Patricia Seybold's Office Computing Group

Editor-in-Chief
Judith S. Hurwitz

INSIDE

UNIX IN THE OFFICE

Guide to Open Systems

Vol. 7, No. 1 • ISSN: 1058-4161 • January 1992

Downsizing with Open Systems

Can Unix Symmetric Multiprocessing Systems Meet MIS Requirements?

By Mitchell I. Kramer

IN BRIEF: The availability of computer platforms based on symmetric multiprocessing technology has allowed many applications that previously required mainframe-class processing to be deployed on smaller, more cost-effective platforms. The vast majority of these platforms use Unix as their operating system, making downsizing and open systems almost synonymous. The ability of these platforms to compete with mainframes will depend on their ability to meet user requirements for reliability, availability, serviceability, and manageability.

Report begins on page 3.

Taking Stock

Looking ahead to 1992

NOW THAT 1991 is officially behind us, it is time to look ahead to see just where the Unix and Open Systems industry is headed and what its prospects are for the future. It would be easy at this point to look upon the malaise and shifting alliances of the computer industry and despair. But despite the problems that will take years to solve, I am optimistic. I see the light at the end of the tunnel because of several factors:

- Vendors are listening to users.
- Users are becoming more active participants in the requirements process.
- Standards-based frameworks are slowly beginning to mature.
- Applications portability has a chance of becoming a reality.
- The computer industry is coming out of a painful adolescence (albeit slowly).

Let me now explain what I mean by each of these factors.

VENDORS ARE LISTENING TO USERS. It is unfair to say that system suppliers have ignored their customers. However, they never really seriously considered letting customers set their future direction. Sure, they were willing to listen to customers' input about what functions should be added to the next release of a product, but users were never smart enough to set strategic directions. That is beginning to change with the maturation of the computer industry. Suddenly, it's a buyers' market, and vendors are now making sure that they hear what users want for their future technology. No vendor would magnanimously decide to implement an open systems strategy unless users demanded it.

USERS ARE BECOMING ACTIVE. In the old days, users complained only after product introduction. Traditionally, they felt they were at the mercy of their system suppliers. The political health within information systems organizations often depended on the fortunes of these companies. With the advent of open systems, users are becoming bolder and more active. They have also grown beyond the point of simply telling vendors their priorities for standards and functionality. Instead, they are beginning to articulate the complexity of the problems and issues they face within their organizations. Helping vendors to understand this complexity is crucial to finding meaningful, pragmatic solutions.

STANDARDS-BASED FRAMEWORKS ARE EMERGING. Users have learned painful lessons about what happens when they begin to implement proprietary technology frameworks in areas such as CASE (spelled AD/Cycle). They realize that it is the same as being locked into one vendor's proprietary operating system or hardware. Emerging development frameworks are gathering enough momentum to offer new hope for standards-based infrastructures.

APPLICATIONS PORTABILITY MAY BECOME A REALITY. The maturation of Posix and the initial interest in new technology such as the Open Software Foundation's DCE and ANDF technologies offers users hope that portability will become a reality in our lifetime.

THE COMPUTER INDUSTRY IS MATURING. It has been a painful few years for the computer industry. No, the pain is not over, but the fact that standards are emerging quickly and finding wide acceptance means that the industry will reach a new stage where innovation at last will be able to flourish.

UNIX IN THE OFFICE

Editor-in-Chief
Judith S. Hurwitz

MCI:
JHurwitz
Internet:
jhurwitz@mcimail.com

Publisher
PATRICIA B. SEYBOLD

Analysts and Editors
JUDITH R. DAVIS
DAVID S. MARSHAK
RONNI T. MARSHAK
JOHN R. RYMER
ANDREW D. WOLFE, JR.

News Editor
DAVID S. MARSHAK

Art Director
LAURINDA P. O'CONNOR

Sales Director
RICHARD ALLSBROOK JR.

Circulation Manager
DEBORAH A. HAY

Customer Service Manager
DONALD K. BAILLARGEON

Patricia Seybold's
Office Computing Group
148 State Street, 7th Floor,
Boston, Massachusetts 02109

Telephone: (617) 742-5200 or
(800) 826-2424
Fax: (617) 742-1028
MCI: PSOCG
Internet: psocg@mcimail.com
TELEX: 6503122583

Unix in the Office (ISSN 0890-4685) is published monthly for $495 (US), $507 (Canada), and $519 (Foreign) per year by Patricia Seybold's Office Computing Group, 148 State Street, 7th Floor, Boston, MA 02109. Second-class postage permit at Boston, MA and additional mailing offices.

POSTMASTER: Send address changes to *Unix in the Office*, 148 State Street, 7th Floor, Boston, MA 02109.

Downsizing with Open Systems

Can Unix Symmetric Multiprocessing Systems Meet MIS Requirements?

Mainframe Minds Meet Open Systems

The computer industry is witnessing a remarkable phenomenon: Unix-based open systems are being considered as alternatives to mainframes.

- With symmetric multiprocessor SMP architectures, open systems can now deliver the processing power of mainframes for about a tenth of their cost. These systems are primarily Unix-based.
- The pressures of competition and of the current economic recession are forcing corporations to run ever more efficiently and productively. IS executives are being asked to find cost-effective alternatives to the next mainframe upgrade. SMP open systems are just such alternatives.

MAINFRAME POWER ON THE CHEAP. The economies inherent in standards-based hardware and software systems have been significantly enhanced through SMP. Systems with the capability of delivering hundreds of MIPS can be acquired for tens of thousands of dollars vs. millions of dollars for a mainframe.

Corporate management has heard the news. Corporate executives are directing IS to consider SMP systems as an alternative to upgrading mainframes. If the power of SMP systems can be used to run mainframe applications, then literally millions of dollars can be saved.

Downsizing: Offloading Work from Mainframes

Using systems smaller than mainframes to offload mainframe processing is called *downsizing*. These smaller systems could be PCs, networked PCs, minis, or supermicros. Downsizing is not a new idea. It's been going on for years.

Twenty years ago, user departments purchased their own Digital Equipment Corporation PDP minicomputers and, later, VAXs, when IS organizations were not able to be responsive to their needs in either function and/or performance to their needs. The advent of the personal computer 10 years ago and of PC LANs a few years later accelerated downsizing to include the desktop. Today, Unix-based, open systems are the focus of a lot of downsizing activity. They offer high compute power, large data capacity, and attractive pricing.

In the past, most downsizing activity was undertaken by line-of-business organizations. The motivation to downsize occurred outside of IS, in response to IS, and while IS continued to develop and support the mainframe.

Mainframe Minds Meet Open Systems

The latest downsizing trend is coming from within IS. Facing the next upgrade in the regular three-to-five-year cycle of upgrades as well as facing a large backlog of applications to be developed, IS is being asked to do what users have been doing for 20 years: run the work on less expensive systems. The downsizing process is the same, but the scale is orders of magnitude larger. This inside-out downsizing requires considerable change, and considerable change means high risk. IS has traditionally been a change-resistant, risk-averse group.

Multiprocessing

Multiprocessing systems are systems that bring the power of multiple processors to bear against a workload. Tightly coupled multiprocessor systems share buses, memory, and I/O. SMP systems are tightly coupled multiprocessor systems in which all processors are peers, all processors can perform any portion of the workload, and all processors have the same CPU power.

Symmetric Multiprocessing

SMP systems have been around for nearly 20 years, with mainframes among the first. SMP open systems, which combine very high performance with the economies of standards-based hardware and software, are the new multiprocessing attraction. They provide a wide, scalable range of computing capacity and power. Scalable systems can address the requirements of many workloads and can provide for nondisruptive growth of these workloads over time.

