Patricia Seybold Group

Editor-in-Chief
Michael A. Goulde

INSIDE

EDITORIAL

Page 2

***Cautious migration** from today's architectures to tomorrow's is the common wisdom. Perhaps a more radical change would be preferable. Although the short-term costs would be higher, over the long haul, the benefits might significantly outweigh the risks.*

ANALYSIS

Page 34

***Borland** is seeking to ensure that its stable of database products is not left out in the cold as users tie distributed databases together. Along with IBM, WordPerfect, and Novell, it has announced its Integrated Database API (IDAPI), which will add support for non-relational databases to support for the SQL Access Group interface and will be available on multiple platforms.*

OPEN INFORMATION SYSTEMS

Guide to Unix and Other Open Systems

Vol. 7, No. 12 • ISSN: 1058-4161 • December 1992

The Unix Data Center

Fact or Fiction?

By Michael A. Goulde

IN BRIEF: Whether a company's objective is to downsize applications or to introduce open systems into the data center, Unix is increasingly being considered for a role that, in the past, it has not played particularly well. Within the last two years, however, the robustness of Unix systems has improved, and third parties have been introducing data center management applications for Unix that make migration from a proprietary IBM MVS environment to Unix a much more realistic possibility. But there are still shortcomings, not the least of which is that many data center products for Unix are first generation.

Report begins on page 3.

Which Path to Progress?

Incremental vs. Revolutionary

SUPPORTING LEGACY SYSTEMS is generally included as a key consideration in plans for an open systems migration strategy. At issue are not the mainframe computers, nor the disk farms attached to them, nor even the licensed software that runs on those systems. The value contained in the legacy systems is the business analysis that has been captured in the application code and in the information which that code collects, accesses, and manages.

Supporting a legacy system involves maintaining the application code and the data as much as it does maintaining the physical system on which that system is running. Maintenance of the code, often Cobol or other equally inefficient languages, is an extremely costly process. But, without it, the entire system is in jeopardy.

The cost becomes increasingly less defensible as alternatives emerge that are faster, cheaper, and easier to support and maintain. New software development tools allow applications to be built in months instead of years. Existing applications become a living functional spec for a replacement system.

A number of different strategies have emerged for gradual migration of older applications to contemporary technology. Some of those strategies are discussed in this month's feature article. They often call for careful incremental change over what is in use today. Replacement systems are phased in, changing just a few variables at each step. This cautious approach is defended as a necessity for preserving access to the knowledge managed by the legacy application while the infrastructure is being slowly replaced.

Caution is both understandable and reprehensible. Understandable because senior managers don't rise through the ranks by being wild-eyed innovators or by taking risks. Reprehensible because the incremental approach to progress makes it difficult if not impossible to introduce a scope of change broad enough to fully capture the benefits of open systems.

Open systems, networked distributed computing, and object-oriented software are combining to force radical changes on platform architectures, network infrastructures, and software development tools and processes. These changes cannot be responded to with timid incremental adjustments in approach and technology. Distributed applications require different network infrastructures, network management, system management, software licensing and distribution, application development tools, end-user support systems, budgeting, maintenance, and more.

Even encapsulating legacy systems—that is, placing a shell around them that hides their underlying structure from contemporary applications—constrains designers from making the most of contemporary technology. They still have to deal with unique security requirements, access methods, and communications protocols that are not up to snuff for today's environments.

In a practical sense, introducing change on a revolutionary scale requires an analysis that shows not that there is zero risk, but that there is a high probability of achieving benefits that far outweigh any risk. In reality, there will be problems. But, if the gains are great enough, the problems will be tolerable. Incremental change doesn't occur without problems, either. Granted, the potential for problems is greater with a revolutionary shift, but that can be managed. The steps to revolutionary progress are taken toward a clearly defined goal, purposefully and with massive planning and preparation. The key is to not leave anything to chance.

OPEN INFORMATION SYSTEMS

Editor-in-Chief
Michael A. Goulde

MCI:
MGoulde
Internet:
mgoulde@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.
STANLEY S. DOLBERG

News Editor
DAVID S. MARSHAK

Art Director
LAURINDA P. O'CONNOR

Sales Director
PHYLLIS GUILIANO

Circulation Manager
DEBORAH A. HAY

Customer Service Manager
DONALD K. BAILLARGEON

Patricia Seybold 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

Open Information Systems (ISSN 0890-4685) is published monthly for $495 (US), $507 (Canada), and $519 (Foreign) per year by Patricia Seybold 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 *Open Information Systems*, 148 State Street, 7th Floor, Boston, MA 02109.

The Unix Data Center

Fact or Fiction?

Unix: A Step toward Open

Traditionally, the corporate data center has been marked by two characteristics—it is mainframe based, and those mainframes have the IBM logo on them. IBM's dominance of the data center defined the rules of competition: plug compatibility at price/performance better than IBM's. Whether the hardware was an IBM 3090 or an Amdahl machine didn't matter. The operating system was still MVS, and thus all applications and all support software ran in a consistent environment, no matter who supplied the hardware.

The Center of the Universe . . .

Until the 1970's, the data center environment was the center of the computing universe. Then came the minicomputer challenge, led by Digital Equipment Corporation's (DEC) VAX. In spite of the encroachment of these "mid-range" machines, the data center lived on, and "IBM connectivity" became a key requirement that DEC, as well as the rest of the minicomputer vendors, had to meet. The next challenge came in the 1980s with the arrival of the personal computer. Although the PC liberated users from dependency on the data center for some of their processing, users couldn't escape the fact that the vital data they required were still under control of the data center mainframe.

Until the late 1980s, real alternatives to the IBM mainframe or compatible data center were nonexistent. Overall system capacity and performance aside, the robust MVS operating system has historically been able to thwart efforts to displace it. Digital's attempts to create an alternative data center environment with clusters of VMS-based VAX computers have been largely unsuccessful, and PC LANs have lacked the capacity and the robustness to ensure data integrity and system reliability.

. . . Is No Longer at the Center

However, over the past 18 to 24 months, the invincibility of the IBM mainframe-based data center has begun to lessen, due to the confluence of five factors:

- The price/performance, as well as the absolute performance, of microprocessor-based alternatives to IBM mainframes has led cost-conscious customers to consider them as replacements, or at least supplements, instead of automatically accepting the annual or biannual mainframe upgrade.
- Application architectures are shifting from centralized, host-based environments to distributed and client/server environments.
- The downsizing craze, driven both by inexpensive microprocessor MIPS and dirt-cheap PC LAN MIPS, has led customers to reassess their true mainframe requirements.
- Vendors of Unix-based systems are supplying enhanced system software products that are far more robust than earlier versions, offering features, reliability, availability, and security that approach IBM standards at a fraction of the cost.
- Vendors of data center products ranging from utility programs to tape silos, which had previously only been available on MVS or VM systems, are porting those products to Unix systems.

Unix: A Step toward Open

What was a risk taken only by the lunatic fringe in the mid-1980s, i.e., building a data center on Unix, is now considered to be a relatively safe option. Or is it? Are we just being inundated with hype and hyperbole once again? Although the potential benefits from switching to Unix in the data center are large, ramifications have to be well understood before any moves are made.

Who Is Switching?

Three Perspectives on Unix Data Centers

Not all customers who are considering converting to Unix or building a data center around Unix-based systems are doing so for the same reasons. In some cases, the move to Unix is strictly aimed at driving down costs; in others, it is to meet strategic technology objectives; for still others, it is to meet tactical objectives. There appear to be three general categories of users who are interested in employing Unix in data center applications: mainframe users who want to reduce costs, mainframe users who have committed to open distributed computing, and the long-standing Unix user.

MAINFRAME USERS REDUCING COSTS. One group looking at Unix data center alternatives is the traditional "True Blue" mainframe customer, who is keenly interested in potential cost savings but not in giving up host-based, centralized computing or a coaxial network of thousands of 3270 devices. These customers want the benefits of Unix-based systems but with minimal changes to their environments.

In the past, it had been axiomatic that data center managers seldom had to cost-justify their periodic mainframe upgrades. In today's cost-conscious environment, managers who haven't carefully considered the alternatives to a mainframe upgrade, which can range from $500,000 to several million dollars, will probably need to begin considering employment options instead. Watching capacity utilization creep up 1 or 2 percent a month toward 100 percent is unnerving. Nevertheless, these customers have become cynical by comparing the cost of a unit of mainframe processing capacity versus the cost of similar capacity on many multi-microprocessor-based systems. And if that hasn't caught their attention, then the differential in software licensing fees between mainframe class systems and multiprocessor systems often does. When the differences in the cost of operations, license fees for layered software, cost of maintenance, and communications costs are added in, there is often no rational choice but to switch. However, these customers also recognize that the system cost is only a portion of the cost of switching. The cost of application conversion, staff training, and user training also has to be taken into account. Added to that is the potential risk of introducing radical change in an environment that has been running smoothly and, more importantly, predictably for many years.

CUSTOMERS MOVING TO OPEN DISTRIBUTED COMPUTING. Once a commitment is made at a strategic level to migrate from centralized, host-based processing to distributed processing, Unix in the data center becomes a key step toward eventually implementing a fully distributed computing environment. It is part of a long-range process of converting and rebuilding a proprietary infrastructure into one based on open systems.

This group of users has decided to migrate both applications and data from their mainframes onto networks of powerful, Unix-based symmetric multiprocessor (SMP) systems. These customers have long-term plans to re-engineer their applications using distributed, object-oriented tools and to adopt open frameworks for system and network management. However, they also realize that they cannot achieve their goals overnight or in one step. So they begin their migration by moving applications from an IBM or compatible host to Unix hosts to start the slow process of developing distributed applications and distributing the data along with the applications. By moving to Unix, perhaps with Unix versions of the same databases and applications, they can focus on getting system and network management issues resolved without having to deal with wholesale application changes. The move to

Who Is Switching?

Unix in the data center becomes the first of a great many steps down a long migration path. Once that step has been taken successfully, the move toward implementing a fully distributed environment can begin in a controlled, coherent fashion.

LONG-TIME UNIX USERS. The proponents of distributed processing who have been using Unix for many years to support a broad range of applications now want to extend Unix into the data center. This group discovered a need for more rigorous systems management and tighter control over data security and integrity for mission-critical applications than traditional Unix provided, but they do not want to sacrifice Unix to achieve it.

This group of users has been using Unix in technical applications in engineering and manufacturing departments. In the case of many telecommunications companies, Unix has been at the core of the business for many years. In manufacturing, extensive networks of Unix workstations and servers have been deployed, and second-generation distributed applications are being installed and used. The level of knowledge about, as well as the comfort with, Unix has led these fans to decide to run their order entry, accounting, financial, and customer service applications on the same operating system they have been using to design and build their products and to run "The Network." To them, putting Unix to work in the corporate data center is simply a logical extension to what they have been doing all along. Aside from consistency, the cost savings both in systems and in maintenance and management more than offset the cost of migration. Now, the same personnel who have been supporting Unix all along can also support the data center.

These categories cut across industries and across company size. There are large companies, like many of the regional telephone companies, that have been long-time Unix users, and cost-conscious businesses of all sizes that have been loyal mainframe customers.

Unix Data Center Strategies

Although the users above who are considering Unix for their data centers tend to fall into one of the three categories, there are actually far more options for putting Unix to work in the data center. That these approaches have different tactical and strategic implications is evidence of the flexibility that Unix has to offer to the data center and to large-scale applications. Unix can play a role in a number of different strategies for migrating away from a dependency on the mainframe. Unix systems can coexist and interoperate with applications running in IBM mainframe environments, they can replace IBM mainframe-based data centers with products that have comparable functionality but very different architectures, or they can support a completely different application architecture, such as one which is fully distributed. The richness of possibilities for employing Unix as an alternative in the data center is reflected in the number of buzzwords the industry has created to describe them all, including downsizing, rightsizing, smartsizing, client/server, distributed computing, surround, mainframe alternatives, and alternative mainframes. From the data center perspective, however, Unix is either supplementing, substituting, or succeeding the traditional mainframe.

Mainframe Application Processing Offload

Mainframes are designed primarily as application processors. The large memory spaces and multiple partitions allow large numbers of users to be supported with applications that have large data sets and provide the users with fast response times. In the normal course of business, new applications are developed to run against these data sets, and the data sets themselves grow and more users are added. As a result, excess capacity gradually erodes, and, in the face of the danger of running out of capacity, an upgrade has to be considered. Mainframe upgrades are a frequent and expensive proposition. Depending on the starting point and the ending point, the cost of an upgrade can range up to several million dollars. If capacity requirement estimates are off, then the danger exists of budgeting either way too much or way too little.

SELECTING APPLICATIONS. In some circumstances, upgrades can be avoided by pulling other applications off the mainframe and running them on more cost-effective platforms. The first to go are applications classified as personal productivity, easily placed on inexpensive personal computer LANs and still tied to the mainframe through terminal emulation. Next to go are applications which provide generic services that can easily be transferred to other platforms, such as scheduling and electronic mail. A surprising amount of capacity can be freed up in this manner, although it still leaves open the possibility in the future of having to upgrade the mainframe platform on which the strategic applications, such as a reservations system or order processing system, are running.

MANAGING BUDGET GROWTH. Expanding capacity using relatively low-cost Unix systems has an added benefit. The risk of budgeting too much or too little is minimized because an incremental unit of increased capacity costs so much less for a Unix system than for a mainframe. As a result, steadily increasing capacity requirements can be matched much more closely with increased capacity than is possible with the typical mainframe upgrade.

BEWARE HIDDEN COSTS. Moving existing applications carries the potential risk of high costs in converting production applications to new versions running in new software environments by personnel who may be unfamiliar with the new environment. Often the costs involved can easily wipe out the initial savings in hardware acquisition.

Surround: Complementing Mainframe Applications

"Surround" is a strategy of building new applications on Unix systems, as opposed to migrating or offloading applications. Instead of introducing these new applications onto the mainframe, the surround strategy calls for bringing new applications up on Unix-based SMP systems. The broader range of development tools, better support for distributed application development, lower cost of development platforms, and lower cost of deployment platforms all make this alternative very attractive.

Using a surround strategy requires that certain issues be faced. These include the extent to which the new applications need to interface with existing mainframe applications, the degree of interoperability required between the new systems and the mainframe systems, and the advisabilty of introducing an entirely new infrastructure for these new applications. The new infrastructure will often include a new network, a new security environment, a new system and user management environment, and a new application development environment.

The potential cost savings of the surround strategy make it an attractive option. Since the alternative requires doing the same new development on a mainframe, the analysis will virtually always come out in favor of doing the new application on the Unix system.

Mainframe Coexistence

Mainframe coexistence, the strategy of blending into existing IBM mainframe environments, is exemplified by IBM plug-compatible vendors, like Amdahl, that have mainframe products but are delivering Unix-based solutions. The objective is to provide a high degree of compatibility in applications and communications environments so that systems can be virtually plugged right into existing IBM data centers as upgrade replacements and cost-effective alternatives.

