DRAFT

JANET DAVID -
THANKS FOR YOUR HELP TO DAVID STEWART. I THOUGHT YOU MIGHT BE INTERESTED IN WHAT THE (NEARLY) FINAL PRODUCT WOULD LOOK LIKE.
ROBERT SHELTON

Tandem and Telecommunications

A Natural Partnership

September, 1991

DRAFT

Tandem and Telecommunications

Introduction

The telephone is an indispensable part of the modern age. It is woven into the fabric of daily life for most residents of the world's industrial nations. The availability and reliability of telephone service have made it an essential element of business. Deals are proposed, meetings arranged, contracts negotiated, and orders committed via the telephone. Users expect to be able to lift up a receiver, hear dial tone, dial a number and either be connected to the called party or receive a busy tone, 24 hours a day 7 days a week.

As telephone service has improved, expectations for service have increased. One of the ways public switched networks have responded to increased demand is with innovative technology. Telecommunication providers (telcos) have long recognized that moving from analogue to digital technology allows increased traffic while reducing maintenance costs. Recently, customer expectations have been raised by competition. For example, in many countries telephone bills contain no individual call detail information. With the introduction of competition, these same telcos are now faced with providing detailed billing.

The emergence of a global economy has only increased competition among telcos. Even those network providers which are protected by monopolies feel the pinch of competition as multinational companies use by-pass technologies to search for the best deal in telecommunication services. This has driven telcos to emphasize the values of service, productivity, flexibility, and innovation. In this quest, public switched network providers have harnessed computer technology to serve virtually every aspect of their enterprise. This in turn has caused telcos to search for computer systems that demonstrate the same levels of reliability and availability.

Tandem computers is a natural partner with telcos because:

1) The architectures of the telephone network and Tandem computer systems are similar in design concepts and operational standards;

2) The types of functions that Tandem computer systems perform are analogous to the types of functions that telephone networks perform;

3) Tandem is committed to, and is a world leader in the delivery of continuously available online computer systems.

Tandem computers are currently deployed by many telecommunications providers, including the top 29 telcos and PTTs in the world. These telcos have found that Tandem computer systems provide continuously available, and reliable processing for their most critical business functions.

Telephone network architecture and Tandem computer systems

The parallels between the design of the public switched network (PSN) and Tandem computer systems are remarkable. Both support real time distributed processes. The availability and reliability of the public switched network has long been a hallmark of telephone service. Tandem computers systems deliver continuously available service 7 days a week 24 hours a day (7 X 24).

Tandem computer systems achieve the availability and reliability standards of telcos by using sophisticated hardware and software technology. These systems are designed for continuous availability and prevent a single component failure from disrupting operations. Tandem design is comparable to the engineering of telecommunication networks, that rely on redundancy and automatic reconfiguration in response to failure on the network. Parallel processing is at the heart of Tandem computer's fault tolerant design. Tandem's use of parallel processing ensures efficient computing and thus the best price performance for machines in its class.

As telecommunications providers enter the information processing arena a key competitive factor will be database capability. Tandem architecture assures the integrity and security of the data processed. Tandem offers NonStop SQL, a high performance industry standard relational database management system. NonStop SQL is the only commercially marketed, distributed relational database that provides continuous availability. For telcos with applications such as calling card and 800 and 900 services this type of database capability is essential.

Tandem has demonstrated a commitment to open standards-based systems, and participates in a variety of standard's groups including CCITT and ANSI. Supported standards include, Open Systems Interconnection (OSI), the UNIX System V operating system, Signalling System #7, and Integrated Services Digital Network (ISDN).

Since Tandem computer systems match the architecture and operational standards of the public switched network, Tandem computers are the ideal platforms for telco applications.

Tandem computer systems and the telephone networks perform similar functions

The architectural resemblance is not surprising when you consider the similarity of the functions that the two types of networks support. Both the PSN and Tandem computer systems provide an environment where customers send and receive information by electrical means in real time. For the purposes of this paper this type of design will be called a message based architecture. The primary elements of this design are the sender and receiver of information. Terms such as "originating office" and "terminating office" ,or "calling party" and "called party" reveal the message based nature of telecommunications network.

The type of processing involved in traditional computer systems are very different from the types of processing involved in telephone networks. Initially, commercial computers were stand-alone machines that performed batch functions such as check processing, inventory control, or complex numerical calculations. Communication between these computers was usually limited to the transfer of data files. Surprisingly, most general purpose computers are still using this architecture. They are not able, therefore, to take advantage of the availability and reliability of messaged based architecture. Stored program telephone switches, on the other hand, were distributed processors that supported the sending and receiving of electronic signals (messages) in real time.

Since its inception in 1974, Tandem Computers applied the telephone system model to the design of commercial computer systems. The results were general purpose, multi-processor computers using a message based operating system with a fault tolerant architecture. Tandem computer systems are composed of one or more machines that

support the sending and receiving of information in real time. These machines can be distributed over great distances and connected in a network. The term generally used for this type of computing is online transaction processing (OLTP).

OLTP has a wide range of applications in business from securities trading to health maintenance. Telcos may justly lay claim to the invention of OLTP,with the development of call processing and common channel signalling. Tandem's unique contribution, however, was to recognize that a messaged based architecture in a general purpose computer could provide solutions to a wide variety of business problems. Since the introduction of its first fault tolerant system, Tandem has provided computer systems that provide continuously available, cost effective, innovative and flexible solutions to mission critical applications.

Tandem systems are deployed in telcos worldwide

Probably no single technological trend is having greater impact on the telcos than the intersection of telecommunications and computer technology. The convergence of telecommunications and computers is driving the use of a messaged-based architecture from the central office into all areas of the telecommunications business. Telephone companies are finding that the application of messaged-based architecture can provide a significant advantage. It provides increased levels of services and operational capability that telcos can offer their customers. For these reasons the top 29 telcos and PTTs have implemented Tandem computer systems.

Tandem Computers have successfully integrated messaged-based systems in a variety of telecommunication venues outside the traditional basic telephone service arena. These include the intelligent network, business support systems, value added services, and network operations support systems. In order to meet the 7 X 24 availability and reliability standards of the public switched network, the elements that make up the network must operate at the same level. Computer systems that are used to offer telco services must measure up to the 7 X 24 standard. Tandem computer systems are designed for continuous availability 7 days a week 24 hours a day.

Besides the 7 X 24 reliability and availability standard that they provide, Tandem computer systems appeal to telcos because they can offer significant improvement in price/performance, integration, and real-time responsiveness. This improves telco service levels. The linear expandability of Tandem systems offer a significant competitive advantage when introducing services, because the cost of initial deployment is significantly lower. This means that a telco can buy just the processing power needed to introduce a service and expand the system easily and cost effectively as the market for the service grows. Linear expandability means that as processors are added to a Tandem system, or systems are added to form larger and larger networks, the processing power of Tandem systems grows linearly. That is, four processors do twice the work of two, and eight processors do four times the work of a two-processor system.

Tandem computer provide responsive, productive, flexible, innovative service oriented systems

The use of computers in telecommunications has been increasing since the introduction of the stored program controlled switch in the early 1960s. An early factor in automation was the growth of traffic on the telephone network. In the United States, for example, market forecasts predicted that without automation the number of operators

required to handle the growing volume of calls would have exceeded the U.S. population by the early 1980s. Another reason for the increasing use of computers in telephone networks was the desire to provide as high a level of service as possible, while keeping the cost of providing that service at a minimum. The increasing availability and decreasing cost of computer components has spurred the convergence of telecommunications and computer technology.

The emphasis on providing a superior service at lower cost has become increasingly important to the network provider. The changing face of government regulation and the restructuring of telcos from monopolistic utilities to competitive business entities plays a role in this development. Bandwidth, the major commodity of the telco, is becoming more abundant and less and less expensive thanks to fiber optics technology. Because of competition, the decrease in cost to the carrier must be passed onto the customer. As price becomes less of a differentiating factor, competition will come down to service, productivity, flexibility, and innovation in each part of the telecommunications industry. Tandem Computers Incorporated provide systems that support these business requirements.

Customer Service

The increased use of high technology in all areas of the economy have caused customers to become more demanding of the services they purchase. A service is judged by its price, reliability, availability, and responsiveness. Besides the continuous reliability and availability that they provide, Tandem computer systems offer significant advantages in terms of price/performance and responsiveness.

The linear expandability of Tandem systems offer a significant price/performance edge when introducing services. A telco need only invest in a Tandem system large enough to support the initial offering of a new service. It can then offer that service at a price which is competitive and yet achieve a return on the investment faster. As the market for that service grows, Tandem architecture allows expansion of the system easily to meet increased demand and without interruption to the service.

Tandem's messaged based architecture enables telcos to implement responsive applications that are online and perform in real-time. The high-volume transaction processing capability of Tandem systems allows applications which can be tailored to customer needs while keeping pace with peak activity. The Tandem Call Application Manager (CAM) utilizes these capabilities to support call centers.

Calling centers are often the first point of customer contact and thus are critical elements in a company's image and marketing strategy. The Call Application Manager provides a software development platform for call center applications on Tandem NonStop™ systems linked to private branch exchanges (PBXs) or automatic call distributors (ACDs). The Call Applications Manager receives incoming-call details, such as the automatic number identification (ANI) or the dialed number information service (DNIS), and delivers the information to a host application in a simplified format. Using this data, the application builds customer-specific screens, which CAM then delivers to the call center agent as the call arrives. If the call must be transferred to another location, CAM software manages screen transfer so that the new agent receives all information gathered up to that point. The Tandem Call Applications Manager improves teleservicing efficiency through faster call handling, increased agent productivity, fewer agent errors, and prompt, personalized service.

Integration of systems is emerging as a significant factor in providing better service. Application integration can provide a seamless interchange of information between systems responsible for various aspects of telecommunication service. Integration cuts down on customer frustration by not requiring the customer to provide information more than once. It cuts costs and improves information accuracy by reducing the amount of data entry. And responsiveness is improved by allowing different systems to share up-to-date information.

Tandem systems allow telcos to integrate services whether they reside on different Tandem computers, workstations, or other manufacturers' hosts through a variety of networking and connectivity techniques. The foundation for Tandem's strength in distributed processing is the EXPAND™ networking software which gives users a single view of distributed applications and data anywhere in the network. To support high speed communications on local networks of its NonStop systems, Tandem supports the Fiber-Optic Extensions (FOX) product line with an aggregate bandwidth of 5 megabytes per second.

Open and flexible enterprise networks are supported with products such as Ungermann-Bass's Net/One system and Access/One network delivery system. Ungermann-Bass is a wholly owned subsidiary of Tandem Computers, Incorporated. Net/One interconnects devices from all leading computer vendors, while Access/One uses ordinary twisted-pair telephone wire to connect local users. The Access/One system was the first to integrate both local area network (LAN) and wide area Network (WAN) applications into a single platform. Ungermann-Bass installed the first fully operational Fiber Distributed Data Interface (FDDI) network in 1989.

A wide variety of networking and connectivity standards are supported by Tandem including the Transmission Control Protocol/Internet Protocol (TCP/IP), the System Network Architecture (SNA), and the X.25 protocol. Tandem is also compatible with LAN standards including NETBIOS, Ethernet, and Token-Ring. Tandem is committed to supporting the international Open Systems Interconnection (OSI) standards.

Productivity

Productivity is important to any telecommunications provider that must compete on the cost/benefit of services. Telcos are finding that one way of improving service while reducing cost is the introduction of messaged based/transaction oriented systems. Recently a Regional Bell Operating Company (RBOC) implemented a service negotiation system on Tandem computers and achieved significant improvements in productivity.

Service negotiation is the means by which customer needs are ascertained and the required service provisioned. Since customer expectations for quick and informed responses to their requests are growing, the service negotiation system must provide real time access to a myriad of customer related data. One such system is GemStar implemented by Pacific Bell. GemStar is a Tandem based Customer Contact Platform which allows Pac Bell service reps to assign a customer a telephone number, validate the customer's address, verify the customer's credit, provide pricing quotes on various services all in the course of a single customer contact session. The Gemstar system provides the service rep with information regarding product availability, product fit and suggests proposals to make to the customer regarding other products based on the customer profile.

The productivity gains that this system has provided are impressive. The Tandem based GemStar system has reduced customer contact time by 10%, and reduced service order errors by 15%. It has also cut training costs considerably by reducing initial training by 33% and continuation training by 40%. And it has decreased the need for printed reference material by 60%.

Flexibility

As competition grows, the flexibility of service offerings will be decisive. Telecommunication providers must fashion services to the requirements of individual customers or market segments in a timely fashion. A unique and powerful set of Intelligent Network products, which can support the delivery of flexible and timely services, is offered through Tandem's wholly owned subsidiary Tandem Telecommunications Systems Inc(TTSI). TTSI-NET products include a service control point (SCP), a service management system (SMS), and the applications development tool set of the service creation environment (SCE).

The service control point and the service management system will be discussed in detail later in the document, but the SCE bears special attention as a facility for rapid service deployment. The service creation environment has a set of software development tools which facilitates the design and development of network based applications such as a Line Information Database, 800 Service Database, or Calling Card Validation Service. The SCE allows application engineers to design databases, and data validation logic for services, without the need for complete knowledge of network protocols or system architecture. Tandem supports application development tools such as the SCE to ensure that services are specified and delivered in response to the requirements of customers or the needs of market segments in record time.

Innovation.

In order to differentiate themselves and accommodate shifts in the market place, network providers will need to offer innovative new solutions to customer problems. One solution that has been developed to meet customer needs is Integrated Services Digital Network (ISDN). ISDN offers secure, flexible, and cost effective access for voice and data communications equipment to the public network for both circuit and packet applications.

The Tandem ISDN Communications Subsystem (ICS) provides a platform for development of applications that can take advantage of the communications, network control, and call control benefits that ISDN offers. Tandem ICS is a combined hardware and software product that provides either a Basic Rate or Primary Rate connection between a Tandem host and the public switched network. With the delivery of the ICS platform Tandem has demonstrated its commitment to team with telephone companies, switch vendors, and third party software providers to develop ISDN applications that meet real needs and provide competitive advantages.

Tandem's pioneering development of messaged based/transaction oriented computers makes it an ideal platform for telco information systems which emphasize service, productivity, innovation and flexibility. The continuous availability and messaged based design of Tandem computer systems ensures responsive, real time service. Tandem's architecture allows competitively priced introductions of new, untried services while providing a clear path for future expansion of the system to meet demand. And Tandem

continues to demonstrate its commitment to innovative and flexible solutions to business problems with its support of standards and its development of systems to support new technology such as SS#7 and ISDN.

Tandem responds to the needs of the Telco Industry

For the purposes of this document the architecture of a modern telecommunications business will be seen as having five components:

1) the network fabric (switching, transmission),
2) business systems,
3) value added services,
4) the intelligent network,
5) network operations support systems.

Except for the network fabric layer, an area in which it does not compete, Tandem has processors installed in all components of the network architecture. Network providers choose Tandem because Tandem's design fundamentals parallel the architecture of the the public telephone network. These fundamentals include real time message processing, linear expandability, 7 X 24 availability, data integrity, distributed processing, and a commitment to standards.

Business Systems

Telecommunication business systems are those manual and automatic systems that support the business functions of a telco. They include: service negotiation, inventory, billing, customer support and repair. Many business systems use old batch oriented, mainframe resident programs that were built to support a largely analogue network and are now hard pressed to support digital features and services. Because these systems represent the primary customer interface, they need real time responsiveness. Customers demand control of, and access to network information. Business systems must provide this information or serve as a conduit to it. This in turn requires that business systems are knit into the overall architecture of the network and that they share with network elements the requirements of availability, reliability, data integrity, and distributed processing. Tandem systems are well positioned to support these types of requirements.

A business system example is the use of Tandem computer systems by Jutland Telephone. Jutland Telephone (JT) is Denmark's second largest telco serving more than one million customers with annual sales in excess of $700 million. As with many other telcos, JT is confronted with a changing business environment as Europe moves toward a single market. To meet this increase in competition, JT chose Tandem computers as the platform on which to develop a series of integrated, real-time applications. These applications include service order entry, billing, sales and order processing for large corporations, directory assistance, and operator services support.

The service order entry system of JT automates the ordering of new services or equipment for customers from initial customer contact to installation of the service or equipment. The system interfaces to directory assistance and billing to insure that information is updated in real time. The system maintains information on each subscriber line. This Tandem based system represents a marked improvement in

efficiency and productivity, since almost 60 percent of the changes of the telephone subscribers connections are now made automatically.

JT's billing system has resided on a Tandem platform since early 1988. The system provides summary billing for a company with many subscriptions. Services provided by different JT divisions are listed on a single invoice for those business that request a single bill. A Tandem Optical Storage Facility, which provides compact storage and improved access, is used by JT to store a required five year billing history for each of its customers. All of this translates positively to the bottom line, as the 6 million invoices the JT prepares annually are done sooner, with less handling and reduced postal costs.

Operator services and directory assistance systems are both key telco business systems that benefit from Tandem's design fundamentals. Volt Delta's Operator Services System is a product that can provide telcos a competitive edge. This application offers sub-second retrieval speeds from a multi-million listing database. Retrieval features include alternate spelling searches, automatic secondary searches, emergency number searches, and frequently called number caching. This reduces the average work time resulting in a significant cost savings.

The Volt Delta Operator Services System provides direct access to customers who have an ongoing need for current telephone listing information as well as an electronic directory assistance option to support network access. Other features include automatic interception of calls to disconnected numbers, customized messages, split referrals,and automatic call completion using the announced telephone number. The system also supports automatic billing for collect and third party calls using interactive voice response technology to offload calls from operator service positions. Because it is implemented on Tandem NonStop computer systems, the Operator Services System delivers continuously available, reliable service.

Value Added Services

Value added services use available switching and transmission facilities and then add features to increase the value of the transmission. Although 90% of a telephone company's current business is spoken calls, it is the other 10% of the business--data calls and value added services--which is the fastest growing market. Value added services require a computer which supports a messaged based architecture and provides continuous availability and reliable processing power. They include voice, electronic mail, and facsimile messaging, Electronic Data Interchange (EDI), data services and gateway services. Tandem processors are ideally suited for value added services because of their fault tolerance, distributed processing, and data integrity.

The linear expandability of Tandem systems is a key competitive factor in deploying new services in untested markets. A provider can offer a new service with minimum investment. As demand grows and operational requirement are better understood, the service provider can invest in more processors, confident that Tandem's architecture guarantees a linear increase in processing power This is a significant competitive advantage.

Toyo Information Systems, Ltd's Fair-Way Intelligent Facsimile Transmission service is an example of a value added service which takes advantage of Tandem's continuous availability and distributed processing. The Fair-Way service, which is implemented in Japan, provides a number of enhanced facsimile transmission (FAX) services. The basic

service includes store and forward capabilities that allow a FAX to be broadcast to multiple locations at user specified times. The system automatically retries busy and does not answer (DNA) destinations until successful. Fair-Way also supports simultaneous transmission to multiple locations. At the customers direction, the service can also stop receipt of unwanted FAX messages.