SMP Makes Downsizing Comfortable

Cost motivates downsizing to SMP open systems, but the similarities in application environments are the key enablers. Because SMP open systems can potentially run large, multiuser, data-intensive mainframe applications the way they run on mainframes, downsizing to such systems feels comfortable and appears feasible to IS. It's an approach that minimizes change and, therefore, minimizes risk.

Downsizing involves two frames of reference: IS managements' consideration of a cost-effective alternative to the next mainframe upgrade, and the open SMP vendors' hopes of satisfying their requirements. The specific elements are:

- Open SMP vendors must understand the characteristics of mainframe applications.
- IS management must define criteria for selecting SMP open systems and identify those that most closely meet those criteria.
- Both groups must be aware of areas where SMP open systems' capabilities do not meet IS requirements. Vendors must improve those capabilities.
- IS management must develop an approach to downsizing. This necessitates understanding how to specify and implement SMP open systems, what applications to develop for them, and how to maintain coexistence between legacy applications running on mainframes and SMP open systems.

Mainframe Application Processing Characteristics: The Requirements for Downsizing

SMP open systems must provide an environment supporting several key characteristics in order to allow large applications to be downsized from mainframes. A description of these characteristics follows.

Business Functions. Mainframe applications implement corporate business functions. They include order entry, inventory control, general ledger, accounts payable, accounts receivable, and payroll. In the extreme, these applications are mission critical. We're all familiar with reservation systems and automatic teller machines. Whatever the level of

importance or sophistication, mainframe applications affect the bottom line by generating revenue or by managing costs.

Centralized. Mainframe processing is centralized. There are one set of application programs and one set of corresponding data. Users make terminal connections to the mainframe to access the applications that execute within the mainframe. Data is stored on disks locally attached to the mainframe. There are a single point of system management and a single point of control.

Large Numbers of Users. A large mainframe installation may support thousands of users who, combined, work with dozens of applications. These applications are composed of millions of lines of code, most of which is Cobol, and require extensive resources to support and maintain .

Multiuser Applications. Mainframe applications are multiuser applications. At peak loads, it would be reasonable to expect that anywhere from 50 to 5,000 concurrent users are accessing a single mainframe application. The type of processing that is done can often be characterized as online transaction processing (OLTP).

Lots of Data. Hundreds of gigabytes of data are available online to mainframe applications, and terabytes of data are stored on offline media. Data is the most critical mainframe resource. Data integrity and availability are essential to the proper processing of mainframe applications. Data is stored in flat files accessed in batch mode by Cobol programs as often as it is stored in relational database management systems (RDBMS) accessed by applications written in fourth generation languages (4GLs).

Databases. Data is organized in databases of various sorts. In mainframe databases, data is often accessed by several related or even unrelated applications concurrently. However, a single mainframe application would require online access to many gigabytes of data. Several times that amount of data would be managed between online and offline media for the proper administration of a mainframe application. The offline data consists of backups of the online data, journals of online database activity, logs of online processing activity, and RDBMS and user program libraries.

High I/O Rates. With so many users accessing so much data, the rate of data access can be extremely high. Within a large mainframe environment, it is common for hundreds of I/Os per second to be performed as data is accessed, searched, changed, and stored again.

Performance. Mainframe users have definite performance expectations. Response time to users of online applications must be kept to a few seconds for all but the most complex transactions. Overnight batch production schedules must be met in order to start the online systems for the next business day.

Based on business needs, the priorities, and therefore the response time requirement, for applications vary. The hardware system needs sufficient CPU and I/O resources for the size of the workload, and the operating system must have the capability to manage a complex, heterogeneous workload.

Operation. Typically, mainframe installations run a three-shift operation. Online systems run for one-and-a-half to two shifts, and the rest of the day is used for batch processing. Batch processing consists of report writing, reconciliation or clearing, and backing up data.

Availability. Availability is a critical requirement for mainframe applications. Because their execution is directly related to the bottom line, when they're not executing, the company isn't doing business.

To achieve high availability, mainframe operating system software is designed to recover from a broad range of hardware, software, operational, and environmental failures. The mainframe mentality is to do whatever is necessary to keep the system going: Applications must remain available to online users, and no work or data can be lost, even if that means maintaining redundant systems.

RAS. Mainframe availability is achieved through reliability and serviceability functions, and mainframe vendors are quick to emphasize the RAS (reliability, availability, and serviceability) features of their systems. Good reliability means long mean time between failure (MBTF). Serviceability refers to the repair capabilities of a system once it has failed. Good serviceability means short mean times to repair (MTTR).

Users want systems with good RAS, that is, systems that stay up by maximizing MTBF and minimizing MTTR.

Reliability is largely a hardware technology and manufacturing process issue. Serviceability is more interesting. It relates to the capabilities of systems to minimize MTTR from hardware, software, operations, and environmental errors and failures from within. Serviceability strategy involves anticipating all possible error conditions, collecting data about them when they occur, and providing recovery routines that are automatically invoked to restore an acceptable processing environment.

System Management

System management requires a set of software tools and management processes intended to provide IS personnel with the means to run their systems to meet availability and performance requirements. Areas of mainframe processing addressed by system management include performance, chargeback, and operations.

PERFORMANCE MANAGEMENT. Performance management tools abound on mainframe systems. Performance is measured and monitored in real time, and performance data is tracked historically. Systems must be adaptable to installation uniqueness and workload dynamics, allowing I/O configuration to be easily changed. Also, the operating system scheduler should be tunable to meet specific workload requirements.

CHARGEBACK. Accounting and chargeback are important in mainframe applications. IS organizations provide a computing service to multiple departments in the corporation. Systems are usually procured based on their justified usage by several corporate departments. Accounting and chargeback schemes ensure that each department gets its fair share of computing resource, that each corporate entity gets its fair share of IS support services, that the cost of the system is correctly amortized, and that future systems are procured in a timely and cost effective manner.

OPERATIONS. Operations requires the following:

- Monitoring the status of system hardware components for utilization and for error conditions. Thresholds can be set, and exceeding thresholds can result in the automatic execution of rules-based procedures and commands.

- Automating the operator console function.
- Managing batch production workloads by job schedulers.

Establishing a Baseline

The capabilities of the various open systems that are available and the positioning of those systems by their vendors are extremely varied. Not all open systems are appropriate for downsizing. We'll attempt to establish a minimum standard of system capabilities and characteristics to identify those which are.

As a minimum, today's open systems that can be considered for downsizing applications can be generally characterized as being Unix-based, supporting open systems interface standards, configurable as multiprocessing systems, and supporting some common enhancements to Unix disk management. Such systems are very well suited to a variety of applications, of which downsizing is but one.

OPEN SYSTEMS. The use of standards-based hardware and software interfaces in open systems makes them less costly to manufacture and support than proprietary systems. The use of RISC microprocessors in their design further contributes to their lower cost. This cost difference is the basis for our interest in them as alternatives to mainframes.

We'll be somewhat flexible and practical in defining *open*. An open system in this context is built around commonly used processors and I/O devices. However, it is the emphasis on some level of compliance with commonly accepted de facto and de jure standards that makes these systems open. Although all of these systems are Unix-based, in the near future, non-Unix systems will fit this description as well.

SYMMETRIC MULTIPROCESSING (SMP) SYSTEMS. SMP allows systems to scale smoothly through a wide range of processing capacity and power to a high end in the hundreds of MIPS. Binary compatibility across a vendor's entire range of systems is now assumed.

It is important to IS management, after one application is downsized to an open system, that growth in the usage of that application or in the number of applications run on that system can be accommodated without disruption.

There are many uniprocessor systems that could accommodate a single mainframe application, but further growth can be accommodated only with new technology or with additional systems. Mainframe application workloads are not easily split across multiple systems. Room for growth in a single, centralized system environment must be provided, or open systems become a short-term solution not worth the initial effort.

