

*Consolidation and Recentralization with the
Ultra™ Enterprise™ 10000 Server
New Deployment Strategies For Improved Manageability And
Lowered Costs*

Technical White Paper



© 1997 Sun Microsystems, Inc.—Printed in the United States of America.
2550 Garcia Avenue, Mountain View, California 94043-1100 U.S.A

All rights reserved. This product and related documentation is protected by copyright and distributed under licenses restricting its use, copying, distribution and decompilation. No part of this product or related documentation may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX[®] and Berkeley 4.3 BSD systems, licensed from UNIX[®] Systems Laboratories, Inc. and the University of California, respectively. Third party font software in this product is protected by copyright and licensed from Sun's Font Suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

TRADEMARKS

Sun, Sun Microsystems, the Sun logo, Ultra, Ultra Enterprise, Gigaplane, and Gigaplane-XB are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd. SAP and R/3 are trademarks or registered trademarks of SAPAG.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.



Please
Recycle

Contents

1. Introduction	1
2. The Need for Simplified Management of Distributed Computing	3
3. How the Ultra Enterprise 10000 Addresses Needs in the Data Center.	7
Manageability through Dynamic System Domains	8
Consolidated Servers for Cost Reduction	8
Capacity Redeployment Based on Business Cycles	10
Server Consolidation for Improved Availability	11
High Availability and Serviceability Features	12
No Single Points of Failure	12
Dynamic Reconfiguration	13
Hot Swappable Components	13
Domains Eliminate the Leading Cause of Server Downtime	13
Bringing High Availability to UNIX® for the Distributed Applications	14
Capacity for Large-Scale Applications	14

High-Volume OLTP	14
Large-Scale Decision Support	14
Capacity for Today and Tomorrow	16
Capacity for Maximum Configurability	17
The Ultimate Single Server Solution	19
4. Application Examples	21
Application Example 1–Server	
Consolidation for Large-Scale SAP R/3 Deployments	22
Database And Application Server Consolidation	23
One Platform For Production,	
Development, And Training	23
Application Example 2–Recentralization of	
Corporate Business Intelligence Services	24
Fat Clients	24
Network Bottlenecks	25
Application Recentralization	
with the Enterprise 10000	25
Application Example 3–Mainframe	
Database Consolidation for Data Warehousing	28
Data Warehouse Deployment Issues	28
Database Consolidation with the Enterprise 10000	29
Data Mart Consolidation	29
Customer Data Warehousing with the Enterprise 10000	31
5. Conclusion	33
6. References	35
Gartner Group	35
META Group	36
Sun White Papers	36

Corporations are deploying an expanding array of information technology to improve competitiveness, exploit new market opportunities, and streamline operational costs. Computing styles have migrated from mainframe host-based to PC-based client/server, with an emerging trend toward thin client/web-based styles.

Key issues facing IT organizations today include:

- How to reduce the operational costs of distributed computing
- How to manage increasing amounts of corporate data and support rapid user community growth
- How to make IT a strategic corporate weapon.

The Sun Ultra™ Enterprise™ 10000 server is a new class of systems designed to meet these and other IT challenges. The Enterprise 10000 is the only system which marries the ability to provide large-scale distributed user support with centralized management features to ensure improved service levels, while keeping costs under control.

The Need for Simplified Management of Distributed Computing



Distributed computing has brought great benefit to corporations, but it has also created a tremendous challenge for IT organizations—how to manage complex distributed environments. According to Gartner Group, with the trend towards IT budgets migrating from the central IT organization into the business units, the central IT organization is increasingly being forced to manage highly distributed systems without the benefit of planning for deployability, manageability, and rapid expansion (Gartner Group, *The Third Generation of Client/Server Computing*, C/S: R-210-106, 26 September 1995)

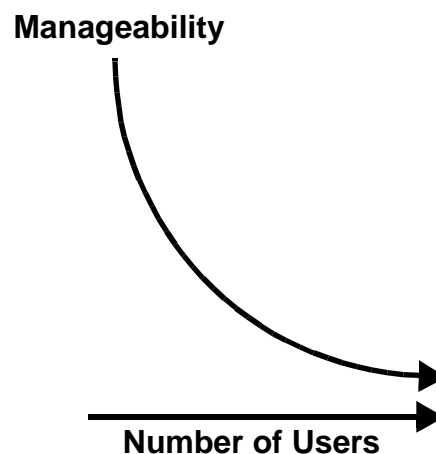


Figure 2-1 The manageability of distributed systems decreases dramatically as user communities grow. Source: Gartner Group, *Riding the Whirlwind: The Transformational Effects of Managing Distributed Systems*, NSM: R-MGT-111, 26 September 1996

In many organizations, the almost spontaneous growth of distributed applications has led to administrative complexity that is approaching unmanageable levels, and resulted in systems implementations with precarious stability. The need for centralized management is clear.

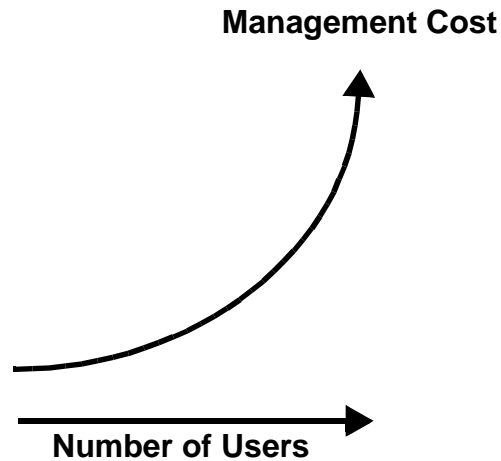


Figure 2-2 The effect of diminished manageability in distributed systems is rapidly rising management costs due to increased complexity. Source: Gartner Group, *Riding the Whirlwind: The Transformational Effects of Managing Distributed Systems*, NSM: R-MGT-111, 26 September 1996

As Gartner Group notes:

“ . . .The number of servers and network change requests in many environments has become unmanageable, so many enterprises are looking to reduce the complexity through centralization. . . . IS managers are used to a more centralized approach in dealing with hardware and staff, and they are seeking ways to reduce costs and improve service, while establishing some control over distributed environments.” (Gartner Group, *Minimizing LAN Costs, Improving Support, Yield Big Dividends, Point-To-Point*, 26 January 1996)

The Ultra Enterprise 10000 server offers a means to reel in operations costs and mission risk by providing a consolidated and centrally-managed platform for distributed applications.