Another value added service which is gaining in importance is Electronic Data Interchange (EDI). The exchange of electronic information for buying, moving, selling or paying is being adopted by both large and small corporations to streamline operations, cut costs, and improve customer service. Tandem computer systems provide the continuously available distributed processing power with the data integrity that is required of EDI applications. EDI utility software offered under the NonStop EDI program provides message management, security, and translation software for medium and large enterprise EDI systems.

The Message Way system provides flexible addressing capabilities that permit users to separately schedule and control when messages are delivered to trading partners. XWay performs translation of in-house files to EDI standard formats, and EDI documents into in-house formats. XWay supports both EDIFACT and ASC X12 standards. Both of these systems benefit from Tandem's linear expandability. As volumes increase, the capacity of the system can be increased in a cost effective manner without service or tracking interruptions.

Globalization of the economy is significantly affecting the telecommunications business. Competition is increasing, while opportunities are expanding. As companies look beyond their national borders to increase business and cut costs, they need communication tools that allow them to keep in constant contact with distant locations. Tandem computers are used by Sprint International to provide just such a capability with its electronic message handling service--Telemail 400.

Telemail 400 lets business people send, receive, or file messages from any terminal, personal computer, communication word processor, or telex device on their network. The system supports the X.400 protocol, ensuring businesses that they can connect to public and private messaging systems around the world. Telemail 400 acts as a hub for communications connecting dissimilar messaging systems and offering international communications. By implementing the service on Tandem computers, Sprint International takes advantage of Tandem's message based architecture to provide continuous availability and distributed processing power.

Intelligent Network

The goal of the Intelligent Network is to improve the delivery, cost, and variety of telecommunication services. It does this by separating network control and routing functions from the switching equipment that establishes and maintains voice and data connections. This has driven the institution of an open network architecture for both the hardware and software elements of the network. Most of the software that will control the network will reside on processors that sit beside the switching and transmission elements of the network . These processors must support a continuously available message based distributed architecture. These processors must also provide a high degree of data integrity. Because Tandem was a pioneer in the development of fault tolerant, transaction based, distributed processors, it has become a leader in the deployment of the Intelligent Network processors.

The Intelligent Network provides two major advantages for the telecommunication provider. It allows for increased operational efficacy with the use of Signalling System #7 for out of band signalling to support look ahead routing and call set-up functions. The Intelligent Network also provides a rich vein of new services including calling card validation and software defined network features such as 800 and 900 services, and customer network management features. Many of these applications are based on online database systems resident on service control points (SCP).

Tandem Computers Incorporated, through its wholly owned subsidiary Tandem Telecommunications Systems Inc. (TTSI), was the first company to successfully implement a service control point in a major telco's network. The implementation is based on Signalling System #7 and is used for calling card validation and call forwarding. TTSI-NET offers three Intelligent Network Products a service control point, a service management system, and a set of application tools, the service creation environment (SCE).

The service control point provides the online database architecture required for high-speed query support and call processing logic. These features facilitate the interactive communications queries that are generated in an Intelligent Network application. The concurrent operations of multiple enhanced services such as credit card validations, and programmable call forwarding is also provided by the service control point. The system provides SS7 network management and connection control functions as well and Transaction Capabilities Application Part (TCAP) support. The service control point takes full advantage of Tandem's message based architecture and continuous availability to provide interactive call processing treatment as prescribed by the Advanced Intelligent Network architecture.

The service management system is designed to allow diverse applications to be incorporated into the architecture of the Advanced Intelligent Network. It provides communications facilities that can interact with the products of multiple network element vendors to offer flexible solutions to network operations problems. A wide range of data management functions that include update processing, and network data validation can be supported by the service management system. Interfaces to new and existing business support systems, and customer control systems are supported. Both TTSI-NET's service management system and service control point are built on Tandem NonStop computer systems. These systems combine state-of-the-art Very Large Scale Integration technology and messaged based architecture to provide high performance, high reliability and low service cost.

Network Operations Support systems

Network operations support systems monitor, measure, control or report on the status of the network and its elements. Many systems in this environment share the requirements of continuous availability and reliability, since these systems provide for the operation and control of the network in real time. As telecommunications customers become more sophisticated they require greater and greater control over the services that they purchase. Telcos expect to offer a whole range of customer network management features in the coming years. This in turn will require that the operational support systems achieve the responsiveness of other customer oriented message based systems. The reliability and availability of Tandem computers joined with their message based architecture make them an ideal basis for network operations support systems.

The NMCC Network Management Systems by GTE takes advantage of Tandem's architecture to provide reliable and continuously available network operations functions to independent telephone companies and businesses. The NMCC system offers both information management functions and network control functions. Information that is entered or changed in one portion of the system is automatically updated in all other portions of the system. The information management functions include support of service order entry and status reporting, network facilities management, repair administration, inventory management, directory assistance, and billing and cost accounting. The network control functions comprise support for network alarm monitoring, traffic administration, centralized network interface for intelligent network devices, call detail record keeping, and electronic office mail. The cost benefit that the integration of systems provides was achieved through the use of Tandem's database integration and connectivity tools.

Universal Measured Service (UMS) developed by GTE, is a billing support system that provides an efficient, cost effective, and automatic means of collecting, and processing both switch recorded and operator handled customer usage information. UMS is adaptable to all major switch vendors formats. The Tandem based UMS Billing Intermediate Processor electronically collects local billing information based on usage from central office switching equipment. It then performs extensive verification on raw switch data and translates the data into a standard call record format. Information is aggregated for each billing telephone number and group accounts. Finally this information is passed to a mainframe computer over electronic communication links. Because the information that UMS processes represents a major revenue stream for a telco and because of the dispersed nature of switch sites, the continuous availability, distributed processing, and data integrity of Tandem computer systems ideally suit the requirements of the UMS system.

The Basic Operations System (BOS) developed by US West, NewVector Group, Inc. is a UNIX based cellular network management and surveillance system that runs on Tandem's Integrity S2 System. BOS is a modular application that support a number of cellular network management functions. BOS can be configured by users and scaled to their needs in an integrated, open system environment. BOS supports the display of network conditions, the analysis of network failures, and the ability to correct network failures through an interactive user interface. Security in the Basic Operations System includes the ability to restrict access to data, remote systems, and BOS commands. Tandem's Integrity S2 was ideal as the BOS platform because of its its fault tolerance and support for a wide variety of data communication protocols.

Standards

Tandem Computers Incorporated has long demonstrated a commitment to industry standards as a means to improve service, increase productivity, and ensure flexibility and innovation in system development and maintenance. Tandem has implemented applications that support such telco industry standards as SS#7,and ISDN. It actively participates in the development of the International Open Systems Interconnection (OSI) standards. A wide variety of networking and connectivity standards have been implemented on Tandem computer systems including: Transmission Control Protocol/Internet Protocol (TCP/IP), System Network Architecture (SNA), and the X.25 protocol. Tandem is also compatible with LAN standards including NETBIOS, Ethernet, and Token-Ring.

Tandem recognizes that open system standards can support the goals of productivity by allowing customers to benefit from system interconnectivity and application portability. For that reason, Tandem introduced the Integrity S2 system, which combines fault tolerance with an enhanced, but fully compliant, UNIX System V operating system based on RISC technology. Tandem did not change the industry-standard interface for UNIX, it simply made it's NonStop UX software more reliable. Running standard UNIX applications on Integrity S2 hardware with NonStop UX ensures the continuous availability and the data accuracy of the applications.

Tandem is committed to a strategy of providing open application access to the unique advantages of Tandem NonStop systems via industry standard interfaces. Tandem recently announced that it will support the Portable Operating System Interface (POSIX) and X/Open Portability Guide (XPG3) standards. Applications adhering to these standards can run on a variety of different systems supporting the POSIX and XPG3 interfaces. The goal of these standards is to increase the ease of porting programs and programmer skills. This is illustrative of Tandem's ongoing support to open systems development; it ensures that customers are furnished an environment where applications are truly portable. This translates to reduced software development and maintenance costs and greater productivity.

Conclusion

The restructuring of telcos from monopolistic utilities to companies that must compete in the free market to sell their services have caused significant changes in the telecommunications business. Telecommunications, perhaps more than any other industry, is played on a global level. Even those network providers which are protected by monopolies feel the pinch of competition, as multinational companies use by-pass technologies to achieve cost efficient telecommunications services. Competition coupled with abundance has driven down the profit margin of bandwidth, the major telco commodity. All of these factors have caused the telco to emphasize service, productivity, innovation, and flexibility. This emphasis has spurred the convergence of telecommunications and computing technologies. Computers are used in ever increasing numbers to meet the needs of telcos and their customers.

Tandem Computers is a natural partner with the telcos in providing innovative and flexible solutions that enhance service and productivity. Tandem's design fundamentals parallel the architecture of the public switched network. These fundamentals include a fault tolerant, message based architecture that supports distributed processing with a high degree of data integrity. The Tandem architecture is linearly expandable which ensures cost efficient deployment of systems with a clear path for expansion. The inherent strengths that stem from these design criteria make the Tandem computer solution a natural fit to many of the telcos mission critical applications were continuous availability and fault tolerant processing is required.

file
Intelligence Report

March 1990

In This Issue

Issue Number 65

DEC VAX/MicroVAX Planned & On-Order Systems

(1,500 US Site Interviews/Month)

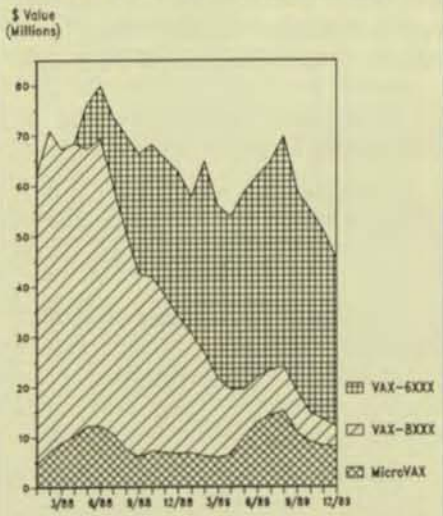

The VAX-6XXX series is rapidly displacing the VAX-8XXX series in VAX user purchasing plans. -DE

COMPUTER INTELLIGENCE
3344 North Torrey Pines Court
La Jolla, California 92037
619.450.1667

European Mainframe Market

Karen Landis

Like the US market, competition among mainframe vendors is equally fierce in the European community. Marked by mergers and acquisitions throughout the 1980s, the US mainframe market has diminished to a handful of major players. Consolidation would appear to be the name of the game in Europe as well. After bowing out of the mainframe market in late 1989, Nixdorf recently announced plans to merge computer operations with Siemens AG.

Currently, the majority of installed European mainframes are IBM-supplied. However, IBM may possibly face stiffer competition selling to Fortune 1000 sites in the European community than in US markets. The Mainframe Market Share graph (page 4) details IBM and competitors in both Europe and the US. As evidenced by this graph, IBM's share of the European market is six percent less than that enjoyed in the US.

IBM's smaller share can be attributed to several factors. Not only must IBM face competition from old familiar companies like Amdahl, Bull, Unisys, et al., but also from several players not familiar to US markets, such as Comparex and Siemens.

Comparex Information Systems GmbH began operations in 1987 as a joint venture between Siemens and BASF, offering mainframes which are plug compatible to IBM sys-

(See Europe on page 4)

Tandem - Looking Strong!

Ralph Busch

While the stocks of most computer vendors are rated either "hold" or "sell" by the stock investment analysts, a number of analysts have been rating Tandem as a "buy." Some analysts have even rated it as an "aggressive buy." What's going on here? Why does Tandem rate so highly while some of the giants in the industry, like IBM and DEC, are rated much less favorably? A look at the **Financial History** graph (page 2) shows that Tandem has achieved a steady pattern of sales growth over the last six years. Many analysts feel that they are on the road to even more dramatic growth in the 1990s.

The Financial History graph shows that the percent of Net Income to Sales has been relatively healthy, growing over ten percent in 1987. The percentages were

(See Tandem on page 2)

LAN Use at AS/400 Sites

David Eulitt

Over the last few years, local area networks (LANs) have grown to be extremely popular with personal computer users. The many benefits of LANs include the ability to share resources (such as printers and tape drives) and data (such as programs and files). More sophisticated LANs might include midrange and/or mainframe computers which would allow PC users to access computing and storage services on the larger systems.

Computer Intelligence divides LANs into two classifications: *PC LANs* and *General-Purpose LANs*. A PC LAN, as its name implies, is a LAN consisting primarily of PCs. Other equipment can also be attached to PC LANs, such as file servers, office equipment, bridges,

(See LANs on page 5)

Financial History - Tandem

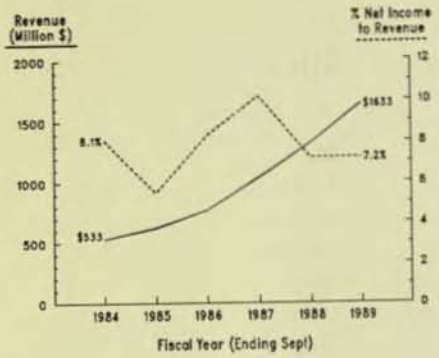

Industry Profile

General Purpose Systems
Tandem - US Sites

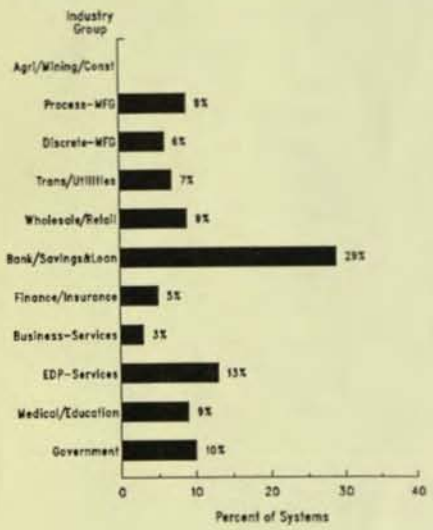

Tandem

(Continued from page 1)

down somewhat in 1988 and 1989. Heavy expenditures for the development and marketing of new products announced near the end of 1989 influenced expenses. These products include enhancements to the CLX entry level NonStop systems and the new top-of-the-line Cyclone system. Net income should improve this year as these products begin to ship, and the manufacturing learning curve benefits the current mid to high range NonStop products.

By looking at Tandem's current product line and industry orientation, we can understand how they have achieved their success to date. The Tandem - General Purpose Systems table provides some historical and performance perspective of Tandem's products. Tandem was founded in 1974, and as the table shows, the first products were installed in 1976. All of these general purpose systems are members of the NonStop family. One of the key features of the NonStop family is that they are fault tolerant. That is, if a fault occurs in one processor, another processor will continue processing the user's applications while the fault is being corrected. The offending processor can be removed from the system, repaired, and returned to the system without interrupting the application processing. Note that the single processor CLX systems are not fault tolerant.

Fault tolerant processing is especially important to customers who require continuous systems availability, high transaction throughput and system expandability in on-line transaction processing (OLTP), real time, and on-line and support environments. The banking industry is an example of an industry that has these types of requirements. We all know how irritating it would be to run out of cash and go to an on-line teller machine that is out of order because a computer is down.

Given that banks do large volumes of OLTP and require fault tolerant processing, it is not too hard to see why they have been prime targets for Tandem systems. A look at the Tandem - Industry Profile graph shows that 29% of the Tandem system base is in the Bank/Savings and Loan industry segment. The majority of the remainder of the base is

Tandem - General Purpose Systems

Model	Number of Processors	Date of Installation	Performance MAX TPS*	Comments
NonStop	2-14	1976	--	First Model Shipped
NonStop II	2-16	1981	32	Migrated to TX or VLX
TXP	2-16	1983	72	Migrating to VLX
EXT 10	2 or 4	1986	5	Former Entry System
EXT 25	2 or 4	1986	15	Former Midrange System
CLX 600	1-6	1987	15	New Entry System
CLX 700	1-8	1989	30	New Midrange System
VLX	2-32	1986	200 +	Mid-to-High Range
Cyclone	4-16	1989	600	Top-of-the-Line

* TPS = Transactions per Second

COMPUTER INTELLIGENCE
3344 North Torrey Pines Court
La Jolla, California 92037
619.450.1667

Outside Intelligence

Territory imbalance due to industry concentrations or absence of prerequisite hardware or software quickly demotivates a salesforce. Demotivation turns into higher turnover in the territory, lack of sales, missed product quotas and worst of all a lack of trust by the salesforce in sales management.

Company's Market Share

Your market and territory share greatly influences quota attainment of a product/service. This is the easiest of all factors to assess. You should be able to determine this from your own company records. The additional benefit of analyzing this data is you may determine that you do better in a specific geographical area or in a particular industry even if you are selling cross-industry software.

Market share is relevant to sales only from the viewpoint of determining what potential remains in the marketplace. Geography may be extremely important when it comes to assessing territory share and potential. For example, your service or product may sell better in the large cities such as New York, Boston, Chicago and Los Angeles rather than smaller cities such as New Orleans or Jacksonville, Florida. This would dramatically affect the salesperson with a territory that includes the South.

Years ago I worked for one of the major software companies which obviously obtained it's largest revenue share from new sales. However, no analysis was ever done on what our market share was versus the competition, how fast the market was growing or slowing for specific products and so forth. The company didn't really have to worry about all of that because it was doing so well at the time.

Quotas were derived from the amount of revenue needed to obtain the annual growth Wall Street wanted to see. The revenue needed was divided by the number of salespeople and quotas were established. Of course, senior salespeople were assigned a larger quota due to their experience. No analysis was done with respect to the competition's share of the marketplace or how competition might affect individual territories. How do you think the salesperson felt about this method of arriving at quotas when they were meeting stiff competition or their territory was saturated with their product?

Competition's Market Share

Analysis of the competition's market and territory share, although painful at times, will benefit your company and salesforce dramatically. It benefits the company because once you have determined the competition's share, added that to your own and then subtract it from the marketplace total you have what potential remains in the marketplace. This assumes that prospects can convert to a different vendor (YOU, of course!).

Without this analysis you are guessing at your company's bottom-line because most companies obtain a majority of their revenue from new sales. This one factor alone will allow you to establish realistic revenue targets for the company. It may even indicate to your company that the market is losing potential and you must take preventive action for the company. It may be time to speed up development of new products, acquire new products or enter new markets.

Winning Forecasts

When forecasts are real the company benefits. Stockholder and banker expectations are not built up beyond a reasonable level. Support staff levels are consistent with needs. This helps avoid lay offs or under utilized staff situations. Capital equipment and facilities are also controllable to a reasonable level. Company profitability clearly suffers when forecasts are wish lists.

Realistic revenue targets benefit the entire sales organization because it makes quotas attainable and motivates the salesforce. Motivation comes when you explain that you have taken the time to set realistic quotas based upon facts, not guesses. Attainable quotas this year makes next year's quota increases, which are inevitable, more palatable and believable to the salesforce.

Realistic Forecasts & Fair Territories

At times it is hard to face reality, and its even harder to confront others, such as company executives and financial partners, with numbers that do not reach desired growth objectives. By using CI's market intelligence to quantify and understand your market, you can set your territories using a formula that everyone, even your representatives can live with. Wouldn't it be wonderful to eliminate squabbles over territories in 1990!