DATA STORAGE. In parallel with the movement toward SMP, there has been considerable activity in the area of data storage. The Unix file system has been enhanced or even replaced. Many vendors offer or plan to offer capacity, performance, and availability functions such as disk-mirroring, disk-striping, disk-spanning, and disk arrays. These capabilities are assumed in the downsizing baseline. At the same time, standards for distributed file systems that allow data to be accessed across a variety of vendors' systems are a key component of open systems.

Adding Value

Many systems exist that meet our minimum downsizing criteria. Systems that go beyond the baseline to address downsizing-specific requirements are more interesting. We've selected five as examples.

The Requirements for Downsizing

DOWNSIZING IS A NEW OPEN SYSTEMS OFFERING. Positioning SMP open systems for downsizing is a new development. The requirements for downsizing have been addressed by a few vendors who are just now bringing products to market. Therefore, we are dealing here with product plans, systems designs, and the results of early experience.

FIVE EXAMPLES OF DOWNSIZING SYSTEMS. All of the following vendors offer SMP systems supporting open system standards with the enhancements described above. All have invested in the development of downsizing-specific features and functions. Their approaches and investments have some similarities and many differences. Listed alphabetically, the vendors and their products are:

- Control Data Corporation 4680 Information Server
- Hewlett-Packard 9000 Series 870
- NCR Corporation System/3000 Model 3550
- Pyramid Technology MISserver S Series
- Sequent Symmetry 2000

SMP Vendors Targeting Downsizing

Control Data Corporation

Control Data Corporation (CDC) knows mainframe applications. While IBM has dominated the commercial mainframe market, CDC has been the traditional choice for engineering and scientific mainframe applications. The company has had more than 20 years of experience designing systems for large numbers of users that required high compute performance and the management of large amounts of data. They also have a fair level of experience with symmetric multiprocessing and parallel processing.

CDC has used all its experience in the design of the 4680 Information Server, which is targeted for downsizing data-intensive applications. The product was first introduced in January 1990, and the company increased its multiprocessor to four-way in October 1991. The 4680 is among the few systems available today that use the MIPS R6000 processor.

Hewlett-Packard

Most people associate Hewlett-Packard (HP) with high-performance workstations, manufacturing systems, and laser printers. It has traditionally held a small position in several niche markets with its proprietary MPE minicomputer. Certainly, the company has enjoyed success in all those areas. After a somewhat painful migration to its PA RISC architecture in the late 1980s, HP has been quite effective offering both Unix and proprietary systems for commercial processing. In December 1991, the company expanded its 9000 Model Series 800 servers with the high-end Series 870. HP seems to understand the requirement to create a mainframe systems environment within its Unix-based systems in order to enable downsizing.

The HP 9000 Model 870 Series is composed of four models ranging from a uniprocessor to four-way SMP. The systems are based on HP's PA RISC architecture currently running at 50MHz.

HP has developed a strong and comprehensive downsizing program. The Mainframe Alternative Program provides five downsizing alternatives:

- Provide identical applications
- Convert
- Replace
- Rewrite
- Coexist

HP has built partnerships with software companies and systems integrators to implement these programs. The relationship with Computer Associates (CA) will provide all of CA's IBM MVS tools on the HP/UX operating system, and, later, on OSF/1.

NCR

NCR has been offering Unix-based, multiuser systems since 1981. The company recently shipped its 100,000th Motorola 680X0-based Tower system. As one of the old BUNCH, NCR has also been a mainframe vendor. This experience shows in the System/3000, which combines excellent price/performance with the systems capacities and sophistication required for mainframe applications.

The Model 3550 is the most powerful, tightly coupled system in the System/3000 family. Built with Intel 80486 processors running at 50 MHz, the system is aimed at large, multiuser applications. The system has only recently become generally available, after completing Beta testing.

Pyramid Technology

Pyramid seems to understand both mainframe and downsizing requirements. It has even named its product the Corporate MIServer. Pyramid has been offering Unix-based, multiuser systems to the commercial market for 10 years, and SMP systems for 6 years. The company was among the first to recognize the potential in downsizing, and has been developing features to address its requirements and marketing them to mainframe users for several years. Pyramid is using this product experience and marketing success to continue to develop its product and to penetrate the mainframe market.

MIServer systems offer between two- and twelve-way symmetric multiprocessing using 33 MHz, MIPS R3000 processors. The DC OSx operating system offered with the MIServers has been adapted to support mainframe-style application processing.

Sequent Computer

Sequent was an early pioneer in offering symmetric multiprocessing Unix systems. The company seems to understand the requirement for creating a mainframe systems environment within its Unix-based systems in order to enable downsizing. Recent announcements include a "mainframe level" backup capability, improvements to disk availability, and a transaction monitor.

The Symmetry 2000 is Sequent's latest family of Symmetric Multiprocessing systems. The product line was announced in January 1991 and features systems built around the 50 MHz Intel 80486 ranging up to 30-way multiprocessing.

Downsizing with Sequent is distinguished by its relational database management capabilities. The appeal of relational database management is combined with the hardware price/performance and SMP experience Sequent has. The company is partnering with companies like Novell and Oracle to facilitate the implementation of applications on its Symmetry systems.

SMP Vendors Targeting Downsizing

Within the open SMP offerings of these vendors, there are similarities and contrasts in capabilities and in approach. CDC and NCR are longtime computer manufacturers. Over the years, both companies have offered a variety of systems for a range of applications, and they have survived many computer industry cycles. Pyramid and Sequent have always been only Unix system vendors. CDC, HP, and Pyramid all use RISC microprocessors, while NCR and Sequent use Intel CISC microprocessors.

Whither IBM?

IBM is in a tough position in the age of downsizing. After all, it all but owns the mainframe industry. IBM does not view downsizing to Unix-based systems very positively because of the potential for erosion of the mainframe installed base. Even then, if IBM customers were to insist on downsizing, they would likely be steered toward the IBM AS/400 as the right solution.

On the other hand, short of dramatic price cuts on mainframe systems, there is nothing IBM can do to prevent customers from downsizing. Although the company does not currently offer an SMP version of its RS/6000 system, the RS/6000 could be an excellent downsizing platform, and IBM certainly understands mainframe application requirements. Until now, the RS/6000 has been positioned for engineering and scientific computing, although over half of RS/6000 customers actually use it as a multiuser system.

We expect IBM to position the RS/600 for downsizing later in 1992. The downsizing market is predominantly composed of longtime IBM mainframe users, and the easiest sale is a sale to someone who is already doing business with you. An aggressive and agile IBM could capture a fair share of downsizing—but is IBM really agile and aggressive?

What about Digital?

Digital does not currently offer a system that meets our baseline downsizing criteria. The company claims some success with the new VAX 6000 as a downsizing platform, but it does not offer mainframe-class performance, although it does have price/performance advantages over IBM mainframes.

Digital's Alpha technology, planned for future systems, would seem to offer high-performance open systems capabilities. Combining Alpha with Digital's longtime VAX experience and its impressive open systems distributed computing capabilities should result in a formidable offering. The first systems based on Alpha RISC technology are expected to be announced in time for DECworld in April.

Downsizing Experience

Enough companies have had experience downsizing to allow us to draw some preliminary conclusions about the positive and negative aspects of moving to SMP open systems. Overall, downsizing is very feasible, but considerable work must be done in the areas of robustness and system management, and more hardware is required than anticipated to achieve desired levels of performance.

The Good News

The good news about downsizing is:

- Performance of online applications on SMP meets the requirements of mainframe users.
- The capacity of SMP open systems is sufficient to support large numbers of users accessing multiuser applications which, in turn, access large amounts of data.

- SMP systems have, so far, been a cost-effective alternative to mainframes for mainframe applications.