The widely proclaimed cost-of-ownership advantages of LAN-based systems over the mainframe have not been borne out in many corporations because the superior price/performance of server hardware has been overtaken by increased support costs. Research into the cost of computing for the mainframe versus PC/LANs highlights the fact that the support costs are the major determinant of Total Cost of Ownership (TCO) in distributed environments.




Figure 2-3 Distributed systems have a very different cost profile than centralized systems due to the high management and support costs for distributed computing. Source: Gartner Group, *Riding the Whirlwind: The Transformational Effects of Managing Distributed Systems*, NSM: R-MGT-111, 26 September 1996

The recognition of the hidden costs associated with distributed systems management has been heralded as a reason to “return to the mainframe” by some. But realistically we know that the cost basis and inflexibility of mainframe systems combined with the lack of availability of new applications continue to drive an increasing trend towards open systems. What is needed then, is a way to keep the good attributes of each approach without the respective downside.

The ideal case would be a centralized computing resource that could enjoy the power and simplified management of the mainframe, while providing the breadth of application availability and client/server support necessary for distributed computing.

The Ultra Enterprise 10000 server marries the ability to provide large-scale distributed user support with centralized management features to ensure improved service levels, while keeping costs under control. It is the first open system to offer the means to provide multiple operational environments in a single managed system to meet the consolidation and recentralization needs of IS.

How the Ultra Enterprise 10000 Addresses Needs in the Data Center



The Ultra Enterprise 10000 system brings mainframe capabilities to the open systems world and sets a new standard for large-scale server manageability, availability, and performance.

- *Manageability* for distributed computing

Multiple user or application environments in a single system with Dynamic System Domains

Single administration point for multiple domains to simplify management complexity and slash administrative costs

System resources can be transparently reallocated from one department or application domain to another at a moment's notice to meet unplanned demands or cyclic needs

- *High Availability* features for new mission-critical distributed applications

No single points of failure

System dynamically reconfigures after failure to remove failed components

Domains fully isolate users in a single system — other domains are not affected by a system crash in one domain

Systems are repairable and upgradeable while the system is running.

- Industry-leading *Capacity* for the largest applications

Capacity to run the largest DSS queries or support thousands of OLTP users with a single system.

An Enterprise 10000 system can be upgraded to over 60 times its base performance in a single cabinet, without powering the system off.

Manageability through Dynamic System Domains

Dynamic System Domains provide benefits similar to the mainframe feature known as “logical partitioning.” Domains provide both physical and logical partitioning of the system.

With domains, you can easily isolate different user populations or applications, but also dynamically reallocate your computing resources based on business needs, thereby providing you with better performance and more effective resource utilization for lower cost of ownership.

A single Ultra Enterprise 10000 system can be dynamically divided into as many as five smaller discrete servers and back. Any software errors and many hardware failures in a domain are confined to that domain and will not affect the rest of the Enterprise 10000 system.

Most importantly, you can administer all domains from a single point. Here are some examples of how domains benefit business applications.

Consolidated Servers for Cost Reduction

There is widespread industry interest in techniques for consolidating LAN servers to reduce support costs. According to Gartner Group research findings, “Significant cost savings can be achieved by consolidating LAN servers and centralizing certain support functions.”

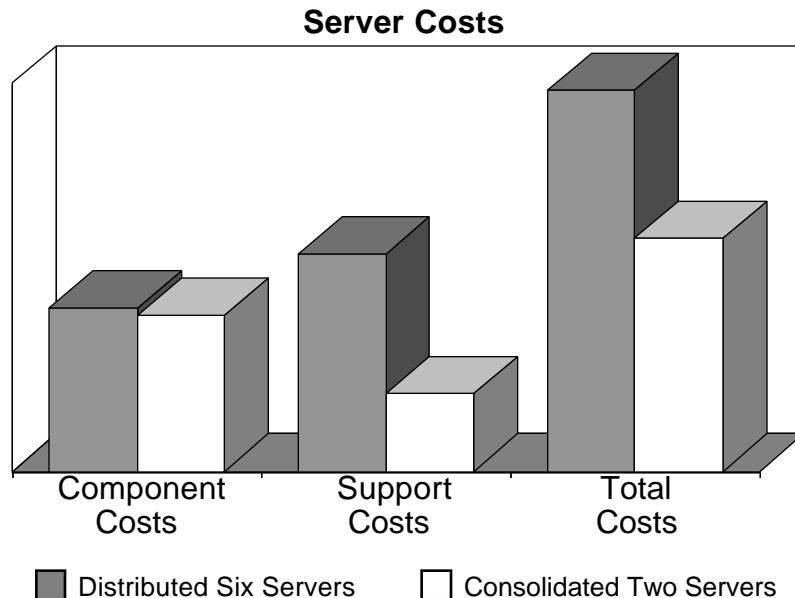


Figure 3-1 Support costs can be dramatically reduced by even modest server consolidation. Source: Gartner Group, *Minimizing LAN Costs, Improving Support, Yield Big Dividends, Point-To-Point (PTP)*, 26 January 1996

Gartner's findings indicate that the benefits increase for larger consolidations, as well.

“In a significantly larger server consolidation—moving from 32 distributed servers to six consolidated servers—the three-year savings total more than \$3 million.” (Gartner Group, *Minimizing LAN Costs, Improving Support, Yield Big Dividends, Point-To-Point (PTP)*, 26 January 1996)

The Ultra Enterprise 10000 server can play a unique role in consolidation—whether it be servers, applications, and databases. Because the Enterprise 10000 domains support consolidated management, and the Enterprise 10000 system capacity allows it to take on the processing burden that usually dictates multiple servers. Distributed client applications can be recentralized on the Enterprise 10000 to lower client hardware and management costs, while improving performance levels. There is also an industry-wide need to consolidate separate databases, especially in data warehousing environments.

Capacity Redeployment Based on Business Cycles

Domains can be employed to balance seasonal departmental processing requirements across the enterprise without maintaining excess capacity and underutilized assets. By using domains to reallocate system resources, a single Ultra Enterprise 10000 server can supply the peak needs of several different departments.