Mainframe Replacement

Replacement strategies characterize vendors of hot RISC as well as CISC microprocessor-based symmetric multiprocessing boxes (See *Unix in the Office*, Vol. 7, No. 1, January 1992). These Unix-based systems offer price/performance that often betters the traditional mainframe by an order of magnitude. The strategy is less one of coexistence with existing hardware than one of supporting existing applications as they are migrated from the mainframe environment to the new Unix environment.

Alternative Architectures

While alternative architectures are not the focus of this article, there are other camps that say that data centers are irrelevant and that data center functionality needs to be spread

throughout the enterprise. Distributed solutions depend on distributed management architectures, distributed control, and distributed technologies. The reality is that the products to support this level of application distribution are not completely available today.

Data Centers Are Built with Hardware and Software

High-Capacity, High-Availability Systems

Data centers require systems of high capacity that are highly available, minimize planned outages, and strive to eliminate unplanned outages. Availability must be designed into the system as well as the software, and that availability extends beyond processors to front-end processors, terminal servers, storage controllers, and storage subsystems. Systems may approach availability with duplication of components, functional modularity designed in, easy recoverability, and designs that support concurrent maintenance. However it is achieved, the objective is to have systems that stay up and remain available.

A slightly different perspective is that it is not systems that need to be available, but services. This implies that a network of distributed systems can provide redundant sources for important services, with immediate failover should the primary provider of a service fail. This can be accomplished either through clusters of machines or simply an N+1 configuration where the number of systems necessary to provide a service plus one additional system are configured.

Disaster Recovery Broadens the Scope

Reliable systems, or even redundant systems at the same site, aren't much good if the building burns down or an earthquake strikes. Therefore, another important data center issue is disaster recovery. Being able to recover from a disastrous event that destroys a data center requires strategies for maintaining backup systems that comprise hardware, software, and data off site. Again, this is somewhat beyond the scope of this article. However, one issue that Unix data centers pose to service bureaus like Comdisco (Rosemont, Illinois) is that, unlike the homogeneous environment of IBM MVS data centers, providing disaster recovery support to multiple flavors of Unix implementations can be a real headache. If standardization of operating environments helped no one else in the industry, it would help the hot- and warm-site disaster recovery service providers.

Components of Data Center-Capable Systems

Defining Data Center Requirements—OSI

Exactly what should a Unix-based system be able to do in order to qualify as data center capable? Going through the software inventory of an IBM data center is not sufficient, because the functionality of that software is closely tied to the specific nature of the hardware in many ways. A somewhat better approach is the OSI Systems Management model, which defines five areas as components of systems and network management:

- **Performance**, which deals with control and tuning of system resources in order to monitor satisfactory service levels and with planning for future system resource needs
- **Configuration**, which focuses on the control of system and network device configurations
- **Accounting,** which tracks system resource utilization
- **Security**, which defines how to protect the system, network, and its components from unauthorized intrusion or surveillance
- **Fault Management,** which provides the ability to quickly identify, diagnose, and recover from system and network problems

Components of Data Center-Capable Systems

X/Open Extensions for OLTP

Online Transaction Processing (OLTP) is a key application in the data center. X/Open has recognized that Unix lacks certain important capabilities needed to perform acceptably for OLTP applications. Therefore, it has supplemented the work that ISO has done to define system management requirements for OLTP. X/Open published an "OLTP Reference Model," which identifies four areas where standard Unix had to be enhanced in order to function effectively as a database host:

- Support for critical processes (to minimize paging and scheduling overhead)
- Better control over process priorities
- The ability of processes to override preemption
- Positive notification of disk I/O completion

These enhancements have been addressed to some extent in releases from Unix Systems Laboratoties (USL) and further addressed in specific implementations from vendors competing in this market.

Patricia Seybold Group Model of Data Center Functionality

Exploring this area in more depth, the Patricia Seybold Group (PSG) has identified 14 areas, beginning with basic operating system kernel design and extending outward to the policy level, that are required to create a data center environment running Unix which has all of the capabilities needed for commercial applications.

1. Kernel Optimizations

PREEMPTION FOR RESPONSIVENESS. In order to meet data center requirements for reliability, availability, and performance, vendors have had to make many modifications to the Unix kernel. Some have concentrated on supporting asynchronous I/O, while others have tweaked scheduling and task preemption. For example, standard Unix downgrades the priority of a task as it continues to run for an extended period of time. Kernel processes are generally allowed to run without preemption, affecting system responsiveness to user programs. Some processes, however, are critical and must be allowed a high priority run to completion. Many vendors have modified the priority scheduling algorithms of their kernels in order to support this.

PROCESSOR AFFINITY. In multiprocessor configurations, it is often necessary to allow a task to run on a particular processor in order to avoid having to move a lot of data around to different processors' caches. Again, kernel modifications have provided support for this critical capability.

MAINTENANCE OF APIs. The key for kernel modifications is to prevent them from affecting the higher-level programming APIs. This is important to ensure that all applications and services which have been developed to open systems standards continue to run unchanged, no matter what has been done to the kernel.

SUPPORT FOR LARGE MEMORY SPACES. Systems need to be able to be configured with very large addressable memory spaces. Sorting a database that is 20 or 30 GB requires very large memory if that sort is going to happen in a reasonable amount of time. Configurations less than 500 MB will probably have performance problems in most data centers.

Patricia Seybold Group Model of Data Center Functionality

PSG Model of System Management Requirements

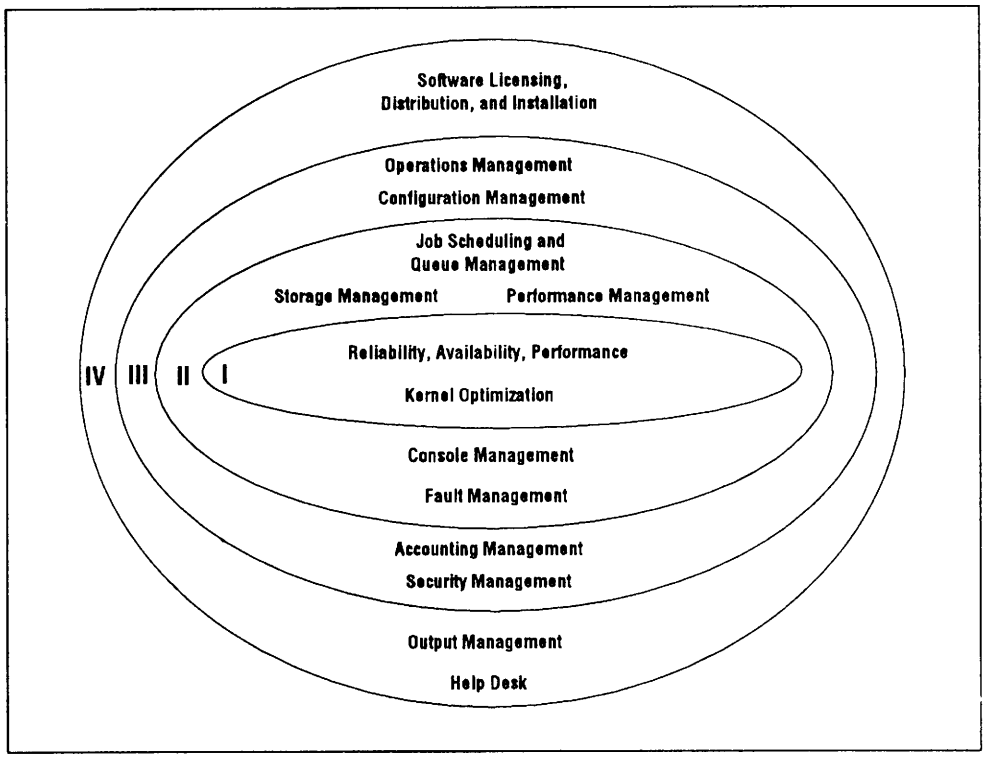

Illustration 1. Requirements for an integrated system management environment. Level I requirements are base technology, not visible to the administrator or user. Level II requirements are system level, but are managed by the administrator. Level III requirements are application related, but are seen only by the administrator. Policy plays a significant role at Level III. Level IV requirements impact users as well as administrators and involve policy as much as technology.

2. Fault Management

Capabilities in fault management include automatic problem detection, automatic notification, automated diagnostics, and failure recovery. The ability to run lights out or to run 24 x 7 (24 hours a day, 7 days a week) depends very heavily on fault management. Many vendors have service offerings that supplement the various notification mechanisms they have built into their systems. Few can support other vendors' systems yet because of the differences in interfaces. With DME, that should begin to change.

3. Configuration Management

Large data centers, with thousands of users spread over many sites and hundreds of gigabytes of data being backed up and archived almost constantly, require capabilities built into the system as well as tools to manage their configurations. Tools for assisting capacity planning, for managing the installation of new equipment, and for hardware and software updates as well as for all of the tasks involved with network configuration and management are critically important.

4. Security Management

While security has to be built into the system all the way down to the kernel, the management interface to those security functions has to be functionally rich in order to handle the task of supporting large numbers of users, files, and applications. Security management must facilitate security for user accounts and logging activity for audit purposes. It should provide policy-based security so that the same entries don't have to be made hundreds of times for members of the same security group.

Patricia Seybold Group Model of Data Center Functionality

5. Accounting Management

Chargeback of costs from the data center is standard in virtually all companies. The system must maintain extensive statistics on usage and must support overhead allocation, split charges, discounting, and credits based on a wide variety of parameters. Even within companies that are simplifying their chargeback, for example, on the basis of average usage, those statistics still need to be maintained.

6. Operations Management

In many ways, the data center is the heartbeat of the enterprise. The operations staff needs proper tools, including those for workload planning, scheduling, and execution. Facilities for logging system events must be provided so that operations can be evaluated and monitored, and the log must be able to be searched and analyzed as well as archived for historical analysis.

7. Performance Management

The system has to provide for the collection of extensive performance data, including user response time, per-user resource utilization, per-terminal resource utilization, per-process resource utilization, file system I/O, and many other system parameters. Facilities should be provided for both analyzing the performance data and optimizing system performance.

8. Software Licensing, Distribution, and Installation

Although some may consider software distribution and installation to be a part of either operations management or configuration management, it is so critical to proper system operation that it deserves its own category. It includes the capability to push upgrades out to distributed users from a central licensing source as well as to allow users to pull installations from a server. This category also includes tracking software licenses, not only for the data center, but for all products that are supported centrally and used throughout the enterprise.

9. Job Scheduling and Queue Management

While Unix provides "`cron`" and "`at`" utilities, data centers require much more sophisticated job scheduling and queue management capabilities. Running a job according to a date/time event is easy. Running a series of jobs that have sequencing requirements because one job relies on the output of another is difficult. Scheduling has to be able to be event based, so that, in these cases, a job runs only if another job has run and has run successfully.

10. Console Management

Console management interfaces should provide an integrated environment for managing all aspects of system performance and operation. Remote console is becoming increasingly important for 24 x 7 operations and lights out operations. The ability to manage multiple systems from a single console and to store, forward, and respond automatically to system console messages is critically important.

11. Output Management

Output management capability, including devices such as printers, plotters, and fax servers, is a data center requirement. Users, as well as administrators, should have the ability to view the status of their jobs and to start and stop their jobs. Also required are queue management capabilities for administrators, including the ability to start and stop queues, re-prioritize jobs in the queue, delete jobs in the queue, and transfer jobs from one queue to another. Included in this area are the ability to archive and retrieve queues for future output, and the management of report distribution. Forms management enabling data to be merged with pre-defined forms at print time is also an important aspect of this area.

12. Storage Management

The system must provide integrated, system-level support for large-capacity mass storage, 500 GB as a minimum, and control of backup and restore, both attended and unattended. The ability to span physical disks with logical volumes and to manage those logical volumes is key. Related capabilities are the ability to mirror drives to protect data and other data integrity features, such as RAID support. Ideally, backup and restore should be online, i.e. while the system and applications are operational. Backup and restore over the network is also important, allowing centralized backup of distributed systems. Support for automatic retrieval and hierarchical storage management helps optimize the use of online, near-line, and off-line devices, making certain that required data are available as needed.

13. Help Desk

Help desk facilities provide support for the interaction between users and systems managers. Help desk capabilities help maintain logs of problem reports, track how those problems are being addressed, and when and how they are resolved.

14. Reliability, Availability, Serviceability

Enhancements to the operating system prevent many system outages caused by defective hardware or software, thereby providing reliability. The extended file system and online peripheral diagnostics provide availability by reducing the amount of system downtime required for operational maintenance. Improvements in binary maintenance support tools and error detection capabilities enhance serviceability of the operating system.

Caveat: Customer Requirements Will Vary

Not unlike the automobile manufacturers' caveat that actual gas mileage achieved by an automobile owner might vary from the EPA ratings depending on usage, customer requirements for a data center functionality will vary. The specifics will depend on the combination of the requirements they currently have, the extent to which those requirements are currently being met, any additional requirements they develop post-migration, and any requirements they are willing to sacrifice in order to achieve benefits such as cost savings. We believe that the areas in which customers will be least forgiving are likely to be backup/restore, security, and job scheduling. These are areas in which Unix has traditionally been weak and mainframes strong, and they are likely to be key differentiators until neutralized either by the Unix vendors or by third party management software providers.

Nine Traditional Weak Points of Unix

Why is this mass migration to Unix just happening now? In the past, Unix has had a deservedly bad reputation in terms of its ability to fit the bill as a robust, commercial-quality operating system. Unix was correctly viewed as an operating system developed by hackers, used by hackers, and only suitable to hackers. Corporate data centers do not run on hacked code. Data center managers were justified in dismissing Unix as inadequate for their purposes for reasons that fell into nine categories: System Administration, Security, Workload Management, Performance Monitoring Tools, Backup Facilities, Storage Management, Accounting Tools, Print and Spooling Management, and General Robustness.

Labor-Intensive System Administration

Administration of a Unix system generally required hand-editing entries in system configuration files. This was prone to error and very time-consuming, and training system administrators was difficult. Most commercial systems now offer some kind of command structure for system configuration, but there are still inconsistencies across them.

Gaps in Security

Shortcomings in security included the lack of auditing facilities and the inability to set security policy. The fact that the superuser account could access any resource on the system also created problems for data center security. Another concern was that passwords are encrypted in a publicly readable file. Theoretically, a user could copy an encrypted password into a personal file and use it to gain unauthorized access.

Scheduling Missing from Workload Management

One of the deficiencies of Unix that data center managers mentioned most frequently was the lack of adequate batch-scheduling facilities. The Unix "`cron`" utility, which can run a program or shell script at a designated time, lacks the conditional scheduling and exception-handling facilities necessary in a data center.