In essence, this good news is a proof of concept. Open systems are viable platforms for downsizing mainframe applications in terms of function, performance, and cost.

The Bad News

However, the bad news about downsizing is:

- Few tools are available to convert the applications and data. Conversion has largely been a reprogramming effort.
- Many applications that run on SMP open systems require mainframe coexistence and access to data that is managed on the mainframe. Data are extracted and downloaded nightly, a lengthy and tedious process. Coexistence takes a lot of work.
- Database application tuning and problem-solving on SMP systems is difficult because serviceability tools don't yet exist. Suddenly, it's 1974 again. Broken database chains and pointers have to be found and fixed by hand. More often than not, the solution to a problem is a database restore.
- Systems break too frequently. They seem very sensitive to load and aren't as robust as mainframe systems. Too often, the solution is a re-boot and a database restore, resulting in lost work at best, and lost data at worst.
- The system isn't managed like a mainframe; discipline is sometimes lacking. For example, in one company, a new version of the database system was installed prior to the start of the online day without testing or planning. It didn't work. After considerable online debugging, the old version was reinstalled. Online users spent a few hours without access to their application.
- Hardware resources required to support the workload, such as processor power and disk capacity, have been greater than anticipated.

The Good News about the Bad News

Many of the "bad news" items described above come about as a result of a characteristic response to system problems:

"Re-boot and try it again. (Maybe it will go away.)"

That response indicates inexperience with production workloads, operating systems that are more easily and quickly refreshed than repaired, and a lack of tools to expedite problem-solving.

Using mainframe terminology, the bad news refers to missing serviceability features and some system management tools and disciplines: problem management, change control, and performance management.

The best news about the bad news is that open systems vendors are addressing these areas within their products. Serviceability and system management have been implemented in a few systems and have been planned for many others.

Open SMP Approaches to Serviceability

The Importance of Serviceability

Inevitably, systems go down. They crash from hardware problems, software problems, operational problems, and from acts of God. System outages are expensive. During the time that a system is down, business functions are down and the users of the system are idle. The time to diagnose and repair the condition causing the crash must be kept to a minimum.

Among CDC, HP, NCR, Pyramid, and Sequent, NCR is the one that, far and away, provides the most comprehensive serviceability capabilities.

NCR

NCR has invested significantly in building mainframe class serviceability into the System/3000 Model 3550. Reliability, availability, and serviceability—RAS—are even included in the name of the operating system: Unix SVR4 MP-RAS. The system's features are:

- Environmental conditions such as temperature are monitored. An alert system notifies the operator when thresholds are exceeded, allowing corrective action prior to system or component failure.
- Lazy panics have been reduced in the Unix operating system. A panic is a condition from which the operating system cannot recover. The so-called lazy panics are conditions from which recovery was possible with some operating system modification. NCR has done those modifications.
- At boot time, components that fail diagnostics are automatically made unavailable.
- The Journaling File System provides fast recovery from system outages and minimizes the amount of lost work by logging changes to disk files. In the event of a failure, the last backup is combined with the log of changes to restore the disk.
- An optional Power Backup System keeps the system up through power failures of 15 seconds or less.
- A dedicated diagnostic processor is built into the system.
- Logging and analysis are performed for recoverable hardware errors. The log can be used to identify system components that require maintenance or repair.
- Remote support capability allows NCR support personnel to access the system to isolate problems.

CDC

CDC offers a mainframe approach to serviceability that is not quite as comprehensive as NCR's. The 4680 Information Server includes the following:

- Logging and analysis are performed for recoverable hardware errors. The log can be used to identify system components that require maintenance or repair.
- In what CDC describes as "hardening the kernel," Unix panics have been reduced.
- Panic dumps are automatically analyzed to extract the error log data for hardware errors. This capability shortens the time required to diagnose the cause of the panic.

- Diagnostics can be run online. Maintenance and repair actions can occur in parallel with production processing.

HP

HP also offers a mainframe approach to serviceability. The 9000 Series 870 includes the following:

- Logging and analysis are performed for recoverable hardware, software and environmental errors. The log can be used to identify system components that require maintenance or repair.
- HP has also "hardened the kernel." Unix panics have been reduced in the HP/UX operating system.
- The system supports battery backup for continued processing across power outages of 15 minutes or less.
- The system will automatically restart following a power failure.
- Diagnostics can be run online. Maintenance and repair actions can occur in parallel with production processing.

Pyramid

Pyramid offers good, mainframe level serviceability. The MIServer S Series includes the following:

- At boot time, components that fail diagnostics, including processors, are made unavailable.
- Logging and analysis are performed for hardware errors. The error log can be used to identify components that require maintenance or repair and the nature of the action required to restore proper operation.
- Diagnostics can be run online.
- Remote maintenance can be performed by Pyramid service personnel. The error log can be analyzed to assist in determining the cause of problems.
- Pyramid has also "hardened the kernel." Unix panics have been reduced.

Sequent

Fewer serviceability capabilities are offered by Sequent on the Symmetry 2000 systems than on the systems of the other sample vendors. The system includes the following:

- The Symmetry Series is equipped with a dedicated diagnostic processor. Diagnostics can be run in parallel with online operation.
- The diagnostic processor may be activated via modem, allowing Sequent service personnel to perform problem diagnosis remotely.
- Sequent has reduced panics in its Unix operating system.

Conclusions About Serviceability

Some impressive serviceability functions are provided. Most mainframe-like are the remote support capabilities, the options for power backup and processing through power failures, and the Unix panic reductions.

Software serviceability should be emphasized. The system should become more and more automatically self-repairing. Panic reduction is a great start, and efforts there should continue. Future efforts should push this error recovery back through transaction monitors and database management systems to individual applications.

Error logging is another serviceability area deserving continued attention. HP's use of the log for software and environmental errors as well as for hardware errors sets a good example. As more events are logged and the volume of data increases, automated log analysis tools will speed problem isolation and determine repair actions.

The SMP Open Systems Approach to System Management

System Management presents a challenge to open systems vendors. IS management wants the same capabilities in open systems as it has today in mainframe systems, capabilities like performance management, automated operations, and chargeback. This is not an unreasonable request, but systems management is a 20-year-old discipline on mainframe systems, and open systems are in their infancy in addressing mainframe-style applications. Open systems vendors are not in a position to deliver mainframe-style system management today. The questions to be answered are clear: What will open systems vendors provide, and when will they provide it?

CDC

CDC has added one-and-a-half million lines of code to its operating system in its efforts to implement what it calls "production capabilities." These capabilities include:

- User and project chargeback.
- Tape handling that includes 3480 support and support for ANSI-standard tape labels.
- Mainframe style error logging.
- Automated backup function.
- Data archive and retrieval.
- Significant workload management capabilities. All capabilities make mainframe-style processing more efficient and responsive.
- A batch job management system which provides for time-scheduled processing and manages batch tape allocation.
- Swapping, which minimizes high paging activity common with large numbers of users.
- Real-time process scheduling functions, which include fixing process priorities, process binding to one or a group of processors, and affinity scheduling — scheduling a process on a processor if the processor is available when the process is. These scheduling functions try to gain a performance advantage for a given process and a throughput advantage for the system. The strategy is to schedule a process on a processor that may contain data relating to that process in its cache because they were previously scheduled on that processor.

The data management capabilities of the EP/IX operating system are significant. All simplify the operational aspects of manual activities prone to human (operator) failure. Tape labeling prevents using the wrong tape, so support of ANSI tape labels is especially important. The number of tapes required to support mainframe database applications can be very large, and the order in which the tapes are mounted is critical to data integrity.

One gets the impression from speaking with CDC developers and product managers that significant efforts have been made to design mainframe system management capabilities into the 4680. All the features described above are available today.