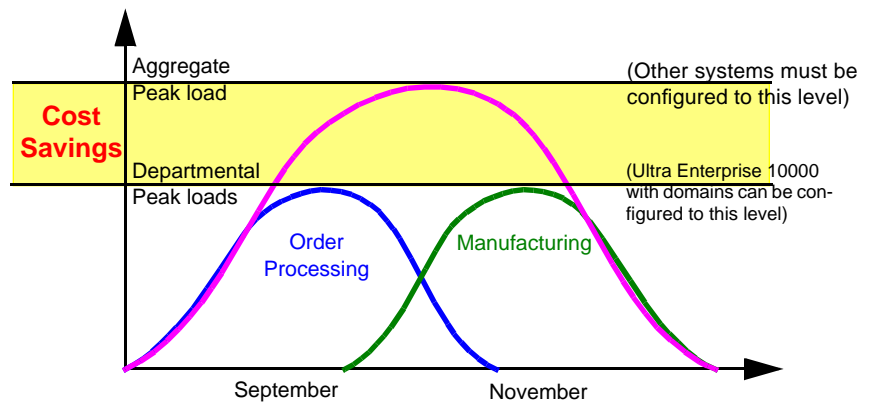


Figure 3-2 The ability to redeploy processing capacity based on business cycles lowers overall hardware expenditure requirements.

Many businesses have cyclic processing requirements that vary by department. However, in many cases, each department has stand-alone systems with idle capacity during non-peak periods. These needs can be mutually met with the Enterprise 10000 system domains. If, for instance, the peak load for order processing occurs in September, extra processors can be assigned to that domain to meet the needs of that department.

Later in the fall, as manufacturing organization needs increase, some of the resources previously assigned to order processing can be reallocated to meet that department's requirements. By using dynamic attach/reattach in conjunction with domains, this migration can be completed while the system is operational in a way that is transparent to both organizations.

Server Consolidation for Improved Availability

While the administrative merits of consolidation are well understood, many do not recognize the impact of multiple servers on combined service availability. It is a common practice for IT organizations to deploy dedicated servers for many services, such as file and print, to name a few. With the increasing complexity of corporate application environments users must rely on many services provided by multiple servers. This means that there are single server dependencies which can affect the ability of the end user to perform his normal work. This data from Gartner shows how severely a distributed multiple server solution can degrade availability based on the availability of a single server, especially when several servers are required for a given user activity.




Figure 3-3 Individual server availability may significantly impact service delivery in distributed application scenarios. Source: Gartner Group, *Strategies to Increase LAN Availability*, NCI: R-LAN-106, 23 September 1996.

If, for example, five servers are required, each with 99% availability, then the combined user availability falls to less than 95%.

The Enterprise 10000 improves availability over a multiple server solution while providing additional reliability, availability, and serviceability features to further improve overall application availability.

High Availability and Serviceability Features

Many IT organizations are looking for ways to deploy distributed OLTP applications, but are concerned about the availability issues with LANs. The Ultra Enterprise 10000 offers features which allow IT to deploy distributed mission-critical OLTP applications while preserving high service levels.

According to Gartner Group,

“There are four high-level business trends that are pushing enterprises into providing high-availability computing or 24x365 operation. These trends are:

- The globalization of enterprise operations and thus enterprise networks
- The increase in the use of remote computing for mobile and home-based employees
- The competition in many business segments, which has driven many enterprises to provide around-the-clock customer service
- The attempt to ensure service stability for large numbers of distributed sites.”

(Gartner Group, *Strategies to Increase LAN Availability*, NCI: R-LAN-106, 23 Sept.1996)

To meet these needs, the Enterprise 10000 server provides the highest level of availability for open systems—without resorting to complex architectures like clusters or expensive fault tolerant systems. Availability levels of 99.95% can be achieved in a single system. The Enterprise 10000 achieves this with an architecture designed for high availability.

No Single Points of Failure

All system components down to the power cords are redundant, ensuring that no single hardware failure will prevent the system from doing useful work. Hardware failures will typically cause the system to crash, but it will quickly recover, thanks to journaling file systems and fast recovery features, and continue to operate in a degraded configuration.

Dynamic Reconfiguration

Dynamic Reconfiguration is the process by which the system removes failed hardware from the system configuration and recovers to continue operation. For example, if a processor fails, it will note that failure by a process known as blacklisting, and return to operation without that processor. All other processors, and memory, and I/O controllers sharing the same system board are unaffected. Even in the unlikely event of a system backplane failure, the Ultra Enterprise 10000 can recover by dynamically reconfiguring out a portion of the system interconnect and rebooting in this new configuration.

Hot Swappable Components

All processors, memory, and I/O controllers can be hot swapped, to either remove them or add them, while the system operates. The dynamic reconfiguration feature allows activity on a system board to be quiesced while the rest of the system continues to perform useful work. Memory data is migrated to other boards, I/O is rerouted or transparently suspended, and processing is moved to other CPUs so that the board can be powered down and removed for service or upgrade. This means that the only downtime due to hardware failure is the system restart which dynamically reconfigures out the failure. Later, service personnel can use dynamic reconfiguration and hot swap to replace the component without loss of availability or additional downtime.

Domains Eliminate the Leading Cause of Server Downtime

Dynamic System Domains provide an additional benefit to minimize downtime. According to the Gartner Group, the most common cause of server downtime is planned hardware and operating system maintenance (Gartner Group, *Strategies to Increase LAN Availability*, NCI: R-LAN-106, 23 Sept.1996). The domains feature of the Ultra Enterprise 10000 system allows concurrent maintenance and production with full software isolation, with significant benefit to system availability—without resorting to purchasing two separate systems.

Bringing High Availability to UNIX[®] for the Distributed Applications

Organizations may deploy multiple servers to help enhance overall application availability, because if one server goes down, another one can help support the workload. However, this drives up complexity and operational support costs. The Ultra Enterprise 10000 provides high-availability features in single server for simplified management.

Capacity for Large-Scale Applications

High-Volume OLTP

Until now, mainframes have been the systems of choice for supporting high volume, mission-critical OLTP. Open systems have delivered high MIPS, but haven't always provided the balanced system architecture and I/O capacity to perform well with high-volume transactional processing. They have also lacked mainframe-class availability features. The Enterprise 10000 server was designed to provide a balanced architecture for the largest OLTP challenges.

With 32 low-latency SBus I/O channels and up to 12.8 Gbytes/sec of system bandwidth, the Enterprise 10000 can deliver tens of thousands of I/O operations per second to meet the largest OLTP needs. What this means is that the Enterprise 10000 can run large enterprise applications with a single database server and avoid the database partitioning headaches that many corporations face.