Crude Performance Monitoring Tools

Being able to support thousands of users and large data sets and to run business-critical applications, often in near real-time, requires the ability to monitor, tune, and optimize system performance. Little can be done to tune the system and plan workloads without the tools to monitor a wide variety of system performance parameters on a fine-grained basis. Mainframes have required highly proficient monitoring tools in order to squeeze out every last bit of performance from very expensive resources. Unix, on the other hand, has traditionally lacked these facilities, partly because they needed to be implemented with

specialized hooks provided by interfaces to the hardware which Unix could not support in its generic distribution and partly because Unix system resources are a lot less precious than mainframe resources.

Lack of High-Speed, Automated Backup

Backup facilities within Unix are often slow and poorly automated, requiring almost constant manual supervision. This is not acceptable in a data center that has hundreds of gigabytes of data to back up every day. Also, support for large-scale backup devices, such as tape silos, hasn't been available. Unix also has no concept of hierarchical storage management.

Limited Storage Management

Unix has had inherent shortcomings in managing large storage configurations and spanning physical disks with logical volumes. Growing and shrinking file systems across volumes has been a nightmare. System crashes could force file system checks on reboot that could last for hours.

Lack of Fine-Grained Accounting Tools

Data center resources are generally charged back to users on the basis of some measure of resource utilization. Unix lacks adequate accounting facilities to enable charging for usage with any real degree of granularity. For chargeback purposes, different measures of resource utilization will be more relevant in some applications than in others. For instance, the amount of storage used may be relevant for large databases, but, in an application with a high volume of transactions, the number of I/Os may be more relevant. The collection of chargeback information is related to the point above regarding performance monitoring.

Hard to Manage Printing and Spooling

The spooling facility within Unix is crude, and queue management is virtually inaccessible to users. Unix lacks the concept of a consistent imaging model or a page definition language. Its formatters are cryptic and difficult to use, and control over partial printing is virtually non-existent.

Questionable Robustness

Unix implementations generally assumed that there was a system manager available at all times who could do everything from rebuilding the kernel when a new device was installed, to reallocating disk space, to restarting the system after a panic. Many types of relatively ordinary system events (such as a file system becoming full) could easily bring down an entire system.

Why Bother with Unix? —It Has Matured

With that background, why even bother discussing the possibility of using Unix in the data center? Simply put, Unix has grown up. Many past deficiencies and idiosyncrasies have been addressed in the current releases of Unix and related technology from the various sources of supply, including Unix System Laboratories (SVR4.X), Sun Microsystems (Solaris, NFS), and the Open Software Foundation (OSF/1, DCE, DME). Unix system vendors have made major improvements in their Unix products that address still more of Unix's weaknesses. Finally, numerous Independent Software Vendors (ISVs) have seen the opportunity for Unix system management utilities presented by the market and have jumped on the bandwagon. Through a combination of better technology, system vendor enhancements, and third-party products, robust, commercial-quality Unix seems to have become a reality.

It is interesting to note that many of the third-party applications are reincarnations of similar products marketed by the same vendors on IBM mainframe platforms. In spite of the perceived power of MVS, the robustness of the mainframe data center still relied very heavily on layered software products from IBM as well as from third parties. The market for third-party data center applications was a significant market. Today, as the market for mainframe software slows down, these ISVs are flocking in droves to Unix.

Economic Benefits of Leveraging Open Systems

Technology Leverage

Unlike proprietary approaches, Unix-based data center systems have the advantage of being able to draw technology from throughout the industry. Vendors are not dependent on being able to develop their own storage management technology, tape backup technology, performance monitoring tools, nor on virtually anything else, for that matter. As exemplified by Computer Associates' (Long Island City, New York) CA-Unicenter and General Atomics' (San Diego, California) UniTree, data center capabilities can be added to systems by porting third-party software. These approaches have the added benefit of offering consistent management across heterogeneous systems supported by each product.

Benefits for Users

Consistent open-system interface standards implemented on platforms from a multitude of vendors have measurable benefits for software developers: the cost and effort of delivering a product on multiple platforms is significantly reduced. But the real winners here are the users because they no longer are faced with the absence of a particular software application on their platform of choice. The ultimate benefit of open systems is freedom of choice for users, and that benefit is being manifested today in the choice of system management products for Unix.

The Ultimate Impact of DME

Looking into the future, to 1995 and beyond, we see that customers will be able to buy products that fit into OSF's Distributed Management Environment (DME) framework. System and network management applications, running on top of this distributed, object-oriented framework, will allow customers to mix and match management applications in different functional areas provided by different vendors, all of which will share a consistent, integrated management interface. The experience Unix vendors are gaining providing data center class systems today will undoubtedly improve the quality and breadth of DME-based products in the future.

Software Vendors Pursue the Opportunity

When a company's business depends on a market that has not only stopped growing but is actually beginning to shrink, it time to rethink the business. That is exactly what the vast majority of software vendors who have depended on a healthy growing mainframe market have begun to do. Whether their focus has been in database management, applications, or system management tools, they have begun the task of making their products available on Unix.

For most software vendors, porting an application from MVS to Unix means a complete re-write of the application in C, using the MVS application as a functional specification. In the course of adapting their ports to Unix, the most typical strategy for software developers seems to be to use as much of the POSIX interface specifications as possible. While this doesn't have a visible benefit to users, it has tremendous benefits to the ISV. Once the first port to a Unix environment has been completed, the time necessary to complete subsequent ports to another POSIX-compliant system is dramatically reduced. The ultimate benefit for users is that the selection of software products available will be far greater than it would have been had each machine required extensive work in order to complete a port. The benefits for the software developer are faster time to market, lower costs of development, and, if the product sells, increased profitability.

Products Require Channels

Just having products ported to new platforms will not be sufficient for the mainframe data center vendors to succeed in the Unix market. New channels of distribution will have to be developed. For the most part, layered data center management products have been sold through direct sales. The environment vendors were selling into was well known by themselves and by their customers—IBM MVS on 3090 mainframes. Now that these

Software Vendors Pursue the Opportunity

products are going to be sold for new, unfamiliar, Unix-based platforms, these ISVs will have to establish relationships with the system vendors as well as with the systems integrators, who are going to be involved in making all the pieces work together. These relationships have already begun to be announced, and we expect to see more as the market evolves. In part, this will be made possible because the Unix system vendors are less likely to have competing products than did IBM in the mainframe market.

Universal Availability Should Spark Planning

It seems to be a safe assumption that virtually all software that is currently running only in the MVS environment but is specific to the IBM hardware or operating system environment will be available on Unix within two years. Therefore, those who begin planning today to migrate their data center operations to Unix will have a much broader selection of products available to them by the time they are actually ready to bring systems into production. In the interim, applications will be ported and/or developed, performance will be characterized, and policies and procedures for managing the new environments will be developed. Large users in particular are in a position to establish strategic relationships with data center software providers to help guide their efforts in bringing product to market. That work can enable those ISVs to ship products that meet customer requirements.

Vendors Search for the Right Strategy

Choosing the Right Approach.

Vendors of Unix systems hoping to enter the data center have a number of different strategies from which to choose. For the most part, these strategies parallel the different ways in which Unix systems are being used in the data center. There are several considerations along which the strategies of the various vendors will differ:

- Scope. Will the vendor try to provide a complete desktop-to-data-center "Single Logo" solution? And will this solution consist of a single scalable architecture, or will there be one architecture for the desktop and another for the data center?
- Relationship to mainframe. Will the vendor try to coexist with existing mainframes, replace them, or ignore them?
- Computing style. Will the vendor introduce a different style of computing, e.g. distributed and peer-to-peer as opposed to host-based?

Finding the Balance between Breadth and Depth

The vendors offering the broadest scope are those with the broadest product lines. Hewlett-Packard, Sun Microsystems, and Data General offer single scalable architectures from the desktop workstation to data center-class machines. The advantage these vendors can offer is that applications can be developed on smaller systems and deployed on larger, binary-compatible systems. The disadvantage is that they may be less focused on data center solutions and their resources may be spread too thin to adequately support difficult migrations for users. Amdahl, Pyramid Technology, and Sequent Computer, on the other hand, make no pretense about delivering desktop-to-data-center solutions. Instead, they provide the application and data management piece and integrate that with commonly available desktop pieces. They specialize and concentrate, which can bring a lot of expertise to bear, but they do not offer one-stop shopping.

Attitude toward Mainframe Coexistence

Some Unix vendors are confident that they can immediately supplant IBM mainframes. Others believe that to be folly and that a complementary role is more appropriate for the time being. Of all the vendors, Hewlett-Packard is probably the most aggressive about being a mainframe replacement, although that is not its only approach. On the other hand, Sequent tends to be the lowest keyed in regard to the mainframe, believing it has a more important role to play as the anchor of a new, distributed, client/server architecture. The other vendors, with the exception of Unix mainframe vendor Amdahl, tend to fall in between.

Vendors Search for the Right Strategy

Computing Style—Centralized or Distributed

Although Unix is, by nature, at the foundation of distributed computing, some vendors emphasize that aspect of their strategy more than others. We believe that Sun and Sequent are the hardest core adherents to distributed computing, while Amdahl and Pyramid represent a more centralized approach. However, all the vendors intend to track and support emerging standards for open systems, so even if distributed open systems are not a central part of their strategy, they will certainly be able to support them.

Selected Vendor Profiles—Can We Buy It?

For the purposes of this report, we have selected six vendors who have explicit data center strategies for their Unix systems, including Amdahl, Data General, Hewlett-Packard, Pyramid Technology, Sequent Computer, and Sun Microsystems. This article is not an exhaustive review of every vendor who sells data center solutions. It is a sample of options that are on the market or will be coming to the market soon, selected to present an overview of the range of capabilities that Unix can deliver to the data center. Some hardware is mainframe class; other platforms are RISC SMP; and others are Intel SMP. The focus is on support and management capabilities, not on database, data management, or transaction processing. Future products are included if they will be generally available within the span of time it normally takes an organization to make a purchase decision, place and process the order, and take delivery.

Amdahl—The Mainframe Alternative

Nothing in the concept of a Unix-based data center dictates that a microprocessor-based multiprocessor architecture is the only one that is acceptable. Unix has not been a popular mainframe operating system because, historically, mainframes have been synonymous with IBM mainframe operating systems, particularly MVS. In addition, the limitations Unix has had scaling upwards in memory capacity, supporting large storage configurations, and in the number of users that could be supported constrained its usefulness as a mainframe operating system. Those limitations are largely disappearing, however, making Unix a much more viable option.

Amdahl Corporation (Sunnyvale, California) has understood the challenge of running Unix on large systems for some time. As early as the mid-1980s, Amdahl customers in the telecommunications industry, long-standing Unix adherents, wanted Unix in their data centers. Amdahl provided a native Unix operating system on its mainframes back then, and its experience with those customers led the company to evolve and expand its Unix support to the point where, in its most recent year, Unix systems represented approximately 15 percent of sales and an even higher proportion of its profits. In fact, Unix plays a large role in Amdahl's own data center as well as in its research and development support.

Amdahl UTS

Amdahl has been working on running Unix natively on the IBM 370 architecture for a long time. In the mid-1980s, the company was developing a new operating system, Aspen, for its next-generation mainframes when it realized that putting the enhancements in development into Unix instead of developing another proprietary operating system made much more sense. The product has evolved to the point where the most recent release of Unix for Amdahl's IBM 370 architecture-compatible mainframes is based on USL's SVR4.1 Extended Security (ES) release (See *Unix in the Office*, Vol 7, No. 2, February 1992). Its previous release, UTS2.1 Service Level 4, was based on SVR3.1. UTS Release 4 (UTS4) is planned for early availability to customers in the first quarter of 1993. General availability is expected for the second quarter, when additional functions, including improved network connectivity through Amdahl's 4655 and 4745 front-end processors, will have been added. This will allow UTS4 to support existing networks of 327X terminals, asynchronous devices, and X.25 networks.

BASIC FEATURES OF UTS. UTS runs in native mode on Amdahl's 5890, 5990, 5995A, and 5995M series of IBM-compatible mainframes. These machines range from uniprocessor up

to eight-way multiprocessor configurations. At the high end, these systems support up to 2GB of main memory and 11TB of magnetic storage (DASD, which stands for Direct Access Storage Device, to mainframe aficionados). A new asynchronous disk and tape I/O capability has enabled Amdahl to further increase its application throughput over previous releases. UTS also offers a fair-share scheduler that allows users' workloads to be partitioned into groups, with specific percentages of processor capacity apportioned to each group, similar to Pyramid's FairShare. This scheduler is in addition to the standard System V time-sharing and real-time schedulers.

Backup And Restore. The UTS tape management system provides access to IBM-labeled tapes. In addition. REELbackup and REELlibrarian, products from Sceptre (Ann Arbor, Michigan), are available for additional tape management capabilities. Not all of the management capabilities that are available on Amdahl's "compatibility" systems are available on UTS because some of that functionality comes from MVS. Amdahl's direction, however, is to provide similar capabilities either through its own enhancements or through third parties.

Tape Management. UTS offers client support for the Storage Technology Corporation's (Louisville, Colorado) STK4400 Automated Cartridge System, a fully automated storage and retrieval system. It adds flexible tape management operations and enhances tape selection for users, incorporating standard tape management extensions that support MVS and VM volume and header labels. Volume labels can be reviewed and a volume selected from a system scratch list. The tape extensions to the tape daemon support an API and a log-file interface to an external tape management system.

Storage. UTS4 will support the IBM Enterprise Systems Connection (ESCON) architecture with full compatibility with the IBM 3990 Model 3 provided by the Amdahl 6100 Storage Processor. It is designed to operate transparently in an IBM Data Facility Storage Management (DFSM) subsystem environment. Support is included for tape silos and for the Open Systems Expansion Features of Amdahl's 6390-3 disk drives. The 6390-3 drive increases the configurable capacity in UTS operating environments to 108 GB per string, using 3.39 GB modules under the Open Systems Expansion Feature. This increases the total addressable storage to 11TB.

UTS supports UniTree Release 1.0 from General Atomics (San Diego, California) for a cost-efficient storage management environment to provide improved data reliability and integrity. It is a distributed hierarchical storage management system that provides an information management solution for large enterprises.

Security. UTS 2.1 customers can convert their UTS systems into secure UTS/MLS systems by licensing and installing a security features tape, which is priced separately from UTS.

Language Support. No mainframe environment could function without Cobol. Amdahl provides Cobol, based on MicroFocus Cobol/2, on UTS. It also supports popular extensions included in RM/Cobol, DG Interactive Cobol, IBM OS/VS Cobol, IBM VS Cobol II, and Microsoft Cobol. Its intermediate code is fully portable across platforms. C++ is supported, operating as a preprocessor for the Amdahl C compiler, generating C source code tailored for the compiler.

Enhanced System Management. UTS backup and restore facilities schedule and automate the execution of backups, and support incremental and online backup. These facilities provide a consistent operator interface independent of the local file systems supported. UTS Release 4 supports multiple file systems, including Amdahl's extended file system (EFS); the SVR4 system expanded for 4K block size; the Berkeley fast-file system; a secure file system with access control lists; a boot file system; remote file system, RFS; and Network File System, NFS.

