprise Resources Planning center.

14 The SD benchmark consists of eight dialog steps. Components of the program include creating an order with five line items, changing the delivery, and creating the invoice. These results produced by the programs provide objective statements about how a system performs and allow customers to compare different computer systems.

In real deployments, the response time (or “think time”) is unpredictable, and the particular workload may be constrained by the performance of other SAP modules. The time an average user needs in order to process the information received from a transaction before making another request of the database naturally vary tremendously from user to user and application to application.

Application Server Performance (ServerBench Benchmark)

The scaling characteristics of applications can vary from processor architecture, operating systems, and system design. To understand scalability better, it is necessary to understand the causes of these variations. For example, different programming languages can be used to piece together what appears to be the same application on entirely different code bases, implemented differently for different operating systems.

ServerBench is an industry-accepted benchmark developed by ZD Benchmark Operation. This benchmark measures performance of application servers in a client/server environment. Specifically, the benchmark program consists of transactions that are the basic unit of work that the client uses to stress the server. Each transaction consists of the work request a client sends to the server and the response the client receives. Each client measures how long each transaction takes and how many transactions take place. The client then calculates the Transaction Per Section (TPS) number. ServerBench uses an individual client TPS number to produce the overall TPS number for the server.

Hardware vendors and independent software vendors can download ServerBench from Ziff-Davis on-line ZDNet(TM) services. Readers may want to learn more by examining the ServerBenchFAQ .

Exchange Server Performance (LoadSim Benchmark)

No generally accepted benchmarks exist for electronic mail systems that are analogous to the TPC-C benchmarks for online transaction processing. Therefore, Microsoft developed a tool that generates performance data called Microsoft Exchange Use Load Simulation utility, or LoadSim. LoadSim simulates a client-user load on an Exchange Server. Its purpose is to enable a single Windows NT Server computer—called a LoadSim client—to simulate multiple Microsoft Exchange client users. The current version of LoadSim provides simulation of MAPI protocol clients, as well as support for POP3, NNTP, LDAP, and HTTP client protocols.

The LoadSim profile governs the operation of LoadSim users. This profile controls multiple factors, such as how long a LoadSim “day” is, how many e-mail messages to send in a day’s time, and how many times to open and read existing e-mail messages. A LoadSim test is considered valid when the 95

percent of the test cycles have a response time of less than one second and no leftover work items exist in the Exchange Server queues.

LoadSim has limitations because it is a simulator. First, LoadSim results do not properly address when users are logging on and off. During a test, all users log on sequentially. Once this process is complete, users begin their tasks. If users typically log on and off multiple times during the day, the Windows NT logon can have an impact on server and network rates of utilization, especially when many users log on simultaneously. Second, the response times generated by LoadSim are based on server response to clients and do not account for strictly client-side actions such as rendering rich text once a message is received. Despite these limitations, LoadSim is very useful in understanding relative performance issues.

Web Server Performance (SPECWeb)

There are no Web server benchmarks as widely accepted and as thoroughly defined by the Transaction Processing Council. However, Standards Performance Evaluation Corporation (SPEC) develops suites of benchmarks intended to measure computer performance.[15] Specifically, they developed SPECWeb96, the first standardized World Wide Web server benchmark that measures basic Web server performance. Though not perfect, this benchmark is a standardized test, which numerous vendors and users have embraced.[16]

SPECWeb96 targets Web server systems designed to handle millions of hits per day and hundreds or thousands of hits per second. The benchmark program uses a file mix developed after analyzing the server logs from several Web service providers. The SPECWeb96 benchmark measures HTTP GET commands, the command used to retrieve a Web page. Results are reported in terms of operations per second.

Table 3 – SPECWeb96 Workload Distribution

File Size	Percent of Workload
Less than 1 KB	35%
1-10 KB	50%
10-100 KB	14%
100 KB – 1 MB	1%

Unfortunately, this benchmark does not contain several elements central to Web application processing. For example, the benchmark only measures static page

[15] The three groups are Open Systems Group (OSG), High Performance Computing Group (HPC), and Graphics Performance Characterization Group (GPC).

[16] SPECWeb96 results are not subject to the same rigorous reporting and independent auditing as TPC benchmarks. SPEC reviews the results, but no warranties are made with respect to accuracy or veracity of the data.

serving. CGI calls, script execution, and security event handling are not part of the test. Generation of dynamic content would require the use of CGI or script execution. In addition, tests are based on the HTTP 1.0 protocol. Also excluded from tests are HTTP 1.1 Keep-Alive packets, which can dramatically improve HTTP server performance. Consequently, measurements overstate the performance of an HTTP server when CGI calls are made, scripts are executed, and security mechanisms are utilized. Similar to other benchmarks cited, SPECweb may be more useful to assess relative performance than it is to assess absolute performance, because the demands of Web sites are highly variable from enterprise to enterprise.