Large-Scale Decision Support

Open systems have often been the choice for data warehousing initiatives because of low entry prices, inexpensive DASD, and the plethora of decision support software applications available. The 64 UltraSPARC[™] CPUs and 6.4 Gbytes/sec of I/O bandwidth in the Enterprise 10000 server allow it to move massive data volumes and support the processing needed for complex queries. It can handle the largest queries and outperform exotic architectures like MPP, clusters, and NUMA, but without the implementation risks inherent in these unproven and more complex technologies.

The Ultra Enterprise 10000 server has demonstrated linear scalability for DSS queries with popular databases and linear scaling for data mining applications from leading vendors. And the Enterprise 10000 server can support the largest Data Warehouses—into the 20Tbyte range today.

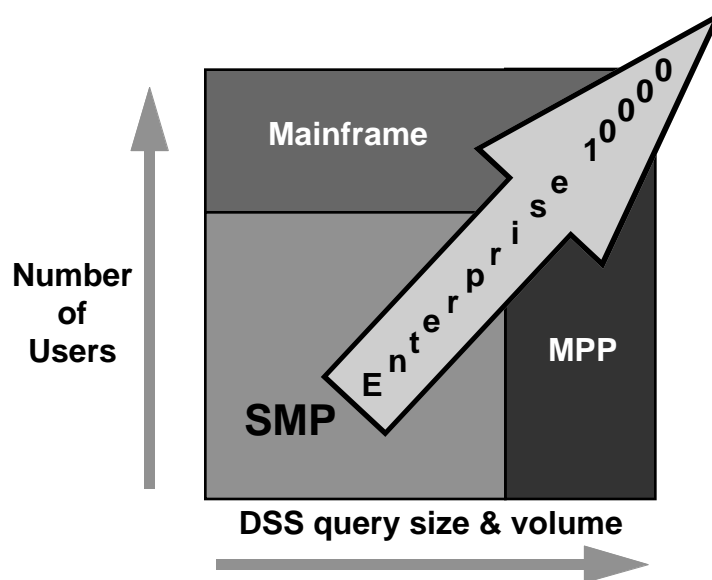


Figure 3-4 The Enterprise 10000, with its production-ready SMP architecture, can be configured for the largest OLTP and DSS problems that have previously been the realm of costly mainframes and exotic MPP architectures.

Most importantly, the Enterprise 10000 is a general-purpose system. It handles either transactional or large query applications with equal ease. In fact, with domains, you can configure both OLTP and DSS databases on the same system. This flexibility means that your investment is preserved because the Enterprise 10000 can easily accommodate changing processing needs as your business applications evolve.

Capacity for Today and Tomorrow

New strategic applications like enterprise resource planning are driving up client/server OLTP transaction volumes. Companies have deployment plans which call for hundreds, even thousands, of users to come online in the next 12 to 18 months. The Ultra Enterprise 10000 server allows you to start small and then upgrade the system as your needs increase, all the way to 64 processors, 64Gbytes of memory, and 64 I/O controllers in the same chassis. No other shared memory system has this capacity. With other systems you have to pay for additional system chassis, power supplies and hardware infrastructure anywhere from two to six times as often as with the Enterprise 10000. With the Enterprise 10000, your investment is protected, and hardware cost of ownership is lower than alternative systems.

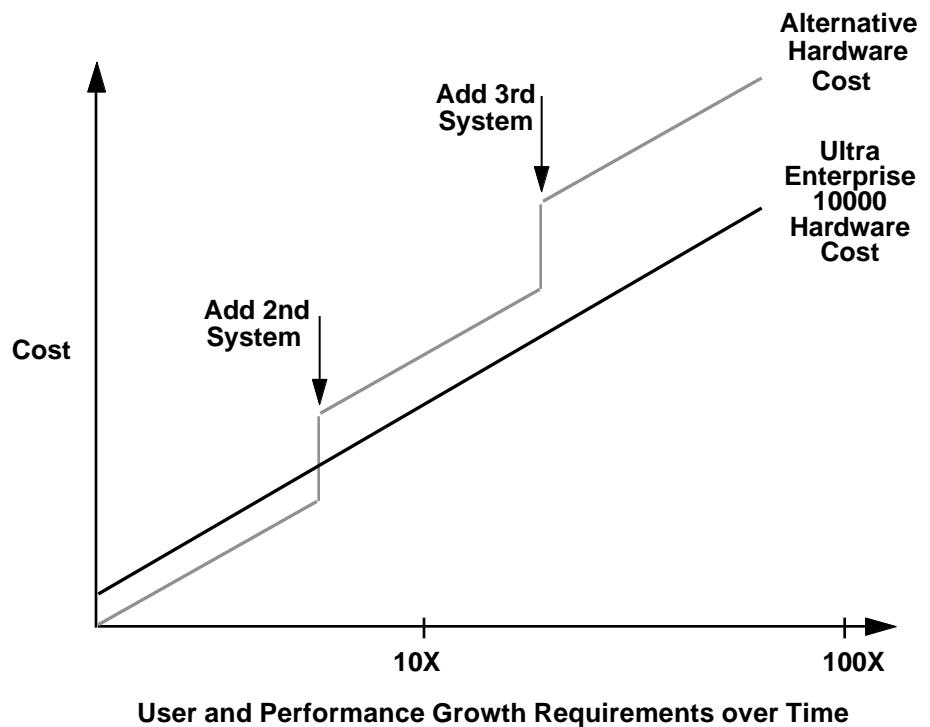


Figure 3-5 Hardware costs for the Enterprise 10000 are lower than other SMP systems which are forced to resort to multiple system solutions as user loads and performance demands increase.

Nor do you need to struggle with the complexity and inefficiency of administering loads and data across multiple servers. In addition to lowering your hardware costs, you lower your operational costs, which studies have shown rise dramatically with each additional system you own and manage.