John Grant is the VP of Marketing for Indisy Software, Inc., which markets an enterprise-wide communications system that connects LANs to IBM mainframe computers and provides applications such as electronic mail across LANs, mainframe and minicomputers.

Indisy Software
6100 Wilshire Boulevard
Suite 1150
Los Angeles, CA 90048-5102
213-933-8393

Outside Intelligence

Market Intelligence and Sales Territories

$$\frac{\text{Total Number of Sites}}{\text{Total Number of Salespeople}} \neq \text{Number of Sites per Territory}$$

John Grant
Indisy Software

Does this sound familiar? Divide the total number of sites by the number of salespersons and arrive at sales territories. However, is this really a method? When assigning territories shouldn't you analyze these three key components:

- What is the product/service potential for the entire marketplace?
- What is your company's share of the market?
- What is the competition's share of the market?

Without answers to these questions using the equal quantity of sites method for territory assignment is analogous to using a crystal ball.

Proper use of market intelligence to assign sales territories enables sales management to make more accurate sales forecasts and establish realistic revenue targets thus providing an environment of success for each salesperson.

Sources of market intelligence on specific hardware, software and services are available in the computer marketplace today. You can't pick up a trade journal without seeing a graph on IBM or DEC regarding hardware, the number of PCs sold compared to last year, or how local area networks are cutting into minicomputer and mainframe sales.

Besides keeping up with the trade press you can purchase published reports and database information on specific markets, or even commission studies of your own. CI's Market Intelligence System is my primary source for market intelligence information.

Once you have access to market intelligence you need to arrive at a formula for setting territories. The formula for each company should be unique, based on various factors that are largely out of your control. Before arriving at a method, or revamping your current method, examine each of these three components.

Product/Service Potential

Consider whether your marketplace is growing or shrinking. A shrinking market may necessitate more sites to close a given revenue amount than a growing market. A growth market may offer the potential for multiple placements per site, whereas the mature market may yield only an occasional new placement. You also need to know the total size of the market with regard to specific prerequisites such as hardware or software. For example, assuming you have an AS/400 banking application, the number of IBM AS/400s installed vs. the number being shipped to banks on a monthly basis would be important for estimating the entire market potential for your product.

I have worked for companies that assumed the salesperson had enough accounts to meet prerequisite hardware/software requirements in their territory. This assumption didn't affect the salesperson when everyone had 750 accounts per territory. However, as the salesforce grew and territories shrank this assumption began to negatively impact the salesforce. Some salespeople started to complain vehemently concerning the lack of qualified prospects in their territory. A review of the territories indicated that little potential existed for a specific product because there weren't enough accounts that meet the prerequisite requirements. After this incident the salesforce was suspicious and they all started taking time away from selling to analyze their territories. Management should have performed a detailed analysis in the first place and avoided the lowering of sales morale and productivity. Imagine the consequences of this assumption had the company been selling only one product in a competitive market.

Product/service potential is especially important when first introducing your product to the marketplace. As mentioned in the above example, geographical or industry concentrations may create an imbalance when assigning territories based strictly on total number of sites. If your product is favored by manufacturers, the downtown legal and banking centers, although dense in end-user site count, may not produce many viable prospects.

In mature markets product/service potential is even more important and easier to assess due to the wealth of information available. Unless there is a good reason why your product can displace a competitor's the potential of a territory may be very little due to the penetration by all vendors. Based on an analysis of all territories you may decide you have to offer special incentives to prospects. For example, a discount on your product if the prospect replaces your competitor's product with yours or additional dollars toward quota attainment for the salesperson if they are able to replace a competitor's product.

in five industry segments that each have between nine and thirteen percent of the total.

One other important characteristic of the Tandem installed base is that a relatively high percentage (43%) of Tandem sites also have an IBM/PCM mainframe installed at the same site. Fifteen percent of Tandem sites also have a DEC VAX system installed, and eight percent have an IBM midrange system installed. The significant factor about this high level of co-residing with other vendor's systems is that there are considerable requirements to exchange information between these unlike systems. Tandem has met this requirement with a number of networking systems that facilitate exchanging data between systems from multiple vendors. Their acquisition of Ungerman-Bass has further strengthened Tandem's networking capabilities. A group of six new network management products recently made available by Tandem significantly enhances their ability to perform network management functions in a Tandem/IBM network.

Tandem's financial history shows that they have been quite successful selling fault tolerant OLTP systems. However, competition is picking up in this market. Besides IBM and Stratus, who have been in this market for some time, other vendors like Bull, DEC, HP and Unisys are going after this market.

Despite this competitive onslaught, Tandem and the financial analysts feel that the market for Tandem products is expanding. Joint efforts with third-party software developers and system integrators to develop what Tandem calls "Application Environments (AEs)" for the manufacturing, finance and retail/distribution markets will fuel the expansion. TIME (Tandem Integrated Manufacturing Environment) is available now and the other two AEs will be available late this year.

Another way to expand the fault tolerant market is to offer a UNIX based system for users (certain government agencies and telephone companies for example) who have standardized on the UNIX operating system. Tandem introduced the LNX Series in 1984 that ran the UNIX operating system, but it was not very successful in the marketplace. The LNX experience apparently provided training for the much more aggressive UNIX effort that was announced in January 1990. Also, the acceptance of UNIX as a standard has progressed significantly since 1987.

"Application Environments" and UNIX should provide incremental business for Tandem. Additionally, Tandem and some of the financial analysts are predicting that there will be a major move by organizations whose businesses are becoming ever more time-critical to what they are calling "On-Line Enterprise Computing (OLEC)." Where OLTP is generally restricted to a few key applications within an enterprise, and may be run on a separate system from the system running the production data processing functions, OLEC will combine all of the data processing for the enterprise on one, on-line, interactive, secure, distributed system.

OLEC moves what had been somewhat of a niche market for processing transactions against a specialized database into the mainstream. Clearly, OLTP provides a good base upon which to make this move, and Tandem has major strengths in this market. With the announcement of the new Cyclone systems, Tandem will have the computer power required to move into the OLEC era. As an example, Tandem claims that a 4 processor Cyclone equalled a 4 processor IBM 3090-400S in an OLTP benchmark.

If Tandem is successful in moving into this new world of OLEC, the future could indeed be bright.

Ralph Busch, Industry Analyst, researches the General Purpose Computer market for CI.

Top Ten PC Models Sold By US Specialty Stores Channel During December 1989

PC Brand and Model

1. IBM PS/2 Model 55sx
2. IBM PS/2 Model 30 286
3. Compaq Deskpro 386s
4. IBM PS/2 Model 50z
5. IBM PS/2 Model 70
6. Compaq Deskpro 286e
7. Apple Macintosh SE
8. Apple Macintosh Plus
9. Apple IIGS
10. Epson Equity II+, III, III+

CI's StoreBoard reports that IBM PS/2 models held four out of the top 5 spots in December 1989. StoreBoard monitors the US specialty store channel for CI. -JS

COMPUTER INTELLIGENCE
3344 North Torrey Pines Court
La Jolla, California 92037
619.450.1667

Mainframe Market Share

US & European Fortune 1000 Sites

US

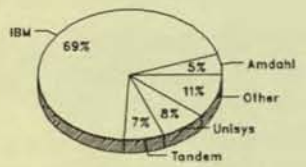

% of Systems

European

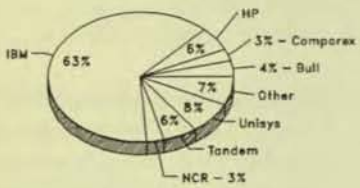

% of Systems

IBM AS/400 Planned & On-Order Systems

(3,500 US Site Interviews/Month)

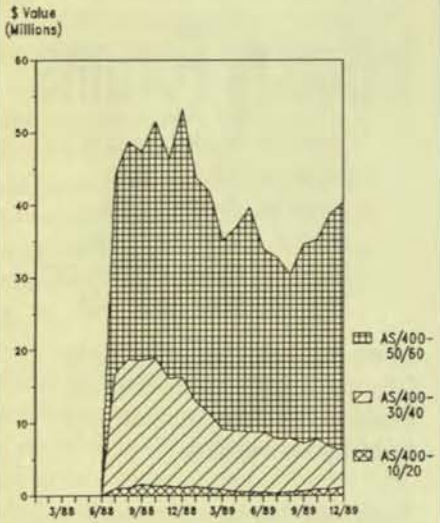

IBM is hoping that the recent enhancements to the AS/400 models 10 and 20, and the introduction of the model 25, will improve sales of low-end AS/400 systems. -DE

COMPUTER INTELLIGENCE
3344 North Torrey Pines Court
La Jolla, California 92037
619.450.1667

Europe

(Continued from page 1)

tems. After only three years in operation, Comparex holds a three percent share of CI's European mainframe database.

Siemens, on the other hand, has been in the data processing field since 1954. Though not directly compatible to IBM systems, the Siemens BS3000 operating system allows these systems to run IBM and compatible software with little or no alteration. Two percent of the mainframes in CI's European database are supplied by Siemens. A comparison of systems based on performance from both of these manufacturers and IBM is offered in the **IBM and Compatible Mainframes** chart (below).

As indicated in the IBM and Compatible Mainframes chart, both of these companies offer a wide range of alternative and competitive systems to IBM. Along with BULL, Hewlett Packard and other familiar competitors like Amdahl, Hitachi Data Systems and Unisys, these European companies appear to be giving IBM a run for their money.

Data used in this article is based on 4,000 US sites and 800 European Fortune 1000 sites from the February release of CI's Computer Installation Data File.

Karen Landis, Industry Analyst, researches the IBM/PCM mainframe market for CI.

IBM and Compatible Mainframes Ranked By Performance

IBM	Comparex	Siemens
3090-600J		
3090-600S		7-500/120S
3090-500J		
3090-500S	8/98	
3090-400J		
3090-400S	90-8	
3090-380J	8/96	
3090-300J		7-500/120P
3090-380S		
3090-300S	90-6	7-500/120I
3090-200J	8/95, 8/94	
3090-280J		
3090-200S	7/90-22, 90-4, 8/93	7-500/120F
3090-280S	8/92	
3090-250JH	7/90-11, 90-3, 8/91	7-500/90T
3090-250J	7/88MP	
3090-180J		
3090-250S	7/90-2	
3090-180S	8/90	7-500/90I
3090-170JH	7/78MP, 7/90-1	
3090-170J		
3090-150JH	7/88	7-500/90F
3090-150J		
3090-170S	8/89	
3090-150S	8/87	7-500/90D
3090-120J	7/78, 8/85	
4381-92E		
3090-120S, 110J	7/75	
3090-100S	7/73, 8/83	
4381-91E, 23		7-500/60S
4381-90E, 22	7/71, 7/72, 8/81	7-500/60F
4381-21	7/68, 7/69, 8/80, 7/63, 7/65	7-500/60B, 60D

LANs

(Continued from page 1)

etc. A General-Purpose LAN, however, is a LAN that contains at least one system larger than a PC (such as a minicomputer or a mainframe) as a node on the LAN. For example, a Token-Ring LAN consisting of four PS/2s and a file-server is classified as a PC LAN. But add an AS/400 and you've got a General-purpose LAN.

LANs are gaining popularity at US IBM midrange sites. The **LAN Penetration - US IBM Midrange Sites** graph shows that in general the larger the system at the site, the greater the use of LANs at IBM midrange sites. Only seven percent of System/34 sites and eight percent of System/36 sites have LANs, versus 19% of System/38 and 18% of AS/400 sites. The graph also shows that PC LANs greatly outnumber general-purpose LANs. In most cases these midrange systems are not included as a native node on the network.

Although an equal percentage (4%) of System/38 and AS/400 sites have general-purpose LANs, the AS/400 advanced networking features clearly establish it as IBM's preferred midrange networking platform of the future. A review of the LAN Industry Penetration table shows that LAN use at US AS/400 sites differs significantly by industry group. For example, 29% of the sites in the medical/education segment have an installed LAN, versus only 8% of the wholesale/retail sites. Other industry groups with a higher than average LAN penetration include discrete manufacturing (22%), agriculture, mining and construction (21%), bank/savings and loan (21%), finance/insurance (19%), and process manufacturing (19%). General-Purpose LANs are most prevalent at bank/savings and loan (9%), business services (6%) and discrete manufacturing sites (6%), while PC LANs are predominant in medical/education (25%), discrete manufacturing (18%) and agriculture, mining and construction sites (16%).

LAN Industry Penetration - US AS/400 Sites

Industry Group	All LANs	GP LANs*	PC LANs
Agriculture, Mining and Construction	21%	5%	16%
Bank/Savings & Loan	21%	9%	13%
Business Services	16%	6%	11%
Wholesale/Retail	8%	2%	6%
EDP Services	15%	5%	11%
Finance/Insurance	19%	4%	15%
Government	18%	3%	15%
Medical/Education	29%	5%	25%
Discrete Manufacturing	22%	6%	18%
Process Manufacturing	19%	5%	15%
Transportation/Utilities	13%	3%	10%

*General Purpose LANs

It's evident that the use of local area networks at AS/400 sites differs by industry group. As the **LAN Penetration - US AS/400 Sites** graph (page 6) shows, LAN use also differs dramatically by the AS/400 model installed at the site. Only 11% of the AS/400 model 10 sites have an installed LAN, compared to an amazing 43% of the model 70 sites. Clearly, the reason for the wide range in overall LAN penetration is caused by the wide range in PC LAN penetration (from nine percent at model 10 sites to 35% for model 70 sites).

Since only four percent of AS/400 sites currently have a general-purpose LAN, on the surface it appears that PC-to-AS/400 connectivity is quite low at AS/400 sites. However, there are other methods of PC-to-host connection than a direct peer-to-peer LAN link. The **PCs Linked to Host System** (page 6) graph shows that 33% of PC LANs at AS/400 sites are linked to a host system via a remote link, such as a gateway. Seven percent of PC LAN sites plan such a link in the future, and 60% have

LAN Penetration

US IBM Midrange Sites

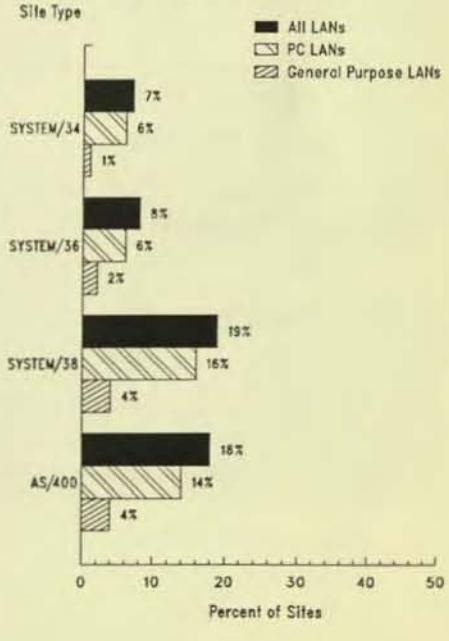

COMPUTER INTELLIGENCE
3344 North Torrey Pines Court
La Jolla, California 92037
619.450.1667

LAN Penetration

US AS/400 Sites

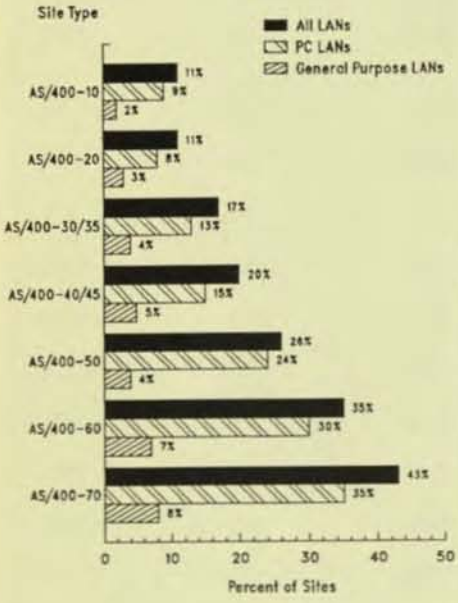

PC LANs Linked To Host System

US AS/400 Sites

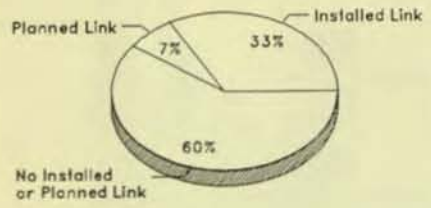

Percent of PC LAN Sites

COMPUTER INTELLIGENCE
3344 North Torrey Pines Court
La Jolla, California 92037
619.450.1667

no link and no plans to install one within the next year.

Many other PCs are connected to AS/400s via other means, such as 5250 terminal emulation. Of the AS/400 sites that don't have a direct PC-to-AS/400 link via a general-purpose LAN, 65% have at least one PC communicating in some way to the AS/400. Overall, 32% of the PCs at these sites communicate to the AS/400, usually via 5250 emulation.

We have examined several ways in which PCs communicate with AS/400 systems. Although a direct PC-to-AS/400 connection via a general-purpose LAN is currently the least common method of connection, it will be much more common in the future as the vision of client/server computing becomes a reality. Many software vendors will rewrite their applications to take advantage of the client/server architecture, which combines the ease of use and friendly user interface of PCs with the computing and storage resources of minicomputer and mainframe systems. Part of an application may run on the PC, and part on a larger system, but it will be transparent to the user. All this will be made possible by the local area network.

David Eulitt, Industry Analyst, researches the DEC VAX and IBM Midrange markets for CI.

Two Years of Fortune

Suzan Sachdeva

Multi-site Fortune 1000 corporations are using more and more computing power to automate the way they do business.

Each new release of a mainframe, mini, or personal computer introduces a new price/performance level to the market. This means that each time a firm installs another system they're getting more bang for their buck than ever before.

Analysis of data extracted from CI's Market Intelligence System reveal notable changes in corporate America's computing. The average multi-site Fortune 1000 corporation had installed 8.6 mainframes, 63 minis, and 2,287 personal computers at year-end 1989. These figures, when compared with those from two years earlier, show a significant change in computing power over a short time period. At year-end 1987, multi-site Fortune 1000 corporations had installed an average of 8.9 mainframes, 34 minis, and 982 personal computers.

As you might expect, personal computers account for the lion's share of growth on a unit basis. The average number of installed personal computers grew 133% from 1987 to 1989. Minis experienced an 85% growth rate, while mainframes experienced a three percent decline.

These trends take a somewhat different route when each industry is looked at individually. The **Average Number of Installed Systems Per Corporation** table (page 7) highlights the differences in installed products for two time periods at multi-site Fortune 1000 corporations in the US. As the table illustrates, the use of mainframes, minis and personal computers did not change in a uniform manner from 1987 to 1989.

First, the work-horses of American computing, mainframes, present a mixed picture. Within this category, the average

number of installed systems per corporation varied from a decrease of 25% (Life Insurance) to an increase of 27% (Diversified Services). Differences in how these industry segments conduct business contributes to this variation. For example, participants within the Life Insurance industry are more dispersed and prone to distributed processing than many other industries. In addition, the increase in power available with systems often times reduces the corporation's need to continue to add mainframes.