SCALABILITY WORKLOADS

Based on customer feedback, analyst reports, and ISV usage data, Microsoft believes that Windows NT-based solutions can address the overwhelming majority of OLTP, application, Web, and messaging requirements in the marketplace today. As reported by SAP, over 80 percent of their installations are 400 simultaneous users or less, which is a fraction of the throughput that Windows NT supports today. This is especially significant because SAP is a very demanding and complex application, that exercises many aspects of the operating system and the underlying hardware platform.

This section examines widely available data on a set of specific workloads to better illuminate the significant and growing capabilities of Windows NT Server, especially Windows NT Server, Enterprise Edition.

TPC-C Benchmarks

When analyzing TPC-C benchmarks, it is important to remember that an average user of a system, benchmarked using the TPC-C workload, will issue about 1.2 transactions per minute. Therefore, a system that can produce 12,000 transactions per minute is capable of supporting roughly 10,000 simultaneous users in this workload.

The earlier section, *Scalable Hardware Platforms with Windows NT*, highlighted the recently posted TPC-C benchmarks from Compaq Computer. The test configuration was a 4-way ProLiant 7000 system with Pentium II Xeon 400-MHz microprocessors. With Windows NT Server 4.0, Enterprise Edition and Microsoft SQL Server 7.0, Enterprise Edition, the system recorded 18,127.40 TPC-C throughput and $26.10 per transaction. The closest, comparably configured Ultra Enterprise 450 from Sun Microsystems had only two-thirds the performance (11,559.6 tpmC), and more than double the cost per transaction ($56.60 per tpmC).[17]

Statistics for 8-way systems *available today* are equally impressive. For example, an 8-way NetServer LXr Pro8 system from Hewlett-Packard delivers *better overall performance at one-third the cost* of an Ultra Enterprise 6000 from Sun Microsystems with 8 250-MHz UltraSPARC processors.[18] The operating system and database for the Hewlett-Packard server was Windows NT Server 4.0, Enterprise Edition and SQL Server 6.5, Enterprise Edition.

More generally, Table 4 shows the Top 10 TPC-C Results by Price/Performance. As of September 15, Unisys reports the system with the lowest price/performance—$25.49 per tpmC. The benchmarked configuration used Microsoft SQL Server 7.0 or SQL Server 6.5, Enterprise Edition on Windows NT

---

17 Clock speeds for Intel Pentium II Xeon and Sun UltraSPARC II were 400 and 300 MHz, respectively.

18 Sun Microsystems has recently posted benchmark results for the Ultra Enterprise 3500, which was configured with 8 336-MHz microprocessors rather than the 250-MHz processors used in the E 3000 benchmark.

Server 4.0, Enterprise Edition. Each Top 10 system used SQL Server and Windows NT Server, Enterprise Edition.

The most cost-effective UNIX-hosted system is a Compaq AlphaServer running Sybase, and delivering a mere 7,023 tpmC at $41 per transaction. The overwhelming majority of UNIX-based benchmarked systems have costs per transaction that range from $50 to $150.

In addition to their UNIX/RISC benchmark data, Compaq demonstrates excellent throughput using Windows NT and SQL Server – over 18,000 transactions per minute at $27 tpmCs. Although the largest and most expensive UNIX implementations can deliver more throughput than the largest Windows NT systems, the price of UNIX relative to Windows NT is painfully high for the overwhelming majority of the marketplace. For example, an Ultra Enterprise 6000 from Sun Microsystems with 16 processors offers approximately 23,000 TPC-C but at a cost of $119 per transaction. For enterprises that are fundamentally transaction driven, this tremendous premium in cost per transaction places UNIX users at an ongoing, daily competitive disadvantage.

Table 4 – Top 10 TPC-C Results by Price/Performance

Company	Price/ Perf.	Throughput	Total Cost	Database Software	Operating System
Unisys Aquanta QS/2 Server	$25	18,154	$462,803	SQL Server 6.5, Enterprise Edition	Windows NT Server 4.0, Enterprise Edition
Compaq ProLiant 7000-6/400-M1	$27	18,127	$473,203	SQL Server 6.5, Enterprise Edition	Windows NT Server 4.0, Enterprise Edition
Unisys Aquanta QS/2 Server	$27	17,700	$463,958	SQL Server 6.5, Enterprise Edition	Windows NT Server 4.0, Enterprise Edition
Compaq ProLiant 5500-6/200	$27	11,748	$312,680	SQL Server 6.5, Enterprise Edition	Windows NT Server 4.0, Enterprise Edition
Acer Altos 19000 Pro4 Pentium Pro 200 MHz	$28	11,072	$301,663	SQL Server 6.5, Enterprise Edition	Windows NT Server 4.0, Enterprise Edition