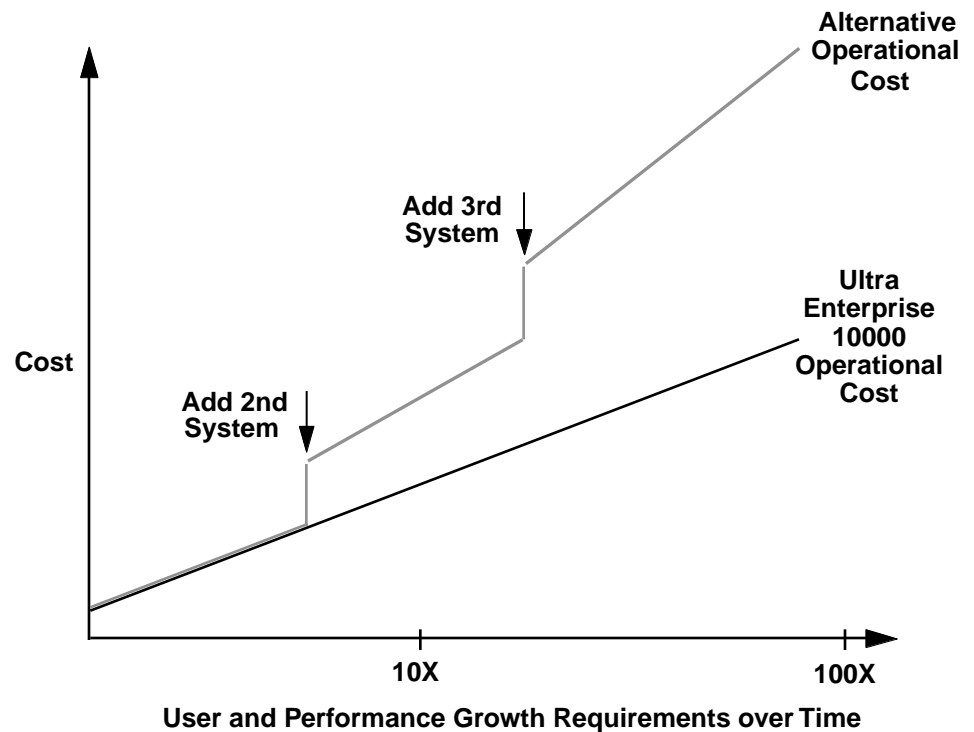


Figure 3-6 The operational costs for the Enterprise 10000 remain much lower than other SMP systems which are forced to resort to multiple system solutions as user loads and performance demands increase.

Capacity for Maximum Configurability

Most servers are designed to meet a small to mid-sized range of requirements. They are general purpose systems that meet the needs of a broad market audience. The Ultra Enterprise 10000 is different. It was designed specifically to meet the needs of the enterprise for large-scale problems.

For example, the expansion capability of the Enterprise 10000 system is designed for full scalability in all areas. Because of the way each system board is built, with independently upgradeable CPUs, memory, and I/O channels, you can upgrade your systems in a balanced way. Unlike other systems, where capacity is limited by slot trade-offs, you can configure an Enterprise 10000 system to full capacity: 64 processors, 64 Gbytes of memory, and 64 I/O controllers, at the same time.




Figure 3-7 Unlike other systems which have slot limitations, the Enterprise 10000 is designed to support the combined maximum configuration of CPUs, memory, and I/O boards.

In fact, every part of the hardware scales as the system grows. As you increase the power of the system, the capacity of the interconnect grows, as does the capacity of the power supplies, and so on.

The Ultimate Single Server Solution

With 64-way everything, the Ultra Enterprise 10000 server has the capacity that is needed in the data center. Whether it is used as a single enterprise server, or morphed into several powerful smaller servers, the Enterprise 10000 can provide the transparent flexibility that is needed for today's data center challenges.

Application Examples



The three application examples that follow illustrate the capabilities of the Ultra Enterprise 10000 server to meet different business and application needs. The first uses a typical implementation of SAP™ R/3™ to discuss the features of the Enterprise 10000 that support server consolidation in the enterprise, specifically the ability to consolidate multiple database servers and application servers.

Another application example explores the rapid growth in business intelligence applications and the burden they are placing on IT infrastructure. This example highlights the use of the Enterprise 10000 to recentralize DSS applications, by moving processing off multiple client machines onto a larger server system.

The third applications example investigates the challenges of consolidating mainframe databases in deploying customer data warehousing applications. It further details how the unique architecture of the Enterprise 10000 can provide the flexibility to meet the changing needs of the rapidly growing data warehouse in the enterprise.

Application Example 1–Server Consolidation for Large-Scale SAP R/3 Deployments

Many businesses are migrating to large packaged enterprise resource planning applications like SAP R/3 which provide a comprehensive view of the business to all departments and help streamline operations. The compelling cost advantages of packaged applications over custom applications, along with the year 2000 issues and the desire to integrate all business units are fueling the rapid adoption of these applications.

However, with conversion costs in the millions of dollars, and the risk associated with rearchitecting the enterprise, implementations must be carefully planned to ensure their success. Because of its complexity and the number of users in large corporations, SAP R/3 requires powerful systems with high transactional query capacity. And businesses also need high reliability and availability for this mission-critical application. The capacity, partitionability, and high availability of Ultra Enterprise 10000 can uniquely meet the requirements of SAP R/3 deployments in large corporations.

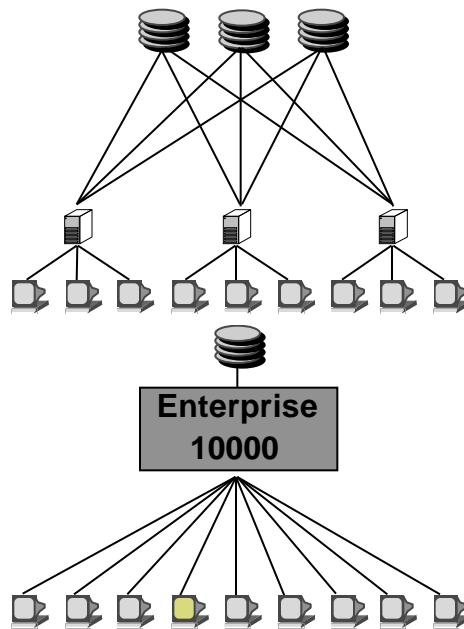


Figure 4-1 Some corporations are forced to partition both SAP R/3 database and application servers to meet user loads. The Enterprise 10000 can act as a consolidated database and application server with excellent performance and simplified administration.

Database And Application Server Consolidation

While most large R/3 installations are not yet fully in production, current data shows that database server capacity will be an issue for many corporations. No current system has demonstrated the ability to support more than a few hundred actual users without resorting to some form of multi-system architecture. Companies are being forced to segment the database server across multiple platforms to meet the user load requirements. The risk associated with this is an additional burden that IT organizations would rather avoid. Application server partitioning is a lesser technical issue, but drives up the overall management complexity and costs of the deployment over a consolidated alternative.

