UNISYS
imagine it. done.

white paper

The Emergence of a Windows Mainframe

Dr. Michael Salsburg

growth

Technology that was once considered the embodiment of the personal computer has emerged as a powerful, scalable, high-performance force in the mainframe market—one that offers levels of familiarity and standardization that are both compelling and cost effective. This paper tracks the effects of CPU commoditization on the evolution of mainframes since their early proprietary days to their current standardization on Microsoft Windows, and presents a hardware and software solution for providing traditional mainframe characteristics that is packaged in a surprisingly non-traditional manner.

Table of Contents

Introduction

No doubt you have an intuitive definition of the term, "mainframe." At one time, the word described a system's physical characteristics—a complex collection of highly specialized electronics, bundled in monolithic boxes and used as the main central repository for processing information. Today, the term mainframe simply describes a monolithic computer complex used for large, commercial business applications, one that is associated with centralized rather than distributed computing.

Mainframes embody a number of attributes that are inherent in robust information technology such as high availability, sufficient capacity to handle maximum demand, security, and manageability. Most successful IT organizations realize these valuable attributes at a minimal cost by establishing standards for development, deployment and execution within an environment consisting of one or more mainframes.

The Mainframe Era

In the early days...

The world's first large-scale, general-purpose computer was developed at the University of Pennsylvania by J. Presper Ekert and John Mauchly and the first commercial version, the UNIVAC, was brought to market by a predecessor Unisys company, Remington Rand. IBM came to the computer market later but, with an already pervasive presence in business offices with their typewriters, became the dominant force in computing with a huge percentage of the overall computer market share, consisting primarily of large mainframes. There were other mainframe vendors, but they could not encroach upon that market share for a number of reasons.

One explanation rests in the overall involvement of many third parties in the "value chain," the various bits and pieces that add value to the basic hardware and thus provide a complete solution for the customer. The IBM value chain included many large software companies that provided turnkey software solutions, along with a bevy of consultants to apply the hardware/software to provide complete solutions for large data centers.

IBM hardware was proprietary, but there was a large commodity market consisting of makers of peripherals, supplies and all of the other components that are needed to make a computer environment successful. As more vendors participated in the ecosystem, prices dropped for several reasons: widespread computerization, the growing variety of competitors and the high volume that allowed for cheaper manufacturing. This ecosystem continued to grow, with success begetting success until IBM dominance seemed perpetual.

In general, standardization offered compelling economies of scale both on the product procurement side and in the area of skilled personnel. And, in essence, most mainframe computing was standardized on the IBM approach.

The Open Systems Era

And then came UNIX.

Though the general IBM ecosystem flourished through the commoditization of various supporting components, the main central electronic complex (CEC) was strictly proprietary, with IBM dictating the prices of new machines. The first wave of disruptive technology for mainframes emerged in the form of UNIX systems.

UNIX showed all the signs of being “open” in that many vendors would manufacture hardware that would collectively support a single operating system. The promise was that, as various vendors achieved higher performance, the IT manager could incorporate new, cheaper, UNIX based hardware into the complex yet continue to support a homogeneous environment. UNIX was supposed to become the new, standard operating system. Thus a new value chain also emerged, with a new ecosystem to provide all of the pieces needed to establish a complete solution.

As the ecosystem grew, cost continued to drop and raw performance continued to climb. These two trends were not independent. There is a causal relationship between lower cost and higher CPU performance. With the lowering of cost, the adoption of computerization becomes more widespread. In the case of UNIX, computerization first spread to the departmental level. It was no longer a corporate-level decision to buy another computer system.

Once again, the commoditization of a system created a “snowball” effect, with momentum growing as unit sales climbed and prices dropped. The higher sales volume created more competition and provided more R&D dollars to invest in more efficient and higher performing CPU architectures. In retrospect, UNIX was not quite ready for the “big time” when it was first adopted for corporate (as opposed to academic) deployment. But the hardware, operating system and ecosystem matured until UNIX became the solution of choice for new data center applications. By 1999 it was widely accepted that UNIX could be used to implement a large, data-intensive computing environment.