Company	Price/ Perf.	Throughput	Total Cost	Database Software	Operating System
Dell PowerEdge 6100	$29	10,984	$324,623	SQL Server 6.5, Enterprise Edition	Windows NT Server 4.0, Enterprise Edition
Compaq ProLiant 3000-2P	$31	8,228	$253,569	SQL Server 6.5, Enterprise Edition	Windows NT Server 4.0, Enterprise Edition
Unisys Aquanta HS/6 Server	$32	13,728	$441,289	SQL Server 6.5, Enterprise Edition	Windows NT Server 4.0, Enterprise Edition
Compaq Digital Server 7105	$33	11,358	$370,812	SQL Server 6.5, Enterprise Edition	Windows NT Server 4.0, Enterprise Edition
NEC Express5800 HX4100	$33	12,105	$396,588	SQL Server 6.5, Enterprise Edition	Windows NT Server 4.0, Enterprise Edition

Source: TPC – September 15, 1998

In some workloads, with appropriate software, clustering two or more loosely coupled computers can extend the throughput of a system significantly. For example, a Compaq submitted a TPC-C benchmark running Windows NT Server 4.0, Enterprise Edition with Tandem ServerNet Cluster software. This system of six
4-way 200-MHz Pentium Pro microprocessors delivered 27,383 transactions per minute, the highest throughput of any tested Windows NT-based system, at a cost of $72 per tpmC.[19] By comparison, an IBM SP2 cluster running AIX delivered more than 57,000 tpmC, but at twice the cost per transaction ($148 tpmC). The key point is that Windows NT can scale both in SMP and in clusters, with the appropriate software. While SMP is more cost effective, higher absolute levels of performance can be achieved with well-selected clustered components. This is why Microsoft continues to invest in both improving SMP scalability and cluster scalability.

[19] Sun Microsystems offers a similarly configured system, Ultra Enterprise 6000 with 24 250-MHz UltraSPARC microprocessors, which delivers only 13 percent higher performance, but at a 50 percent higher cost per transaction.

This benchmark data charts scalability on along one dimension—system size. Another important dimension to understand is how scalability increases when adding processors to a single host. This is important because these numbers reflect the cost to users of scaling up a system over time.

Using audited TPC-C benchmarks for a selection of Compaq ProLiant machines, performance increases at near-linear rates as processors are added. At the same time, cost per transaction decreases in the larger SMP systems, with a four-processor system delivering over 11,000 transactions per minute at a total aggregate cost of only $40. A uniprocessor ProLiant 2500 with a single Pentium Pro 200 -MHz processor provides 2,502 transactions per minute at a three-year total cost of $72 per transaction per minute. Figure 2 illustrates this point:

Figure 2 – Cost per TPC-C per minute

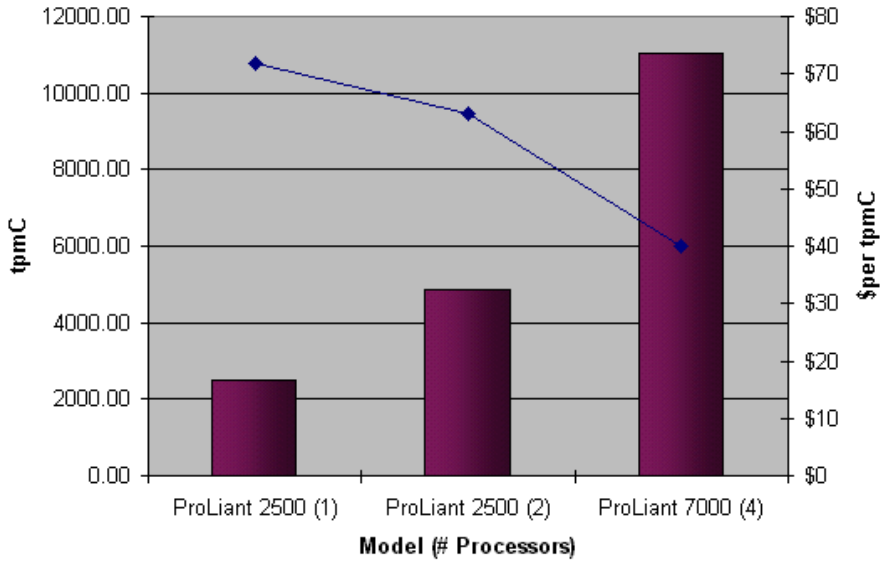

Source: TPC

Taken together, this analysis of audited benchmarks suggests that Windows NT Server, Enterprise Edition-based systems provide excellent scalability by system and per processor. Furthermore, the cost per transaction decreases at greater than linear rates when adding processors to individual systems.