The Ultra Enterprise 10000 enables corporations to consolidate application and database servers. For mid-sized corporations, both application and database servers can be consolidated into a single Enterprise 10000 system. For the largest corporations, a single Enterprise 10000 database server and separate application server can be implemented to still avoid database partitioning. The Enterprise 10000's capacity enables you to consolidate many application servers onto a single physical machine. And this capacity also means that your R/3 deployment can easily grow to support many more simultaneous users than before.

One Platform For Production, Development, And Training

You can also avoid the need to maintain separate platforms for development, training, and production, by deploying development and training domains on the Enterprise 10000. And development system crashes cannot impact the production domain. Since training is not always needed continuously, its system resources can be combined to assist development of production uses as appropriate, as can development domains.

Application Example 2—Recentralization of Corporate Business Intelligence Services

Businesses have recognized that rapid responses to changes in business climate and competitive pressures are key to success. One of the strategies employed is to empower a broader user community to analyze data and make decisions as needed with business intelligence applications like DSS, EIS, OLAP, and data mining systems. Recent Gartner Group findings expect the deployment of business intelligence applications to continue to increase dramatically in U.S. corporations.

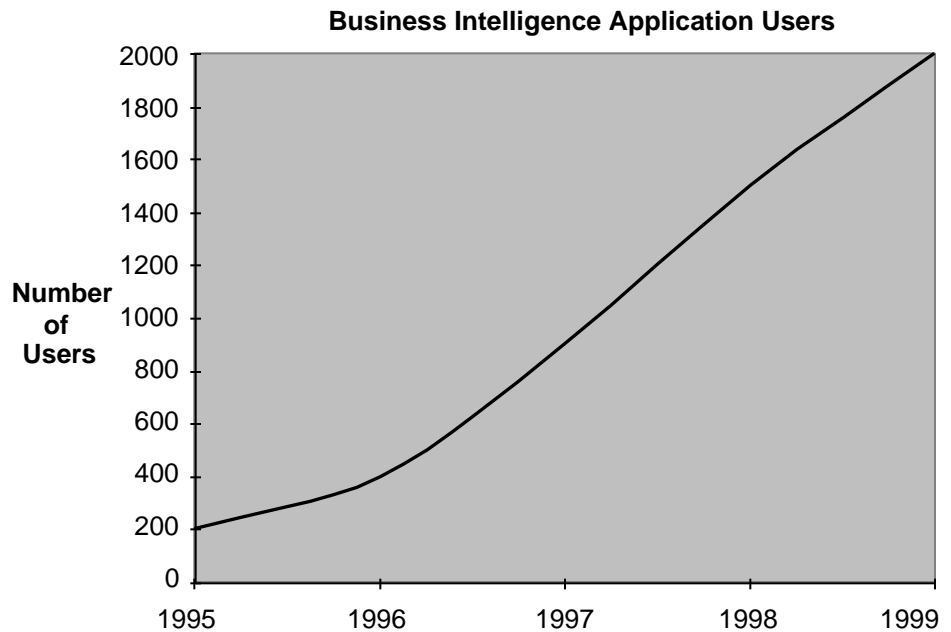


Figure 4-2 Research findings project dramatic increases in the number of business intelligence application users in corporations Source: Gartner Group, *Business-Intelligence Deployment Strategies*, OIS: R-BI-131, 26 September 1996.

Fat Clients

As a result, many companies find themselves deploying more and more complex distributed applications on PCs and small servers. Business intelligence applications require a great deal of processing power and are data-

intensive, driving up the costs of support and requiring frequent CPU and memory upgrades. Even low-cost packaged OLAP application acquisition and support costs are high when multiplied by an increasing number of clients.

And often, the analysis complexity is artificially limited by constraints on the power available on the user's desktop. With client processing and storage limitations, the user rapidly learns not to ask the most meaningful questions to address his business needs, but rather to ask questions that can be answered in a reasonable amount of time based on his available processing resources.

Network Bottlenecks

Business intelligence applications return large result sets to the client, sometimes resulting in network bottlenecks, increasing the amount of local storage required, and exacerbating the already overwhelming network backup requirements.

Application Recentralization with the Enterprise 10000

The Ultra Enterprise 10000 server has the capacity to support the recentralization of DSS applications for better performance at reduced cost. This is implemented in a three-tier client/server application architecture.

A three-tier client/server architecture is comprised of data, business logic, and presentation layers, differing from the traditional two-tier model which does not differentiate between the application logic and presentation layer. The three-tiered model has gained increasing favor for its manageability and scalability. Many current business intelligence applications support this architecture.

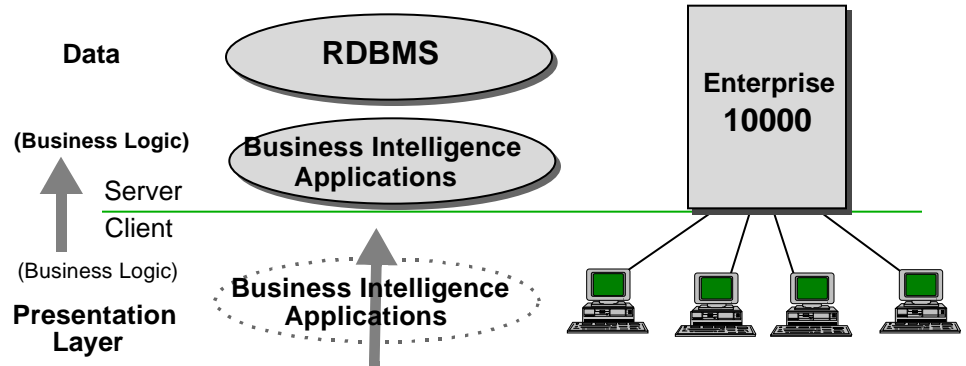


Figure 4-3 The Enterprise 10000 can enable recentralization of client applications for improved user performance and reduced costs to IT.

The key to the recentralization is the migration of the business logic tier to the server. The Enterprise 10000 system logically acts as the upper two tiers by providing both the database and application logic, while maintaining the presentation layer on the client. By retaining the database and application logic in a single server, bulk data no longer needs to flow over corporate networks, benefiting all network users.

Recentralized applications can also exploit the parallel processing capacity of the Ultra Enterprise 10000 system to rapidly scan large quantities of data and analyze result sets. Faster processing allows the end user to more rapidly perform ad hoc queries based on the results of his previous query. The ability to quickly iterate through multiple queries based on new information is key to successful use of decisional applications.