HP

HP offers good system management capabilities. The company's relationship with Computer Associates has improved system management function significantly for downsizing IBM MVS applications. These functions include:

- Significant workload management improvements to the operating system.
- A batch job management system.
- Open/Spool, which provides network-wide spooling capabilities.
- Real-time process-scheduling functions, which include fixed priorities, process binding, process exclusion, and preemption control.
- Software tools for system performance monitoring and capacity planning. An application performance monitor is planned.
- A high-volume, automatic disk-backup system.
- As a result of HP's relationship with Computer Associates, all the system management tools offered by CA in the IBM MVS mainframe environment are available under HP/UX and will carried over to OSF/1.

NCR

NCR has developed and packaged the Open Systems Administrator (OSA) to provide systems management capabilities across the System 3000 family. The functions of OSA are comprehensive and quite impressive. These functions are:

- Backup and restore facilities.
- Utilities for real-time performance tuning.
- The Hardware Manager, which displays the hardware available with the system and the current status of each device.
- The Batch Manager, which creates and manages local and remote batch jobs.
- The RAID Manager and the Volume Manager, which support disk arrays, virtual disks, mirroring, striping, and spanning.
- Although performance management is not an OSA capability, it is a system management function that NCR plans to implement with workload management functions within its operating system. Capabilities include process-binding, gang-scheduling, and affinity-scheduling.

The SMP Open Systems Approach to System Management

At this writing, OSA is in the late stages of Beta testing; general availability is scheduled for 1Q92. Missing from OSA is chargeback functionality. NCR has stated that chargeback will be implemented—sourced from a third party—but that there is no schedule for its implementation.

Pyramid

Pyramid Technology has recognized the importance of system management. The company has been providing tools under its DC OSx Unix system which are familiar to mainframe users. These tools are:

- 3480 tape support.
- Significant workload management improvements to the operating system.
- A job management system that provides job scheduling, management of batch jobs, and spooling.
- Real-time process scheduling functions that include process-binding, process exclusion, preemption control.
- A performance monitor that collects and displays data on the usage of CPUs, memory, and I/O.

Sequent

Sequent, too, understands the importance of robust system management. However, current capabilities are described by the company as "rudimentary." Sequent recognizes the requirement to improve and has begun to do so.

In October 1991, the company announced ptx/Backup, ptx/SVM, and ptx/transACTION, all aimed at enhancing Sequent's support for commercial environments. ptx/Backup is a "mainframe level" backup system that simplifies the administration and reduces backup time.

Sequent would like to find third parties to implement performance management tools within its Dynix operating system. The company is talking to software developers who have built the tools for MVS, and who have the expertise and experience necessary to build the tools effectively for Unix.

Conclusions About System Management

There are two key points to be made about systems management in open systems:

1. Open systems vendors all recognize the importance of systems management to mainframe applications processing and to IS management. Vendors have invested and plan to continue investing significant resources to provide system management capabilities.
2. We are only in the development phase of system management for open systems. Many offerings are not generally available. What is available is not yet complete. Many more offerings are planned. We're off to a great start, and we're certainly heading in the right direction.

Performance Is More than MIPS

System balance, raw MIPS, raw CPU power are not enough to give good performance to mainframe applications. More important than the processing power of SMP open systems is balancing the performance of these systems among their CPU, memory, I/O, and bus

components. It's essential to organize systems within which work flows smoothly, avoiding potential bottlenecks.

Good mainframe application performance needs both CPU and I/O resources and a system that balances them. Given a number of high performance CPUs, performance balance can be achieved through caching, large memory, wide I/O buses, fast I/O devices, distributed I/O processing, asynchronous I/O processing, and combinations of these approaches.

Table 1 lists the performance-related specifications of CDC, HP, NCR, Pyramid, and Sequent systems.

Without benchmarks, it's difficult to draw any hard conclusions about performance balance from these specifications, but, because the range of capabilities is so wide, some inferences are obvious. For example:

- Both HP and Pyramid have large caches. The larger the cache, the less the requirement to access memory and I/O, the less the bus traffic. On the other hand, with many processes accessing wide areas of memory and making many I/O requests, cache data could be frequently invalidated.

- Sequent supports the largest number of processors and implements a single bus architecture that has the smallest system bus bandwidth. The smaller the bus bandwidth, the greater the effect on performance of memory and I/O access.

- System organization varies from one to three buses. So?

- System bandwidth varies from 80 MBps to 400 MBps. Is one too big? Is the other too small?

While processor organization and CPU power of SMP systems are reasonably similar, system organizations and I/O capabilities are remarkably different. One should conclude that there are workloads which could yield very different performances if run on each of our sample systems. Unfortunately, we don't have these workloads available. With them, we could produce benchmarks and compare the systems. Without them, we're left to deal with specifications. However some theoretical comparisons are possible.

A PROBLEM IN COMPARING SYSTEMS

"Theoretical" is really an appropriate term in the comparison of different systems. While processor types, disk interface, memory technology, and the like are similar, the means of interconnecting these components into a system differ significantly among vendors. Bus structures and protocols are proprietary. Device drivers are, too. On the other hand, these differences are not so great as to preclude any comparison, and, as shown in Table 2 below, can at least lead us toward some conclusions. The ratio of the different bus bandwidths to the maximum number of processors that can be configured provides a measure of theoretical bandwidth per processor.

Processors use the system bus for memory references not satisfied from cache. Memory references can be minimized by large caches or by workloads with tight processing kernels and little I/O, e.g. compute-intensive, scientific processing. Commercial, data-intensive mainframe applications make frequent and wide memory reference and do lots of I/O. High system bandwidth per processor is important in supporting mainframe workloads.

Performance Is More than MIPS

Comparative System Specifications

Vendor	CDC	HP	NCR	Pyramid	Sequent
Product	4680	9000 Series 870	3550	MIS-12S	Symmetry 2000
Number of processors	1-4	1-4	2-8	1-12	2-30
Processor used	MIPS R6000	PA-RISC	Intel 80486	MIPS R3000	Intel 80486
Processor speed	60 MHz	50 MHz	50 MHz	33 MHz	50 MHz
MIPS per processor	56	53	40	25	40
System organization	system bus I/O bus	system bus I/O bus	system bus I/O bus	system bus memory bus I/O bus	system bus
Cache size (KB/-processor)					
Primary	80	1024	256	128	512
Secondary	512		4096		
Memory interleave	8-way	4-way	4-way	4-way	4-way
System bandwidth (MB/sec)	240	220	400	80	80
I/O bus bandwidth (MB/sec)	180	40	160	40	42
Channels	24	2	16	16	14
Controllers (max)	24	12	16	16	7
Disk drives/ controller (max)	1	5	7	4	8

Table 1

System bus bandwidth for CDC, HP, and NCR appears more than adequate. In fact, there appears to be sufficient bandwidth to support either greater levels of multiprocessing and/or faster processors.

Pyramid and Sequent seem to be lacking in bandwidth when configured for the maximum number of processors. Note that both these systems are mid-1980s designs, while the others are 1990 systems. It could be that some updating is needed.

Ratio of Bus Bandwidth to the Maximum Number of Processors

CDC	240 : 4 = 60
HP	220 : 4 = 55
NCR	400 : 8 = 50
Pyramid	80 : 12 = 6.67
Sequent	80 : 30 = 2.67

Table 2

The system bus bandwidth-per-processor ratio for the Sequent Symmetry system is misleading. While most implementations use the system bus for the flow of both data and interprocessor control, Sequent provides a separate bus for each function. Its System Link Inter-Connect (SLIC) bus carries control information. So, while the ratio for Sequent appears low, it is understated; the ratios for the other systems are overstated. Unfortunately, it's impossible to quantify the differences.