SAP Sales and Distribution (SD) Users Benchmark

Since Scalability Day, Windows NT Server-based systems have made tremendous gains in absolute performance in all major benchmarks developed by the leading vendors of ERP solutions. For example, in August of 1997 Compaq supported

1,011 SD Users with a 4-way ProLiant 7000 6/200 database server. In this benchmark, Windows NT Server 4.0, Enterprise Edition was the operating system and Microsoft SQL Server 6.5, Enterprise Edition was the database server.

In July 1998, using a server system based on Pentium new Xeon 40-MHz microprocessors, Compaq boosted its best SD users number by 126 percent to 2,288 SD Users. The operating system was the same; the database server was Microsoft SQL Server 7.0.[20]

Unlike the Transaction Processing Council, SAP does not publish price/performance statistics for benchmarked system configurations submitted by vendors. However, Aberdeen Group has worked with SAP and major RISC/UNIX and Intel/Windows NT Server vendors to assemble price/performance statistics for the first time. Their report, *Price-Performance of UNIX and NT Systems: Reading the SAP R/3 Benchmark Tea Leaves*, can be viewed at Compaq's password protected Web site, ActiveAnswers for SAP R/3.

Table 5 – SAP R/3 Performance Benchmarks

Date	Hardware Vendor	Server Model	Database	SD Users
8/97	Compaq Computer	ProLiant 7000 6/200 4-way Pentium Pro 200 MHz	SQL Server 6.5, Enterprise Edition	1,011
8/97	Compaq Computer	AlphaServer 8400 10-way Alpha 440 MHz	Informix Online 7.23	2,001
8/97	Data General	Aviion 6600 6-way Pentium Pro 200 MHz	Informix Online 7.20	1,300
12/97	Hewlett-Packard	NetServer LxrPro8 8-way Pentium Pro 200 MHz	Informix Online 7.23	1,670
7/98	**Compaq Computer**	**ProLiant 7000 6/400 4-way Pentium II Xeon 400 MHz**	**SQL Server 7.0**	**2,288**

On September 2, Hewlett-Packard announced 2,700 SAP SD benchmark users. The tested configuration was a NetServer LXr 8000 system with 4 Pentium II Xeon 400-MHz processors. The benchmark used Oracle version 8.0.4 and Microsoft Windows NT Server 4.0, Enterprise Edition.

To gauge the validity of the benchmarks, it is important to investigate the raw

[20] The same system reported equally impressive gains on Baan and PeopleSoft benchmarks. Specifically, this configuration achieved 3,232 BAAN Reference Users (BRUs), which is more than 10 times the best SQL Server 6.5 benchmark. SQL Server 7.0 also produced the highest performance of any Windows NT-based system to date in the PeopleSoft Payroll benchmark—41,328 checks per hour.

data and understand how the vendors in most cases had to tune and modify their systems in order to achieve maximum performance on the benchmark. For this reason, systems vendors like Compaq and Hewlett-Packard have established "competency centers" to help customers plan for and deploy critical business systems. A competency center can help customers set up testing and development environments to observe and measure the behavior of a particular application in the Windows NT Server environment.

For example, at Compaq's SAP Competency Center in Houston, Texas, a high-level sizing might merely count the modules expected to receive heavy use, along with the expected number of simultaneous users.[21] Data derived from Compaq's SAP competency center, along with other customer surveys indicates that relatively few SAP installations require support for more *concurrent* interactive users than can be supported by readily available Windows NT-based systems that are shipping today. In fact, Windows NT Server-based systems host more than half of all SAP R/3 deployments worldwide today.[22] Windows NT Server can now support 90 percent of the new R/3 customers—all but the very largest installations, according to Steve Rietzke (SAP America). Further, Rietzke points out that 80 percent of current SAP deployments today actually support 400 or fewer concurrent users, well within the performance range of a well configured Windows NT Server system.

For example, Pennzoil, an SAP customer, supports 600 concurrent (2,000 named) users in production today, using a 4-way SMP Pentium Pro server, using Windows NT Server and SQL Server 6.5. With a Pentium II Xeon or Alpha-based system, as well as the Enterprise Editions of both Windows NT Server and SQL Server, performance could potentially increase by at least 50 percent.

ServerBench Benchmark

The ServerBench benchmark simulates an application server environment to show the performance improvements on similarly configured systems that vary in the number of installed processors. The systems cited here are Hewlett-Packard 200-MHz Pentium Pro systems running Microsoft SQL Server 6.5, Enterprise Edition on Windows NT Server 4.0, Enterprise Edition. The results show that an eight-processor system delivers almost eight times the throughput of a single processor system.