Unfortunately, one of the key goals, standardization on a single operating system, did not come to pass since different hardware and software vendors introduced different versions of the UNIX operating system, thus negating its openness. As a result, IT was again forced to “lock” into a proprietary environment: a specific flavor of UNIX. The ecosystem that once provided “open” values was thus splintered into many pieces. A software company that supported UNIX was ultimately forced to dilute their R&D into support for as many as 11 different types of UNIX!

Enter Windows

The new contender.

Over the past few years, most IT professionals have felt the effects of CPU commoditization. Anyone who has replaced a personal computer within the past year knows that they can now get from two to five times more performance at a lower cost than two years ago. As this paper is being written, a 2.2 GHz system with 256 MB of memory, a 17-inch monitor and 30 GB of disk can be bought for less than $600.

Such low prices are possible because, as more and more CPUs are purchased, revenues can be plowed back into R&D to help fit even more transistors into even less real estate. The higher volumes drive down manufacturing costs per unit, causing memory prices to drop and CPU cycles to climb. On the operating system side, Windows had already become the standard and pervasive operating system for both home and business desktop workstations.

In the year 2000 the trend was to use small and inexpensive computer systems in distributed networks and "scale out" business applications to meet increased demands. In other words, when more computing capacity was required, IT departments simply deployed more and more small systems.

Let's go back to the personal workstation scenario for a moment and use it to better understand what scaling out really looks like. Imagine that, as you do more and more tasks, you deploy a number of different workstations in your office. Each workstation handles a different application, such as word processing, spread sheets, slide presentations and Internet access. Besides thoroughly challenging your sanity, this scale-out approach will be sub-optimal for a number of reasons.

Information sharing among those applications on different workstations would involve transfers across a network, not within the internal memory complex of a single computer. As your application workload required changes, you may have to re-provision many or all of the various workstations to accommodate them. And, while reducing the reliability of your workstation environment overall, such a setup also makes managing the "system" more time consuming and costly.

Obviously, the tactical scale-out approach needed rethinking if it was ever to become a successful overall strategy.

Here is an interesting quote from a white paper published in 1999 by DataBase Associates International[1], a company since renamed, Intelligent Business Strategies:

> *This paper has demonstrated that a S/390 server running OS/390 and DB2 UDB for OS/390 is a powerful combination for supporting a large data warehousing system. The only real, viable alternatives to S/390 are large UNIX, AS/400 and niche hardware parallel servers. Microsoft Windows NT has yet to demonstrate the scalability and functionality required for such a warehousing system.*

1 DataBase Associates International, Inc., "IBM System/390 and DB2 UDB: A High Performance Data Warehouse Server," May 1999 http://www.databaseassociates.com/pdf/IBMos390%20Version%202.pdf

Enter Big Windows

The industry-standard mainframe.

Early on, Unisys had a unique vision of the tremendous benefits of integrating the industry-standard Windows operating system with commodity Intel processor components to create an entirely new kind of mainframe environment. Unisys envisioned a single, large-scale Windows mainframe for the data center on which many different applications and very large databases could be deployed and administered from a central location. This Windows mainframe would not only be extremely cost-effective, but it would also "scale up" to accommodate growing applications and databases.

In the fall of 2000, after a long collaboration with Intel and Microsoft, Unisys announced the industry's first 32-processor system based on industry standards. The company's ES7000 enterprise server offered a unique, Intel based architecture that took full advantage of Microsoft's new enterprise software, Windows 2000 Datacenter Server.

The availability of this system marked the beginning of a new and innovative mainframe technology that would, within three years, finally bring the industry around to labeling UNIX a proprietary operating system. Since Unisys ES7000 systems have gained wide acceptance, a number of other vendors, including HP, IBM, NEC and Fujitsu, have entered the large-scale Windows and Intel systems marketplace.

Windows standardization.

Though the overarching rationale for standardization on Windows is reduced cost, there are a number of mainframe values that were envisioned by Unisys and the mainframe engineers who designed the system. Before data centers were splintered off into distributed, heterogeneous systems, mainframe environments took a single, homogeneous approach to computing. There was a wealth of skill and experience from which to draw. As IT staff moved from one data center to another, they brought along their knowledge and best practices. The operational model of a well-run data center was, as a result, consistently improving throughout the industry. And as Windows becomes a standard in data centers, we can expect this same type of cultural change. As best practices are institutionalized, we will see more independent software vendors develop tools to further simplify focused high-end, Windows data center administrative tasks.

