Consolidation in the Enterprise

Mike Malgeri

Technical Report 22
Version 1.0
March 1998

Document 300-TC050, V1.0, 980409

Auspex Systems, Inc.
2300 Central Expressway
Santa Clara, California 95050-2516 USA

Phone: 408/566-2000
Fax: 408/566-2020
Email: info@auspex.com
http://www.auspex.com

TABLE OF CONTENTS

INTRODUCTION **1**

Corporate Challenges 1

Client/Server Trends 1

1 FILE SERVER REQUIREMENTS **2**

Reliability 2

Scalability 2

Performance 2

Data Management Capabilities 3

Support for Shared NT/UNIX Environments 3

2 AUSPEX IN THE FILE SERVER ROLE **4**

APPENDIX A: INDUSTRY REFERENCES TO DATA CONSOLIDATION **5**

APPENDIX B: WHY CONSOLIDATED SYSTEMS ARE MORE RELIABLE THAN DISTRIBUTED SYSTEMS **6**

Introduction

Corporate challenges are steering IT managers towards consolidation of critical enterprise data. Consolidation helps meet these challenges by increasing reliability, scalability and performance when accessing network data. Data management and support of mixed UNIX/NT environments is facilitated. Auspex has helped hundreds of companies consolidate their data. This paper describes the motivations for consolidation and how Auspex uniquely satisfies corporate objectives through consolidation while offering continuous data access in distributed computing environments.

Corporate Challenges

Today's corporations are faced with many challenges, each impacting the success or failure of the enterprise. Some of these challenges include:

- *Global Enterprise*

  In order to tap new sources of revenues, corporations are venturing beyond their traditional customer base, establishing sales, distribution and manufacturing facilities worldwide. Successfully meeting the needs of diverse markets presents technical, cultural and political challenges.

- *Supply Chain Integration*

  Corporations are forming strategic partnerships with suppliers and customers in order to achieve synergy across the entire spectrum of functions from defining customer requirements to delivering and maintaining finished products.

- *Customer Requirements*

  As competition increases, companies are forced to offer higher quality, feature rich, cost justifiable products approaching the "envelope" in performance and functionality. Terms such as "a market of one" have been coined to designate a trend toward rapid product development to meet increasing customization requirements.

- *Strategic Use of Information Technology*

  Information has become a competitive weapon. Operational data must be continuously available, and corporations are discovering that new opportunities can be found by mining for untapped markets revealed in their data warehouses.

- *NT/UNIX Integration*

  IT executives recognize the distinct advantages of NT and UNIX and are required to create environments leveraging the strengths of each technology.

Meeting the challenges described above depends heavily on IT's ability to create an effective computing, networking and data management infrastructure centered on data consolidation.

Client/Server Trends

As Microsoft, system vendors and ISVs simultaneously battle and form alliances, two client/server trends are clear. The first is the distribution of processing power, and the second is the consolidation of data. Microprocessors provided users with deterministic CPU cycles to complete their work. As processor power increased, developers answered the call for greater functionality by creating user friendly, graphically driven applications, increasing the trend toward desktop processing. Unfortunately, as users retreated to their desktops, their data followed. Individual productivity was achieved at the expense of corporate synergy. Reliability, security and data management benefits of mainframe computing were lost in the trend.

Today, thanks to standard APIs at all levels of computing, users can select a variety of hardware and software to meet their processing and functionality requirements. Interoperability facilitated advances in network and data management technology prompting users to rediscover the benefits of consolidated data storage. The next two sections focus on the features of high performance, reliable file servers that allow distributed processing while maintaining the benefits of data consolidation.

1 File Server Requirements

Although high capacity file servers have many attributes, there are core features required to build an enterprise data and networking infrastructure.

Reliability

Corporations and industry analysts estimate the "cost of downtime" in order to determine the impact of unreliable hardware and software on revenues and profitability. DataQuest estimated downtime costs range from $88K/hour for companies in the transportation industry to $6.4M/hour for companies in the financial industry[1]. In many cases, such as in manufacturing companies, the actual number is understated if it does not consider lost market share resulting from missed product delivery schedules.