Not appearing in this graph is Hewlett-Packard's most recently reported 1,655 ServerBench TPS using a 4-way HP LXr 8000 with Pentium II Xeon 400-MHz processors.

---

[21] Hewlett-Packard has a similar organization, SAP-HP Competence Center.
[22] *Business Week*, September 14, 1998.

Figure 3 – ServerBench Peak Throughput per Processor

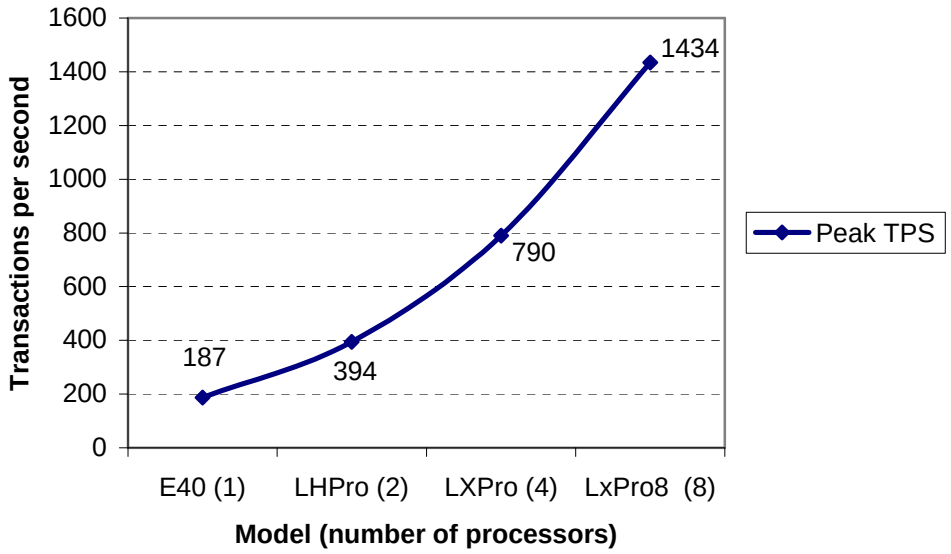

Source: Hewlett-Packard

Microsoft recognizes the likelihood of continuing rapid improvement in microprocessors and system interconnect speeds, and has engineered Windows NT Server to be able to take advantage of these improvements as they occur.

There is excellent scalability in Windows NT for both Alpha RISC and IA 32 architectures, which reflects Microsoft's commitment to both types of systems. For example, throughput measurements of Digital Alpha Server 4100 systems appear in Figure 4. Dynameasure/SQL from Bluecurve generated the load in this particular test. (Dynameasure/SQL is a capacity and performance analysis test suite.[23] It uses a simulated order/entry application to produce these results, and can be modified to reflect customer-specific database and SQL OLTP workloads so that customers can tune it to more fully meet their specific needs.)

[23] Dynameasure/SQL is available in two configurations: Dynameasure/SQL Standard Edition and Dynameasure/SQL Professional Edition.

Figure 4 – Microsoft SQL Server throughput per processor

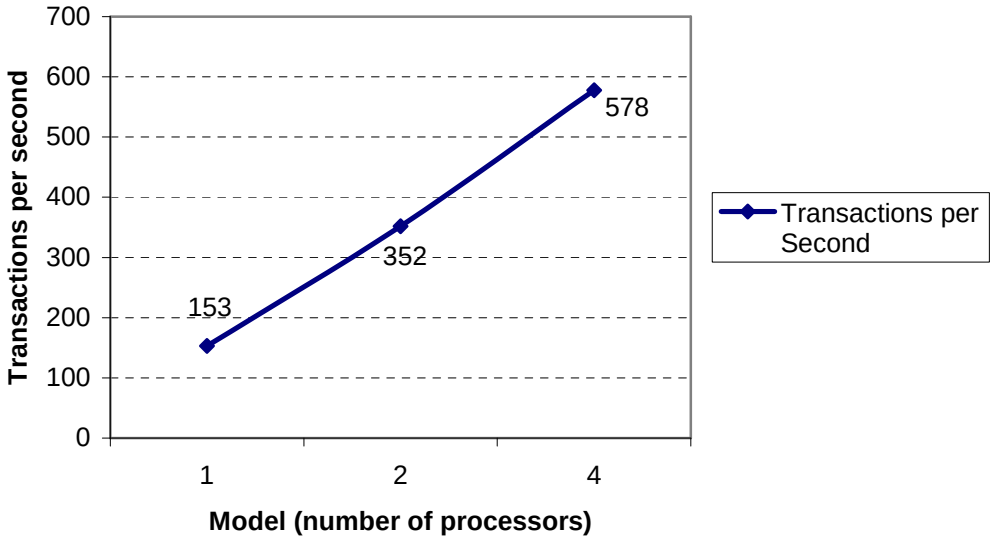

When taken together, these charts indicate that Microsoft SQL Server 6.5, Enterprise Edition with Windows NT Server 4.0, Enterprise Edition are able to take advantage of additional resources in a nearly linear fashion in the presence of increasing client workloads.