Second, for systems classified as minis, the growth rates varied from 40% (Transportation) to 177% (Diversified Financial). The Industrial industry group, which includes large manufacturing comapnies as diverse as aerospace, mining and cosmetics, had close to the average growth rate with 82%. Unlike the mainframe category, the average number of installed minis per corporation increased across all industry groups. These systems not only provide a cost-effective means for computer resource sharing, but also lend themselves well to departmental computing within diverse environments.

Third, in the case of personal computers, the growth rate varied from a low of 48% (Life Insurance) to a high of 230% (Diversified Services). Diversified Services, which includes, in part, corporations that participate in the health care, entertainment, and business services industries, far outpaces all others in its use of personal computers. The significant change seen here suggests that corporate America views personal computers as a cost effective tool for increasing worker productivity.

A review of the table overall shows that several industry groups experienced a decline in the average number of installed mainframes. Conversely, the average number of minis and personal computers has grown. This observation suggests industries that lend themselves to distributed processing are opting to enhance their computing arsenal with smaller, less costly systems.

The trends discussed thus far are not confined to installed systems. A look at the multi-site Fortune 1000 corporations shows the average organization plans to

(See Fortune on page 8)

Multi-Site Fortune 1000 Corporations (US)
Average Number of Installed Systems Per Corporation
Year-End Figures

Fortune Industry Group	Mainframes		Minis		PCs	
	1989	1987	1989	1987	1989	1987
Banking	8.2	8.3	17	8	2059	730
Diversified Financial	9.0	7.7	25	9	2816	1239
Diversified Services	9.8	7.7	43	25	1813	549
Industrial	9.1	9.9	93	51	2643	1178
Life Insurance	4.9	6.5	10	6	2130	1444
Retail	8.7	9.2	20	9	1474	643
Savings & Loan	2.7	2.4	3	2	497	156
Transportation	6.4	6.4	14	10	1175	383
Utility	11.6	12.4	87	42	3142	1435
All Industries	8.6	8.9	63	34	2287	982

Note: CI classifies High Performance Workstations as mini computers

Multi-Site Fortune 1000 Corporations (US)
Average Number of Planned Systems Per Corporation
Year-End 1989 Figures

Fortune Industry Group	Mainframes	Minis	PCs
Banking	2.4	2	232
Diversified Financial	2.6	3	314
Diversified Services	2.6	4	132
Industrial	2.4	9	189
Life Insurance	1.8	1	178
Retail	2.6	5	92
Savings & Loan	0.8	1	38
Transportation	1.5	2	160
Utility	2.6	8	301
All Industries	2.3	4	185

Note: CI classifies High Performance Workstations as mini computers

COMPUTER INTELLIGENCE
3344 North Torrey Pines Court
La Jolla, California 92037
619.450.1667

Intelligence Report is published monthly by:
Computer Intelligence
3344 North Torrey Pines Court
La Jolla, California 92037
619-450-1667

Editors
Paul Hairopoulos
Nancy Pantone

Fortune

(Continued from page 7)

install 2.3 mainframes, four minis, and 185 personal computers. As with the installed systems, differences surface when these corporations are evaluated according to their Fortune industry group. The table **Average Number of Planned Systems Per Corporation** (page 7) displays the average number of products planned by multi-site Fortune 1000 corporations within each Fortune industry group. These figures suggest that firms classified within the Utility group will continue to increase computing power at a faster pace than others.

Clearly, corporate America is using computers to more effectively automate the way it does business. The addition of each new system brings more computing power as systems become stronger and able to do more with each generation. This data shows the market for minis and personal computers grew rapidly from 1987 to 1989. Mainframes, on a per unit basis lost ground. However, a look at the growth of MIPs supports the hypothesis that the availability of stronger, more powerful mainframes effects the number of installed systems. In either case, it is clear computers will continue to leave their footprints within multi-site Fortune 1000 corporations.

The year-end 1989 averages reflect interviews with 26,900 sites associated with 949 of the Fortune 1000 ranked based on 1988 financial figures. The year-end 1987 figures are based on interviews with 21,400 sites associated with 934 of the Fortune 1000 ranked based on 1986 financial figures.

Suzan Sachdeva, Industry Analyst, researches Demographics and National Accounts for CI.

Compaq Leads in Dealer Satisfaction

JoeAnn Stahel

In November 1989, 206 US computer specialty stores responded to StoreBoard/Computer Intelligence's survey regarding satisfaction with manufacturers. The manufacturer receiving the highest total score was Compaq Computer with an average total score of 34.3. On a scale of 1 to 10 for each category, resellers rated Compaq the highest of all vendors for Service and Support (9.0) and among the top four for Product Reliability (9.1). In Product Availability, Compaq scored 8.4 while Pricing was rated at 7.9.

Atari, with a score of 26.4, was rated lowest in total satisfaction, scoring especially poorly in the area of Service and Support.

Over-All Satisfaction with Manufacturers

November 1989 Ratings by US Computer Specialty Stores
(A score of 40 was the highest score possible)

Manufacturer	Total Scores	Manufacturer	Total Scores	Manufacturer	Total Scores
1. Compaq	34.3	6. ALR	31.3	11. Toshiba	29.8
2. Hewlett-Packard	32.7	7. Arche	31.2	12. AST	29.8
3. NEC	32.5	8. Epson	31.1	13. Hyundai	29.5
4. IBM	32.3	9. Apple	30.7	14. Leading Edge	28.0
5. Everex	32.3	10. AT&T	30.0	15. Commodore	27.1
				16. Atari	26.4

JoeAnn Stahel, of StoreBoard/Computer Intelligence, researches the computer specialty store channel.

SOFTWARE MANAGEMENT STRATEGIES

Tandem Business Information Center

Software Management Strategies *is published by Gartner Group for subscribers to the Software Management Strategies program.*

September 5, 1985

Strategic Analysis Report
M. Braude

Vendors' Links to IBM -- Superiority or Vulnerability?

Management Summary

The theme of Gartner Group's April 24-30 Conference, which was held in Chicago, was "Integrating Workgroup/Departmental Computers: The Critical Middle Tier." As indicated during that conference, many users are looking to departmental computers not only to bridge the functionality gap between large mainframe and personal computers, but to add value to those two environments. Any products capable of fulfilling that role must clearly possess very rich communications capabilities.

In our presentation entitled "Vendors' Links to IBM -- Superiority or Vulnerability?", we analyzed these communications requirements and carefully considered the potential of some leading vendors to satisfy them. This report includes, but is not limited to, the material covered in that presentation.

To establish the necessary knowledge base, this report first explains, with historical perspective, how and why IBM is revamping and extending SNA, and then specifies the role of each of the key components. It then analyzes the strategies of three major vendors -- Digital Equipment Corp. (DEC), Tandem and Wang -- for building interfaces to those evolving components. As will become clear, it is an essential part of their strategies to acquire technology from a series of small software "boutiques" that have sprung up, and we view this situation as extremely positive for both the industry in general and the vendors of departmental computers in particular.

This report summarizes the capabilities of those "boutiques," or independent software vendors (ISVs).

Among the strategic key issues considered are:

- How can outside vendors provide an acceptable level of communication with IBM PCs and hosts compared with IBM itself?
- How is IBM likely to revamp and evolve SNA and perhaps obsolete the interface products of other vendors?
- How open is IBM being with SNA specifications? Will "catch up" ball be acceptable from other vendors?

- To what extent are the capabilities of other vendors' communications architectures really the important issue, with linking to IBM products required, but of secondary importance?

- Which vendors are likely to be most successful with their communications products?

Among the conclusions reached are:

- IBM is revamping the SNA infrastructure, using LU6.2 to replace all the other Logical Units. This does not mean that all communication devices in the field will quickly become obsolete -- LU6.2 implementations normally have a base component that will emulate the older Logical Units. Also, this transition will not occur quickly.

- While IBM may not be particularly concerned with the fallout from this evolution, neither is it primarily motivated by competitive or marketing considerations. The changes are needed to simplify SNA and to give it some of the higher connectivity it needs.

- Vendors of departmental computers are well aware that, to be players, they must handle these new SNA technologies; and they are practical enough to buy these software capabilities from others. This maturity indicates a welcome change in their philosophies.

- While SNA is not really open, there is no real barrier to these vendors' coexistence with the extensions and changes to SNA.

- DEC and Tandem are likely to enjoy success in developing these technologies. DEC's strength is its aggressiveness in announcing and developing the new technologies, while Tandem's is in its well-developed base SNA capabilities.

- Those evaluating the purchase of equipment that is to maintain SNA communications (with an IBM architecture environment) must ensure the capability of the vendor to supply these communications technologies.

When analyzed this way, the area -- communications among different levels of the machine hierarchy -- becomes reasonably transparent, and there is an effective set of standards for technology that is not only understood but also soundly implemented.

CONTENTS

Introduction

The intent of this Strategic Analysis Report is to analyze how well the vendors of departmental systems -- including IBM -- fit in with and accommodate the strategic data streams that IBM has defined and evolved. These data streams (DCA, LU2, LU6.2 etc.) were defined to meet needs perceived by certain IBM development facilities at certain points in time, and then were evolved, enhanced and built upon as customer and internal requirements changed. Therefore, it is not necessarily the case that all of IBM's midrange offerings work well with these data streams. IBM will continue to install support for these data streams on most, if not all, of its hardware products (note that the 3820 laser printer announced in February 1985 is an LU6.2 device), and will continue to use those data streams as the infrastructure for the functionality required by its customers -- increased connectivity, network recovery, network management, distributed database, etc.

The next part of this report will analyze the strategic directions of these data streams, and will discuss the impact of those directions on users and vendors. As part of that discussion we will analyze IBM support of those data streams -- in other words, how IBM has implemented those facilities on its own hardware products. We will see further indications that IBM is very software bound, and that it is all the software developers can do to reasonably support advances in hardware.

There will be discussion of what several of the more advanced vendors of departmental computers -- we will label them "OEMs" for brevity -- are doing to live with these data streams. We will emphasize that some of them are using ISVs to provide the technology needed. The day is past when any one vendor, probably including IBM, has enough design capability and development resource to provide all the technologies needed by these vendors for them to grow the way that they want to. This situation is giving rise to several alliances -- a trend that is accelerating. This report will profile the major players.

Extensions and Revamping for Systems Network Architecture

Figure I is a high-level view of how an OEM machine -- from DEC, Wang, Hewlett-Packard (HP), Tandem, Data General (DG) or even IBM -- would fit well within an IBM environment. The overall structure of the hierarchy is well known -- according to Gartner Group Comtec data, the share of the mainframe market belonging to the IBM architecture (IBM itself plus the plug-compatibles) in 1984 was 77 percent, up from 75 percent in 1983. IBM was very strong on the desktop (PC and clones), that architecture winning 55 percent of that market. Consequently, most observers are conceding to IBM the mainframe and the desktop markets. But IBM clearly is not dominating the midrange; in fact its market share in this range last year was 21 percent, down from 22 percent in 1983. This trend may continue, particularly if the OEMs follow up on and continue the announcements that will be discussed here.

Most, if not all, OEMs have bought the fact that they must support SNA. While their performance in delivering product, particularly timely product, has historically left something to be desired, at least they now understand the technical requirement. A problem for the OEMs is that IBM has introduced some higher-level layers, or extensions to SNA, such as DIA and SNADS, that can be difficult and unclear. In addition, IBM is doing significant things with LU6.2, and these will be discussed later. It is

our position that the OEMs must build upon their base SNA support and handle these data streams, if they want to be serious players in this market.

This report will not spend much time on PC software. We use the term "PC support" for the type of facility announced by DEC in April. DEC's facility -- DECnet-DOS -- provides the PCs with a full capability to participate in the VAX architecture instead of having PC virtual diskettes maintained on a larger computer. It also provides means to convert files between PC-DOS and VAX/VME format. Tandem has announced a similar facility.

Figure II represents the "old" SNA. By "old" we mean the components that most users either have installed or are in the process of installing, and whose next stage is not yet fully available to them. While these components/data streams are not yet perceived as obsolete, it is our contention that they are.

Figure I

OEM Machines in IBM Hierarchy

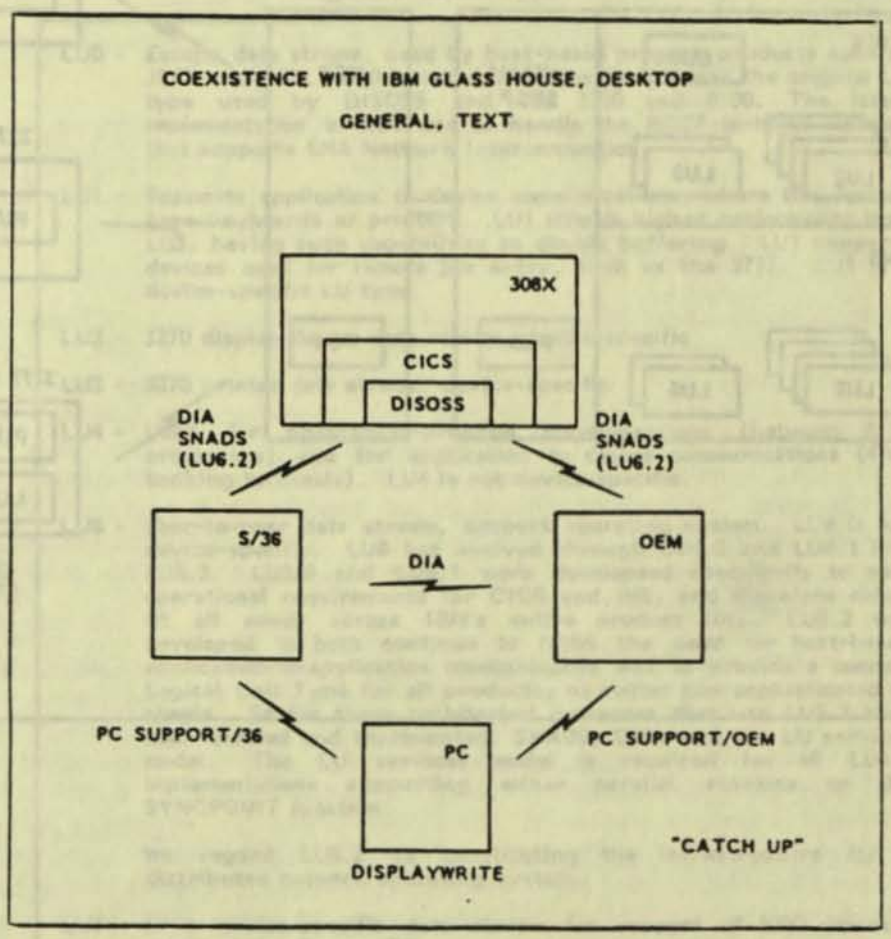

Figure II

Old SNA Components

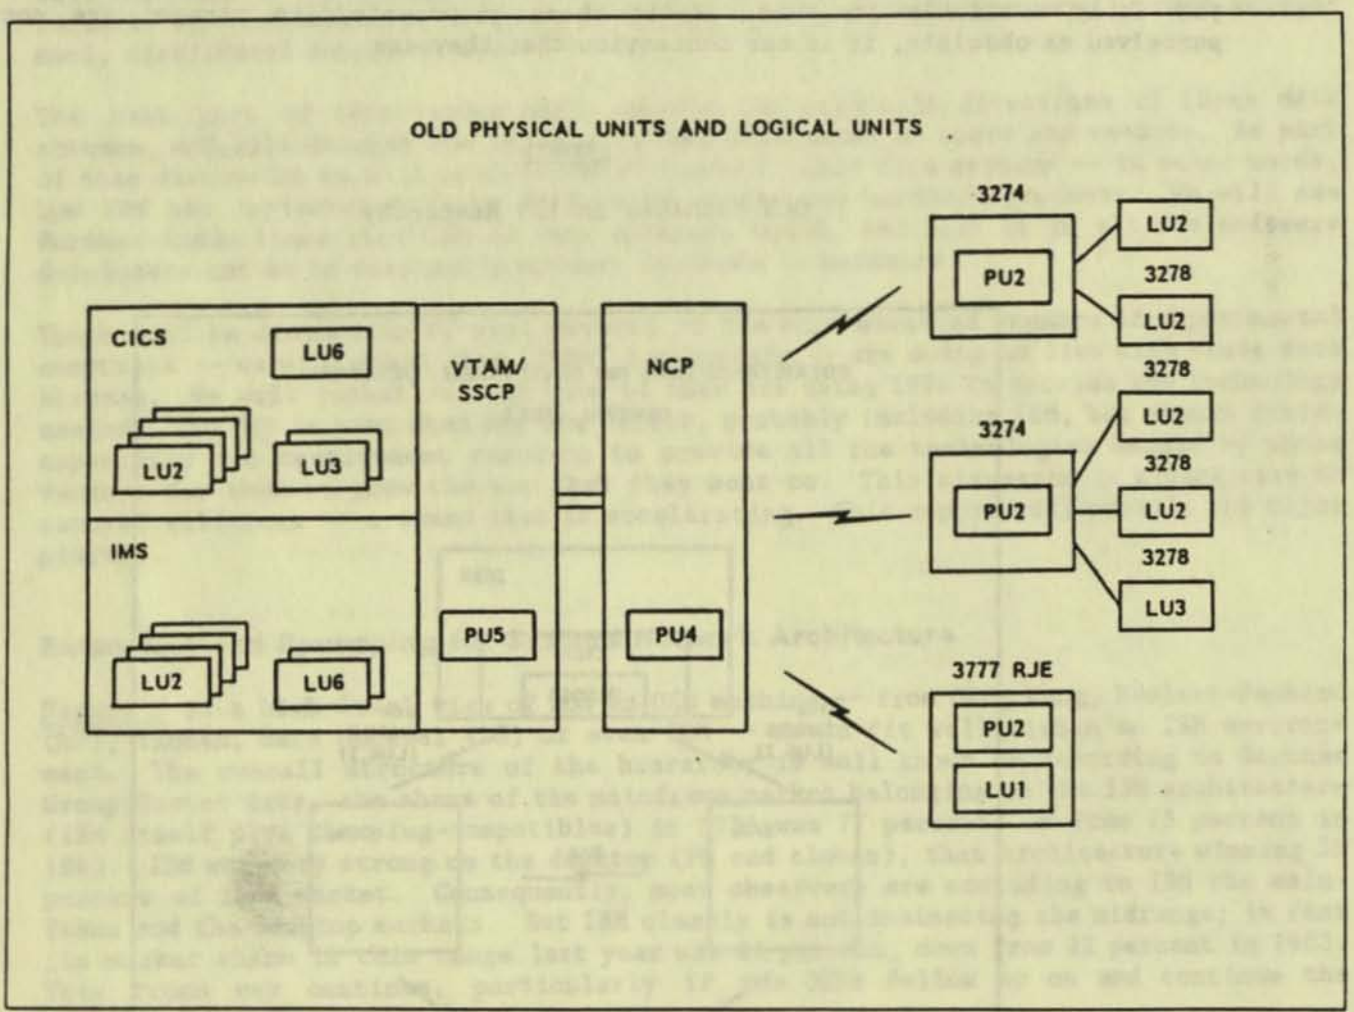

IBM's Basic Building Blocks -- LUs

An exploration of SNA technology in general and Logical Units and data streams in particular will aid understanding of this report. IBM's technical definition of a Logical Unit (LU) is that it is an end user's "port" into the network; it is the means by which an end user accesses the network in order to communicate with another end user, where an end user means either a terminal or an application program. It is much more useful to think of an LU either as a data stream -- a set of rules for encoding data -- or as a set of coded control characters together with protocols that govern data exchanges between end users. This report will identify Logical Units with data streams. Sometimes, however, we will refer to LUs also as protocols. So LU2 will denote the 3270 interactive data stream, rather than a particular 3270 terminal. This perspective will make the entire presentation much more clear, and in fact is more true, architecturally. With LU6.2, IBM may clearly separate the protocols from the data stream. However, we will not analyze that possibility here.