Software, hardware and service effect fileserver reliability. A vendor must be dedicated to continuous improvement. It is required that these processes permeate a vendor's organization to ensure that reliability is built into the product and a top priority of every department. System vendors should prove their commitment to reliability by justifying reliability claims with empirical data. Buyers should not accept theoretical reliability estimates calculated from hardware MTBF ratings of a system's components. This does not represent what may be experienced from a system deployed in a mission critical application.

Data consolidation is a proven storage strategy for meeting reliability and data management goals[2]. An approach using many smaller, distributed servers, creates network problems through cross mounting, where one "downed" server can affect the entire group. Multiple small servers increase data unavailability, as measured by downtime[3]. Also, the competitive advantages organizations gain through data sharing among cooperating departments and partners, require a consistent view of information. Data consolidation is the best way to achieve this view. Multiple servers must use cross mounting to create this view. This decreases reliability and creates system hardware, software and data management problems which are almost completely avoided by deploying a highly reliable, consolidated file server. In a distributed server environment, various firmware levels on disks and NIC cards and different patch and revision levels on operating systems create configuration management challenges. Incompatibility problems can cause significant downtime and are difficult to debug.

Consolidation and reliability are complementary. Consolidation makes an information system more reliable, and a consolidated system must be extremely reliable because many people in an enterprise depend on the data it serves.

Scalability

Data growth rates are keeping pace with improving processor power. Many companies report 50-100% yearly increases in storage requirements. "Forklift" upgrades of critical network file servers, requiring many hours of downtime, are not practical when capacity increases are needed. File servers must accommodate data growth and support new or additional network interfaces with minimal disruption in service. It's not enough to rely on replacing current disks with higher density drives, as data growth rates are higher than drive density improvements.

Performance

As developers continue to enhance the functionality of applications, desktop and networking performance play an increasing role in user productivity. When IT managers implement data consolidation, users of powerful

---

[1] DataQuest 9/96

[2] See Appendix A for industry Quotes and articles regarding the importance of consolidation on reliability, data management, security and cost.

[3] See Appendix B for further discussion.

desktop systems should not be constrained by slow disk and network I/O. A fileserver must be capable of sustained saturation of its network channels with rapid response for all disk operations. Performance of network file service is characterized by sustainable peak throughput, low response time under various load conditions, and sustainable peak file transaction rates. High scale consolidation requires special software and hardware architectures to optimize I/O flow. Auspex Functional Multiprocessing® Architecture (FMP) meets these requirements. General-purpose systems with architectures better suited for compute, database or application processing are ill suited to meet the performance demands of consolidated network file service.

Data Management Capabilities

IT managers protect corporate data. Reliability and security must be achieved while maintaining flexibility. System administrators are required to perform frequent and fast backups. Backup is difficult because of high data capacities and the need to access data on a 7x24 hour basis. There are often no "windows" in which a file system can be taken offline in order to create a consistent (non-corrupted) backup image. A solid backup strategy is crucial for the protection of enterprise data. FastBackup addresses the narrow window that mission critical environments have and provides the capability of backing up live data on terabyte systems in a 10-hour shift.

A file server must support different projects and users with separate, sometimes opposing needs. The ability to flexibly configure filesystems around needs is extremely valuable. Support for multiple filesystems provides the benefits of data consolidation while simultaneously permitting sharing or separation of projects that have independent data. It may also be necessary to relocate filesystems in the event of disk upgrades or expand them in response to project requirements. An administrator should be able to perform this task online so that users, who are working on the original filesystem, can transparently access their data from the new location without knowledge of the move or expansion.

Failure of any one filesystem should not affect users of others. This requires that a file server provide file system isolation and hot plug-ability of failed drives. In addition, multiple RAID levels are needed to support the various reliability and performance requirements of the user community. Critical application binaries may require the highest performance and reliability, justifying the added expenditure of a RAID 1 filesystem. "Temp" directories used as high performance "scratch space" may not need redundancy and administrators can cut costs by deploying RAID 0 filesystems for this purpose.

A consolidated data server, with built in capabilities for fast backup, online filesystem relocation and expansion, filesystem isolation and multi-level RAID, is the best tool for meeting demanding data management needs. Multiple small, distributed fileservers add data management complexities such as the need to schedule local and/or network backups, slow backup solutions, inflexible RAID options and, in some cases, a requirement to deploy ONLY one large filesystem per server.