LoadSim Benchmark

As important as transaction processing is, it is not the only workload of interest to IT organizations. Enterprises are also properly concerned with messaging, workflow, and collaboration projects. LoadSim provides insight into these workloads. Results indicate that Microsoft Exchange version 5.5, Enterprise Edition can sufficiently scale messaging and collaboration systems as the number of processors increase in SMP systems.

Figure 5 demonstrates a 50 percent increase in the number of simultaneous users supported when the number of CPUs doubles in Compaq server models 3200, 5200, 7105, and 7305, which are Intel-based systems in the formerly-Digital product line. In this product branding scheme, the 3200, 5200 and 7100 series are Pentium Pro-based (light gray bars), while the 5300 and 7300 series are Alpha-based. This demonstrates that for mail server workloads, the scalability of Windows NT Server and Windows NT Server, Enterprise Edition is quite positive for both microprocessor architectures.

Figure 5 – Microsoft Exchange 5.5 User Load

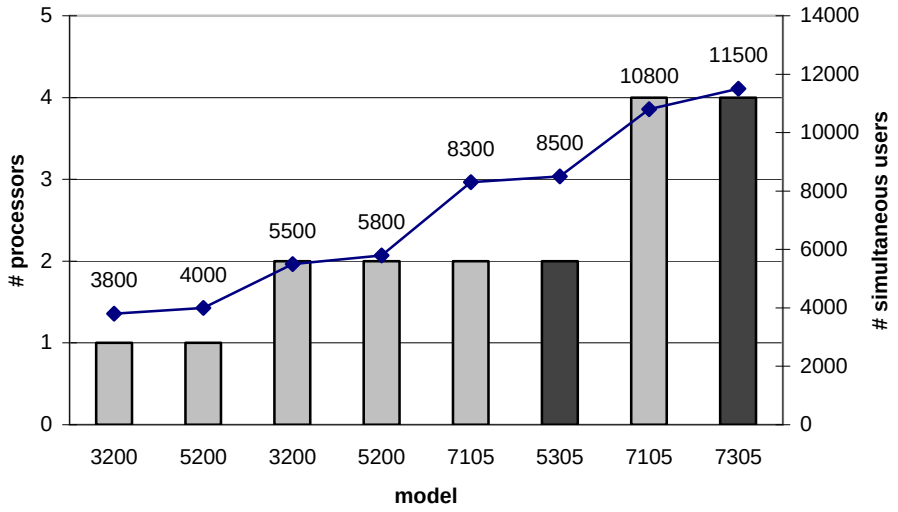

Source: Digital

Table 6 presents recent LoadSim results published by a variety of original equipment manufacturers (OEM). All results used Microsoft Exchange Server 5.5 on Windows NT Server 4.0, Enterprise Edition. Different configurations are identified. A full description of each configuration is available at: http://www.microsoft.com/exchange.

The response time score represents a 95th-percentile score of the measured test run. The score is expressed in milliseconds (ms). For example, a response time score of 1000 ms or less is considered an acceptable response time for e-mail users utilizing Exchange Server's MAPI protocol. The tests were completed in accordance with Microsoft Exchange Server OEM Benchmarking Policy Guidelines version 1.1.

Table 6 – LoadSim results

Vendor	System	Response Time	Users	CPUs
Digital	Server 5305	454 ms	4500	1
Digital	Server 5305	608 ms	8000	2
Compaq	ProLiant 6000	336 ms	9000	4
Digital	Server 7305	509 ms	10,000	4
NCR	4300	422 ms	10,000	4
Hewlett Packard	NetServer LXr Pro8	344 ms	14,000	8
Compaq	ProLiant 7000	1000 ms	15,000	4

Today, some of the world's largest corporations depend on Microsoft Exchange Server. Boeing, for example, hosts 130,000 Exchange users, and up to 1,000 users per server with capacity to spare. General Electric uses Microsoft Exchange as its single worldwide messaging system. It currently supports 125,000 Exchange users and up to 1,200 concurrent users per server. Texaco, Hoechst Marion Roussel and Dell Computer Corp. also have large-scale Exchange deployments.[24] For a technical discussion of Texaco's implementation of Microsoft Exchange Server, read the deployment case study. It covers a range of topics, including: Migration Strategies, Windows NT Domain Topologies, Microsoft Exchange Site Architecture, Microsoft Exchange Connector Architecture, and Mail-Enabled Applications.