Figure III is a complete listing of the SNA Logical Unit types. The important ones for this discussion are LU2 and LU6.2, and we will summarize both of them.

Figure III

Logical Unit Types

LU0 -	Escape data stream, used by host-based program products such as JES2 and DSX. LU0 is not device-specific. It was the original LU type used by DISOSS and the 3790 and 8100. The latest implementation of LU0 was to handle the NCCF-to-NCCF session that supports SNA Network Interconnection.
LU1 -	Supports application to device communications, where the devices have keyboards or printers. LU1 affords higher performance than LU3, having such capabilities as double buffering. LU1 supports devices used for remote job entry, such as the 3777. LU1 is a device-specific LU type.
LU2 -	3270 display device data stream, device-specific
LU3 -	3270 printer data stream, device-specific
LU4 -	Used for program-to-program communications (between 8100 processors) and for application to device communications (4700 banking terminals). LU4 is not device-specific.
LU6 -	Peer-to-peer data stream, network operating system. LU6 is not device-specific. LU6 has evolved through LU6.0 and LU6.1 into LU6.2. LU6.0 and LU6.1 were developeed specifically to meet operational requirements for CICS and IMS, and therefore didn't fit all needs across IBM's entire product line. LU6.2 was developed to both continue to fulfill the need for host-based application-to-application communication and to provide a common Logical Unit Type for all products, no matter how sophisticated or simple. So far three architected processes that use LU6.2 have been defined and implemented: SNADS, DIA, and the LU services model. The LU services model is required for all LU6.2 implementations supporting either parallel sessions or the SYNCPOINT function. We regard LU6.2 as constituting the infrastructure for a distributed network operating system.
LU7 -	Is a device-specific data stream for support of 5250 devices attached to a System/36.

The first defined Logical Unit types were Logical Unit types 0, 1, 2 and 3. LU2 is used for communications between a 3270 terminal and an application program. This data stream includes commands for such field attributes as color, blinking and modifiability. LU3 is just a presentation layer modification of LU2 for 3270 printers, running in what is called 3270 data stream mode. For example, there are commands for page eject and new line, and the read commands are invalid. LU1 is basically a richer, high-performance data stream for printers that might have keyboards. It is synonymous with the so-called SNA character stream (SCS), which includes facilities such as double buffering.

Remote Job Entry (RJE) stations such as 3777's are supported by the LU1 data stream. Printers such as 3278's and 3289's will support either LU1 or LU3, depending upon the microcode in the 3274 controller to which they are attached.

Logical units 1, 2 and 3 are "master-slave" data streams. One component -- the host -- is the Primary Logical Unit and controls the interaction (session) between the two entities, so that it alone can build sessions and tear them down, initiate error recovery, etc.

Logical Unit 0 is really an escape data stream; it features the absence of constraints. It is used by those who want to build their own types of data streams and device-to-device sessions, and it has been used to build application-to-application communications, particularly for specific device types. We would not encourage the use of this facility.

Potentially, an organization with a great deal of skilled manpower and very unique requirements may find LU0 appropriate. IBM itself has used it for writing program-to-program protocols for specific uses, such as certain JES facilities. We tend to regard LU0 as a predecessor facility to LU6.2, and one that will tend not to be used further as LU6.2 becomes fully developed.

Similarly, LU4 is seen as an intermediate step to LU6.2. These sessions may be either master-slave or peer-to-peer. IBM used LU4 before LU6.2 was ready for applications/devices sessions that needed more flexibility than LUs 1, 2 and 3 were able to provide. LU4 was used with 5520 and 8100 peer-to-peer communications, and to support the 6670 printer. Like LU0, LU4 should not see much further use. Nothing is completely impossible, however, and there is no doubt that, somewhere deep in the bowels of some development laboratory, someone is contemplating use of LU0 or LU4 because there is something about LU6.2 that he does not like. Perspective on LU6.2 should be based on a knowledge of the "new" SNA components, and should be seen in the context of what is really happening within SNA.

Figure IV shows the "new" SNA components. There are two obvious differences between Figure II and Figure IV. First is the connectivity, indicated by arrows, between the "S/36" and "3274" boxes. This should not be a surprise; LU6.2 is commonly identified with the direct connectivity of departmental computers and cluster controllers to one another. The second difference is the absence of LUs 1, 2 and 3 from Figure IV. Those data streams have no place in the "new" SNA. We strongly believe that LU6.2 will replace LUs 1, 2 and 3. This succession will obviously have very strong competitive effects on the OEMs, who have only comparatively recently delivered support for the older data streams. This of course does not mean that 3270's will become instantly obsolete. LU6.2 consists of sets of functions (see Figure IVa), of which the base set

is supported by every product and the others are optional -- indeed, IBM has called them option sets. The base set normally includes support for LU2, or 3270 data streams. However, the August announcement of LU6.2 support for the Series/1 did not include support for the older Logical Units.

Figure IV

New SNA Components

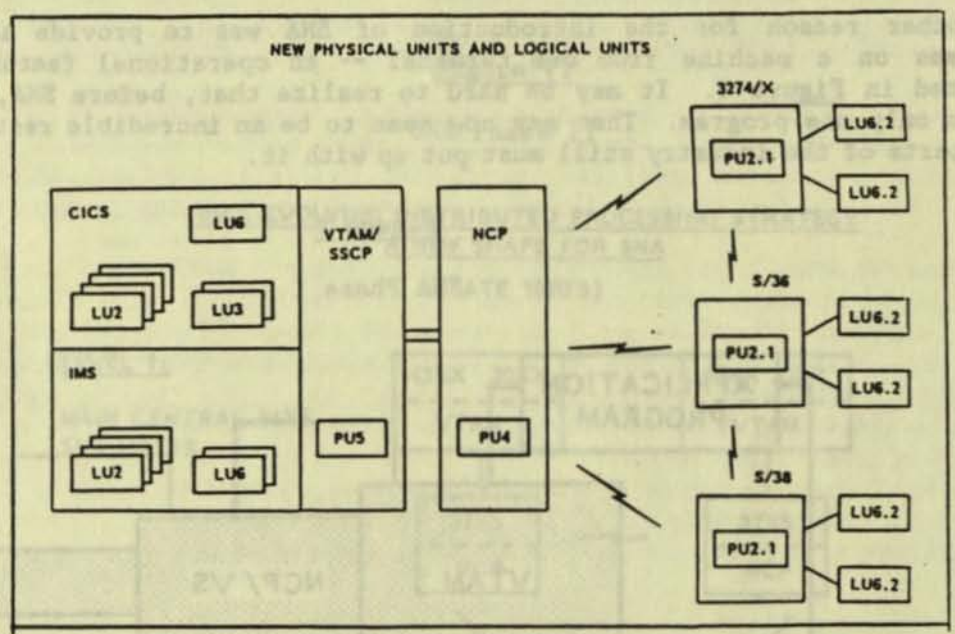

Figure IV A

LU6.2 Function Sets

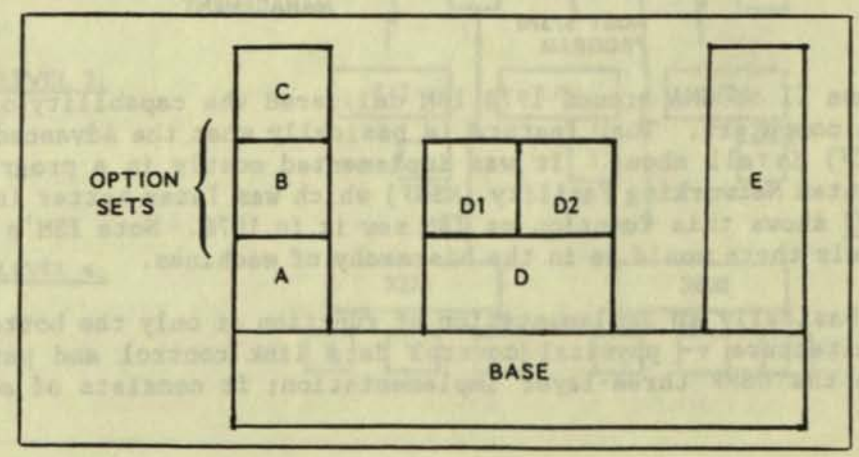

Later in this report, we will discuss the functional necessity for that replacement. We view this infrastructure replacement as "Phase III" of SNA. IBM is not stressing that, nor calling attention to it in any way. Why is that? What were Phases I and II?

Phase I was the 1974 introduction of SNA, which IBM made for two good reasons. One reason was to remove barriers to customer growth and internal economies. By IBM's own count, there were 35 telecommunications access methods and 15 link control procedures at that time, and these supported only about 200 communications products, including software. Clearly, this was a disaster area, confusing for customers and expensive for IBM. One objective of the initial SNA function, therefore, was to consolidate and simplify a large number of access methods and devices.

The other reason for the introduction of SNA was to provide access to multiple programs on a machine from one terminal -- an operational feature for customers, depicted in Figure V. It may be hard to realize that, before SNA, a terminal could access only one program. That may now seem to be an incredible restriction, although some parts of the industry still must put up with it.

Figure V

SNA Phase I

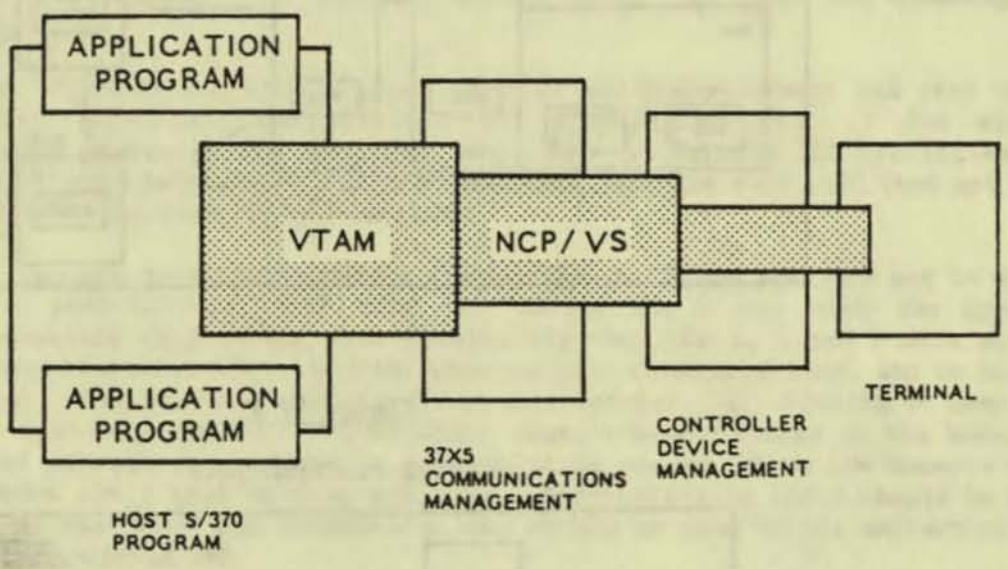

With Phase II of SNA around 1978 IBM delivered the capability of a terminal to access multiple computers. That feature is basically what the Advanced Communications Function (ACF) is all about. It was implemented mostly in a program product called the Multi-System Networking Facility (MSNF) which was later better integrated (into VTAM). Figure VI shows this function as IBM saw it in 1978. Note IBM's uncertainty about how many levels there would be in the hierarchy of machines.

MSNF is basically an implementation of function at only the bottom three layers of the SNA architecture -- physical control data link control and path control. LU6.2 is built on the MSNF three-layer implementation; it consists of additional function in

layers four (transmission control), five (data flow control) and six (presentation services), Figure VII summarizes these layers.

So we have set the stage for Phase III, which takes the connectivity still further (Figure IV). With Phase I IBM provided connectivity within a mainframe. Phase II made connectivity across multiple mainframes possible, and now Phase III provides connectivity directly across cluster controllers and departmental computers. We want to show that LU6.2 is much more than a protocol that provides more connectivity.

Figure VI

SNA Phase II

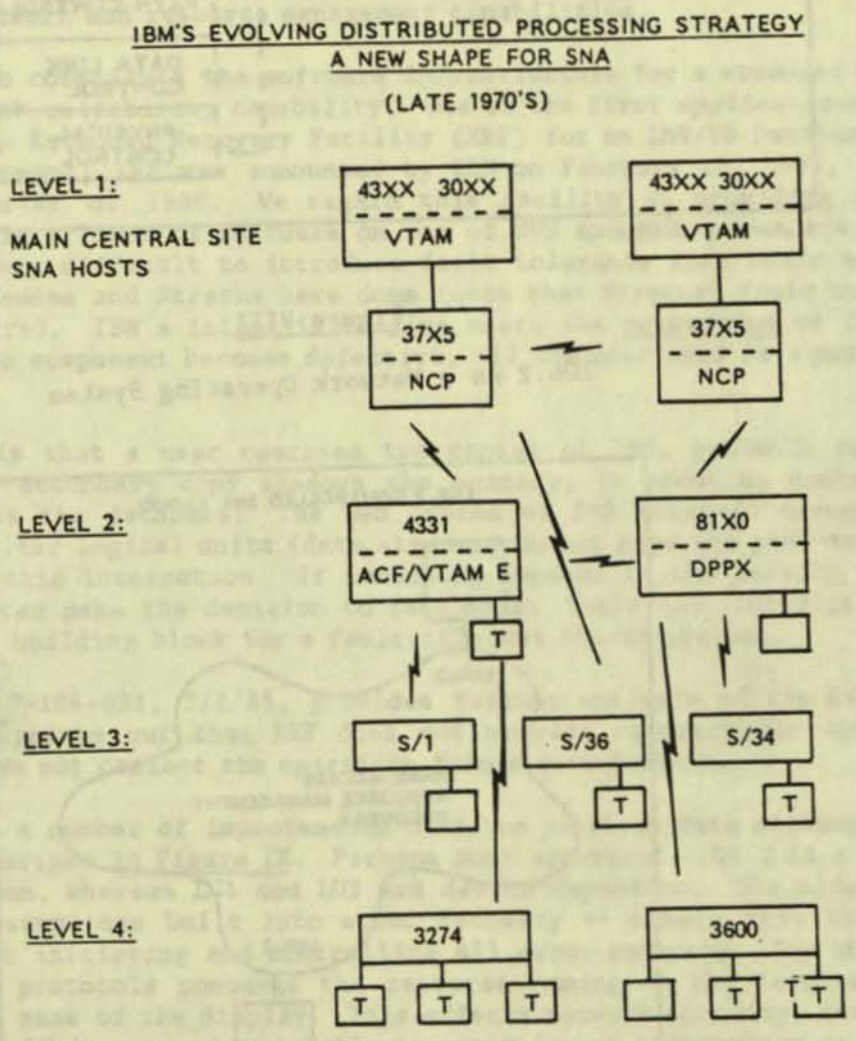

Figure VII

IBM's Architecture and Layering

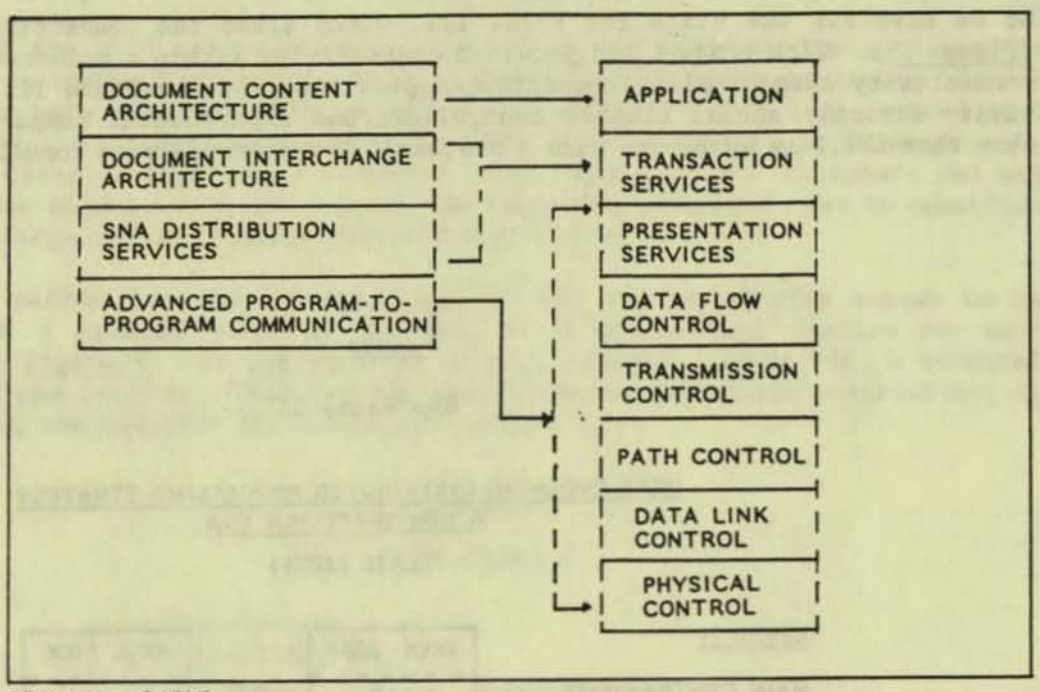

(Courtesy of IBM)

Figure VIII

LU6.2 As a Network Operating System

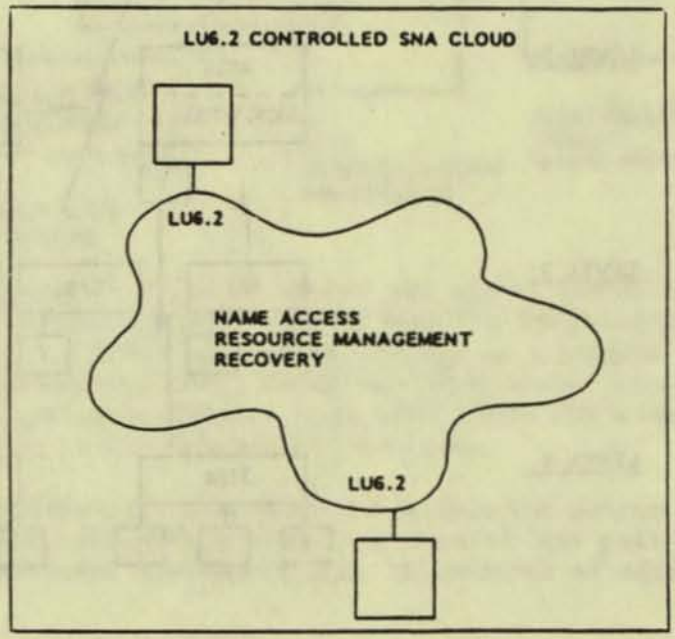

In fact, LU6.2 is a Network Operating System. What does an Operating System (OS) do? Two of its primary functions are resource management and error recovery. Key to the resource management function is the fact that LU6.2 (support) provides a session manager which allocates temporary ownership of a pool of communications sessions among competing applications. The design assumption is that the applications are transactions with shorter lifetimes than the communications sessions. For example, as can be seen from Figure VIII, CICS transactions or IMS transactions at the host might be sending many print jobs to the 3820 printer. If the LU1 or LU3 data streams (protocols) were used, the host would have to build up a new session for each print job and then tear it down. That scenario would cause tremendous overheads -- not just tne obvious CPU load, but also a lot of traffic over the communications lines. Not everyone is aware of the number of messages (packets) that must be exchanged by nodes when a session is being activated between them. By using the session management facilities of LU6.2, one session is maintained, and it handles the repeated application requests.