Simplification of these tasks has two important by-products. First, operational errors are reduced. (Gartner analysts often cite 40 percent of unplanned downtime in Windows data centers as attributable to operator error.) Second, as administration becomes much easier, there is a new focus on managing the application itself in the form of workload management, security management, business continuance and the ability to quickly adapt to changes. These tasks are more difficult to accomplish in distributed, heterogeneous environments.

Engineering innovations.

Edison once described genius as one percent inspiration and 99 percent perspiration. It may look simple enough on a white board to draw a 32-CPU Windows based system, but there have been (and continues to be) a number of much-less-visible technical achievements that have contributed to the final outcome. The realization of a homogeneous Windows data center, running on a 32-CPU Symmetric MultiProcessing (SMP) platform required a lot of ingenuity and hard work.

Manufacturing and delivering a reliable, enterprise-quality system using 32 Intel processors requires a solid working relationship with Intel, which provides building blocks of up to four processors. Combining these building blocks with 64 GB of memory, 40 PCI slots and an efficient method of moving data from memory to the individual processors is the subject of another paper. It should suffice to say that there are many ways to tackle the problem, with each method requiring its own level of cost/risk analysis. The key here is that Unisys spent a great deal of time and energy on innovating and developing unique, patented technologies to create the internal architecture of its ES7000 enterprise servers that allows 32 processors to quickly and efficiently access all of the system's memory in parallel. Yet these systems still rely on a single commodity supplier to continually improve raw CPU performance.

A critical mainframe attribute is availability. Four years ago many questioned the reliability of a "Windows mainframe" due to the negative perception of workstation reliability. However, the Windows mainframe concept envisioned by Unisys embodied a radical approach to achieve high availability, essentially infusing the solution with its own culture and expertise in providing highly available mainframes to critical business sectors such as the airline and financial industries. Unisys mainframe engineers worked with Microsoft software engineers first to help them develop the Windows 2000 Datacenter Server operating system and later the Windows Server 2003, Datacenter Edition version, both of which are specifically focused on enterprise-class availability.

For more information about the unprecedented availability of the ES7000 server family, you can access a recent study showing that, on average, customers are achieving the "five nines" of availability required for mission-critical computing environments. Visit http://www.unisys.com/es7/availability for the full report.

All of the hardware infrastructure, from the chipsets that support the CPUs out to the PCI buses, have been engineered to detect and correct errors before they cause disruption. Server Sentinel systems management software oversees the entire ES7000 mainframe complex including both the hardware and operating system, providing self-monitoring and self-healing capabilities that ensure system availability. It delivers these mainframe-class data center server management capabilities for resiliency and business continuance in Microsoft Windows operating environments by simplifying systems management, automating corrective and preventive measures to avoid a system failure and increasing the systems' reliability, scalability and performance. Application Sentinel provides a comparable degree of protection for the application and database layers of an ES7000 solution.

In addition to the CPU and hardware, consider the Windows operating system and environmental software, such as SQL Server. Three years ago the architecture of this software had not evolved to take advantage of large, mainframe-class systems. Many of the critical software routines could not execute concurrently on 32 processors. Unisys engineers had developed considerable expertise in parallel processing through decades of experience in delivering multiprocessing mainframe systems. Working side by side with Microsoft engineers, Unisys mainframe engineers analyzed and rewrote critical code so that processes and threads could execute in parallel in a large multiprocessor ES7000 server complex. Many tools and techniques were developed to focus on the serialized portions of the code and then to provide high-performance locking and thread management.

What about the real world?

Many factors come into play in measuring success. To demonstrate a number of these factors, the reader is referred to a case study about Microsoft's deployment of its World Wide Marketing Database[2]. Some applications, such as those dependent on large databases, cannot be deployed successfully in a scale-out environment. Only a large mainframe-class system such as the ES7000 server can meet the challenges. This case study describes the ES7000 Windows mainframe solution for a large database application. When the study was published in April 2002, the database was "only" 1.8 terabytes in size. Since then, it has grown significantly. Certainly, the previously referenced white paper's author could not have anticipated such an environment.