SPECWeb Benchmark

Having discussed both OLTP and messaging in the earlier sections of this White Paper, it is now time to turn to the third major workload of importance to servers in the commercial marketplace: the Web. As the leading multipurpose platform, support for Internet services is an integrated component to the overall operating system design. Selected SPECWeb96 results from systems running Windows NT Server 4.0 and Windows NT Server's Internet Information Server (IIS) version 3.0 and IIS version 4.0, respectively, appear in Table 7, and demonstrate the significant capabilities of the operating system today. Results indicate the number of accesses per second. IIS has the highest results of any Web server running on Intel processors. It is also faster than select Web servers

[24] Additional case studies can be found at: http://www.microsoft.com/exchange/showcase/papers/casestudies/default.asp?A=3&B=3

such as Sun Web Server 2.0 running on comparably configured Solaris/SPARC operating system.

For example, Third Quarter '98 SPECweb96 Results report 3,151 pages access on a Aquanta QS/2 from Unisys running IIS 4.0 on Windows NT Server 4.0. The best results posted by Sun Microsystems were posted in the Third Quarter '97 SPECweb96 Results. The system tested which, posted results of 2,906 pages accessed, was Ultra Enterprise 450 with Sun Web Server 1.0.

Table 7 – SPECWeb96 Results

Time	Server System	CPUs	Result
Third Quarter 1998	Unisys Aquanta QS/2	4	3,151
Second Quarter 1998	HP NetServer LH 3/400	2	2,192
First Quarter 1998	HP NetServer LXe Pro 6/200	4	1,878
First Quarter 1998	HP NetServer LXr Pro 6/200	2	1,525
Fourth Quarter 1997	HP NetServer LX Pro 6/200	4	1,733

Although IIS 4.0 is not quite as fast as the most costly RISC/UNIX server systems, Microsoft has made impressive performance gains—81 percent—in the last 12 months, according to benchmark results. Furthermore, the April 13, 1998 issue of *Internet Computing* published a study comparing the performance and features of nine Web servers on five different platforms. The *Internet Computing* Net Best Winner award went to IIS 4.0.[25]

Recently Mindcraft, Inc., an independent benchmarking company, published results to a series of performance tests on Web servers that is broader than SPECWeb96. Mindcraft used WebBench 2.0 to simulate SPECWeb98, which is a newer Web benchmark test that SPEC is now defining. Ziff-Davis developed WebBench 2.0 for testing static and dynamic requests. In SPECWeb98, 70 percent are static HTTP requests and 30 percent are dynamic HTTP requests. SPECWeb98 requires 6,408 files totaling over 900 MB to support a request rate of 1,600 HTTP requests per second.

Mindcraft wrote ISAPI and ASP programs to test the potential of IIS dynamic request handling. Because the Sun Web Server (SWS) supports only CGI to handle dynamic requests, that interface was used in SWS tests. Table 8 summarizes the results.

[25] For more details about this study and performance and feature results, see http://www.zdnet.com/icom/zdlabs/power.servers/.

Table 8 – Mindcraft Results

Web Server	Peak RPS	Peak Throughput	Price/ Performance
IIS 4.0/Compaq ProLiant 3000 (2 x 333 -MHz Pentium II, 512 MB)	1,337 (ISAPI) 694 (ASP)	17.1 MB/S 8.8 MB/S	$9.17/RPS (ISAPI) $17.65/RPS (ASP)
SWS 1.0/Sun Ultra Enterprise 450 (2 x 296 -MHz Ultra SPARC, 512 MB)	336	4.5 MB/S	$94.75/RPS

The first column in Table 8 outlines the configuration of each system. IIS tests were executed on a Compaq ProLiant 3000 running Windows NT Server 4.0. The ProLiant used two 333-MHz Pentium II and 512 MB of memory. SWS tests executed on a Sun Ultra Enterprise 450 running Solaris 2.6. This system had two 296-MHz UltraSPARC processors and 512 MB of memory.

The second column shows the peak requests per second (RPS) supported by the Web server. ISAPI and ASP results are reported for IIS. ISAPI modules run in the same process space as IIS, while ASPs are HTML pages that invoke scripts and components. A new process is created every time a CGI program handles a dynamic request. ASPs can service about half the number of requests per second of an ISAPI module, but twice as many requests as a CGI program.

Column 3 reports the total number of bytes that each Web server sent to test clients. Again, results show that ASPs offer half the throughput of ISAPI, but twice the throughput of CGI.