MIGRATION FROM UTS2.1 TO UTS4. UTS4 represents a major upgrade from the earlier release, which was System V.3 based. Amdahl is providing various options for migrating to the new release, including a binary compatibility feature that allows existing UTS 2.1 applications to run under UTS4 without change. Migration tools are available at no charge to UTS2.1 customers to modify application source code before making the transition. Modified applications can continue to execute in the 2.1 environment until transition is complete. The nature of the migration is no different from any of the other V.3 to V.4 migrations that the Unix community has faced. Migration issues are encountered mostly when applications write to interfaces close to the hardware.

STANDARDS COMPLIANCE. UTS4 will comply with SVID Issue 3, POSIX 1003.1, XPG3, and ANSI C standards (X3.159-89). It comes with OSF/Motif 1.1, OpenLook Release 4, and X11 Release 4.

CONNECTIVITY. Amdahl uses TCP/IP to communicate over Ethernet, HYPERchannel (from Network Systems Corporation (Minneapolis, Minnesota)), and FDDI networks, and over channel-to-channel connections. Amdahl also resells Network Systems' Ethernet and FDDI controllers. GOSIP requirements are supported with a standard product offering. Support is provided for ONC, including NFS, remote execution, lock manager, automounter, and NIS. Through the 4655 front-end processor (FEP), UTS can support Datakit (a networking scheme popular with Amdahl's highly valued telecommunications customers), X.25, and BX.25 protocols. VCS Fiber 2.1 support for the Amdahl 4655 Datakit FEP allows fiber optic connectivity to the Datakit VCS, which supports over 1,500 virtual circuits on a single 4655 FEP. UTS also supports VCS Host 2.0 for Amdahl's 4635 Datakit VCS customers.

UTS applications can communicate on SNA networks with applications running under MVS using LU6.2 over a channel-to-channel interface. UTS also supports a channel-to-channel interface to MVS Network Job Entry facility, allowing simultaneous peer-to-peer transfer of job and print spool files between MVS and Unix systems. UTS supports Tuxedo Release 4.2 for OLTP applications.

FUTURE DIRECTIONS Amdahl is researching the application of RISC technology to mainframe architectures and applications. The work it is doing in Unix will come in very handy when those products are released, since porting both operating system and applications to a new architecture will be much less an issue for UTS than it will be for MVS.

A future release of the Amdahl UTS operating system will be submitted for branding for XPG4 compliance. Amdahl plans to implement X/Open's Transport Interface, XTI, in a future release. (XTI is a transport service interface that is independent of any specific transport provider. It has been extended to include both TCP and User Datagram Protocol, or UDP.) Future releases of UTS4 will conform to the UK GOSIP Version 4, OSI X.400 MHS 1988 standard, and X.500 Directory Service. Also slated for a future release of UTS will be SNMP agent capability for all elements of the TCP/IP stack.

POSITIONING A UNIX MAINFRAME. Amdahl provides an option for the customer that still requires very large capacities but wants to use the Unix operating system. UTS4 will be a mainstream release allowing applications to be developed on smaller systems of different architectures and then recompiled and run on the mainframe. Running UTS on the Amdahl mainframe allows mainframe processing within an environment of open systems interoperability.

Hewlett-Packard's Mainframe Alternative

Hewlett-Packard Company has been aggressively pursuing a strategy of selling its high-end RISC multiprocessor machines as mainframe replacements for almost two years. This effort initially was led by its HP 3000 line, which runs HP's proprietary MPE/ix, and was later extended to the Unix-based HP 9000 as well. (See *Unix in the Office*, Vol. 7, No. 5, May 1992). It is important to recognize that the HP 3000 and HP 9000 are the same hardware and

differ only in the operating system and software they run. The Unix systems that HP is offering as data center alternatives are the top-of-the-line HP 9000 Model 890 Corporate Business Systems. These systems are currently four-way symmetric multiprocessing RISC machines based on the HP Precision RISC Architecture. In the future, they will be expanded to support 16 or more processors. The Model 890 can be configured to support up to 6,000 users, with 2GB of memory, and 1.3TB of mass storage. HP has measured its systems to provide MTBFs (Mean Time Between Failure) of over three years. Its standard disk storage subsystems exceed MTBFs of 20 years. Clearly, HP is hoping to exploit the renowned reliability of its products in its efforts to get into the data center.

OPENVIEW IS MANAGEMENT FOCAL POINT. HP's umbrella for system and network management is OpenView. Several components of OpenView were selected as parts of the Open Software Foundation's DME, and it is HP's plan to move OpenView to fully support DME as it evolves. HP supplied the OpenView Network Management Server, Network License System, and Software Distribution Utilities to DME, and also worked with Bull on the Consolidated Management API (CM-API). OpenView will also be the vehicle for integrating third-party management applications on HP platforms.

Central System Manager (CSM) sits on top of the OpenView services. It is an interface and integrating platform that is being developed by HP within its open management framework, and it will allow both HP and third-party management applications built on top of the framework to share common screen, event notification, and event-handling facilities. It will also allow users to predefine actions to be taken in response to system events.

PERFORMANCE MANAGEMENT TOOLS. HP provides several performance management tools of its own as optional supplements to the capabilities of HP-UX. HP GlancePlus/UX is an online tool for viewing a snapshot of current system activity. CPU, memory, swap space, and disk utilization can be monitored by the system administrator, allowing the system's performance to be characterized and managed. Problems can be identified and resolved either on the host or from across the network using "`rlogin`". LaserRX/UX is another tool that combines historic performance data collected on an on-going basis from the kernel with graphical analysis software to analyze system performance and activity over time. It can be used to identify system bottlenecks, perform load-balancing, and locate potential performance problems. LaserRX/UX can work across the network to analyze all supported systems.

Building on the data collected by LaserRX/UX is a capacity-planning tool called RXForecast/UX. This product builds on the historical performance data from the current system configuration to forecast performance under future loading. It uses a variety of statistical forecasting tools to perform this analysis, and the documentation takes users through the analysis so they can make informed forecasts of future requirements.

HP PerfView is designed for managing a large number of systems in a distributed environment, including servers and workstations, rather than a few large data center-class machines. PerfView has two components: Performance analysis software based on OSF/Motif runs on HP 9000 Business Servers, and performance collection software runs on both HP 9000 machines and Sun SPARCstations. It uses management-by-exception techniques to spot and resolve both actual and potential problems before they affect users. It is integrated with OpenView to provide a graphical map of the system environment.

OPERATIONS MANAGEMENT TOOLS. HP includes output management and production scheduling under the area of operations control. Print-spooling beyond that offered in Unix is provided by HP OpenSpool, which is a network-based spooling application built on the MIT Palladium Print System Version 1. Palladium Print System Version 2 is a part of OSF DME, and HP expects a straightforward migration in future OpenSpool releases. OpenSpool allows establishment of print priorities, viewing and changing requests; provides form and

font management; and supplies templates for repetitive tasks, single-point administration, and security. It supports SunOS as well as HP-UX. OpenSpool will be integrated in the OpenView management platform under CSM in 1993.

In addition to what HP provides, the database management system from Software AG (SWAG) (Reston, Virginia), Adabase, has extensive output management capabilities, including sophisticated report formatting. These capabilities are being tied into the OpenSpool environment by SWAG.

HP has no offering of its own for production scheduling beyond what is in Unix, "`cron`" and "`at`", and will rely on third parties for the time being. However, CA-Unicenter provides scheduling capabilities comparable to its mainframe products. In addition, solutions are available from AIM Technology (Santa Clara, California), which has Job Scheduler; Software Clearing House, or SCH (Cincinnati, Ohio), which has Qbatch; and Unison (Sunnyvale, California), which will release a product called Maestro in the first quarter of 1993.

HP Task Broker is a load-balancing facility that provides automatic distribution of jobs across a network of systems to optimize the use of available resources. Task Broker automatically directs jobs from a backlogged machine to an idle machine.

STORAGE MANAGEMENT. In the area of backup/restore, HP has two offerings beyond standard Unix facilities for the data center. OmniBack is a network backup management system that provides automated central backup and recovery. It can be configured from any system on the network and can back up any HP-UX, Domain, or SunOS system on the network. It provides unattended backup/recovery with a variety of media, including magnetic tape, DAT, or rewritable optical disk. With the rewritable optical option, OmniBack can do an unattended backup/restore of 2TB of data. OmniBack/Turbo is a superset of OmniBack that increases performance in networked environments and on systems with large amounts of online data. It can provide local backup of up to 25GB per hour in certain configurations. Storage management is also available from Storage Technology, which will provide 3480 tape support by mid-1993 and silo products by late 1993.

Disk management is provided with HP's implementation of OSF's Logical Volume Manager (LVM) within HP-UX. LVM adds disk-spanning capabilities and error-recovery features. HP's MirrorDisk/UX is the mirroring component of OSF's LVM that supports up to three-way mirroring, with online backup while maintaining mirroring. It is transparent to applications and can be brought into operation while applications are running. HP also has RAID arrays available as a data integrity option.

Tape management solutions are primarily obtained from third parties, such as SCH's REELlibrarian, REELbackup, and REELexchange for HP-UX. During 1993, HP plans to provide an integrated storage management product family of its own, integrating backup, archive, and hierarchical storage management, supporting files as well as databases. Multivendor support, online capabilities, faster speeds, advanced media management, larger storage capacity, and file migration will be addressed in HP storage management product announcements during 1993.

CONFIGURATION MANAGEMENT. HP breaks down configuration management into four functional areas: Installation and Update, System Configuration, Network Configuration Management, and PC Management. The installation and update area is addressed by a utility shipped as a part of HP-UX called Distributed Update and Install (DUI). It enables the creation and use of a network server system as a source of update software. System configuration is managed by the System Administration Manager (SAM), which is a screen-oriented system administration environment. Administrators can use SAM to manage users and groups, the file system, printers, Unix-to-Unix Copy Protocol (UUCP), trusted systems,

devices, backup/recover, and auto-configured devices. The release of HP-UX Version 9.0 added a Motif interface to SAM, along with single-system remote administration, logging, tape drive configuration, process management, system configuration, and other functions.

Network configuration management is provided through OpenView and the OpenView Network Management Server. The inclusion of this technology in OSF's DME will help HP become a leader in DME-based solutions. Other products relevant in this area are Network Node Manager and Network License Server.

PC Management is provided through NewWave Office System Services; LAN Manager for Unix; NetWare integration and Native NetWare/9000; and Pacer Software's PacerShare, PacerPrint, and PacerLink for Macintoshes. HP Software Vendor is a product that provides a means of centrally controlling and distributing PC applications in a LAN environment. It provides everything for PC software management, including installation, push/pull software updates, and license-tracking.

In mid-1993, HP will release Software Distribution Utilities (SDU). These will manage the creation, installation, and removal of system and application software packages. SDU will be a part of DME as well.

ACCOUNTING MANAGEMENT. HP currently relies on Unix utilities for accounting management. Disk Quotas, used to limit the disk space usage of users on a per-file-system basis, is standard with HP-UX. Otherwise, HP relies on third-party products, such as CA-Unicenter, to provide applications that generate chargeback information. HP will work with third parties to continue to provide software for system accounting on the OpenView platform.

SECURITY. HP-UX is designed for C2 security, and a B1 version, called B-Level Security (BLS), is available. Security enhancements to HP-UX include protected password database, access control lists, and auditing. Security is administered from the SAM, enabling the administrator to audit events, such as file creation, deletion, open and close; process operations; logins and logouts; and others. BLS incorporates components of trusted software technology from SecureWare, Incorporated (Atlanta, Georgia). Demax's (San Mateo, California) SecureMax is also available for distributed HP-UX environments.

A STEP AHEAD WITH CA-UNICENTER. HP was the first Unix vendor to announce CA-Unicenter for its systems. The HP 9000 has served as the reference port for CA-Unicenter, and, as a result, CA-Unicenter for HP 9000 systems is available now in an early shipment program and will be generally available soon, giving HP a six- to nine-month lead over the other systems to which CA-Unicenter is being ported. CA-Unicenter will be integrated within OpenView in those areas where it makes sense. In a similar fashion, aspects of HP functionality will be integrated into CA-Unicenter. One example might be the integration of OpenSpool with CA Report Distribution Manager.

RAXCO: ANOTHER KEY HP SUPPORTER. Raxco Incorporated (Reston, Virginia), a major provider of VAX/VMS data center tools, signed a joint marketing agreement with HP to make some of its products available on the HP 9000 Series 800 machines. Raxco has over 10,000 DEC VAX customers to whom it can introduce HP. The company's Security Toolkit allows managers to assess system security, monitor users for suspicious activities, and protect data from unauthorized use or tampering. BACKUP.UNET is an automated backup, restore, and media management utility for multivendor Unix networks. PRINT.UNET manages printer resources throughout multivendor Unix networks. One print spooler can be used across all Unix platforms.

Sun Microsystems Takes SPARC to the Data Center

The fall announcement of the SPARCcenter 2000 was Sun Microsystems' (Mountain View, California) most significant foray into data center-class hardware. Its earlier high-end products had been used to offload data center applications, but these systems were explicitly

designed to offer mainframe-class performance. They represent an entirely new generation of Sun systems, extending the company's capabilities far beyond its previous line of servers. With these machines, Sun hopes to provide desktop-to-data-center scalability with a single architecture. This scalability is being positioned as an advantage for the development of large-scale, distributed applications that will initially supplement, rather than replace, mainframes. However, that does not mean that Sun will not seek to replace mainframes in the future.

Although designed to be configured with 20 processors, the first release of SPARCcenter 2000s will only contain up to eight 40 MHz SuperSPARC processors. Volume deliveries of these systems will begin in April 1993, while shipments of the 20-processor configurations should begin shipping by the end of 1993.

UNIQUE DESIGN FOR SPARCCENTER 2000. The design of the SPARCcenter 2000 systems is interesting to consider because of its departure from earlier Sun designs. The systems have a unique backplane bus designed in conjunction with Xerox PARC research. It is a high-speed, packet-switched, dual bus design that, in addition to providing designed-in redundancy, provides a sustained throughput of 500 MBps. There are ten slots in the backplane for system boards. The first release supports up to five system boards. Four of these system boards may contain up to two processors, yielding the eight-processor configuration. Each system board also has 4 SBus slots for I/O expansion as well as SIMM slots for memory expansion. There is 1MB of external cache configured with each processor in the system, resulting in a total of 8MB of external cache for an eight-processor system. The four processor boards that the systems support have up to 512MB of shared memory each, and a fifth system board, with memory but no processors, can support another 512 MB, for a system total of 2.5GB of memory.

These systems will be able to support faster SPARC processors as they become available. With the current CPUs, an eight-processor SPARCCenter is rated at 864 MIPS. The systems will initially only support approximately 138GB of mass storage. However, by the end of 1993, up to one TB will be supported as a result of adding support for drive arrays.