While the users enjoy faster turnaround because of their ability to exploit the parallel query performance of the Enterprise 10000, the greatest benefit is to IT, as noted by the META Group (1996 Open Computing & Server Strategies META Trends, 6/21/96):

“As C/S applications and deployment styles mature, application control will gradually migrate to servers. By 1997-98, ‘fat server’ will dominate ‘fat client’ as the preferred IT architecture, enabling lower-cost C/S applications, encouraging broader use (1998/99) of clients with ‘skinny’ footprints (e.g., mobile and Web browsers)...”

Centralized application logic enhances manageability. Users no longer require large local data sets for processing, resulting in simplified data backup and recovery. And importantly, now users can readily share query result sets and reports without resorting to hardcopy.

For the key emerging application of data mining, an application domain can be created to perform the data mining analysis without interfering with the normal interactive user query processing. The implementation of data mining and the performance implications for the Enterprise 10000 are detailed in a whitepaper entitled *Scalable Data Mining with Sun Ultra Enterprise Servers*.

Application Example 3–Mainframe Database Consolidation for Data Warehousing

Corporations are increasingly recognizing that data gathered about its customers is one of its most important competitive advantages. Unfortunately, customer data is often kept in many isolated mainframe databases that are organized by product line only, making it impossible to view customer data in a holistic fashion. Without a consolidated view of customer data, it is almost impossible to answer important business questions about the customer base to develop cross-selling or household-based marketing strategies.

Many businesses are implementing customer data warehouses to centralize customer data and enable effective analysis by decision support applications.

Data Warehouse Deployment Issues

Most have opted for a UNIX[®]-based open systems implementation for their data warehouse because of merchant database and DSS applications availability, compelling MIPS and DASD cost advantages, and the superior network connectivity found in open systems.

However, corporations are finding the decision regarding the hardware and RDBMS platform for data warehousing to be relatively simple compared to the complexity of integrating sometimes more than a dozen standalone databases spread across numerous mainframes. While there are clear advantages to creating a consistent enterprise view of customer data for all organizations, recent research has noted that some organizations have experienced difficulty in deploying enterprise data warehouses. The complexity of developing a consolidated enterprise data warehouse, coupled with the dynamic requirements of business users, can obsolete a data warehouse before it is implemented if the development cycle is too long and the result inflexible. It is also difficult to gain funding for a large-scale data warehouse without the ability to demonstrate the ROI, perhaps with a pilot project that targets a specific area of use.

The solution offered by some is to implement small project-specific data warehouses, known as data marts. These implementations have a shorter development cycle and can be tailored to a specific business need, with demonstrable value to the bottom line. Industry observers are also quick to note, however, that multiple data marts suffer from the same management and

disaggregation issues that the data warehouse was designed to correct. A consolidated enterprise view of customer data is in fact necessary to maintain the data mart data over the long term.

The emerging trend is towards a multi-tiered data warehouse architecture which employs data marts supplied from a consolidated database. When implemented in a staged fashion, this strategy offers both the benefit of the low-cost, quick ROI of the departmental data mart, and the consistent enterprise data view necessary to ensure maintainability and long-term value to the overall corporation.

Database Consolidation with the Enterprise 10000

The Ultra Enterprise 10000 server supports the needs of customer data warehousing and data mining applications with its capacity to rapidly analyze massive amounts of data in conjunction with the unique ability to partition and redeploy resources to support a growing and evolving data warehouse architectural implementations.

Data Mart Consolidation

With Enterprise 10000 systems, IT can easily implement a multi-tiered data warehouse architecture that can readily move through the implementation phases from pilot to multi-departmental production.

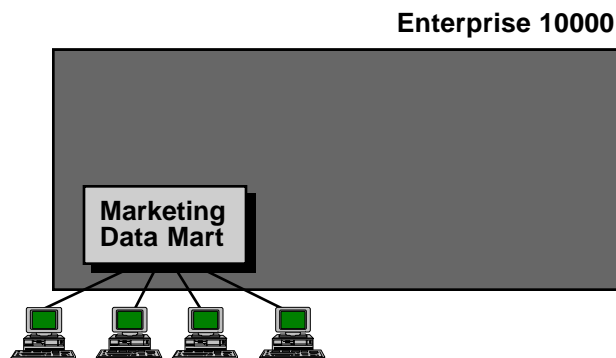


Figure 4-4 Data marts can be easily implemented as an Enterprise 10000 system domain.

The initial pilot project is likely to be a marketing data mart which targets a high profile need within the corporation.

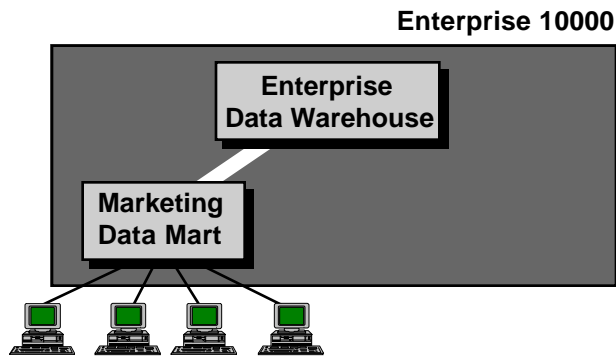


Figure 4-5 Additional tiers in the data warehouse architecture can be incorporated as needs grow.

As the project succeeds, a new domain can be deployed to begin developing the consolidated enterprise data view to support additional data marts. This domain will employ a separate database which will evolve to the primary database for updates from the mainframe product databases.

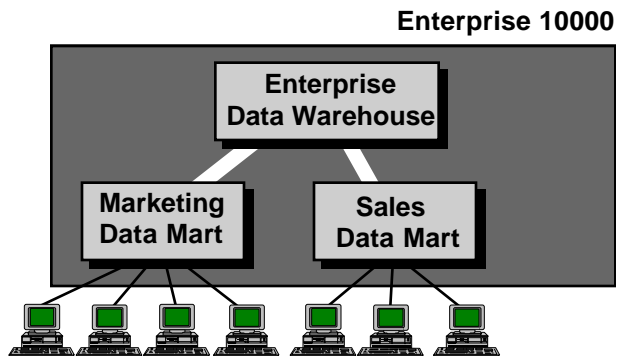


Figure 4-6 Additional data marts can also be easily added.

New and changing requirements can be incorporated easily by implementing the architecture in a phased multi-tiered fashion with the flexibility of domains on the Enterprise 10000 system.




Figure 4-7 The Enterprise 10000 can simplify implementation and management of a multi-tiered data warehouse architecture with an enterprise warehouse that feeds departmental data marts.