In column 4, Mindcraft calculated the street price of the servers and software, and then divided it by the peak RPS. The street price of the Windows NT-based system is $12,260. The Solaris system costs $31,828. Accordingly, the Windows NT, IIS, and ProLiant configuration offers price/performance that is 5 to 10 times better than comparable system from Sun Microsystems. Read Web Server Comparison for more information.

SUMMARY

Companies such as Boeing, Chicago Stock Exchange, Compaq, Dow Chemical, Fidelity, General Electric, Merrill Lynch, Texaco, Saturn, and many others rely on Windows NT Server-based systems to deliver computational, transaction, and I/O capacity requirements in their demanding environments. When hosted on properly configured hardware, Windows NT Server-based systems enable a diverse set of applications to scale to support thousands of online transaction processing users, or mail and messaging client systems on a single, large SMP server. In replicated, distributed and/or partitioned systems, Windows NT Server, Enterprise Edition can support even larger user populations, in some cases exceeding 100,000 users across multiple servers.

This broad adoption notwithstanding, there is every reason to believe that the tremendous scalability gains demonstrated by Windows NT Server, Enterprise Edition-based systems since Scalability Day will continue at a rapid pace. Looking ahead, Microsoft believes that support for Enterprise Memory Architecture in Windows NT Server 5.0, Enterprise Edition and SQL Server 7.0, Enterprise Edition will help Microsoft systems and solutions partners to demonstrate up to 16-way SMP scalability in transaction processing and data management workloads. Microsoft's recent acquisition of Valence Research Inc., developer of TCP/IP load-balancing and fault tolerance software for Microsoft Windows NT Server, lets customers build Internet Web Farms with up to 32 cluster nodes today. This software will be renamed Windows NT Load Balancing Service and will be integrated into Windows NT Server, Enterprise Edition. Web clustering augments SMP scalability allowing Windows NT to be used to power the largest Web sites.

Microsoft understands the increasing demands of the marketplace for scalability. Some of these demands are based on integrating a wider range of stakeholders into Web-based customer service, order fulfillment, and commerce systems. Some of these demands are a function of broader use of advanced data types, including streaming media. Some of these demands are a result of more meaningful integration of worldwide workgroups through video conferencing, distributed authoring, remote learning, and workflow. The eight-fold improvement in throughput that Microsoft has achieved in the last 36 months reflects its vision of bringing the benefits of Windows NT-based computing—its technical features and uniquely cost-effective business model—to as many customers as possible.

However, enterprises do not "live by scalability alone," and so Microsoft's investments in scalability are matched by other enhancements to the operating system. With Windows NT Server, Enterprise Edition 5.0, customers will benefit from higher system availability, advanced PKI and security enhancements, higher performance communications (including ATM QOS services), faster sorting, faster printing, greater manageability, advanced directory services, and the industry's most advanced component services (COM+). In addition to greater scalability, this increased functionality and integration will pay dividends for customers over the long haul in an even wider variety of demanding, mission-critical computing.

FOR MORE INFORMATION

The Transaction Processing Performance Council
777 North First Street, Suite 600, San Jose, CA 95112
http://www.tpc.org

The Standards Performance Evaluation Corporation
10754 Ambassador Drive, Suite 201, Manassas, VA 20109
http://www.specbench.org

Ziff-Davis ServerBench http://www1.zdnet.com/zdbop/svrbench/svrbench.html

Bluecurve Dynameasure http://www.bluecurve.com

"Web Servers ," *PC Magazine*, May 5, 1998.
http://www.zdnet.com/pcmag/features/webserver98/index.html

"New Web Servers Deliver Power Sites," *Internet Computing*, April 13, 1998.
(http://www.zdnet.com/icom/zdlabs/power.servers/)

"A Web Server Comparison: Microsoft Windows NT Server 4.0 and Windows NT Server's IIS 4.0 on Compaq ProLiant 3000 vs. Solaris 2.6 and Sun Web Server 1.0 on Sun Ultra Enterprise 450," Mindcraft, Inc., May 26, 1998.

SQL Server Scalability, A Microsoft White Paper, January, 1998.

Microsoft BackOffice Scalability Site
http://www.microsoft.com/backoffice/scalability

Microsoft Exchange Web site,
(http://www.microsoft.com/exchange)

Compaq
http://www.compaq.com

Digital
http://www.digital.com
http://www.windows.digital.com
http://www.unix.digital.com

Hewlett-Packard
http://www.hp.com

Intel
http://www.intel.com

NCR
http://www.ncr.com

Microsoft
http://www.microsoft.com

Oracle
http://www.oracle.com

Tandem
http://www.tandem.com

Unisys
http://www.unisys.com

Designing Applications with MSMQ, Alan Dickman (Addison-Wesley, 1998)

Porting UNIX Applications to Windows NT, Andrew Lowe (Macmillan Technical Publishing, 1997)
http://www.porting.com