For the sake of completeness, it should be pointed out that multileaving of data streams within a session could be accomplished in the "old" protocols, and in fact it was done in LU1 for 3777 printing. But these facilities were not a suitable base for generalized network and resource management capabilities.

LU6.2 will also constitute the software infrastructure for a stronger network management and network reliability capability. One of the first applications of this aspect of LU6.2 is the Extended Recovery Facility (XRF) for an IMS/VS Database Data Communications environment; XRF was announced by IBM on February 12, 1985, for shipment in the third quarter of 1986. We regard this facility as providing simulated fault tolerance; it is a layer of software on top of MVS and IMS, which are so complex that IBM finds it very difficult to introduce fault tolerance into their architecture, in the way that Tandem and Stratus have done (note that Stratus' fault tolerance doesn't involve software). IBM's intent is to give users the perception of fault tolerance, so that if some component becomes defective, all the user sees is a prolonged response time.

What happens is that a user operates two copies of IMS, probably on two different machines. The secondary copy shadows the primary, in updating control blocks, but does not access the database. The two copies of IMS maintain communications using LU6.2; the earlier logical units (data streams) do not have the peer-to-peer capability to support this interaction. If something happens to the primary version of IMS, the secondary can make the decision to take over. Therefore, LU6.2 is the core of the communications building block for a fault-tolerant IMS/VS system.

Research Note P-104-031, 3/1/85, provides further analysis of the Extended Recovery Facility, and points out that XRF does not provide recovery for all failures; for example, it does not protect the operation from a disk failure.

LU6.2 provides a number of improvements over the previous data streams/protocols, and these are summarized in Figure IX. Perhaps most important, LU6.2 is a device-independent data stream, whereas LU1 and LU3 are device-dependent. The older LUs also have restrictive assumptions built into error recovery -- namely that the primary LU is responsible for initiating and controlling all error recovery. The third restriction of the older protocols concerns the resource naming -- the terminal's LU name is implicitly the name of the display. This effects network security, since the operator at the terminal LU is assumed to be able to enter log-on information on the display.

Figure IX

LU6.2 Functions

EMULATES LU1, 2, 3

PROTOCOL COMPATIBILITY BETWEEN DIFFERENT WORKSTATIONS

STANDARD APPLICATION PROGRAM INTERFACE

PEER-TO-PEER COMMUNICATION

BASE FOR NETWORK MANAGEMENT

TOKEN-PASSING RING CONNECTIVITY

The key point for users and vendors to know is that IBM will evolve the SNA networks in the field out of the earlier Logical Units and into LU6.2. As indicated earlier, LU6.2 is normally implemented with a base set of functions that furnish support for devices now in the field -- e.g., this base set emulates LUs 1, 2 and 3 protocols. Above the base sets, LU6.2 provides option sets which in turn provide the session management functions, making possible the peer-to-peer capabilities. As IBM ships devices that support LU6.2 (see Figure X) -- for example, the 3820 laser printer -- it can gradually introduce function that requires the capabilities of the option sets. This function would probably be in the areas of network recovery and performance. Vendors supporting only 3274 emulation (LUs 2 and 3) will not be able to provide those levels of support, and will be perceived as purveyors of obsolete technology.

IBM's Basic Building Blocks -- PUs

Physical Units are the other important entity within SNA. A Physical Unit manages, monitors and controls the physical resources (such as attached links and adjacent link stations) of a node. An important new Physical Unit (PU) type is currently being introduced -- namely, PU type 2.1. It makes sense to think of a PU as just a hardware/microcode implementation that processes SNA data streams. PU type 2 is in essence a cluster controller such as a 3274; an RJE station could also be described as a PU2. The facilities of a PU2 device are not sufficient to exploit fully the capabilities of the LU6.2 data stream/protocols. To do that, an advance over PU2 capabilities is needed, and IBM has designated that PU2.1 .

The PU2.1 devices have two key capabilities that PU2 hardware lacks. The first is that there can be multiple "upstream" links, as shown in Figure IV. The second is the so-called "Single Node Control Point" facility, whereby the device itself can initiate

sessions and tear them down. Unless PU2.1 and LU6.2 are both operative in a node, it is necessary to involve the host in any session establishment/disestablishment.

The LU6.2 and PU2.1 technologies are the basic building blocks, the raw technology infrastructure, for IBM's future networking technologies. There will be no further implementations of LU4 and probably relatively few of LUs 1-3 and PU types 1 and 2. IBM will be migrating all those to the LU6.2/PU2.1 technologies over the next five years.

Figure IX A presents a comprehensive picture of IBM's layered structure for hardware and communications. Figure X shows the status of support for these technologies within IBM's present product line.

Figure IX A

IBM's Comprehensive Layered Structure

<table>
<tr><td>APPLICATION</td><td rowspan="3">SOFTWARE</td><td colspan="5">APPLICATION PROGRAMS</td></tr>
<tr><td>TRANSACTION SERVICES</td><td>JES</td><td>TSO</td><td>CICS</td><td>IMS</td><td></td></tr>
<tr><td>PRESENTATION SERVICES</td><td colspan="5"></td></tr>
<tr><td>DATA FLOW CONTROL</td><td rowspan="3">FIRMWARE</td><td colspan="5" rowspan="2">VTAM OR TCAM
NCP</td></tr>
<tr><td>TRANSMISSION CONTROL</td></tr>
<tr><td>PATH CONTROL</td><td colspan="2">S/370 CHANNEL</td><td colspan="2">S/370 CHANNEL</td><td>SDLC</td></tr>
<tr><td>DATA LINK CONTROL</td><td rowspan="2">HARDWARE</td><td colspan="2" rowspan="2">S/370 CHANNEL</td><td colspan="3">3725</td></tr>
<tr><td>PHYSICAL CONTROL</td><td>CHANNEL ADAPTER</td><td>COMM. LINES</td></tr>
</table>

Figure X

IBM Support for SNA Extensions

S/370	DCA	DIA	SNADS	LU6.2
(Disoss/370)	Now	Now	Now	Now
System/38	11/85	9/85	9/85	Now
System/36	Now	Now	Now	Now
5520	Now	Now	Now	Now C
IBM PC	Now	Now	No	No
Displaywriter	Now	Now	No	Now C
Scanmaster	No	Now	No	Now C
3820	No	Now	No	Now C
8100/DPCX	Now	Now	SOD	Now
Series/1 RPS	No	SOD	SOD	Now

SOD means statement of direction

C means closed implementation -- supporting specific IBM-written applications.

Extensions to Systems Network Architecture

We now understand how IBM is evolving the SNA infrastructure, and the competitive implications of that great change. Almost as crucial are the extensions to SNA that IBM is building -- SNA Distribution Services (SNADS), Document Content Architecture (DCA) and Document Interface Architecture (DIA).

• SNA Distribution Services (SNADs): SNADS is just what the title suggests -- a delivery system that operates within an SNA network. Why is such a thing needed -- would simple file transfer programs not suffice? Well, the problem is that those file transfer programs are synchronous, i.e., both the sender and the receiver must be active at the same time, which is not always the case in an office environment. One of the two key points about SNADS is that it is an asynchronous delivery system -- an office worker could use SNADS to transmit a document to someone else who is away from his desk. The other key point about SNADS is that it presupposes the existence of LU6.2 -- but LU6.2 certainly does not presuppose SNADS, so the two are not synonymous.

What all the above says is that SNADS is just a store-and-forward delivery system well suited to an office environment, which is of a distributed nature and has no single location that should control everything. SNADS is essential to a distributed environment.

Nothing in the SNADS architecture prevents a network of small processors from implementing a store-and-forward network using SNADS and without involving a large S/370 processor with DISOSS. In fact, Personal Services/36 and Personal Services/38 will include those functions.

• Document Content Architecture (DCA): Right now, document distribution is the application that is served by SNADS. The significance of DCA is that it is a machine-independent data stream for text -- IBM deserves accolades for pioneering this technology. That several hardware products support DCA does not necessarily mean that documents transmitted between them will look exactly the same, for the supporting products may not have the same text capabilities; for example, one of them may not support footnoting. What mutual DCA support does mean is that, to the extent that the two machines have common function, the DCA data stream will be interpreted the same way on both, and the printed documents will look the same.

• Document Interface Architecture (DIA): A frequently made analogy likens DCA to a letter and DIA to the address on the envelope. Using this prosaic comparison, SNADS would be the post office vehicle. But this analog is a bit understated. It would be more accurate to say that DIA is the mailman/sorter. DIA is really a communications program as well as an architecture, and it gives transmission orders to SNADS and other software components. That relationship is displayed in Figure XI.

Figure XI
SNA Extensions

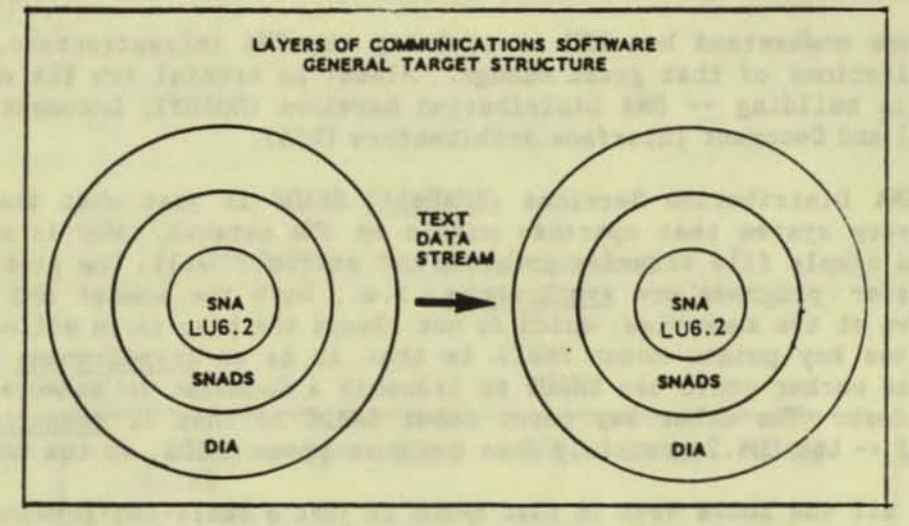

Stated more fully, DIA is a set of defined services performed by peer communications processes. These processes can be people, applications programs or devices, a process being a source or recipient of information.

While it is IBM's intent that DIA use LU6.2 facilities (often through SNADS), DIA function is independent of the LU type; it can and does use LU2 or other facilities. Office systems such as S/36 and the 5520 use DIA to distribute and specify the use of information, and whether the information is transported directly by LU6.2 or mediated by SNADS depends on the type of transaction. For example, if the source and recipient nodes are not attached to the same office system, SNADS must be used. So, between autonomous office systems, SNADS is always used (by DIA). Within an office system, DIA may or may not use SNADS.

Of course, end users are not directly aware of DIA, SNADS or LU6.2, which merely constitute the infrastructure. An end user interfaces to Personal Services, which invokes DIA, SNADS or LU6.2. As part of Personal Services, IBM will deliver the capability shown on Figure XII. Here DIA operates on two PCs. If, for example, there is a Displaywrite 3 document on one PC, DIA will transmit it to the other, using an asynchronous link (or any other) for the communications function.

Figure XII
DIA Operating on PCs

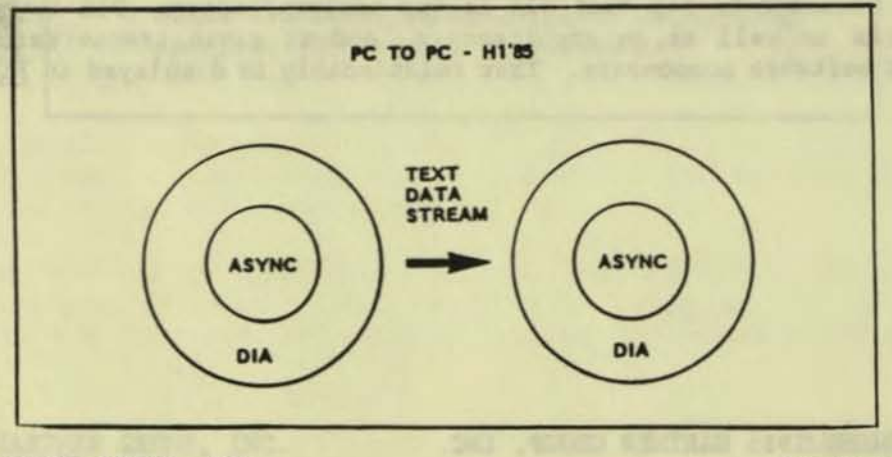

In other words, DIA is independent of SNADS and of SNA. It is not just a conglomeration of coding conventions written on an envelope, and interpreted by some other piece of software with a view to transmitting information around an office. DIA itself consists of several services. Document Distribution Services are concerned with the functions of distribution and delivery of information, discussed above. DIA is usually thought of as a document profile, or a structured form that conveys the characteristics of information on a document, the so-called address on the envelope -- but that is only one part of DIA. It has other services -- Applications Process Services can initiate programs and take other actions. For example, the DIA function set includes such commands as OBTAIN (to request the delivery of documents), DELIVER, LIST (list available documents and their status) and SIGN-ON.

Why are these points poorly understood? Because there is a tremendous education problem here. The inadequacy of IBM's explanations does not mean that the company is in any way nefarious. It arises from the fact that the marketing representatives and systems engineers too are a little confused, not having been trained to explain complicated software to clients. This confusion plays into the hands of IBM's competitors to some extent, but it also creates a barrier to the ability of IBM's customers to grow and to manage their installation.

To understand these three "extensions" to SNA, it is important to relate them to the SNA layering. DCA is purely an application function, and is therefore in the "application" layer.

Figure X shows which of these facilities are operative on particular IBM products. IBM's implementation is incomplete, and most of the pieces will to be delivered during this year. Usability may not be high. IBM has its problems, pursuing what is really a brute force approach to making all these different environments (S/36, S/370, PC) look the same by providing Displaywrite and Personal Services, hooked to DISOSS, as a software front end.

Independent Software Vendors (ISVs)

Independent vendors -- for example, minicomputer suppliers -- may react with concern and trepidation over the expense and time involved in supporting the data streams/facilities described earlier. After all, SNA is not really "open." (For a discussion of this issue see Research Notes K-321-040, 4/1/85 and K-321-074, 6/19/85.) It took over six years for most of the independent vendors to deliver support for the more elementary SNA data streams, which amounted to 3274 emulation and remote job entry support.

These vendors' worries may be addressed by the presence of a new factor in the marketplace -- one that was not present in the 1970's. There are now Independent Software Vendors (ISVs) who claim to provide SNA software technologies, for a price. In fact, there is now a whole infrastructure of small software boutiques to supply these technologies. Their existence indicates a change in the character of the marketplace as user companies realize they cannot (afford to) develop all these technologies themselves. Alliances much be formed, and software expertise can be packaged and sold.

What are these boutique-like companies? We will look first at the ISVs that provide document translation services, which in the context of this report would focus on the

translation of, to and from IBM's DCA of other vendors' text data streams. These text data streams correspond to such products as DEC's All-In-One, Wang's Office and Data General's CEO, as well as to PC products such as Multimate.

Document Conversion Services

Dominating this niche is Soft-Switch Inc., which, founded in 1979 as "Integrated Technologies Inc.," had approximately 75 employees at the beginning of 1985. Soft-Switch Translation Services are designed to support the exchange of revisable form text among the hardware products of various vendors. All translation is really a two-step process -- Soft-Switch has defined a "canonical" format, which is an enhanced version of IBM's Level 3 (revisable) DCA. Figure XIII depicts the process, the first step of which is to translate the text data stream from the originating device to Soft-Switch's canonical format. A second translation is then performed from the canonical format to the (revisable) text data stream of the target device. Figure XIV lists the products handled by Soft-Switch.

Figure XIII

Soft-Switch Document Translation Process

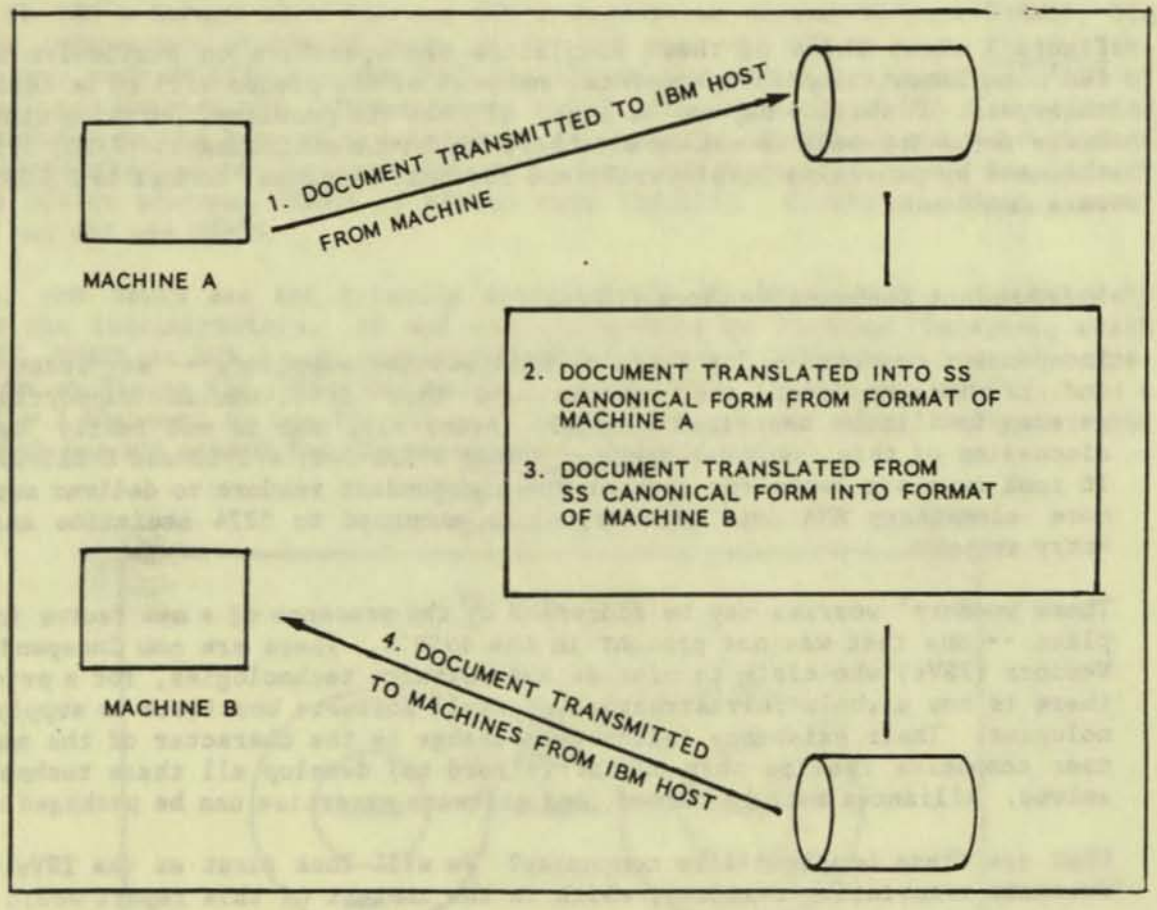

Figure XIV

Products Handled by Soft-Switch[1]

IBM DISOSS

IBM Profs

IBM DCF (Script - under CMS, TSO)[2]

IBM Displaywriter

IBM Displaywrite 2 and 3

IBM 5520

IBM OS/6

IBM 6640

IBM 6670 Information Distributor (Laser Printer)

MultiMate on the PC

NBI OASYS 8, 64, 3000 and 4000

NBI 4000

Wang OIS

Wang VS

Wang WPS

Xerox 850/860

---

[1] Gartner Group has not verified all of these.