Supporting Mainframe-Class I/O

I/O performance is dependent on many factors which are best described by queuing theory and statistics. While a description of the theory is beyond the scope of this article, some I/O performance guidelines are presented (without proof) along with some mainframe-class terminology.

- Devices attach to disk controllers attach to channels attach to the I/O bus(es).
- Buses are faster than channels are faster than controllers are faster than devices.
- Buses are fewer than channels are fewer than controllers are fewer than devices.
- Controllers transfer data at device transfer rates; channels transfer data at multiples of device transfer rates; I/O buses transfer data at multiples of device transfer rates. Device data transfer rates are typically 3 to 5 MBps.
- Increasing the number of controller and channel paths between a disk drive and the I/O bus increases performance. For I/O subsystems of similar component performance, better performance will be achieved in systems with greater numbers of channels per controller and greater numbers of controllers per device.
- During data transfer, both the controller and the device are tied up, and all other devices attached to the controller wait. Minimizing the number of devices attached to a controller minimizes wait time.
- Maximum I/O throughput is achieved by balancing activity equally across I/O subsystem components.

Given these guidelines, some observations regarding bottlenecks can be made about the I/O capabilities of the systems described above.

Performance Is More than MIPS

Performance bottlenecks:

- The I/O bus bandwidth on HP, Pyramid, and Sequent would be completely utilized by the concurrent data transfer to or from eight drives (for 5 MBps drives).
- The channel bandwidth on HP (20 MBps) can support data transfer for eight drives concurrently. Although the channels are balanced with the bus, more bandwidth is required.

Capacity bottlenecks:

- The Sequent system supports only seven controllers.
- The CDC system supports one device per controller.

Commercial, data-intensive mainframe applications perform many I/Os per second against large stores of data. An I/O subsystem with high bus and channel bandwidths and many I/O paths to many devices is required. NCR provides that sort of system today. CDC plans it. HP, Pyramid, and Sequent have some work to do.

Comment

This performance analysis leaves a lot to be desired. It provides no real way to compare systems, just a way to make observations and predictions about them. Minimally, it's shown that CPU performance as expressed in MIPS, MFLOPS, or SPECmarks tells only a small part of the performance story.

The TPC benchmarks can provide the metric to begin system comparisons. TPC-A is an online, transaction-oriented, database benchmark. TPC-B is batch oriented. Both provide system performance in transactions per second, but they are complex and time-consuming to set up and run. TPC results are available for only some of the configurations of some of the systems described here. The vendors are working on producing results, but users must demand them.

User Adoption and Migration

New Development, Not Conversion

Applications developed for proprietary mainframe architectures cannot easily be converted to run on SMP open systems. There are several reasons for this lack of portability.

- Many mainframe applications are system dependent. They use hardware and software features unique to the architecture on which they were built.
- The mainframe applications now in production were developed as long as 20 years ago. Not only might the source code be lost, the developer might even be retired!
- The transaction monitors under which the applications execute are extremely system dependent.
- Mainframe applications access hierarchical and network database structures. Such structures cannot be converted directly or automatically to the relational structures used on SMP open systems. Few tools are available even to do partial conversions.

A better approach to downsizing is to use SMP open systems as the target platform for new applications. Application development is a continuing process. There are always new functions to be automated and old applications to be rewritten.

Comfort, Familiarity in Application Development

Downsizing to SMP open systems has significant advantages in application development. Several important application characteristics remain the same:

- Applications are multiuser.
- Applications access a centralized, local database.
- Familiar programming languages like Cobol are supported on most open systems.
- Application development tools such as CASE and 4GLs used on the mainframe are available on open systems.

Coexistence with Existing Application

Downsizing will require data sharing between the SMP open systems and the mainframe. There will be interrelationships in the applications and data. Three approaches to data sharing are:

1. File transfer
2. Terminal emulation
3. Distributed processing

FILE TRANSFER. File transfer is the simplest approach to coexistence. No programming is required, but there is considerable replication of files. Every night, the changes made to the databases during online processing are extracted from one system and downloaded to the other. The next day, applications access the snapshot of the mainframe databases that contain changes made on the mainframe the previous day. This approach is suitable for reporting and query, but not for transactions and updates.

TERMINAL EMULATION. Using terminal emulation, users log onto applications on the "other" system when they need data controlled there. While no application programming is required, the users must be aware of where the information is and the commands and procedures required to get it. Terminal networks can be complicated. A connection must be defined from every terminal (or terminal server) to each system it will access. The terminal types and protocols used by the mainframe must be supported on the open system. For example, in an IBM MVS environment, the open system has to both support a 3270 data stream and provide the means for 3270 terminal sessions to attach to applications running on the open system. It can't be expected that the mainframe environment will or can be adapted to meet open systems requirements, so it is essential that the open system support the mainframe protocols.

DISTRIBUTED PROCESSING. Distributed processing is the most elegant and most resource-efficient solution to the coexistence requirement. It involves real-time access of data on one system from applications executing on another. Considerable programming is required, but no resources are duplicated.

Some mainframe applications will not be redeveloped, even in the long run. Some databases will stay forever on the mainframe. For these situations, distributed processing becomes the most attractive alternative, although either file transfer or terminal emulation addresses the necessary capability.

DCE. The Open Software Foundation's Distributed Computing Environment (OSF's DCE) greatly simplifies and standardizes distributed processing. Without DCE, very low-level data communications programming and complex process synchronization programming were required on an application-by-application basis to perform distributed processing. Using the DCE's RPC, synchronization services, and other services, distributed programming will be simplified. Most significantly, virtually all mainframe and open systems vendors have committed to implementing DCE.

Downsizing Relational Database Applications: A Very Attractive Opportunity

SMP open systems are the right platform for relational database applications. DB2 on IBM MVS mainframe systems will never provide an effective, cost- and resource-efficient environment for those applications.

As more RDBMS vendors compete the work of "parallelizing" their systems, relational database applications on SMP open systems appears to be the best approach to downsizing. DB2 development projects should provide likely candidates for being downsized, beginning with a pilot. The incremental costs are minimal since staffing is already in place (the DB2 development organization). Relational application development tools and systems exist on SMP open systems. The system and supporting software are the only incremental costs.

The DB2 Story

It's been more than eight years since IBM introduced a relational database system for its MVS system. DB2 was intended as a strategic offering. It was to be the growth path for IMS DB/DC systems and CICS systems using a variety of other database organizations. The bulk of mainframe application development would migrate to DB2. Relational databases would be easy to build, relational applications would be easy to develop. The additional resources required to run this essentially parallel system on IBM's big, serial, asymmetric multiprocessing mainframes would fuel the continued growth of IBM's mainframe business.

Installations have found DB2 hard to learn. Application developers with years of experience on hierarchical and network databases had difficulty thinking relationally. The DB2 learning curve has been steep. Development cycles have been long, and resource requirements for DB2 systems have been immense. Transaction response time was often measured in minutes. Although almost every IBM mainframe installation has licensed DB2, there are very few installations running DB2 for production work today. Some examples:

- A major financial services company planned to redevelop its entire system under DB2 using a CASE tool. Performance and resource requirements are such that each application consumes a 3090-600 (at the time the largest mainframe) system and provides only adequate performance to online users.

- In the wake of the 1989 acquisition of Cullinet Software by Computer Associates, users of Cullinet's IDMS DB/DC system, fearing loss of support, considered DB2 as an alternative. Almost three years later, most installations are still using IDMS. Few are running any DB2 applications.

An Opportunity

In spite of this DB2 experience, mainframe users recognize the benefits of relational systems and would like to implement relational applications. SMP open systems provide the right platform for the following reasons:

- The capabilities of open systems relational database management systems in the areas of application development tools, recovery, and data integrity now approach those of DB2. Products like Oracle and Sybase have reached commercial quality.