Support for Shared NT/UNIX Environments

Companies recognize the need and advantage of successfully integrating NT and UNIX environments. File servers must provide NT/CIFS users, the same data consolidation benefits traditionally supplied to UNIX/NFS users. This shouldn't require that NT users or administrators relinquish tools and security services familiar to their environment. Users of MCAD or ECAD applications that have migrated to NT platforms require the ability to access and share files originally created in a UNIX environment. Also, as they create files on their new NT systems, "collaborative engineering" requires they share files with colleagues that are still using UNIX desktops. In many cases, mixed workstation environments will persist. UNIX desktops will be reserved for users requiring the highest power and functionality, or for backend applications producing content that will be consumed by NT workstations.

Sharing files between producers and consumers of information in their native protocols is essential. However, even in companies where NT will fully replace UNIX workstations, the conversion process must occur over time while ensuring a recovery to UNIX should the transition process encounter problems. NeTservices[4], a

[4] For detailed information regarding this product, please see "Technical Report 19, Auspex NeTservices Delivering "No-Compromise Consolidation of UNIX and Windows NT Data"

software product from Auspex, combines with the FMP architecture to bring continuous data access to shared NT/UNIX environments, or those migrating completely to NT, while meeting the specific needs of both classes of users.

2 Auspex in the File Server Role

Auspex has helped hundreds of companies by meeting or exceeding critical file server requirements, which play a role in IT infrastructure. Auspex's strategy is to provide solutions ahead of problems customers will face. In anticipation of disk and network I/O bottlenecks Auspex invented the Functional Multi-Processing (FMP) architecture to specifically address architectural limitations of general purpose SMP fileservers. FMP has a proven scalability record, growing from a 40 GB/8x10BT maximum configuration in 1991 to 1.8 TB/30x100BT support today. Customers have leveraged their investments in FMP technology by adding disks and network interfaces as needed with near linear scalability and minimal or no downtime for upgrade migrations. In this context, "linear scalability" refers to the fact that system resources, dedicated to moving data, are increased by a factor of X. The I/O throughput of the original system is also increased by a factor of X. General-purpose, symmetric multiprocessing systems (SMP) can not provide linear scaling due to overhead in maintaining memory consistency in a multi-CPU architecture. As more CPUs are added in SMP systems, the incremental performance increase gets less and less.

It is important to note that data and network capacities quoted by Auspex are derived from "real world" configurations. Auspex can provide an extensive list of customers using systems with >1 TB of data, in mission critical environments. Competitors specify multi-terabyte capability with few, if any, production reference sites. Typically, performance and capacity specifications quoted by competitors are based on testing of a unique and highly tailored configuration that misrepresents how customers in real environments would use systems. For instance, some vendors, in order to publish high NFS I/O figures, configure a fileserver with 360 filesystems. Each filesystem is created on one of 360 disks! This is a configuration most system administrators would not deploy.

In response to growing IT environments and 24x7 system availability requirements, Auspex developed high-availability software solutions, such as DataGuard, ServerGuard and FastBackup. Auspex has always recognized the importance of UNIX as well as its potential to cause problems. DataGuard provides IT professionals with the best of both worlds. UNIX can be used to run required systems management software such as NIS, DNS, backup, sendmail, license management and performance monitoring tools, while isolating UNIX panics from the NFS data path. Auspex NetServers will sustain a complete reboot of the UNIX operating system while continuously delivering data to hundreds of networked users. ServerGuard takes this approach a giant leap forward by allowing complete server redundancy, in the event of a catastrophic system failure. Failover to the surviving system is accomplished in a few seconds, with users completely unaware of any problems.

Auspex continues to anticipate the needs of its customers and is focused delivering continuous data access to corporations deploying both NT and UNIX systems. Auspex will bring the same performance, scalability and consolidated data management benefits to mixed NT/UNIX networks that it did for access to shared NFS data.

In support of the information presented in this document, detailed customer case studies can be obtained from the Auspex Web site at www.auspex.com. Learn how companies have successfully deployed consolidated Auspex file servers to help meet corporate challenges in today's global and dynamic marketplace.