In the real world, many different types of applications are expected to run and be administered flawlessly. As mentioned previously, a high percentage of system stops are caused by operator error. In many cases, operator error is the result of the administrator not being familiar with the proper procedure or not having a clear model of how to manage the system. When building a team of administrators, it is very useful to draw from the same operational model. Due to the pervasive deployment of Windows, its operational model is almost second nature within the industry.

2 Microsoft Corporation, "Scaling Up to the Needs of the Worldwide Marketing Database with Datacenter Server," January 29, 2002 http://www.microsoft.com/windows2000/datacenter/evaluation/casestudies/unisys.asp

What's different?

One might wonder about the differentiation from one vendor to another in this brave new world where processors and operating systems are no longer developed by the system vendor. There can be significant innovations in the architecture that is constructed around the CPU chips, including innovations in cache coherence, memory interconnects and IO handling. For example, the 32-CPU ES7000 server uses a unique memory crossbar and cache directory to ensure that memory latencies are kept as low as possible. Low memory latency translates into an increased number of useful CPU cycles. An architecture that does not minimize latency will yield a system where the CPU will waste a majority of its cycles waiting for memory fetches to complete. This is a constant challenge when systems are designed to include more and more memory. The ES7000 system incorporates a patented crossbar technology to accelerate memory fetching and minimize latency.

Another key differentiator is established on the basis of overall cost, including the cost of the initial system and the total cost of ownership. The savvy vendor is now focusing its ingenuity and intellectual leadership on new methods to create a more agile system at lower cost. The high-density rack-mount ES7000 servers show the rewards of focusing engineering time on cost reduction. Unisys experience with its Windows mainframes has shown that many data centers want to start off relatively small (four CPUs) and scale up as needed to meet the growing workload demands. The rack-mount ES7000 servers address this need for agility with packaging that can start with a 4-CPU system that is essentially a quarter of the cost of predecessor ES7000 servers with 8 CPUs. Therefore, the initial system cost is significantly lower, with incremental costs increasing along with the demand for capacity.

Two contributors to the total cost of ownership are availability and the cost of staffing. Both of these factors are significantly reduced through the simple, industry-standard operational model provided by the ES7000 system's Server Sentinel systems management software, which addresses them by providing simplicity within a familiar environment and by greatly reducing the need for operator intervention. The software self-manages the system and optimizes performance when it finds an opportunity. When human intervention is needed, it focuses management attention through a clear, easily recognizable user interface. Its simplicity and clarity reduce operator uncertainty.

Finally, by taking a mainframe approach, as opposed to a distributed approach, staffing costs can be optimized. Fewer systems need detailed attention and administration, leaving more time for strategic and capacity planning instead of tactical firefighting. Since all of an IT staff can focus on a single operational model, they can share their knowledge and expertise. All of this adds to the greater optimization of staffing costs.

Conclusion

This paper has discussed some significant changes in the mainframe platform marketplace over the past decade. The "Windows mainframe" is gaining more widespread acceptance in the market place. The concept of commoditization as a force in elevating the performance of CPUs in a mainframe environment was discussed. Concurrently, the standardization on Windows as the mainframe operating system continues to drive down costs. Since its initial vision, Unisys has continued to innovate and improve on its family of ES7000 enterprise servers to extend the value of Windows into the data center.

Author's Biography

Michael Salsburg, Ph.D.

Michael has been active in the computer architecture and performance evaluation arena for more than 20 years. He received his Doctorate in Mathematics (Probability and Statistics) from Drexel University in 1992 and has been awarded two international patents for performance modeling algorithms and software. Michael has also published dozens of papers and has lectured worldwide on the topic of computer performance evaluation. He was the founder and president of Performance and Modeling Incorporated and is currently a technology director for Unisys.

For more information, contact your Unisys representative.

Or call:

1-800-874-8647, ext. 365 (U.S. and Canada)

00-1-585-487-2430, ext. 365 (Other countries)

In a hurry to learn more? Visit:

http://www.unisys.com/es7/guide

For even more details, visit:

http://www.unisys.com/es7/guide/community

01/06

4126 3781-200