- SMP provides an excellent processing environment for relational applications. Relational systems view data in tables, and many table rows are accessed and manipulated concurrently, in parallel. (The traditional mainframe systems view data serially.) Such parallel processing is best done on parallel processors.

- Additional resources required for relational processing — more disk space and more memory, as compared with hierarchical and network data models — are much less expensive on SMP systems. Even processors can be added cost-effectively.

Approach Downsizing Opportunity

SMP systems supporting open systems standards provide an excellent platform for relational DBMS:

- Most downsizing development will be new development; few Cobol programs will be ported.

- What must be migrated should port pretty easily; the differences among different SQL dialects are less than the difference between SQL and Cobol.

- Mainframe resources can be reclaimed for application processing that can't be converted easily or that requires those resources.

- Open systems can be used where easy success can be experienced.

- Given the difficulty with DB2, much of the mainframe relational activity has involved pilot projects. Downsizing pilot projects is easily accomplished, and chances for success are very high.

The State of Downsizing Today

Downsizing represents a huge potential market for SMP systems vendors, as well as an opportunity for tremendous IS efficiency. Here's the state of downsizing today:

We're Getting Closer...

There is very little actual downsizing going on. Downsizing is a good idea, worthy of serious consideration, and serious consideration is all it's getting—not much action. Only the earliest of adopters are adopting.

From the viewpoint of IS management, serviceability and system management need considerable improvement. From the viewpoint of open SMP vendors, systems are only just being brought to market. For example, four-way multiprocessing on the CDC 4680 was announced and shipped in October of 1991. The HP 9000 Series 870 was announced in November, and its four-way system became available in December 1991. The NCR Model 3550 has just been released from Beta testing. They're attractive downsizing systems, but obviously, they are very new.

THE SERVICEABILITY ISSUE. SMP systems go down more often than users should have to tolerate. We're not talking about fault tolerance but system availability. Too frequently, the response to a variety of software, hardware, and environmental conditions is a system halt. Too frequently, the solution to abnormal conditions is a re-boot.

The State of Downsizing Today

SMP SYSTEMS LACK SYSTEM MANAGEMENT TOOLS AND DISCIPLINES. The level of system management required for mainframe applications is just beginning to be delivered with open systems. Currently, there is a lack of the necessary tools and disciplines.

Taking the Long Road

On the surface, there appear to be many systems suited to downsizing. They all deliver lots of MIPS. However, not all of them deliver total system performance — performance and capacity balanced among all system resources, which is what it takes to support mainframe applications in the long run.

New Development, Not Conversion

Mainframe applications are often written in arcane, proprietary languages; they access data controlled by proprietary DBMSs, and they execute under a proprietary operating system. Very little of the software that runs on mainframes has been ported to open systems yet. Conversion is not feasible, since tools don't exist, and costs are prohibitive because conversion represents a massive reprogramming effort. SMP systems should be used for new development. Long-term coexistence of applications and data on both the mainframe and SMP open systems will be required.

Great Relational Engines

Relational systems have been poor performers and resource hogs on mainframe systems. Installations have had a hard time getting relational systems like DB2 to perform at the level of traditional mainframe hierarchical and network database systems like IMS and IDMS. CPU, memory, and I/O resources required for a DB2 system are much greater than resources required for the traditional mainframe systems.

Open systems are tailor-made for RDBMS. They're cheap, fast, and they process in parallel. There are years of relational experience on them. It makes great sense to move relational processing to SMP open systems.

Next month's *Unix in the Office* will
compare S VR4. and OSF/1
For reprint information on articles appearing in this issue,
please contact Donald Baillargeon at (617) 742-5200, extension 117.

Open Systems: Analysis, Issues, & Opinions

DATABASE WATCH

OnLine 5.0 from Informix Software

This month, Informix Software will fire one of the first shots in the next round of product announcements from the independent relational database management system (RDBMS) vendors. Version 5.0 of Informix-OnLine will add significant base functionality on the database server side—stored procedures, declarative referential integrity, and full support for the SQL89 standard—plus optional enhancements in transaction processing support and storage of complex data.

REORGANIZATION. Another interesting development at Informix is the appointment of Chuck House as senior vice president of product management and product development. Formerly general manager of the Software Engineering Systems Division at Hewlett-Packard, House was in charge of developing HP's very successful Softbench CASE strategy and products. We have always been impressed with House's abilities, and he is already proving to be a valuable addition to Informix.

In December, House reorganized his marketing and development group, taking a business unit approach on both the server/connectivity and tools side. As a result, we expect to see from Informix a broader and more open perspective on working with competitors. For example, the product marketing folks on the server side will begin to view the Informix development tools more as a strong third-party offering, and to work proactively with other tools vendors as well.

On the tools side, Informix is already emphasizing partnerships in establishing an open development environment. An example is the company's OpenCase program, in which Informix expects to have 20 partners by now. Some momentum has been lost here due to HP's delay in shipping the Softbench source code on which OpenCase is built, and Informix plans to work on regaining its initiative here. We also would not be surprised to see Informix adapt its development tools to run against other DBMS engines in the future.

UPBEAT FINANCIALS. The third piece of news is Informix's return to profitability. Revenue for the third quarter, ended September 29, was $48.3 million, an increase of 27 percent over the same period a year ago. Net income was $5.1 million versus a loss of $4.3 million in Q3 1990.

Stored Procedures

A major plus for Informix is the implementation of stored procedures in OnLine 5.0. This is the first step toward offering an intelligent, or programmable, server. Stored procedures are the foundation on which mechanisms such as triggers and event alerters are built, so Informix is establishing the necessary architecture for introducing these features in the future. Stored procedures are also part of Informix-SE 5.0, Informix's server offering geared for smaller, less complex installations.

Stored procedures are those that reside on the server rather than in a client application. Stored procedures offer several major benefits. First, they can be shared by multiple applications, reducing the effort required to both develop and maintain applications. Stored procedures are also a way to achieve consistency in logic and processing, eliminating the need to depend on several developers to implement a function in the same way. Second, they reduce network traffic and improve performance in a distributed environment, since multiple SQL statements no longer have to be sent to the server one at a time. Third, stored procedures are precompiled and preoptimized, which also helps performance.

Informix-OnLine stored procedures can include any SQL statement plus statements to define variables, assign and compare values, and control program flow—if/then/else, for each, while, etc. The stored procedure can also make operating system calls. One advantage Informix-OnLine has over the Sybase SQL Server, which has had stored procedures since its introduction

in 1987, is the ability to define cursors within a stored procedure. (Sybase does plan to support cursors in Version 5.0 of SQL Server.) Informix stored procedures also have the following characteristics:

- Unlimited nesting
- Recursive routines
- Precompilation
- Preoptimization (Even if the objects referenced—e.g., tables or fields—don't exist, the procedure is preoptimized to the extent possible and then flagged as needing further optimization.)

Informix has added tables to the data dictionary for accessing information on existing stored procedures. The dictionary displays the code and can include a description of the procedure as well. If the developer puts a descriptive block in the procedure to describe what it does, the server automatically extracts this description when the procedure is compiled and adds it to the data dictionary.

Informix is also releasing new 5.0 versions of its embedded SQL precompilers for C, Cobol, Fortran, and Ada. Applications written with these products will be able to call stored procedures directly using the "execute stored procedure" command. Informix-4GL applications will also be able to take advantage of stored procedures, but this will require a two-step process initially. The next release of Informix-4GL will include direct support.