[2] Soft-Switch does not support Easy Script, user macros, symbolic variables and GOTO control.

The Soft-Switch software, also called Soft-Switch, runs as a job in the MVS environment; under VM it runs as a service (virtual) machine. Soft-Switch has between 30 and 60 installations. The typical license fee for the installed product is stated to be $75,000.

It is very important for users to understand that the Soft-Switch product cannot be all things to all people -- and this is so because the office products of various vendors are so different from one another. Soft-Switch is an excellent method for translating routine memos and business letters. One installation uses it to gather documents produced in the field on PCs onto an IBM mainframe and then print them. However, the package cannot necessarily handle complex proposals and long documents that include spreadsheets, multiple columns and footnotes and that in other ways fully exploit the advanced features of the environment within which it was created.

Soft-Switch Inc. has marketed its technology to the OEMs, as well as to users, with significant success so far. In fact, most of the major OEMs (with whom a subsequent section of this report deals) including DEC and Wang, have licensed Soft-Switch's technology.

Until recently, Network Applications Incorporated (NAI) was Soft-Switch's principal competitor in this niche. NAI was founded in 1983, and had about 13 employees at the

beginning of 1985. NAI's product, TransNET, provides interchange between DISOSS and Wang OIS and VS systems. The program code resides on the MVS host system for translation and interchange. For library service, a Wang glossary is employed on the OIS or VS system. NAI had announced its intention to support DEC and Data General systems by year-end 1985.

It was announced in May 1985 that Soft-Switch had acquired NAI. Their product lines were competitive as well as complementary -- the basic products implemented similar function; Soft-Switch had implemented support for a wider variety of vendors while NAI had made further progress in developing support for LU6.2 and a DISOSS interface.

There are other vendors in this business. The only other one we will mention is Software Research Corp. (SRC). This firm was founded in 1978 and its employees numbered approximately 121 at the beginning of this year. Document conversion software is only one part of SRC's business, while it is Soft-Switch's area of main concentration. SRC's Docupower has not been sold to end users, but has been licensed to Hewlett-Packard, Lanier and Honeywell.

We will not explore this document conversion niche further. The main point has been made -- technology can be purchased, so that OEMs need not incur a lot of delay and internal expense to develop it.

SNA Mechanisms for Information Transfer

- The Systems Center

This is an ISV that does not at this time sell SNA technology per se. It has positioned itself to attack a specific niche -- SNA mechanisms for information transfer. The Systems Center's goal is to provide a family of products -- currently called the Network Data Mover (NDM) product line -- that can move data, text and graphics information among microcomputers and mainframes in an SNA network and do so in an easy-to-use, standardized manner under either program or manual control. Since the Systems Center is not involved at this time in DCA or LU6.2 support products, it will not be discussed further. More information can be found in Research Note C-706-018, 2/1/85.

- Communications Solution Inc. (CSI)

For some time CSI has had a product line -- Access/SNA -- that consists of packaged support for SNA data streams. In June 1985, it extended the line to include portions of LU6.2 support, with its Access/SNA APPC package. This software uses IBM's Conversational Verb Interface to provide a structure for the interaction of programs across SNA networks, involving hardware from IBM and OEMs. It also supports Control Operator Verbs that activate and deactivate LU6.2 devices. While it supports a base set of protocols, including session allocation and data flow control, it does not yet include all aspects of direct peer-to-peer communications.

CSI also announced an ACCESS/DIA package in June. It is CSI's intent initially to give systems such as PCs and file servers access to IBM office products supporting DIA, and then later to supply ACCESS/DIA to emulate those systems. However, it is still too early to see the full implications of this product.

The entire Access/SNA product line is written in C language and hence can be ported and installed relatively straightforwardly. The publicly announced price for Access/SNA APPC is $275 per unit, and the package will be available in September. Future enhancements will include network control facilities, multiple data-link support and the remaining peer-to-peer communications facilities.

- Rabbit Software Corp.

In June, Rabbit Software Corp. also announced an APPC product supporting the Conversational and Control Operator Verb Interface; the product is scheduled for shipment in the third quarter of 1985. Called "SNA-Plus," and priced at around $1,100, it includes some DIA functions such as file transfer, application services, document distribution and library services. SNA-plus is also written in C language and is designed for Unix environments.

- Systems Strategies Inc. (SSI)

The fourth ISV here is a wholly owned subsidiary of AGS Computers. SSI also markets emulation programs targeting the OEM community, and its software is also written in C language. Its current product line includes SNA/3270 and BSC/3270, which provide, respectively, emulation of LU2 (3274 functions) with SDLC, and bisynchronous data link control. It also markets an SNA/RJE product. These three products can be licensed at the following fees: BSC/3270: $60,000; SNA/3270: $60,000; and SNA/RJE: $70,000. SSI also offers an Applications Program Interface (API) that enables a program on the OEM box to appear as a 3278 terminal; this software can be licensed for $20,000. The company has stated its intent to develop an LU6.2 product, hopefully in early 1986.

It is our understanding that SSI has been involved in supplying its technology to Siemens AG, AT&T, Tolerant Systems Inc., Megadata Corp. and Sequoia Computer Corp.

Communications With DECnet

Interlink Computer Sciences Inc. of Fremont, Calif. has targeted a market niche related to but distinct from those of the ISVs discussed so far. Interlink sells software that, operating on IBM mainframes, provides communications with DECnet. Its first product, IBM MVS/DECnet Gateway, allowed programs in the MVS environment to access data and files on DECnet. Interlink claims to have sold over 40 copies of the software. More recently, Interlink released the IBM VM/DECnet gateway, which provides programs in the VM environment with that file access. However, this software does not yet allow DECnet users to send jobs to the IBM environment, or to use IBM printers. When that capability, which is planned by the end of this year, is provided, Interlink will be about halfway to presenting an alternative to DEC's own SNA gateway facilities, which are described later in this report. The biggest missing pieces are support for LU6.2 and DCA. Interlink also supplies a Network Controller, called the 3711, that, when used with one of the gateway products, enables an IBM host to operate as a peer node on DECnet.

The appearance of these ISVs is no accident. The existence of Soft-Switch, CSI, SSI and the Systems Center demonstrates three things: first, it is the perception of OEMs that these SNA or SNA extension technologies must be supported by their equipment. Second, it is too difficult and costly to perform the development work in-house -- in

other words, SNA is not really open (see Research Notes K-321-040, 4/1/85, and K-321-071, 6/19/85). Third, the "not invented here" syndrome is disappearing.

Figure XIV A

Document Conversion Technology Licensing

	SOFTSWITCH	SRC
AT&T	X	
DEC	X	
WANG	X	
HP		X
TANDEM	X	

Figure XIV A represents some of the technology licensing agreements that have been made.

Connectivity and Networking of Departmental Computer Vendors

In this section we will analyze the achievements of three leading vendors in implementing the technologies described earlier, as well as their internal networking technologies. Digital Equipment Corp. (DEC) and Tandem have been chosen because of their prominence in their "home" market segments -- DEC for scientific processing, Tandem for transaction processing -- and because of their technological achievements. Wang has been selected for its success in penetrating installations with IBM mainframes, and because of the high degree of functionality of its office products.

Digital Equipment Corporation

Digital Equipment Corp. (DEC) and Tandem have the most highly developed networking technologies among the OEMs. In many ways, their networking facilities are more advanced in function than IBM's. For example, DEC has a local networking technology installed at customer sites -- Ethernet -- which has been incorporated into DECnet; whereas IBM has not yet announced its token-passing ring product. Tandem has a distributed processing capability that is probably more sophisticated than anyone else's. However, in fairness to IBM, it must be stated that DECnet Phase IV supports a total of only 64,449 nodes, while SNA has been recently upgraded to address 16 million nodes, a vast increase on its previous capability of 64,000. We see a similar situation with IBM's IMS product compared to Cullinet's IDMS or Applied Data Research's DataCom DB. Of course, these companies are dealing with very different customer bases from IBM. IBM's SNA must always support a tremendous number of terminal networks widely varied in terms of size, diversity of topology and performance. IBM software copes very well with such environments, but SNA's development is constrained by the necessity to carry forward all those customers. It is easier for a vendor to move quickly, and provide flexibility, when he is involved with smaller environments.

We will look first at DEC's progress in connectivity. Its stated mandate for networking is as follows:

- To adopt existing standards when necessary;
- To define or adopt new technologies to fill gaps;
- To migrate to new standards when appropriate;
- To retain proprietary technologies when no standard exists or when there are significant cost/performance advantages.

DEC's fulfillment of these aims can be appraised best by analyzing its technology before and after the April announcement, with which we deal later.

DECnet Phase IV can contain up to 1,023 nodes, and up to 63 DECnets can be linked by servers to interconnect 64,449 nodes. The VAX minicomputers can be linked (see Figure XV) within the architecture by any combination of four means: Ethernet; telephone lines (leased or dialed); X.25 networks; or SNA networks. A VAXcluster has visibility within DECnet as one node.

In carrying out its avowed aim of interconnecting with IBM's environment better than anyone else, DEC has gone so far as to distribute function very similarly to IBM. Just as IBM has positioned a 3725 communications processor to offload low-level operations polling, message formatting, etc. from the host computer, so has DEC set up a separate machine to interface VAX systems to IBM networks. This gateway is a specially packaged PDP 11/24. It makes sense to use a 16-bit processor as a communications front end; there is little to be gained for that function from the power of a 32-bit architecture such as the VAX. Also, practically speaking, DEC may have had in inventory many 16-bit machines, with relatively good basic communications software. In an IBM network, the control of the communications session does not reside in the 3725 -- it is performed by the host (by VTAM). Similarly, the SNA half session in DECnet does not reside in the gateway -- it is implemented in a special software module that must be purchased for each VAX that wants to emulate a 3274. This design may enable DEC to insulate the VAX machine from changes in low-level communications function, and to respond more flexibly as IBM continues to evolve SNA.

At the Intech '84 conference in October 1984, DEC unveiled three products that make use of the infrastructure described above. The most important was an interface to IBM's DISOSS product, and with this announcement DEC became the first office systems vendor to announce such a facility. The DISOSS Document Exchange Facility (DDXF) became available in June 1985; it enables users of DEC's All-In-One office system to access both the distribution and library functions of DISOSS. DDXF involves software on the gateway and on each VAX computer. The announcement specified final form documents only, but DEC recently announced a capability for revisable form translation also, called External Document Exchange. Further, DEC is using a document translation technology licensed from Soft-Switch. Taken together with DEC's pact with Cullinet to develop facilities for All-In-One to access mainframe resident databases (see Research Note K-115-012, 1/16/85), this announcement shows that DEC is pragmatic, flexible and open in building a very high level of connectivity to the IBM world.

Figure XV

DECnet Phase IV

DNA, DECnet, and Ethernet Relationships

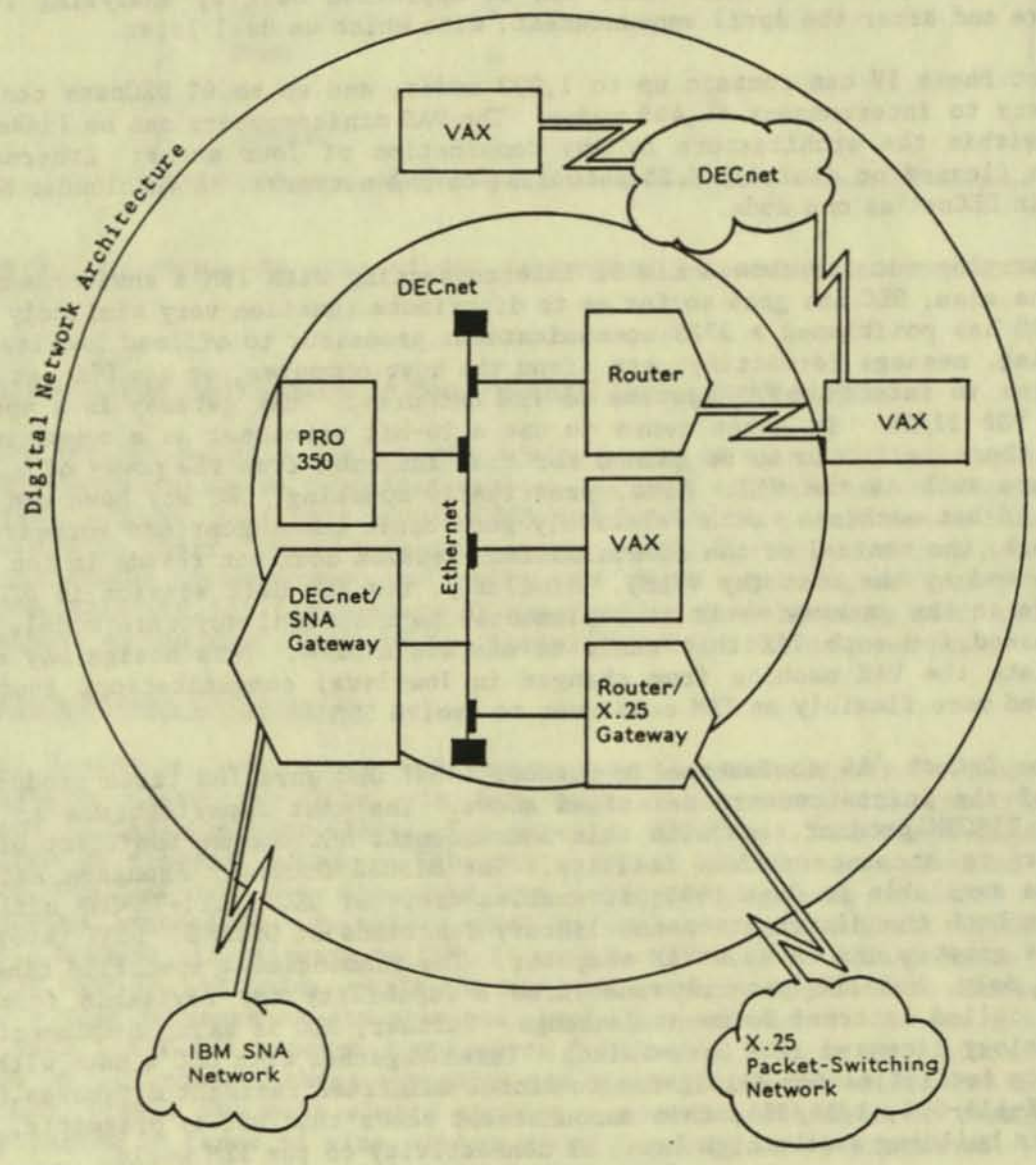

Another important product introduced at Intech was a printer emulator that makes it possible to write information generated on an IBM host on a printer attached to a VAX or MicroVAX system. The interesting thing about this software is that it doesn't rely upon the IBM remote batch facilities -- it will be a more responsible mechanism than the customary emulation remote job entry devices; in particular, it will not require the unfriendly operator interface that is customarily required.

DEC can now connect an IBM cluster controller to its SNA Gateway. Since VAX minicomputers can already attach to IBM hosts, the facility (called the Distributed Host Command Facility, or DHFC) now theoretically enables DECnet to serve as a backbone network connecting IBM 327X terminals to a remote IBM host -- see Figure XVI. What was lacking in this area before April was support within VAX VME for the 3270 data stream, which would enable IBM terminals to access applications on a VAX. As of April 1, there were still some missing pieces, including that mentioned above. The most important was DEC support for LU6.2. Moreover, much needed to be done for PC support. However, with regard to announcements, DEC was further along in IBM connectivity than anyone else.

Figure XVI

Operation of DEC's Distributed Host Command Facility

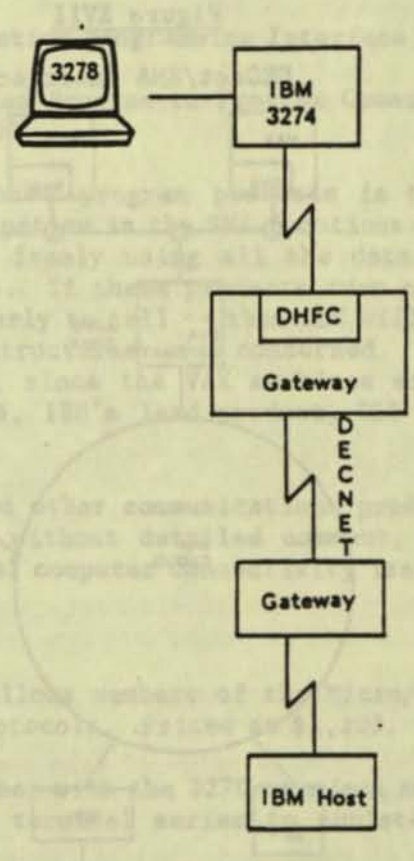

- DEC's Announcements

On April 16, DEC filled in the most important missing pieces. It announced DECnet-DOS software, a product that integrates IBM personal computers with a DECnet network; it introduced three software packages to enhance DECnet's connectivity to SNA; and it offered two protocol emulators providing communication links from DEC's Microl RSX operating systems to the IBM environment.

DECnet-DOS, which costs $500, will allow IBM PCs to acquire the Rainbow's communication capabilities, including file transfer to a VAX and to a Rainbow (through a VAX), and the ability to share VAX peripherals (i.e., disk storage and printers). The significance of this move is that DEC connectivity is now open to the ubiquitous IBM PC. This product should increase significantly the number of PCs connected to VAXs.