ARCHITECTURE FOR THE DATA CENTER. On the software side, Sun's strategy is to provide an architecture, foundation services, and tools in Solaris to which third parties can add value and provide data center services. Sun is looking to a combination of the functionality that it provides in its systems software and third-party products to furnish the necessary functionality in storage management, availability, security, and system administration. Sun has announced an agreement with Computer Associates to port CA-Unicenter to Solaris 2, and another agreement with Legent Corporation (Pittsburgh, Pennsylvania) to port its system management tools. Sun is also aggressively pursuing additional data center partners.

STORAGE MANAGEMENT. Online: Disksuite is a set of utilities that enhance data availability and data integrity. There are also facilities provided within Solaris for online file system growth, high-speed file system recovery, and enhanced file system performance. Online: IPI Dual Port is another utility that allows a dual-ported disk to be switched over in the event of a system failure. Third parties are relied on for RAID support, UPS, and other availability functions. Mirroring is also supported for customers that require that capability.

BACKUP AND RESTORE. Sun provides Online: Backup to do automatic, unattended, high-speed backup as a part of Solaris. It can be used with a sequence of tape drives or with tape stackers. Sun looks to third parties to provide capabilities like scheduling and hierarchical storage management. Sun has deals with Storage Technology Corporation (Louisville, Colorado) and Epoch Systems (Westborough, Massachusetts) to jointly market high-capacity backup and retrieval peripherals for the SPARCcenter 2000. Those products will also be available on other Sun servers.

SECURITY. SunShield is a set of security features and products designed to enhance the overall security of the Solaris 2 environment. Automated Security Enhancement Tool (ASET), part of Solaris 2, is a system administration capability that allows system administrators to audit and manage security of servers and workstations across the network. SunShield also includes C2-level security features and a number of mainframe-type account protection and access control features. ONC+'s secure RPC includes support for multiple authentication technologies, including Diffie-Hellman and Kerberos 4.0, in addition to Unix authentication. Also available as an adjunct product is Data Encryption Standard (DES), which SunSoft is able to make available only in the United States.

SYSTEM ADMINISTRATION. In addition to the backup tool, SunSoft has included a number of management tools in Solaris 2. The User Manager tool is for managing user accounts, including adding users and building directories. The Database Manager tool is used by system managers to review and modify information in the central NIS+ system administration database. Print Manager allows print server configuration and connections to be made throughout the network using a GUI. The Host Manager is a desktop tool that allows the system manager to configure client systems on the network as either standalone, diskless, or dataless clients. Finally, the Software Manager is an OpenLook adminstration application that allows the administrator to browse, install, and remove third-party software that is iconically compliant with the Solaris 2.0 SVR4 Application Binary Interface (ABI) either locally or across the network.

SUNSERVICE 2000. Sun, recognizing the requirement for a high level of support from data center customers, something it was not particularly used to providing, instituted SunService 2000 in conjunction with the SPARCcenter 2000 announcement. This program provides seven-day-a-week, 24-hour-a-day system support and two-hour on-site response. This service guarantees 99 percent uptime on the SPARCcenter 2000. It will be revealing to see the quality of service Sun delivers to very demanding customers under this program.

SUN'S FUTURE IN THE DATA CENTER. Sun still has a way to go from providing large servers to providing mainframe alternatives. In part, this is because Sun's overall approach to distributed computing runs counter to a data center approach. Some of Sun's features are centralized, such as SunNet Manager, NIS+, and JumpStart, but it remains to be seen how much Sun will accommodate its philosophy to data center requirements for centralized management and control.

Data General Prepares Its Data Center Move

Data General Corporation (Westboro, Massachusetts) has been at the forefront in bringing high-performance, low-cost, RISC-based systems to the market. Since it first began shipping the symmetric multiprocessor AViiON servers with DG/UX in 1989, it has positioned those systems as high-end commercial application systems. Now, it is about to embark on a strategic thrust with both a downsizing message and a simultaneous push into the data center.

DG/UX OFFERS ROBUSTNESS. From the beginning of the AViiON development, Data General has been reworking Unix from the kernel outward to achieve what it felt to be critical performance and functionality goals for commercial applications. Although it is an SVR4 licensee, DG proudly points out that no part of its kernel is "AT&T code."

AVIION OFFERS SCALABLE PERFORMANCE. AViiON systems can be configured with up to eight processors on the model 6280, offering a total of 235 MIPS. DG announced a TPC/A (Transaction Processing Council/A) benchmark in December of 239.1 TPC/As at $7,864 per TPC/A. This makes AViiON a price/performance bargain, which is exactly where it chooses to compete. Using DG's CLARiiON disk arrays, AViiON systems can be configured with up to 6TB of storage. With the ability to configure twin SCSI-2 fast connections per array, a total bandwidth of 10 MBps is achieved. In 1993, DG is expected to have support for the new SCSI-2 fast and wide standard, which will give it a bandwidth of 40 MBps. Currently,

AViiONs can be configured with up to 768MB of main memory. When 16Mb memory chips are made available in 1993, that ceiling will rise to 2GB of main memory.

NEUTRALIZING THE APPLICATION ISSUE. DG will use its sophisticated Unix, DG/UX, symmetric multiprocessor designs, and disk array technology as the backbone for launching a full-scale assault on the data center early in 1993. Some in the industry have felt that DG's use of the Motorola 88000 RISC processor in AViiON is a disadvantage in getting third-party software applications ported to its platform. If the applications weren't available, a data center assault would be fruitless. Fortunately for DG, it has not turned out to be much of an issue. The work that the 88Open Consortium has done in developing binary standards, branding programs, and certification test suites has actually made the 88K platform an extremely straightforward port from other Unix platforms. As a result, DG will be able to focus on the support it has from third-party software companies that have the most popular mainframe applications. Dun & Bradstreet Software's Smartstream decision support and Millenium financial applications are already available on DG/UX, as are Cincom's Supra database and Mantis 4GL; Integris' UniKix; VI Systems VIS/TP; Oracle Financials; and Software AG's Adabas DBMS, Natural 4GL, and Network connectivity products. We expect Computer Associates to bring CA-Unicenter, Masterpiece, IDMS, and Datacomm packages to DG/UX as well.

SYSTEM ADMINISTRATION. DG markets OS/Eye*Node, a management integration platform that supports monitoring and control of network devices, computing systems, and applications that support SNMP management protocols. It also supports the integration of system administration tools, such as network backup and restore, through a common GUI. DG also offers a system management menu, both on character terminals and on its GUI interface. The character version of the menuing system is called sysadm, and the GUI version is xsysadm. It translates menu-based input into native management commands, which are then executed. A log file of activities is generated that can be used either for future review or for the creation of automated management scripts.

SECURITY FEATURES OF DG/UX. DG offers both C2 and B1 security options in addition to standard Unix security. Its B1-capable package has been submitted to the National Computer Security Center for evaluation and is targeted at achieving a B2 class evaluation. Both the C2 and B1 packages are offered as modular add-ons to the same base, reducing the possibility of incompatibilities across different security versions.

STORAGE MANAGEMENT. DG/UX has a configurable logging file system and a good disk volume manager. However, these facilities cannot be managed online. DG's file system uses many approaches of the Berkeley Fast File System, including configurable block sizes and self-clustering data organization. DG/UX can support an unusually extensive range of block sizes, from extremely large block sizes, up to 1MB, for large data objects such as images or databases, to extremely small block sizes, down to 512 bytes, to minimize fragmentation. DG/UX employs transaction-oriented logging techniques to improve data integrity and file system reliability. Before updating the file system metadata, the metadata update is entered into a disk-based log. The file structure can then be reconstructed from the log should the system fail, resulting in much faster startup and higher integrity. Fast startup is particularly important when the file system contains hundreds of gigabits of data. It is important to keep in mind, however, that journaling is limited to control structures, not data. DG supplies RAID and disk-mirroring for data integrity requirements. Disk-mirroring can be either two or three way. DG/UX also supports disk-striping for high performance, which allows fine-grained configuration of the stripes down to the level of track-by-track interleaving.

DG/UX supports Control Point Directories, which limit the storage used by all files within a specific directory structure. This is different from the more common disk quotas that limit space allocated in a file system on a per-user basis.

Although DG/UX has a logical disk manager allowing volumes to span physical disks, management has to be done offline. On the other hand, it supports expansion of file systems as well as of volumes, which some other vendors do not. It also supports contraction as well as expansion. Online operation is expected in the next release of DG/UX.

DISTRIBUTED MANAGEMENT STRATEGY STILL BAKING. DG has not taken a position yet on whether it will support one of the available open systems management frameworks, such as OpenView, Tivoli WiZdom, or SunNet Manager, or will throw its support behind DME or USL Distributed Management. The company is a member of both the OSF and Unix International and will take advantage of the complementary efforts of the two organizations in this area. In any case, DG can be expected to migrate its OS/Eye*Node product suite toward integrating with DME as the OSF framework evolves.

Pyramid Technology's Strategy for Survival

Pyramid Technology Incorporated positions its line of MIPS-based MIServer products as complete data center solutions. In fact, Pyramid was among the earliest to implement symmetric multiprocessing, in 1984. Pyramid claims that the largest database running on Unix runs on a Pyramid machine and that Oracle7 running on a MIServer ES produced 645.1 TPC/A transactions, compared to 618.3 on a Sequent S2000/750 and 578.0 on an HP 9000 Model 890. The top of the Pyramid line is the newly-announced MIServer ES-12 system, which can be configured with between 2 and 24 RISC processors and up to 1GB of main memory. The RISC processor used in Pyramid's new machines is a specialized implementation of the MIPS R3000 architecture. Pyramid gets nearly linear scaling across the range of its binary compatible systems.

In the past, Pyramid has relied heavily on its OEM revenue, with customers like AT&T and Olivetti accounting for a significant portion of its business. Following AT&T's acquisition of NCR and Olivetti's new relationship with Digital, Pyramid's OEM business is rapidly shrinking. As a result, the company is scrambling to revamp its marketing and distribution strategy with new channels, new partners, and new strategies.

DATACENTER OSX AT THE CORE. DataCenter OSx 1.1 (DC/OSx), Pyramid's Unix, is designed so that both system and user code can execute on any CPU, and multiple CPUs can execute operating system code simultaneously. This degree of symmetry provides high throughput and predictable response times. It is SVR4 based but has been refined and enhanced to meet data center and transaction-oriented requirements. Pyramid is using some portions of USL's ES/MP code, but Pyramid's DC/OSx operating system has more sophisticated capabilities in the areas of symmetric multiprocessing. For example, SVR4 ES/MP uses round-robin scheduling, which is not efficient in the way it uses processors with large caches. Pyramid has implemented several features that optimize performance in systems with large numbers of users. These include: fixed-priority process control, process binding and exclusion, and multiprocessor affinity. Fixed-priority process control ensures predictable response times for server-based processes; process binding and exclusion allows system administrators to optimize system performance by tuning CPU scheduling in complex environments; and multiprocessor affinity continuously tunes Unix scheduling to maximize efficiency in managing processes.

Pyramid supplied the reference port of SVR4 to the MIPS processor, and it offers this reference source code back to other MIPS systems vendors. DC/OSx is based on this reference port. It complies with the SVID Issue 3 and passes the SVR4 validation suite. It is fully POSIX 1003.1 compliant, and it also complies with the FIPS 151-1 standard. It supports OSF/Motif and is claimed to be XPG-3 compliant.

Among the areas where DC/OSx has been enhanced are its scalability for large systems support, workload management capabilities for commercial and database applications, and other data center enhancements, such as virtual disk and mirrored disk support.

SYSTEMS ADMINISTRATION AND MAINTENANCE. Systems administration functions are integrated within a consistent menu interface that provides the system administrator with access to all the primary Unix system administration functions, including user and device management, network management, job scheduling, and system configuration control. It also includes system-specific management capabilities, such as management of virtual and mirrored disk subsystems.

The interface allows backup and restore operations to be performed using sequences of commands. Services are provided by an integrated backup service that supports a backup history log, online backup, automated backup initiation, and automated processing of restore requests.

The administration facility also has an installation component that provides consistent facilities across software packages, releases, and systems. This facility implements tools defined within the SVR4 ABI for developing and installing add-on packages with consistent procedures.

VISUAL SYSTEM MANAGER (VSM). The graphical user interface provided by the Visual System Manager assists in the administration of networked environments. It is icon-based and simplifies system management by eliminating the need for programming or shell scripts. The first release of VSM provides user and device management; future releases will add clustered system support. VSM will provide a single, federated management platform for multiple Pyramid systems. VSM is standards based, and it is designed with emerging management frameworks in mind, including the Tivoli framework. Pyramid is supplying several key components to USL's Distributed Manager (DM) framework, including the GUI style, Storage Device Management, and User and Group Management products.

LANGUAGE SUPPORT. Pyramid supports ANSI C, C++, Micro Focus Cobol/2, Fortran 77, and ANSI PASCAL languages. A variety of 4GLs is also available from ISVs.

FAIRSHARE. System-wide resource management, control, and allocation are provided by FairShare. This enables administrators to define, establish, and implement a resource allocation policy for the system's use. The user community on a system is grouped, and system resources are allocated and managed on that basis. The allocations change dynamically depending on the load being placed on the system. A wide variety of system resources can be allocated by the administrator, including disk, memory, CPU, process counts, printer usage, and connect time. Histories of usage are maintained and can be used as the basis for accounting and chargeback.

NETWORK QUEUING SYSTEM (NQS). Mainframe-style batch-processing and workload management is provided by the Network Queuing System (NQS). It uses a job queue-based scheduling system to stage and manage batch workloads on single or networked systems. Users submit jobs stating their resource requirements and priority level. NQS verifies submission parameters and schedules execution. Job recovery and restart can be managed by NQS transparent to the application. It includes a scripting facility that can be used to construct dependency-based scheduling options. NQS allows printer management across a network in addition to batch management. Network queuing, which had originally been developed by NASA, is available on other systems that will provide interoperability across supported systems. A client can submit jobs from any compatible system.

PERFORMANCE MONITORING TOOLS. GUI-based tools are provided to monitor system events. A real-time performance monitor tool, "`pmonitor`", provides system performance measurements and real-time tracking of performance metrics against user-defined thresholds. Statistics can be tracked on a per-group, per-user, or process basis. Performance monitoring can be performed on a real-time basis or collected over a period of time for later analysis. Pyramid plans to integrate "`pmonitor`" with VSM.

VIRTUAL FILE SYSTEM (VFS). Virtual File System (VFS) allows different file system types to exist at the same time on the same system, including file systems that have significantly different characteristics and internal formats. Through VFS, DC/OSx supports file systems based on the BSD fast file system, distributed file systems based on NFS, and file systems that provide a mechanism for mapping the address space of running processes.

PYRAMID FASTRAC. FastTRAC is a backup solution for large databases and Unix files. It uses a "tape array" architecture to combine multiple devices into a single virtual tape drive with high throughput and large capacity. FastTRAC backups are up to eight times faster than systems without FastTRAC. In effect, FastTRAC stripes the backup across multiple tape devices. Other approaches address the backup capacity requirement but do not address performance. For example, a 4mm drive backs up 1.2 GB per hour. With FastTRAC using eight drives, Pyramid can do 12GB per hour. The approach could be used with almost any device that has serial SCSI support. Pyramid has also integrated Storage Technology tape silo products as a backup server for a data center.