Enhanced SQL Support

Another new feature of Informix-OnLine 5.0 is full support for the ANSI SQL89 standard, level 2, plus Integrity Enhancement. The ANSI SQL Integrity Enhancement includes the following:

- Declarative referential integrity, which is the ability to define primary and foreign key relationships among the tables in a database and have the DBMS server automatically enforce these relationships. Informix still does not have triggers, but support for declarative referential integrity covers, at least partially, functionality for which triggers are often used. We expect to see triggers in Version 6.0 of OnLine (coming at the end of 1992).

- Default value for a column.

- Check constraints for a column, which are, essentially, field validation options that must be satisfied for every row in a table.

Full ANSI SQL89 support is also part of Informix-SE 5.0.

Online Transaction Processing

Informix is clearly committed to providing state-of-the-art capabilities and performance in the online transaction processing (OLTP) environment. The Informix OnLine database server has been a participant in published TPC-A and TPC-B benchmarks by several hardware vendors, and the company continues to enhance its status in this area.

TWO-PHASE COMMIT. Informix has had support for distributed *queries*, including a distributed query optimizer, for some time. To support distributed *transactions*, the Informix-OnLine 5.0 server now understands two-phase commit (2PC)—that is, a "prepare to commit" statement—and the optional Informix-Star 5.0 software can now automatically enforce 2PC across a network of Informix databases. Among the other independent Unix RDBMS vendors, Ingres, InterBase, and Progress already offer support for automatic 2PC. (Ingres requires the optional Ingres Star software; InterBase and Progress offer 2PC as an integral part of their server architecture.) Sybase has a programmatic 2PC, and Oracle plans to implement 2PC in Version 7.

XA SUPPORT. Informix is the first major RBDMS vendor to offer support for X/Open's XA interface to transaction processing (TP) monitors. Informix-TP/XA 5.0 is an optional add-on that allows Informix-OnLine 5.0 to take advantage of X/Open-compliant TP monitors, such as Unix Systems Labs' Tuxedo System/T, NCR's Top-End, and TransArc's Encina. Informix-TP/XA is a library of C language functions that implement the XA standard. In addition to an XA interface, an important component of TP monitor support is the two-phase commit protocol in the DBMS engine.

Multithreading

By mid 1992, Informix plans to introduce a multithreading version of its DBMS server architecture on Sequent as part of OnLine 5.0. This is an important development for Informix, and will enable the DBMS engine to make more efficient use of system resources

than the current server-per-user architecture does. Informix has developed its own threading facility and describes the capabilities of its planned multiclient/multithreading server as similar to Sybase's recently announced multiserver architecture. The server will be multithreaded on both uniprocessor and multiprocessor platforms. Multithreading will be available on all Informix OnLine platforms with Version 6.0.

Informix is also currently testing a multithreaded version of OnLine 5.0 for Novell's NetWare; this will roll out by mid 1992 as well. It will use the NetWare threading library.

Other Additions to OnLine

OPTICAL DISK SUPPORT. Informix-OnLine/Optical 5.0 is another option that supports third-party optical storage subsystems. Optical disk storage is becoming increasingly important for customers who want low-cost options for storing large volumes of unstructured data (binary large objects, or BLOBs), such as images.

SERVER ACCESS AND DOCUMENTATION. Another addition to both OnLine and SE 5.0 is the inclusion of the interactive schema and query editors as part of the database engine package rather than as part of the tools. Informix has also implemented an online, hypertext, graphical database administrator's guide. Viewing the manual requires FrameMaker or FrameViewer from Frame Technology Corporation.

Availability

Informix-OnLine 5.0 will begin shipping this month on three initial platforms—Sun, HP 9000, and AT&T. Availability on other platforms, including Digital's Ultrix, IBM's RS/6000, ICL, Pyramid, NCR, Sequent, and Unisys will follow.

Interoperability

SUPPORT FOR DRDA. Informix has stated its commitment to supporting IBM's DRDA, which will give Informix a much-needed read/write connection into the IBM mainframe DBMS environment. Access to DB2 data is particularly important. Although Informix already has a gateway to the IBM mainframe, it uses third-party software (Answer-DB from Sterling Software) and is read-only.

Unfortunately, Informix still lacks a connection to the Digital VAX/VMS environment.

Future Enhancements

In its next releases of OnLine, Informix plans to enhance support for very large (mainframe-size) databases, and to implement parallel query processing to speed up long queries on massive amounts of data. The latter will include parallel selects, joins, and sorts with enough smarts in the optimizer to recognize a query that can be parallelized. The focus here will be more on performance than on availability. We also expect to see triggers, a multiclient/multithreaded architecture on all platforms, further optimizer enhancements, remote site backup, and additional distributed database features.

WHAT ABOUT TOOLS? Informix Software's immediate concern after getting OnLine 5.0 out the door will be concentrating on the tools side and shaping up its development environment. There are plans for GUI support, presentation independence, and an object 4GL (a 4GL++, so to speak). We also expect Informix to move the WingZ Hyperscript language to the tools layer and to package and promote it as a powerful development language.

Conclusion

Informix Software is definitely back on track after a difficult time spent grappling with acquisition and financial problems. Its Unix market share is still strong, and the introduction of OnLine 5.0 is evidence that the company is on target technically. Informix is also well aware of the need for *timely* enhancements, not just good ones. Improvements in the tools arena and another major release on the server side by year-end 1992 should keep the company in an excellent competitive position. *—J. Davis*

NOVELL

Novell In the Spotlight

Novell is positioning itself as the Unix equivalent to Microsoft. This became evident when Novell invested in Unix Systems Laboratories (USL) last year, one of the earliest indications that Novell was moving from its traditional role as a LAN-based network operating sys-

tem company to an enterprise systems software vendor. Now Novell has taken two more dramatic steps. First, it has formed a joint venture with USL called Univel. Second, it has formed a strategic relationship with Hewlett-Packard to provide a native version of Novell for HP's RISC servers.

UNIVEL. Why the joint venture? Clearly, Microsoft has shown the industry that it wants to make the network an integral part of the operating system. Therefore, if users purchase the operating systems that Microsoft sells, they will get networking as part of the package. The result will be that they will begin to expect it. If users buy into Microsoft's version of reality, then a company like Novell could be in trouble. Why pay separately for a LAN operating system when it comes as part of the base operating system?

With the new organization, both USL and Novell will try to merge the best of their technologies. USL has been able to capture considerable interest among commercial information technology organizations that are deciding to downsize to Unix-based systems. Novell, on the other hand, is firmly entrenched in the departmental user community. The synergy between the two technologies could be powerful.

RELATIONSHIP WITH HP. Hewlett-Packard has expanded its relationship with Novell over the past few years since its relationship with Microsoft and LAN Manager soured. Clearly, HP understands that being able to offer its customers a native implementation of Novell's networking with its Precision RISC technology could be an attractive offering. Initially, Novell had tried to use its Portable NetWare as a general offering for heterogeneous platforms. However, Portable NetWare never offered the speed of the proprietary NetWare implementation. Novell's new tactic is to partner and thereby customize versions of its networked operating system for specific platforms, such as Hewlett-Packard's 9000 and 3000 platforms. HP and Novell also plan to provide interoperability between their individual software environments, for example, Novell's network management and HP's OpenView. HP will also be merging some of its NewWave technology into the NetWare environment. For example, HP is proposing that the NewWave Agent facility be used to automate many routine networking management tasks for NetWare. HP is beginning to learn that it may have better luck with marketing NewWave if it embeds some of its capabilities into the base system.

CONCLUSION. Novell has always been a sharp marketing company. Its management seems to be nimble enough to react to market shifts and to take advantage of emerging trends before they impact the bottom line. The trend towards corporate downsizing to both Unix and LANs positions Novell—in conjunction with its increasing web of relationships—to capture considerable market share. It will be interesting to watch as the attention of the LAN/operating system world begins to focus on its two titans: Novell and Microsoft.—*J. Hurwitz*