Appendix A: Industry References to Data Consolidation

Many enterprises are developing server-consolidation strategies. Smaller numbers of current technology servers are replacing older technologies. Effectively done, server consolidation can reduce management by up to 40%. It can also reduce bottlenecks, improving storage- and service-related costs. According to Strategic Research (Santa Barbara, CA), while distributed storage costs about $7 per megabyte per year, consolidated storage can cost as little as $2 per megabyte per year. ***Unix Reseller News, June 1997, "Storage Trends for Unix and NT"***

Client-server computing has its origins in the concept that economies of scale can be realized by concentrating the data used by several small (often single-user desktop) computers in one location so that storage can be procured, deployed, protected, and managed as a single asset. ***Unix Reseller News, June 1997, "Storage Trends for Unix and NT"***

As the popularity of LANs with file servers grew, it became apparent that group computing required closer application and data coordination than was possible simply by locating shared files centrally. Network file services rapidly evolved into true client/server computing ***Unix Reseller News, June 1997, "Storage Trends for Unix and NT"***

One business issue is the centralized control of information systems. IT executives must deal with the aftermath of client/server systems and Web sites that have grown uncontrollably…the downside was that many of these systems have grown willy-nilly and companies are now looking to central IT groups to bring order out of chaos. ***Datamation Magazine, July 1997, "Squeezing Profits from IT"***

Why consolidate? One reason is improved manageability: server consolidation can make such labor-intensive chores as software distribution, backup and recovery easier, all while improving system security. By consolidating, users can often lower their total costs of ownership by reducing the need for manpower and other IT management resources, cutting the number of software licenses a company must purchase, and so on. ***Rick Whiting, Client/Server Computing Online, June 1997, "Some Assembly Required"***

Gartner Group's Richmond says [consolidation] should be considered whenever IT cost or operational pressures can be relieved by consolidation. ***Rick Whiting, Client/Server Computing Online, June 1997, "Some Assembly Required"***

Says Lynn Berg, group VP for the Gartner Group, "People want decentralized decision making but centralized support." ***Datamation Magazine, July 1997, "Squeezing Profits from IT"***

"Departmental servers have grown like weeds…," says a January 1997 report from the Yankee Group. "IT managers are discovering that the true cost of departmental servers, including maintenance and support, far exceeds the initial cost of the hardware." ***Rick Whiting, Client/Server Computing Online, June 1997, "Some Assembly Required"***

This concept has been so successful, particularly in the UNIX environment, that is has spawned the file server. Today, file servers are offered as standard computer systems configured optimally for data storage and I/0, and by specialty vendors, such as Auspex Systems (Santa Clara, CA), whose products are highly customized for the single purpose of providing file access to clients. ***Unix Reseller News, June 1997, "Storage Trends for Unix and NT"***

Appendix B: Why Consolidated Systems are More Reliable than Distributed Systems

Assume that an IT manager requires 1 TB of data and is considering either a single consolidated server or an alternative approach of four 250 GB servers. Also assume that all servers in this discussion are equal in terms of system reliability, resulting in 1 hour a year of downtime for each server. This yields 1 hour of downtime for the consolidated solution and 4 hours of downtime for the distributed approach.

One might conclude that total data unavailability is still the same for each approach and can be calculated as 1 TB/hour/year (1x1 TB) for the consolidated server and 1 TB/hour/year (4 x 250 GB) for the 4 distributed servers. This conclusion is incorrect.

The problem arises from the fact that organizations generally want users to have access to most or all of the 1 TB data set from each desktop computer. They also want each desktop to have a consistent view of the data set. In a distributed server environment this requires that every desktop mount filesystems from every server. A mount from one failed server can cause the mounting clients to loose access to some or all of the filesystems it mounts from other servers. This makes the data unavailability number for each down event potentially much higher than 250 GB/hour in the distributed server approach. Use of the “automounter” on client desktops, which drops mounts that have been idle for a certain period of time, can alleviate this problem but it is a partial solution. Client desktops, in the distributed configuration, are almost always mounting filesystems from multiple servers.

Another issue to consider is the fact that problems resulting from mounting filesystems from multiple servers or mounting across routers (a practice which causes slow performance at the desktop) are harder to debug. Therefore, the original estimate of 1 hour per down event for each server in a distributed server environment is underestimated.

As the discussion above indicates, consolidated fileservers are significantly more reliable and provide a much higher level of availability than distributed systems serving the same amount of data and users.