RELIANT FILE SYSTEM (RXFS). Based on the VxFS file systems licensed from Veritas, Reliant File System (RxFS) maintains a transaction log to guarantee the data integrity of every file under RxFS control. It also allows online maintenance of the file system, including defragmentation and file system expansion while the system is in use, and provides fast file system recovery. RxFS also allows online backup and high-performance access for large files.

Pyramid also supports memory-mapped files that allow a file to be mapped into the address space of an application and manipulated as an array. This enables faster application development and more efficient program execution.

STRATEGIC PARTNERSHIPS. In October, Pyramid announced relationships with the systems integrator Integris and IBM hardware lessor Comdisco. These partnerships had a specific focus on data center solutions. Integris's strength in the migration of mainframe CICS applications to Unix will be leveraged in using Pyramid systems as the target platform. The company hopes that Comdisco's expertise in asset management and business relationships with IBM mainframe customers will help generate many leads for Pyramid. Comdisco will also provide disaster recovery services for Pyramid customers as a part of the agreement. In addition to working with Pyramid to migrate CICS applications, Integris will develop, in conjunction with Comdisco, HYPERchannel attachment to IBM mainframes to further facilitate offloading applications onto Pyramid systems while data continues to reside on the IBM host. This innovative solution will allow customers to take advantage of the same IBM security and data integrity features they have been using, while using Pyramid systems for actual application processing.

CA-UNICENTER. CA-Unicenter will be ported to the Pyramid systems, with late 1993 as the target. It may be completed earlier, depending on how the port of CA-Unicenter to Sun's Solaris 2.1 operating system goes, since that environment is the closest to DC/OSx. Going into mainframe shops, CA-Unicenter has a unique position. It provides Federated Network Management of Unix systems side by side with IBM hosts. CA-Unicenter, therefore, offers a significant leverage on the resources that are already in the data center, which is central to Pyramid's strategy.

Sequent Computer Defends Its Turf

Although Sequent Computer's competitors are questioning its commitment to Unix by pointing to its announced support for Microsoft Windows NT, Sequent has no intention of abandoning its core market. Sequent has had success selling into downsizing opportunities and mission-critical applications. It does not, however, propound a strategy of replacing mainframes but of supplementing them as ideal servers for distributed, client/server, business-critical applications. Sequent has been very involved with a group of Unix users concerned with developing standards for large databases on open systems. This group goes

by the name of MOSES (as in "Let My People Go"), which is an acronym for Massive Open System Environment Standards. Sequent is not particularly interested in the massive part of this group's mission, but in large databases, data integrity, and robust, available systems.

SYMMETRY 2000 SYSTEMS. The Sequent product line, called Symmetry 2000, is a series of tightly coupled, symmetrical multiprocessing designs built with Intel microprocessors. The systems support up to 30 i80486 processors and range in price to over $2.5 million. On a high-end Symmetry 2000/750, each CPU has 512KB of cache, and the system as a whole can be configured with up to 1.5GB of error-correcting memory under the 2.0 release of Dynix/ptx. The processors are on a global, synchronous system bus that has a sustained data transfer rate of 53.3MBps. The current maximum storage capacity of a system is 385GB, but that will double to 784GB in the first quarter of 1993 when Sequent implements the SCSI wide specification. Sequent currently uses the SCSI-2 fast interface, which provides a maximum of 10 MBps throughput. The system supports 2GB SCSI-2 disks, 5GB-capacity 8 mm tape drives, and a 3480 tape subsystem. The amount of storage Sequent has supported has been on the low side of the capacity requirements of many large-scale data centers, but it has been more than adequate for Sequent's target markets.

Sequent's operating system, DYNIX/ptx 2.0, is a parallel-enabled version of Unix. It includes support for C2 security, complies with the OSF Application Environment Specification (AES), and has been enhanced to provide online diagnostics.

ENHANCED DATA MANAGEMENT: PTX/CLUSTERS. Clustering is a new capability for Symmetry. It is achieved with a set of related software products that enable Sequent systems to be configured into clusters that provide high availability and scalability while accessing a single, shared database. It offers online disk replacement, shared tape drive support, and support for over 750GB of shared storage. ptx/CLUSTERS supports enterprise-wide consolidation of strategic data and large-scale decision-support applications. It also supports the use of Oracle7 Parallel Server, which is being developed as a part of the joint Very Large Database (VLDB) development between Sequent and Oracle.

QCIC-W HIGH-PERFORMANCE CONTROLLER. QCIC-W is a new SCSI fast and wide controller that supports disk requirements of clustered systems by allowing the use of a variety of devices, including SCSI disks and tapes, fast SCSI-2 disks, and wide SCSI-2 disks concurrently. It also increases the total amount of storage that can be configured on a system. Offering 16 SCSI IDs per channel, the new controller permits more clustered systems per channel without decreasing the number of disks they share.

MAGNETO-OPTICAL LIBRARY SUBSYSTEMS. Sequent's magneto-optical library subsystems combine rewritable optical storage technology with control software to provide fast, straightforward access to large archival databases and file systems. They offer significantly faster data access than offline tape archiving while providing the online capabilities of magnetic disk. Each subsystem stores up to 85GB online. Up to 16 magneto-optical subsystems can be used with each Symmetry 2000 computer, for a total storage capacity of 1,360GB of online information per host. ptx/Jukebox software provides a convenient user interface to the subsystems, allowing use of standard DYNIX/ptx commands and utilities.

PTX/ARGUS. ptx/ARGUS is a new package that provides monitoring, performance analysis, and system management for standalone, networked, or clustered nodes. It provides a graphical display of information and offers intuitive features such as color-keyed alarms for rapid response and decision-making.

ONLINE DIAGNOSTICS. A new feature in DYNIX/ptx 2.0 is Online Diagnostics, the ability to perform a comprehensive analysis of critical system components while the system remains available. A menu-driven, graphical interface enables an operator to fully diagnose

processors, memory, and disks for potential problems. Components can be repaired or replaced during off hours, avoiding costly downtime during peak hours.

PERFORMANCE EVALUATION PACKAGE. Also new to DYNIX/ptx 2.0 is a performance evaluation package that is actually provided by Sequent's Professional Services organization. It provides a detailed system analysis based on information gathered using a set of Sequent-developed performance monitoring tools. The results of the performance evaluation can be used to improve system utilization and as an aid in tuning system performance.

PTX/NQS. ptx/NQS is a new batch execution environment. It provides support for standalone, networked, or clustered nodes. It balances processing loads over groups of systems, controls job priorities, and automatically restarts jobs and queues after network or node failures. ptx/NQS is aligned with POSIX-standard, batch-queuing interfaces. It interoperates with other NQS implementations and can be monitored through ptx/ARGUS.

PTX/BACKUP. ptx/BACKUP is backup management software that supports very large database environments by providing efficient backup and recovery and ensuring high availability for standalone or clustered Symmetry 2000 systems. ptx/BACKUP uses incremental backups that can be quickly performed within a production environment, reducing disruption and providing more recent data for faster recovery. Sequent is currently doing joint development with Oracle to integrate ptx/BACKUP functionality into the Oracle parallel server to be able to offer online and incremental backup. The first phase of this work will improve the backup and recovery performance of very large Oracle servers.

CA-UNICENTER. Sequent will offer the first large-scale symmetric multiprocessing version of Computer Associates' mainframe systems management software. It will be the second platform on which CA-Unicenter ships, and it is scheduled to be available in the third quarter of 1993.

APPC-BASED FILE TRANSFER. The APPC-based file transfer facility is a high-performance connection that operates over Sequent's IBM networking products to enable high-speed file transfers from IBM to Sequent systems.

DISASTER PROTECTION PROGRAM. A Professional Services offering, the Disaster Protection Program is a full suite of Disaster Recovery Programs, including:

- Contingency Planning Service, to assess business risk and plan the means for protection
- Critical Recovery System Protection, to provide fast-recovery, hot-site facilities
- Facilities Program, to provide hardware-ready, cold-site backup facilities
- System Replacement Service, for expedited delivery and installation of replacement equipment.

OPEN SYSTEMS MIGRATION SERVICES. Sequent Professional Services provides comprehensive support in a customer's migration from proprietary technology to a Unix-based open systems architecture or integration of Unix-based open systems into existing environments. Sequent supplies planning, implementation, and training services.

SEQUENT'S POSITIONING. Sequent is interested not so much in replacing mainframes as in using its architecture to supplement mainframes in applications that are designed around more contemporary design centers, such as distributed and client/server applications. Coexistence is a key part of its strategy. Sequent hopes to become a strategic partner with companies migrating to new paradigms of distributed computing.

In some ways, Sequent's product and strategy are not dissimilar to NCR's. However, Sequent has been shipping the large-scale, multiprocessing, Intel-based products much longer than NCR and, therefore, has achieved higher credibility with customers so far.

Computer Associates' CA-Unicenter

While Unix Systems Laboratories may not like the marketing approach of positioning CA-Unicenter as "Making the world safe for Unix" and using a robot super-hero to push the message at trade shows, CA's message is clear. Computer Associates has used its experience in mainframe systems management to bring that functionality to Unix. The concept behind CA-Unicenter is to provide the management capabilities that data center managers had come to expect on mainframes, give them a contemporary graphical user interface, and allow the management of networks of heterogeneous systems. Computer Associates has pledged to migrate CA-Unicenter over time to integrate with OSF's Distributed Management Environment, but it chose to push ahead with development rather than wait for the DME framework and technology to firm up. The company felt the demand was there for its tools and that customers would be willing to make the migration with Computer Associates as the cost of having the tools sooner rather than later.

MANAGING MIXED ENVIRONMENTS. CA-Unicenter is designed to manage mixed Unix environments and mixed Unix and non-Unix environments, including AS/400, VMS, and MVS. CA-Unicenter has the anchor points to participate with all of the environments Computer Associates will support. When the non-Unix Computer Associates management products are enabled in future releases, a data center manager will be able to manage any participating system from any Unicenter console. Security will be the first area to have the heterogeneous Unix capability, followed by storage management. The other areas and other platforms will follow in future releases for each platform.

DISTRIBUTED MANAGEMENT CAPABILITIES. In effect, CA-Unicenter takes advantage of the strength of the Unix distributed environment, giving the data center manager the power of mainframe-class management tools. Although the GUI-based environment of Unicenter may take some getting used to on the part of data center system managers accustomed to character interfaces, the benefits will far outweigh the pain of moving up the learning curve.

The initial platform for CA-Unicenter will be HP-UX, but Computer Associates has announced intentions to port CA-Unicenter to most other commercial Unix offerings, including Solaris 2.1, Pyramid, and Sequent, and OSF/1 can't be far behind. CA-Unicenter functionality is focused on five areas: Security, Control, and Audit (SCA); Automated Storage Management (ASM); Automated Production Control (APC); Performance Management and Accounting; and Data Center Administration. The five areas are unified with a graphical user interface based on Motif, but a command line interface is available as well for administering and monitoring CA-Unicenter.

SECURITY, CONTROL, AND AUDIT (SCA). SCA provides integrated, single-point sign-on coupled with the native security for either a single machine or a network of Unix systems. It provides global enforcement of both user access controls and resource access controls, and sets policy-based security definitions, thereby simplifying security management. SCA manages user registration, user and resource auditing, and monitoring of system integrity. It provides support for enforcing policies for system entry, asset access control, periodic user password change, account suspension, and security for all systems management functions. SCA uses standard Unix administration tools integrated with the security offering to administer user accounts and file access controls. Passwords are kept in a secure, encrypted database instead of a text file. Among the policies it allows to be enforced are preventing users from changing IDs once they've logged on and allowing the Unix superuser authority to be controlled and limited by management.

AUTOMATED STORAGE MANAGEMENT (ASM). ASM addresses the extended data and media management requirements of the complete data cycle, including backup, archive, restore,

transparent restore, recovery, movement, maintenance, retention, and monitoring. It automates virtually every aspect of disk and tape operations. Users can retrieve files they need without help from the systems administrator.

AUTOMATED PRODUCTION CONTROL (APC). APC addresses workload management, including automated workload balancing, batch queue management, console management, and report management. APC organizes and controls production workloads in a flexible manner, adjusting to dynamically changing requirements. It allows the system administrator to specify criteria for job scheduling, including checking whether key preceding jobs have run successfully. Sets of jobs can also be scheduled to run on a wide variety of event triggers. It includes workload policy definition and real-time tracking, as well as workload scheduling. Both calendar- and event-based criteria can be defined in advance to ensure that jobs are run in the right order. This module is central to overcoming Unix's traditional shortcomings in scheduling batch jobs.

The Automated Production Control module includes Console Management, which automates the handling of common console messages from local as well as remote systems. The message/action parameters can alter, suppress, or reply to messages. It can also initiate other actions based on the content, frequency, or certain other characteristics of the messages. The console GUI can be customized to produce site-specific operator console dialogues and can allow consoles to be tailored to specific tasks.

The spool management facility makes it easier to move print files from CPUs to printers and to assess their status and contents. It allows print queues to be viewed, manipulated, and prioritized. Since the spool facility is integrated with the file system, queues can be offloaded and stored on other media, either for archiving, reprinting, or retrieving of archived data.

Report distribution control facilities allow finer-grained control over what is printed and distributed. Different users can get different subsets of reports, and reports can be delivered either in printed form or via E-mail.

PERFORMANCE MANAGEMENT AND ACCOUNTING. The Performance Management and Accounting module reports how system resources are being used, for instance, how memory is being used, which devices are in use, and how processes are running. This is all communicated through graphical windows. In addition to being helpful in tuning system performance, this information can be useful in billing and chargeback functions. Users or departments can be aggregated in straightforward, readable statements. Accounting structures can be customized depending on customer policy.

DATA CENTER ADMINISTRATION. The Data Center Administration module supports help desk functions as well as problem management. The problem manager provides a problem-tracking and -reporting system that can be customized by the user. It facilitates communication so that user requests can be responded to on a timely basis. Problems can be entered into the system manually or automatically when exception events occur. Automatic priority escalation can be set up, and a help desk facility is provided that allows the user to track progress toward solving problems and track vendors' responsiveness in solving problems.

While CA-Unicenter is likely to be an important product, both for the vendors of the systems it supports as well for its customers, a major challenge to Computer Associates will be to track developments in the platforms it is supporting with the product and in the management frameworks and standards which it has committed to track.

Summary and Conclusions

Is the Unix data center fact or fiction? Depending on the stringency of a customer's requirements, it may be more fiction than fact today. But the important point is that the technologies, tools, and standards are in the pipeline that will bring most Unix platforms to parity with IBM mainframe environments within two to three years. Considering the amount of planning and development that any company faces if it decides to migrate from the mainframe, two to three years is within the planning horizon. By the time a user is ready to bring a Unix-based data center into production, all of the capabilities that exist today on MVS should be available on Unix.