Additional data marts can be incorporated by repartitioning the system resources and adding more processing power and storage as needed.

Customer Data Warehousing with the Enterprise 10000

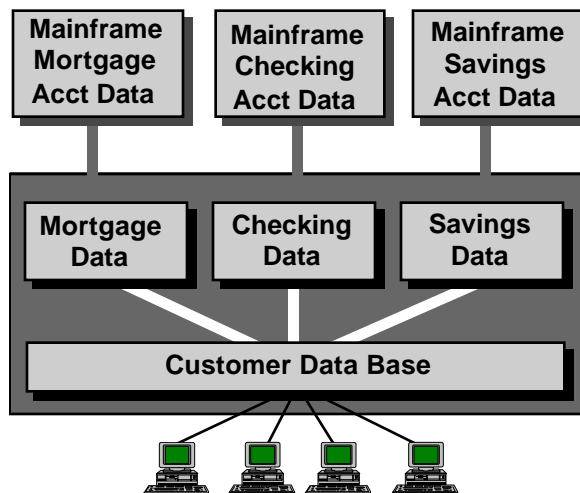


Figure 4-8 Dynamic system domains in the Enterprise 10000 can be deployed as replication points for mainframe data to rapidly implement a customer data warehouse architecture.

Some financial institutions are using domains to consolidate mainframe databases into a customer data warehouse. In this case, they are setting up domains for purposes of database replication and using virtual warehousing applications to present a consolidated view of the multiple databases for customer relationship management applications. The Enterprise 10000 server is more simple to manage and deploy than other virtual warehouse alternatives that require complex middleware residing on many platforms.

Data warehouses must remain flexible and dynamic to be successful and meet ever changing user needs. The Enterprise 10000 server facilitates unique methods of implementation which simplify migration of databases from legacy systems into easily managed data marts and help establish a scalable data warehousing infrastructure that will ensure long term value to the corporation.

Conclusion



Corporate computing styles have migrated from centralized host-based computing to highly distributed PC-based computing over the last decade. New highly-sophisticated strategic applications are being deployed at a dizzying rate to meet the competitive challenges in the marketplace.

However, the rapid expansion and rising costs of IT infrastructure in companies have forced them to look closely at how efficiently their current computing solutions meet their business needs. Companies are increasingly recognizing a compelling need to consolidate and centralize management of systems and data while supporting an increasingly distributed user community.

The Ultra Enterprise 10000 server is a general-purpose solution for high-end computing with the capacity and manageability to reduce the complexity of your computing environment. It can provide you with a highly flexible solution to your business's changing computing needs, both current and future.

References



Gartner Group

Riding the Whirlwind: The Transformational Effects of Managing Distributed Systems, Networked Systems Management (NSM) Strategic Analysis Report, by W. Malik and A. Cushman, September 26, 1996.

Business-Intelligence Deployment Strategies, Office Information Systems (OIS) Strategic Analysis Report, by H. Dresner and B. Lett, September 26, 1996

Strategies to Increase LAN Availability, Network Computing Infrastructures (NCI) Strategic Analysis Report, by J. Barkan and C. Smith, September 23, 1996.

Distributed Computing Platforms Scenario, Distributed Computing Platforms (DCP) Strategic Analysis Report, by G. Weiss, P. McGuckin, and T. Berg, April 2, 1996.

Minimizing LAN Costs, Improving Support, Yield Big Dividends, Point-To-Point (PTP), by D. Cappuccio, January 26, 1996.

Enterprise Resource Planning: Optimizing Business Functions, Inside Gartner Group (IGG), by B. Bond and E. Keller, January 24, 1996.

Why the Mainframe Will Not Be an Enterprise Data Server, Client/Server (CS) Research Note Strategic Planning, by A. Percy, October 16, 1995

The Third Generation of Client/Server Computing, Client/Server (CS) Strategic Analysis Report, by T. Austin and C. Miller, September 26, 1995.



META Group

A Taxonomy of Corporate Data Warehouses, Application Development Strategies, March 11, 1996

Client/Fat Server Computing: Deploying Complex Solutions, not Problems, 1st Quarter Trend Teleconference, Open Computing & Server Strategies, March 6, 1996

Sun White Papers

Scalable Data Mining with Sun Ultra Enterprise Servers, Donna Rubin, January 1997.



Sun Microsystems Computer Company
A Sun Microsystems, Inc. Business
2550 Garcia Avenue
Mountain View, CA 94043 USA
415 960-1300
FAX 415 969-9131

Sales Offices

Argentina: +54-1-311-0700
Australia: +61-2-9844-5000
Belgium: +32-2-716-7911
Brazil: +55-11-524-8988
Canada: +905-477-6745
Chile: +56-2-638-6364
Colombia: +57-1-218-3933
Commonwealth of Independent States:
+7-095-956-5470
Czech/Slovak Republics:
+42-2-205-102-33
Denmark: +45-44-89-49-89
Estonia: +372-6-308-900
Finland: +358-0-525-561
France: +33-13-067-5000
Germany: +49-89-46008-0
Greece: +30-1-680-6676
Hong Kong: +852-2802-4188
Hungary: +36-1-202-4415
Iceland: +354-563-3010
India: +91-80-559-9595
Ireland: +353-1-671-4678
Israel: +972-3-695-6868
Italy: +39-39-60551
Japan: +81-3-5717-5000
Korea: +822-3469-0114
Latin America/Caribbean:
+1-415-688-9464
Latvia: +371-755-11-33
Lithuania: +370-7-298-586
Luxembourg: +352-491-1331
Malaysia: +603-264-9988
Mexico: +52-5-258-6100
Netherlands: +31-33-450-1234
New Zealand: +64-4-499-2344
Norway: +47-2218-5800
People's Republic of China:
Beijing: +86-10-6849-2828
Shanghai: +86-21-6247-4068
Guangzhou: +86-20-8777-9913
Chengdu: +86-28-678-0121
Poland: +48-2-658-4535
Portugal: +351-1-412-7710
Singapore: +65-224-3388
South Africa: +2711-805-4305
Spain: +34-1-596-9900
Sweden: +46-8-623-90-00
Switzerland: +41-1-825-7111
Taiwan: +886-2-514-0567
Thailand: +662-636-1555
United Arab Emirates:
+971-4-366-333
United Kingdom: +44-1-276-20444
United States: +1-800-821-4643
Venezuela: +58-2-286-1044
Worldwide Headquarters:
+1-415-960-1300