Multiple Options Now Exist

It is surprising how many options there are now for running a corporate data center with Unix. Running a Unix-based data center was virtually impossible just a few years ago, and now it is one of the more significant trends in the industry. The really amazing aspect of this transformation is that vendors are accomplishing it with virtually no technology from USL. Nearly all vendors selling Unix for data center-class applications make a big point of how much work they have done rewriting extensive parts of Unix in order to make it data center ready.

Open Systems and the Data Center

The benefits of open systems become really clear when we examine data center solutions. If it weren't for standard public interfaces, the suppliers of the hardware and software technologies that vendors are using to build these highly reliable high-performance systems would have had to build a lot of the technology themselves. This would have added significant cost to their products and led to drastically slower times to market as well.

Although this article was about Unix in the data center, much of the value of bringing Unix into the data center lies in the fact that doing so brings open systems into the data center as well. When that happens, the cost of building heterogeneous, interoperable information systems declines and the speed and flexibility increase dramatically. The other dimension of bringing Unix and open systems into the data center is that it begins to open up the proverbial glass house to the future world of open distributed computing. There is no reason why data and applications everywhere should not enjoy the same degree of integrity, security, and availability as they do when protected by the safe confines of the mainframe.

Implications for USL and OSF

USL and OSF are actively engaged in trying to deliver many of the capabilities required for Unix/Open Systems data centers. The fact is, however, that in the absence of technologies from these two organizations, vendors have built the technologies on their own to meet customer requirements today. Customers are not willing to wait for either OSF's DME or USL's Distributed Management. Vendors will have to continue to deliver working technologies to their customers while weighing the value of migrating those technologies into industry compliance. Customers' primary requirement is for products that work, and then for products that work together. Getting data center management products to work together will have to occur under the umbrella of DME's management framework. However, the process of migrating the products that are on the market today into DME compliance will take at several years at least.

Further ahead, OSF and USL will have to work with data center management requirements as a design center for future technologies, starting with operating systems and extending to the full breadth of management and distributed computing enablers. The reason for this is simple: Customers with valuable, strategic data will accept nothing less, particularly if there are non-Unix alternatives that will provide it, such as POSIX-compliant versions of DEC's VMS, HP's MPE, and others.

Strategies for Proceeding

Customers are in the process today of evaluating strategies and steps for downsizing, rightsizing, upsizing, resizing, or whatever one wishes to call mainframe replacement

Summary and Conclusions

strategies. There are several available strategies that can be combined depending on individual environments:

- **Recompile applications to Unix.** If there is a compiler available on Unix that is compatible with the compiler used in the MVS environment, this is a very inexpensive way to make the transition. Even if the code has to be tweaked to run on Unix, that process will likely result in more maintainable, architecturally neutral code, further benefitting future migrations. Although some MVS C code exists, Cobol and Fortran are the most common and likely candidates.

- **Build new applications on Unix.** Using the distributed development, deployment, and support environment that Unix provides, design and build next-generation, mission-critical applications on Unix. Keep in mind that these applications will probably need to interface with mainframe applications, but also keep in mind that connectivity options are available.

- **Re-engineer applications on Unix.** Applications that have outlived their useful life and need to be rebuilt are good candidates for moving over to Unix platforms. The support provided by Unix for development, combined with the powerful development tools available from third parties, will make the re-engineering job proceed much more quickly than in an MVS environment.

- **Port DBMS applications to Unix.** If your strategic DBMS vendor has moved the DBMS to Unix, and either the vendor's 4GL or a third party's 4GL is available, the port will be straightforward and cost effective. Be careful, however, when specifying requirements for backup/restore, data integrity, and system availablity.

Concrete Steps

There are four steps users can take in order to proceed with a strategy of introducing Unix into the data center:

1. Identify which applications that are currently running on MVS could be easily ported to other environments. This could be either because the language the applications are written in has a compatible version on Unix, i.e. Cobol; the database management system and 4GL they are written in is available on Unix, or the ISV provides the application on Unix, e.g. D&B Millenium.

2. Determine the capacity requirements for those applications in terms of users supported, memory and disk storage, and performance requirements.

3. Assign a migration value to each application, based on the cost of migrating (negative), the risk of migrating (negative), the cost of upgrades for the next five years (positive), the differential cost of software licensing over five years (positive), the differential cost of service and support over five years (positive), and other industry-specific benefits.

4. Rank the applications according to risk, and select the application with the lowest risk and the highest value. Prepare an RFP for vendors to respond to that includes the cost of software migration and full support through the migration process.

Conclusion

It is possible to run a data center with Unix today, providing requirements are clearly spelled out in advance. Very large applications or very large numbers of users requiring sub-second response times may still require a mainframe-class machine. However, that machine could still be running Unix. As long as requirements are spelled out and vendor capabilities are matched to them, most unexpected "gotchas" can at least be avoided. We believe that within three years, the Unix environment will have achieved parity with the MVS environment, and even the most cautious customers will begin to seriously reconsider their most sacrosanct proprietary data center environments.

Companies Mentioned In This Article:

AIM Technology
4699 Old Ironsides Drive
Santa Clara, CA 95054 408 748-8649

Comdisco Disaster Recovery Services
6111 North River Road
Rosemont, IL 60018 708 698-3000

Demax
999 Baker Way
San Mateo, CA 94404 415 341-9017

General Atomics
P.O. Box 85608
San Diego, CA 92186-9784 619 455-3753

Epoch Systems, Inc.
8 Technology Drive
Westborough, MA 01581 800 873-7624

Legent Corporation
Two Allegheny Center
Pittsburgh, PA 15212 412 323-2600

Network Systems Corporation
7600 Boone Avenue North
Minneapolis, MN 55428 612 424-4888

Raxco Incorporated
2440 Research Blvd.
Rockville, MD 20650 301 258-2620

SCH
Three Centennial Plaza
895 Central Avenue
Cincinatti, OH 45202 513 579-0455

Sceptre Corporation
3626 W. Liberty Rd.
Ann Arbor, MI 48103 313 665-8778

SecureWare Incorporated
430 10th St. #N-210
Atlanta, GA 30318 404 876-4840

Storage Technology Corporation
2270 South 88th Street
Louisville, CO 80028 303 673-5628

SyncSort Incorporated
50 Tice Boulevard
Woodcliffe Lake, NJ 07675 800-541-1724

Systems Center Incorporated
1800 Alexander Bell Drive
Reston, VA 22091 703 264-8000

Unison Software
675 Almanor Avenue
Sunnyvale, CA 94086 408 245-3000

Next month's *Open Information Systems* will address
X/Open and the XPG4.

For reprint information on articles appearing in this issue,
please contact Donald Baillargeon at (617) 742-5200, extension 117.

Open Systems: Analysis, Issues, & Opinions

FOCUS: DATABASE ACCESS

Borland Takes the Wraps off Integrated Database API (IDAPI)

In Pursuit of Database Interoperability

Borland International (Scotts Valley, California) has teamed up with IBM, Novell, and WordPerfect to announce an application programming interface (API) that allows client applications to access data from nonrelational as well as SQL-based relational databases. Like Microsoft's Open Database Connectivity (ODBC) API, Borland's Integrated Database API (IDAPI) implements the SQL Access Group (SAG) call-level interface (CLI) specification for SQL-based access to relational databases. However, in addition to accessing the SAG CLI, IDAPI provides an API for accessing ISAM databases, called the Navigational Call Level Interface Client API (NAV/CLI for short). "Navigational" is a term that the group has coined for ISAM databases such as dBase, Paradox, FoxPro, etc.

The IDAPI group will first work to complete the IDAPI specification over the course of 1993 and then will implement it on DOS, OS/2, and NetWare, in addition to Windows. Unix support is likely to follow in short order. In the first half of 1993, Borland will provide Windows and DOS software developer kits (SDKs). IBM will provide SDKs for OS/2 and AIX. Novell will provide support for NetWare and will be working on associated networking technology. WordPerfect will deliver an IDAPI-based InForm application. However, since the NAV/CLI specification will not be completed until then, NAV/CLI will not be included in those SDKs. NAV/CLI-compliant IDAPI SDKs will be delivered in the second half of 1993. The specific form of the other implementations has not been announced as yet.

What Is IDAPI?

IDAPI's client API provides a simplified method for developers to provide users with access to heterogeneous data sources. "Heterogeneity," in this instance, means not just different brands of SQL-based relational databases (RDBMSs), but nonrelational as well. It is designed to allow an application to be written to the SQL CLI, to the NAV/CLI, or to a combination of both. With the appropriate client and server database engines and drivers installed, an application should be able to transparently and simultaneously access and update data from both relational and nonrelational data sources.

IDAPI extends the work of SQL Access Group and other organizations working on data access specifications and standards into the realm of non-SQL and nonrelational databases. This is significant because a massive amount of data are currently being stored in various nonrelational data managers, ranging from IBM's IMS to dBase III and IV. Some estimates suggest that as much as 80 percent of all data stored in databases are in nonrelational databases.

The IDAPI group plans to work within the context of the relevant SAG, X/Open, and ANSI data access standards groups. It is not clear, however, exactly how the IDAPI specification would be incorporated into the broader realm of SQL-based industry standards for database access. Such an effort will be critical if IDAPI is going to have a chance of widespread adoption and support. Otherwise, it will seem to be nothing more than a self-serving specification.

Client APIs and the SAG

Until now, the SAG has been spearheading the effort to develop client-side database APIs. It has defined two common APIs for accessing remote data from desktop clients: an embedded SQL interface, which is more relevant for 4GLs, and a call-level interface (CLI). Both apply ANSI SQL as a lingua franca for accessing databases. Both Microsoft and Borland have been tracking the work of SAG very closely and have included the CLI in their architectures. The difference lies in the availability of ODBC both as a specification from Microsoft and as an actual product in the form of at least one ODBC driver.

The value of standard client APIs can be understood by looking at Illustration 1. The need for each application to have its own driver for access to each database is

inefficient and makes desktop and transport configuration extremely difficult.

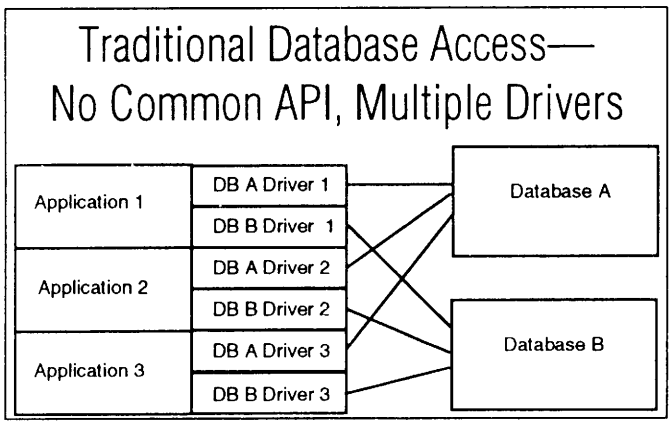

Illustration 1. Each application has to have its own driver for any database it wants to access. This is obviously inefficient and makes desktop and transport configuration extremely difficult.

In addition to having to support each database's unique access method, each application has to be able to contend with different transports for each database. AccessWorks, a proprietary approach to this problem from Digital Equipment Corporation, which was discussed in the June *Open Information Systems* (Vol. 7, No. 6), used Digital's SQL Services as a common client API and DECnet as the common transport as one way of simplifying matters. The value of a common client API is depicted in Illustration 2, which shows how a common API allows multiple applications to share a common set of drivers.

Implementing the SAG CLI

Attention has been drawn to Microsoft's ODBC, mostly because it is the first commercial implementation of the SAG's work. Microsoft chose to use the SAG CLI to solve the problem of providing a common API for database access within the Windows API, which can be used by any Windows application. ODBC is Microsoft's productization of the SAG CLI. There is no reason why other vendors cannot implement the SAG CLI in ways that meet their own product requirements as well. That is where IDAPI comes into the picture.

IDAPI and Its Architecture

The Client API — Application Support Layer

The IDAPI architecture, shown in Illustration 3, is designed to be extensible by adding new drivers, relational and navigational, and by distributing the data model through the requester/responder, independent of the details of implementation.

At the highest level of the IDAPI architecture are two APIs, the SQL call-level interface API (SQL/CLI API) and the Navigational call level interface API (NAV/CLI API). Client applications can be written to one or the other or even both of these APIs. If written to the SQL/CLI API, the application is essentially an SQL application. If written to both APIs, it is considered a mixed application. These APIs constitute the Application Support Layer—an object-oriented abstraction of the relevant database model access methods, whether relational or navigational. This layer comprises a set of functions that manage the state and environment of client applications. The Application Support Layer takes requests for data from the application, parses the queries, and passes the requests on to the appropriate engine, relational or navigational.

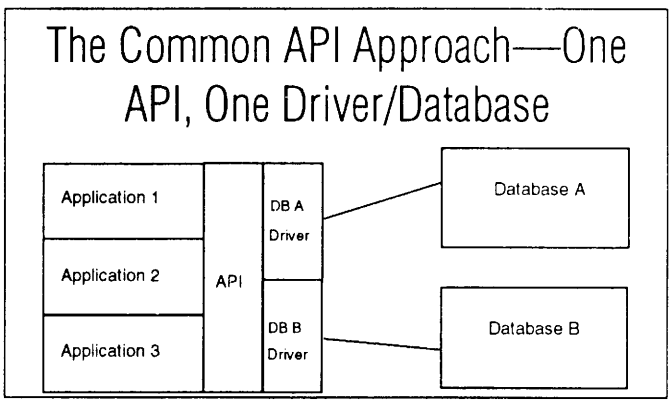

Illustration 2. With a common API, applications are written to a single programmatic interface and can share a driver installed for each database.

This layer handles the initialization of the application and the loading of the appropriate engines and drivers dynamically as required by the client application. It hides all the semantics of the underlying engine and driver from the client application.

Engine Support Layer

At the heart of the IDAPI architecture is the Engine Support Layer, comprising Relational and Navigational Engines and Requestor/Responders. Anyone can provide an engine in this layer, and the IDAPI specification will include different compliance levels that engines can meet. Borland will provide an engine that is essentially the Borland-InterBase Local Engine, on which Borland's desktop data access strategy is built.

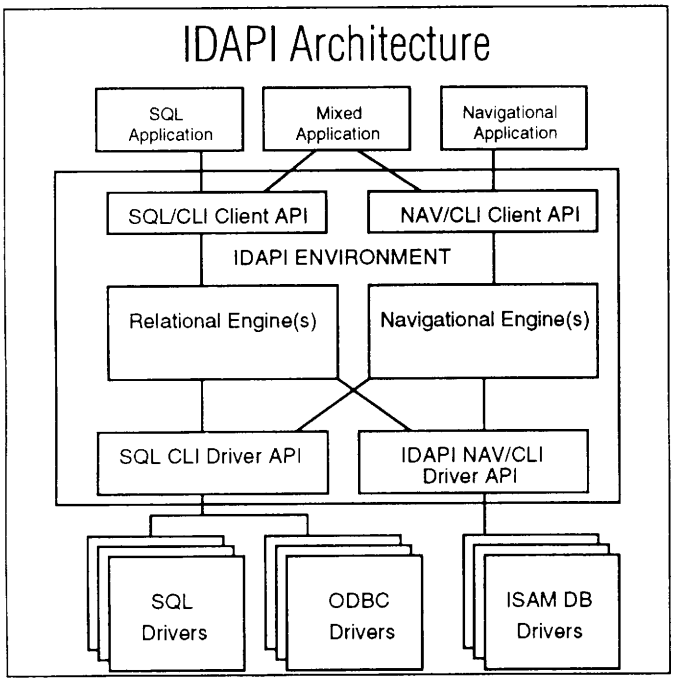

Illustration 3. The IDAPI architecture has multiple components in order to map back and forth across the different data models supported.

Driver Support Layer

The Driver Support Layer provides a consistent API between IDAPI drivers and the Engine Support Layer. This is analogous to the Service Provider interface in the Microsoft WOSA architecture and to the SAG CLI Driver API in the SAG model. There will be analogous APIs within this layer for both SQL and Navigational databases.

Driver Layer

The IDAPI Driver Layer consists of drivers for specific databases that will be provided by the database vendors or by third parties. An IDAPI software developer's kit that will support driver development will be made available. This toolkit will provide services for isolating operating system dependencies as well as for providing services for internationalization and localization.

Relationship among Client APIs

Microsoft's ODBC and SAG CLI

The Open Database Connectivity (ODBC) specification that Microsoft has published is a part of the Windows API. A developer writing a Windows application can write to the ODBC API and automatically have access to any database for which an ODBC driver has been installed on the client machine. It also means that ODBC suffers all of the limitations of the Windows environment as well, such as single threading and lack of memory protection. In the future, that will change somewhat with Windows NT.

The biggest difference between Microsoft's ODBC and SAG's interface is the protocol each uses to transmit requests from clients to servers and results from servers to clients. SAG specifies the International Standard Organization's (ISO's) Remote Data Access (RDA) protocol for this purpose. Microsoft has its own protocol, which is designed to be smaller and less processor-intensive than RDA.

In addition, SAG's interface specifies TCP/IP as the network transport. TCP/IP is important to Microsoft strategically, but the majority of Windows PCs are not using it yet. Therefore Microsoft's ODBC must support the Named Pipes networking interface, which uses NetBIOS as its underlying transport protocol.

Microsoft has already provided an ODBC driver for the Microsoft and Sybase SQL Server products, and Microsoft's new client database product, Access, supports ODBC. As much as Borland likes to criticize Microsoft, in this instance, Microsoft has gotten out early with standards support.

IDAPI and SAG CLI

In the past, Borland had said it that it planned to use the SAG call-level API in IDAPI. However, the SAG API is too limited for Borland's purposes, and that is why IDAPI is a superset of the SAG CLI. Borland is the leading vendor of PC databases, and it needs a desktop API that allows applications to address both remote SQL databases and record-oriented PC databases, or what it is now calling navigational databases. Thus, Borland has described IDAPI as having the ability to process both SQL requests and record-oriented requests from a common set of calls.

In addition, Borland is likely to provide support for binary large objects (BLOBs) and other advanced RDBMS features in the final IDAPI specification. BLOBs are an important technique for storing complex data within relational table structures. As such, BLOBs are key to Borland's support of object-oriented approaches to applications within its database strategy, and they are not currently supported in the SAG CLI.

IDAPI and ODBC

The major difference between the Borland and Microsoft APIs is the role each assigns to SQL. To Borland, SQL is one way to access data; to Microsoft, it is the

only way. All ODBC calls are transformed to SQL, even calls to non-SQL databases, such as IBM's VSAM and Borland's dBase. IDAPI promises to provide both an SQL interface to distributed data and a record-oriented interface more appropriate to data stored in its dBase products.

Also, IDAPI supports access across databases, while ODBC does not support heterogeneous database operations but limits operations to databases of like kind.

IDAPI can be run either in client/server or client/only mode. This is relevant for applications accessing existing dBase and Paradox data and represents a very PC-centric requirement. ODBC assumes a client/server architecture. The IDAPI group plans to have IDAPI on multiple platforms, providing numerous application support options.

Borland is also pledging to offer access to a greater range of database functionality through IDAPI than Microsoft does in ODBC. For example, ODBC does not include access to stored procedures, database programs, or BLOBs. To its credit, however, Microsoft implemented scrollable cursors in ODBC, making it possible to emulate this feature in Sybase/Microsoft SQL Server versions that didn't implement it. Scrollable cursors are a handy feature in a record-oriented environment.

The Ball Is in Borland's Court

Borland, owner of two of the most popular PC database products, Paradox and dBase, has known for some time that it could not allow those products to be made obsolete by SQL-based relational databases. The emergence of standard APIs from SAG and Microsoft threatened to do just that. IDAPI grew out of ODAPI, Borland's original effort at developing a proprietary database API to stay in the game with its desktop databases while leveraging its other database offering, InterBase. Borland opened up its work to its three partners—IBM, Novell, and WordPerfect—and now all four companies are working together on the IDAPI specification. However, the extended time frames and lack of firm details about the IDAPI specification make it hard to take Borland completely seriously as an immediate contender. On the other hand, Borland has lined up an impressive list of supporters in the industry in addition to those participating in the announcement, including Oracle Corporation, Computer Associates, SAS Institute, Cognos, Gupta, Banyan Systems, Frame Technology, Pioneer Software, and others. If these supporters are willing do more than stand on the sidelines and cheer, the Borland-led IDAPI group may be able to lead the industry to developing a nonrelational access standard.

In the long run, data access standards will have to incorporate standard ways of accessing nonrelational data. The development of IDAPI will help to spark interest in extending the work of SAG into nonrelational databases, and SAG is likely to endorse IDAPI's work in that area. The specifics may need to change, however, particularly since there are no specifications yet for anyone to examine. Once there are, the potential for integrating IDAPI with SAG CLI can be more fully determined. *—M. Goulde*

On Your Mark, Get Set, Go! ... to

Take off with the Boston Marathon runners and cross the finish line better prepared to deploy open distributed computing throughout your organization.

Enter Marathon Week and Cross-Participate in 3 Simultaneous Forums!

- **Distributed Object Computing Technology Forum**—*How Can I Use Distributed Object Computing Technology Today to Build Mission-Critical Applications?*
- **Open Systems Forum**—*What Can I Do with Open Systems Today? What Lies Ahead?*
- **Seybold Executive Forum**—*How Do I Design Real Business Solutions That Implements Open Distributed Systems?*

April 18 - April 23, 1993
Copley Marriott,
Boston, Massachusetts

As the technology boundaries begin to merge, so must our Forums. For the first time, Patricia Seybold Group will hold its three annual conferences all in the same week at the same hotel. Our goals are:

- To allow participants to focus on the topics that best meet their business needs. Participants may cross over to attend any sessions of interest.
- To enlarge the group of people with whom you can meet and share experiences.
- To explore how distributed object computing enables open systems.
- To allow participants to enhance their conference learning through business case-focused, hands-on applications of technology.

In addition to the three conferences, participants can visit the New Technology Spotlight, where leading vendors will demonstrate real business solutions.

For more information, call Deb Hay at (800) 826-2424

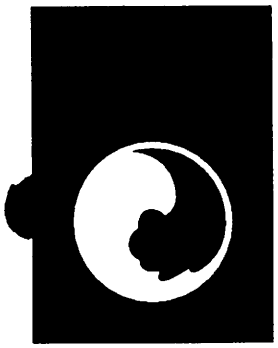

Special Research Reports

from the Patricia Seybold Group

Patricia Seybold's Special Research Reports give you the detail and depth of books ***and*** *the timeliness of magazines.*

New Format!

Unix Relational Database Management *Third Edition*

Vendor Strategies, DBMSs and Applications Development Tools

By Judith R. Davis

Full Report, Third Edition, Available Spring, 1993	$595
Full Report plus Fall and Winter Comparison Matrices	$750
One Comparison Matrix of Your Choice	$200

Vital Information

Complete Details

When You Need It

The third version of this well-respected and popular report continues to cover the marketing strategies of the major Unix relational database vendors: Ask/Ingres, Informix, Oracle, Sybase, Progress, and rising contender Borland/Interbase. Now for the first time, Patricia Seybold Group gives readers the ability to receive the information in this report in a more flexible and timely manner to suit a variety of research needs. Readers can still receive the full report in the spring of 1993, which will include the comparison matrix, and they can now also buy the comparison matrix unbundled in the fall, winter, and spring as we continually update the chart information based on vendor product updates. The chart can be an invaluable ongoing research tool to those readers consistently following this market.

The Full Report—Spring 1993. This includes an in-depth analysis of corporate strategy, product strategy, and positioning for each vendor, a review of product architecture for both the database engine and the application development environment, and the most current product comparison matrix.

Product Comparison Matrix. The report has always included an additional section which is an exhaustive comparison matrix of the features, functions, and architectural characteristics of each product. The matrix includes such categories as file structure, database parameters, user interface capabilities, data types, indexing, forms, architecture type, SQL statements, report writers, application development tools, and more. Readers can now follow the progress of the database products as the vendors make enhancements, via periodic updates of the report's detailed matrix. The next chart update is ready this fall, with updates to follow in the winter and spring.

The Report, the Matrix, or Both. Readers have the option to buy the individual matrices this fall, winter, and/or in the spring. Some readers may choose to buy the fall and winter matrices and also the full report in the spring of 1993. Other readers may choose to wait until the spring of 1993, to buy only the full report which will include the spring matrix. Choose the option that aligns with your database research requirements.

For More Information, Fax (617) 742-1028 or Call (617) 742-5200.

Patricia Seybold's Computer Industry Reports

ORDER FORM

Please start my subscription to:

		U.S.A.	Canada	Foreign
☐ *Patricia Seybold's Office Computing Report*	12 issues per year	$385	$397	$409
☐ *Patricia Seybold's Open Information Systems*	12 issues per year	$495	$507	$519
☐ *Patricia Seybold's Distributed Computing Monitor*	12 issues per year	$495	$507	$519
☐ *Paradigm Shift—Patricia Seybold's Guide to the Information Revolution*	6 issues & tapes per year	$395	$407	$419
☐ *Paradigm Shift—Patricia Seybold's Guide to the Information Revolution*	6 issues per year	$295	$307	$319

Please send me ☐ *Distributed Computing Monitor* ☐ *Open Information Systems*
a sample of: ☐ *Office Computing Report* ☐ *Paradigm Shift—Patricia Seybold's Guide to the Information Revolution*

Please send me information on: ☐ Consulting ☐ Special Reports ☐ Conferences

☐ **My check for $_____ is enclosed.** ☐ **Please bill me.** ☐ **Please charge my subscription to: Mastercard/Visa/American Express** (circle one)

Name: _____ Title: _____
Company Name: _____ Dept.: _____
Address: _____
City, State, Zip code, Country: _____
Fax No.: _____ Bus. Tel. No.: _____

Card #: _____
Exp. Date: _____
Signature: _____

Checks from Canada and elsewhere outside the United States should be made payable in U.S. dollars. You may transfer funds directly to our bank: Shawmut Bank of Boston, State Street Branch, Boston, MA 02109, into the account of Patricia Seybold Group, account number 20-093-118-6. Please be sure to identify the name of the subscriber and nature of the order if funds are transferred bank-to-bank.

Send to: Patricia Seybold Group: 148 State Street, Boston MA 02109; FAX: 1-617-742-1028; MCI Mail: PSOCG
To order by phone: call (617) 742-5200

IOI-1292

Topics covered in Patricia Seybold's Computer Industry Reports in 1991 & 1992:

Back Issues are available, call (617) 742-5200 for more information.

Office Computing Report

1991—Volume 14

#	Date	Title
12	Dec.	IBM/Lotus Relationship—Building a Platform for Communicating Applications

1992—Volume 15

#	Date	Title
1	Jan.	The Groupware Phenomenon—Does It Focus on the Proper Issues?
2	Feb.	Digital's TeamLinks—A Renewed Focus on the Client Desktop
3	Mar.	Requirements for Workflow—What Should We Expect from the Vendors?
4	Apr.	Desktop Multimedia—Moving beyond the Chicken and the Egg
5	May	Borland International—A Database-Centric, Object-Oriented Approach to Desktop Integration
6	June	Apple's Macintosh—Can It Become "the Cadillac of Collaboration"?
7	July	Business Intelligence—A Framework for Data Analysis Applications
8	Aug.	The Quest for Common Mail APIs—Clearing up the Confusion
9	Sept.	BeyondMail for Windows—Epitomizing the Mail-Enabled Application
10	Oct.	Microsoft's Workgroup Strategy—Moving Group Functionality into Windows
11	Nov	Visual Programming—Application Design for End Users and Professional Developers

UNIX in the Office

1992—Volume 2

#	Date	Title
1	Jan.	Downsizing with Open Systems—Can Unix Symmetric Multiprocessing Systems Meet MIS Requirements?
2	Feb.	System V.4 and OSF/1—Matching up in the Marketplace
3	Mar.	Europe's Harness Project—Integrated Technology for an Open, Object-Oriented, Distributed Applications Platform
4	Apr.	The X Window System—Where is Its Future?
5	May	HP's Master Plan—Winning Is Everything in Palo Alto
6	June	Digital's DECworld Gems—Alpha and Accessworks Shine

Open Information Systems

7	July	Integrating Applications in the Real World—Evolution, Not Revolution
8	Aug.	Windows NT 3.1—Microsoft's Bid for Desktop Dominance
9	Sept.	Oracle's Version 7—Can It Leapfrog the Competition?
10	Oct.	Galaxy from Visix—Application Portability Breakthrough?
11	Nov.	European Open Systems Architectures—Europe's Vendors Strike out for Open Distributed Systems

Network Monitor

1992—Volume 7

#	Date	Title
3	Mar.	Smart Hubs—Establishing a Manageable Internet Foundation for Distributed Computing
4	Apr.	Message Express—A Message Platform for Cooperative Processing
5	May	Novell NetWare 4.0—Building toward an Enterprise Distributed Object Computing Environment
6	June	Distributed Printing—Major New Approaches Begin to Relieve One of Distributed Computing's Most Frustrating Problems

Distributed Computing Monitor

7	July	The New E-Mail APIs—Finally, Real Progress toward Mail-Enabled Applications
8	Aug.	Distributed Object Computing—The Merger of Distributed Computing and Object-Orientation into a New Architecture
9	Sept.	Gradient DCE for PCs—PC-DCE Spurs OSF to Make PCs Peers to Unix and Other More Powerful Platforms
10	Oct.	Database Interoperability—A Comprehensive Approach to Database Access for the 1990s
11	Nov	Transarc Encina—Will Distributed OLTP Systems Overrun the Mainframe's Last Stronghold?