It is interesting that DECnet provides functionality that IBM does not yet provide on PCs for the IBM environment -- PCs cannot yet transparently access IMS or DB2 files, for example.

The first of the three software packages is the DECnet/SNA VMS 3270 Data Stream Programming Interface, which enables VMS programs to communicate using the LU2 data stream, with programs operating on an SNA host. This means, among other things, that a person using a DEC VT100 terminal can access CICS/VS or IMS/DC applications, or use TSO (see Figure XVII).

Figure XVII

DECnet/SNA Interface

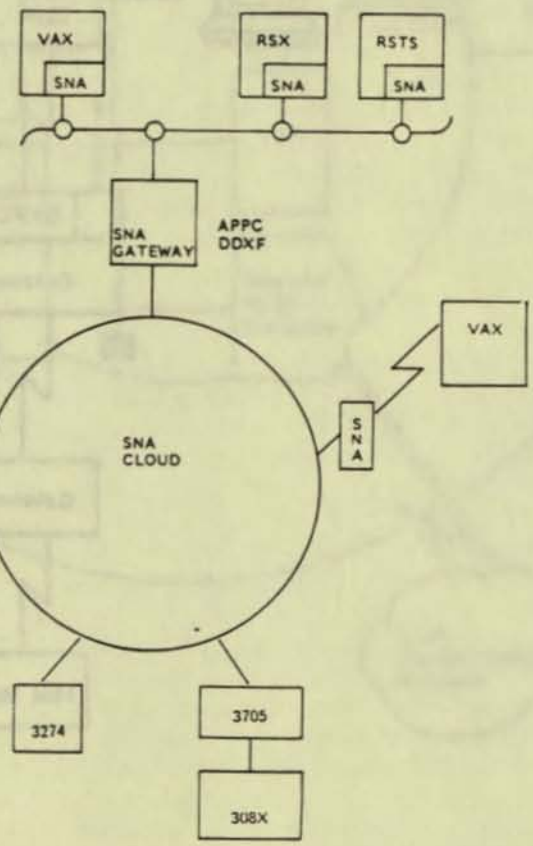

For DECnet programs whose requirements are more complex than the basic function provided by the 3270 Data Stream Programming Interface, DEC introduced the DECnet/SNA VMS Application Programming Interface (API). API makes SNA functionality open to a VAX/VMS program, so that it can exchange data easily. It is really the window DEC is providing to VAX/VMS programs into the LU6.2 environment.

The most important piece of this announcement is the LU6.2 capability, which DEC called the DECnet/SNA VMS Advanced Program-to-Program Communications (APPC)/LU6.2 Programming Interface. Since LU6.2 is being evolved by IBM into the SNA protocol, and in fact LU6.2 can probably be taken as equivalent to all of SNA in embodying the architecture, it is basic for any OEM to accept LU6.2. DEC's APPC enables VMS-based applications for VAX and MicroVAX systems to communicate with IBM host applications, on a peer-to-peer basis, via the DECnet/SNA Gateway. This communication is accomplished by using a set of verbs defined by IBM within LU6.2. There are verb sets in LU6.2 in addition to this one, which is used for program-to-program communication. These additional verb sets will be used for network management and resource control; other verb sets may be defined in the future, because LU6.2 is quite extendable. These verb sets correspond to the LU6.2 "option sets" described earlier.

These three program products support IBM environments: MVS/370, MVS/XA and CICS/VS; they all became available in June, and are priced as follows:

-- DECnet/SNA VMS 3270 Data Stream Programming Interface: $2,500

-- DECnet/SNA VMS Application Programming Interface (API): $1,500

-- DECnet/SNA VMS Advanced Program-to-Program Communications (APPC) LU6.2 Programming Interface: $1,500

The function of these three program products is to provide DECnet processes and elements with full participation in the SNA functions: the ability to attach intelligently and to communicate freely using all the data streams and through all layers including the applications. If these products turn out to be robust and fully functional -- it is still to early to tell -- then DEC will be in an excellent position, so far as networking infrastructures are concerned, to compete in the departmental computer market. In fact, since the VAX machines are more powerful and functional overall than the System/36, IBM's lead product, DEC has some significant advantages over IBM.

DEC has recently introduced other communications products which we will list, for the sake of completeness, but without detailed comment, because they are less directly relevant to the departmental computer connectivity issue.

April Announcements:

-- Micro/RSX 2780/3780 allows members of the Micro/PDP-V family to emulate 2780 and 3780 bisynchronous protocols. Priced at $1,200.

-- Micro/RSX 3271, together with the 3270 terminal emulator utility, allows users of the VT100 and VT200 terminal series to emulate 3270 bisynchronous protocols. Priced at $900.

-- External Document Exchange is a VAX-based software package that permits two-way document transfer, and conversion between a DEC VAX system and a Wang OIS, using the bisynchronous 2780/3780 communications facility common to both systems. Priced at $3,500 for a license, and $5,000 for a license with service.

June Announcements:

-- A multiyear distribution agreement under which DEC will market the Wollongong Group's TCP/IP (Transmission Control Program/Internet Protocol) networking software. Priced at $17,500 for VAX/VMS systems, this package includes media and documentation. Initial deliveries are scheduled for this month. Wollongong supplies this technology for other environments as well -- including IBM's VM operating systems.

-- Product support of the OSI model at the transport layer -- the VMS OSI Transport Service product is scheduled for first quarter 1986. DEC also announced support for the 1984 CCITT X.400 worldwide electronic mail standards as well as the recently revised X.25 packet-switching recommendation, with product delivery during 1986.

-- Upcoming shipment (i.e., during 1986) of hardware and software products to support the General Motors Manufacturing Automation Protocol (MAP). There will be modems, controllers and software designed for the MAP 2.1 version that runs under VAX/VMS and MicroVAX environments.

Clearly, DEC is bringing a great deal of focus and resource to its networking initiatives, and at this point it is as formidable in this area as any vendors of departmental computers, including IBM.

Tandem Corporation

Tandem Corporation is the only non-IBM vendor competitive with DEC in this technology. In fact, Tandem's internal networking technology is probably more advanced than DEC's, although Tandem has not progressed quite as far as DEC in connectivity to the IBM SNA environment.

Tandem's philosophy is that a network must be more than a switching system among "application domains" (an application domain is basically one central electronics complex, or CEC). The "intelligent network" must manage information outside the application domain in which the information was generated. This concept has given rise to three requirements:

1) Networking facilities that can provide interactive as well as store-and-forward, or time-staged, delivery of messages and requests;

2) A database system with strong distributed processing capabilities; and

3) Gateways that implement a wide range of data communications protocols and include format translation services.

Figure XVIII is a conceptual representation of Tandem's networking facilities. Guardian is the operating system. Tandem terms Guardian a "network operating system,"

which may be a bit ambitious. Guardian is message-based and facilitates communications among processes, and was probably designed with more awareness of a networking environment than were IBM's MVS/370 and DEC's VMS.

Figure XVIII

Tandem's Networking Facilities

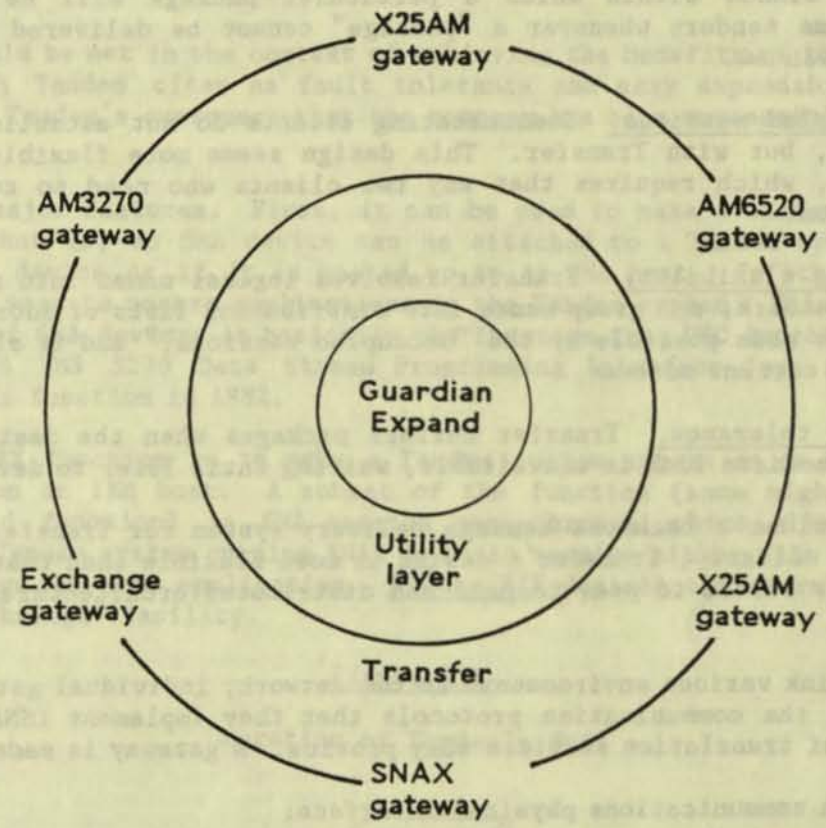

Expand, introduced in 1978, is presented by Tandem as an extension to Guardian. We regard Expand as analogous to a combination of IBM's VTAM/NCP (Virtual Telecommunications Access Method/Network Control Program). Tandem's argument is that because Guardian was implemented in a networking environment from day one, internode communication is a natural extension of existing capabilities; the developers simply introduced routing tables and implemented a packet-oriented communications protocol. The major components of Expand are:

-- The end-to-end protocol;

-- The network control process;

-- The network line handlers;

-- The network routing table; and

-- The network utilities.

Transfer, the component that provides transport and delivery services, can be regarded as roughly comparable to (IBM's) DIA and SNADS operating together. Its most notable features are:

-- Timed delivery of messages or data. A user (program or terminal) can specify a time window within which a particular package will be delivered. Transfer informs senders whenever a "package" cannot be delivered within the specified time window.

-- Decoupled sessions. Communicating clients do not establish sessions with each other, but with Transfer. This design seems more flexible and economical than IBM's, which requires that any two clients who need to communicate maintain a session.

-- Naming flexibility. Transfer resolves logical names into physical addresses in the network, and group names into distribution lists of addresses. This capability is made possible by the "decoupled sessions," and is also more flexible than SNA's current methods.

-- Fault tolerance. Transfer buffers packages when the destination node or some intermediate link is unavailable, waiting until later to develop them.

Expand provides a backbone message delivery system for Transfer, just as LU6.2 does for SNADS. Clearly, Transfer's design is more flexible than that for SNADS, and probably better suited to peer-to-peer and distributed architectures. It was introduced in 1982.

Gateways link various environments to the network; individual gateways differ from one another in the communication protocols that they implement (SNA, X.25, etc.) and in the kinds of translation services they provide. A gateway is made up of three parts:

-- A data communications physical interface;

-- Presentation services, such as document format translation or screen presentation; and

-- An interface to Tandem's network.

Tandem provides gateways to a wide range of systems/protocols, including bisynchronous batch and interactive devices (these are used heavily to communicate with DEC and Wang systems) and Network Systems Corporation's HYPERchannel.

Tandem's SNA gateway is termed SNAX, which stands for "SNA Communications System." SNAX is implemented in software operating on Tandem's processors, and later may be in the Tandem 6100 front-end processor. The 6100 is similar in function to an IBM 37X5, by which a single subsystem can accommodate different protocols, line disciplines and line speeds. However, a 6100 must be channel-attached, and cannot yet function as a remote concentrator. For simplicity, the rest of this discussion will treat SNA as consisting of the software product operating on the Tandem host. This product was introduced in 1982.

The design goals of SNAX are to:

-- allow Tandem systems to operate in SNA networks;

-- allow Tandem customers to preserve their investment in SNA equipment and applications.

The goals should be met in the context of achieving the benefits of the Tandem environment -- which Tandem cites as fault tolerance and easy expandability. It is the perception of Tandem's customers that the company has been successful in meeting these design goals.

SNAX has two major features. First, it can be used to make a Tandem system act as an SNA host -- that is, an SNA device can be attached to a Tandem system, and it will appear to the device as if it is hooked up to an SNA host. In other words, the SNA device can be used to access applications in the Tandem system. This attachability to the OEM host of SNA devices is basically the function that DEC introduced in June with the DECnet/SNA VMS 3270 Data Stream Programming Interface (see page 29). Tandem introduced this function in 1982.

The second SNAX function is to make a Tandem system appear as an SNA device to SNA applications on an IBM host. A subset of the function (some might view this as a discrete third function) is SNA session pass-through, which allows an SNA device attached to a Tandem system running SNAX to be in session either with a Tandem application or with an IBM host application. Figure XIX depicts the operation of this "IBM Session Pass-Through" facility.

Figure XIX

Operation of Tandem's SNAX

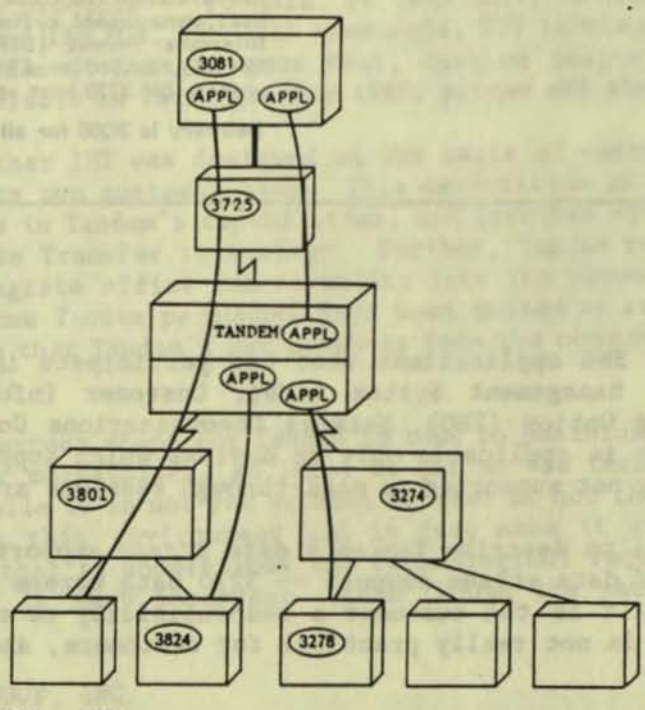

Figure XX

Tandem Software Utilities

Pricing on Tandem Products

	License Fee	Monthly Charge
PC MAIL	$1K-2K*	$100-200
FAXLINK	$12,500**	$100-200
PC LINK	$495/PC***	—
PS TEXT EDIT	$375-500	$50
PS TEXT FORMAT	$375-500	$50

* Free to TRANSFER users with Tandem workstations and IBM PC's using PC LINK.

** Hardware controller and software. Fax machine not included.

*** Free for DYNAMITE workstation.

PS MAIL: (Professional Support Electronic Mail System): Supports Tandem 653X and Dynamite workstations, IBM 327X and asynchronous workstations, and IBM PCs as well as compatibles. Messages can be of any length and either image- or text-based.

FAXLINK: Can use Tandem network, running over leased lines, sending facsimile (fax) messages over leased telephone lines. Fax messages can be "packaged" with text messages using PS MAIL. Uses OCR "mark sense" form or PS MAIL to address to fax recipient. Also acts as a remote printer for text documents; format is switched from ASCII to fax code and printed at another location. Uses CCITT Group III fax machine, which is not included in the price.

PS TEXT EDIT (TEDIT): Supports IBM 327X, Tandem 6520 and 653X terminals, and Tandem's DYNAMITE workstations. Uses function keys. Two windows available.

PS TEXT FORMAT: Can be used with TEDIT or with other Tandem editors. For designing page layouts, selecting fonts and for storing customized layouts.

PC LINK: Allows IBM and compatible PCs to have online access to data integration through emulating Tandem 6530 or 327X terminals. Consists of four software utilities: EM 6530PC for emulating a 6530; IXF/PC to transfer files between Tandem host and PC; PCFORMAT to convert files from the Tandem distributed database management system, ENCOMPASS, into Data Interface Format (DIF), Symbolic Link Format (SYLK), BASIC or ASCII formats; and EM3270 to access IBM 3270 host applications.

Delivery is 3Q85 for all five utilities.

Examples of SNA applications that can participate in pass-through session include: Information Management System (IMS), Customer Information Control System (CICS), Time-Sharing Option (TSO), Network Communications Control Facility (NCCF). Session pass-through is applicable only to devices which support Physical Unit type 2 (PU2). Specifically not supported in pass-through sessions are PU1 devices.

It is useful to describe Tandem's data stream support capabilities. Tandem provides two types of data stream support -- 3270 data stream (LU2) or SNALU, which is really IBM's LU0. It is the customer's responsibility to create the communications protocols. That is not really practical for customers, and in 1984 Tandem introduced SNAX

High-Level Support (SNAX/HLS), which is a general-purpose high-level interface for specifying data streams and communication protocols -- such things as bracket support, request mode, function management profiles and transmission subsystem profiles. This facility would generally be invoked for custom and special-purpose applications.

The Tandem SNA interface capabilities analyzed here (introduced around 1982) are structured similarly to those DEC is providing this year. This product design is a very natural one. Tandem's "SNAX" offering is the most sophisticated interface to SNA operating today, and it provides an excellent base for its implementation of the extensions to SNA.

- Tandem's June Announcements

In June 1985, Tandem addressed this area. It introduced Information Management Technology (IMT), which presently includes five software products, and stated its intent to provide a DCA capability, revisable form and final form, as well as to support IBM's LU6.2. Tandem also indicated that within a year it will provide support for Ethernet -- the IEEE 802.3 standard -- probably initially in the form of interfaces to 3Com's products and AT&T's Starlan. Also within a year Tandem will probably interface to the IBM PC Network. In an unspecified time frame Tandem will deliver support for X.400 (the CCITT electronic mail standard) and IBM's DISOSS.

At the same time as it made these announcements, Tandem announced two other agreements. It signed a multiyear, multimillion dollar agreement with Sytek to remarket LocalNet/20 and IBM PC Network compatible communications equipment. We would expect that Tandem will offer that hardware and software for broadband local-area network configurations involving IBM PCs and Tandem equipment, while the LocalNet/20 technology will be used for terminal-to-host communications within the Tandem environment. The other agreement Tandem signed was with AT&T Information Systems. Under its terms Tandem becomes a licensee for AT&T's Digital Multiplex Interface (DMI), which is AT&T's standard for integrating PBX equipment with computers.

Tandem's five new software products -- PS MAIL, PS TEXT EDIT, PS TEXT FORMAT, FAXLINK and PC LINK -- allow users of IBM PCs, IBM 327X terminals, TTY terminals, Tandem terminals and Group III facsimile devices to move text, data or images through a Tandem network. They will be available in third quarter 1985; prices are shown in Figure XX.

Tandem's presentation is that IMT was designed on the basis of customer requests and will be very popular in its own customer base. This expectation is probably correct, since IMT fills real holes in Tandem's capabilities, and involves not much more than a natural exploitation of the Transfer technology. Further, Tandem regards IMT as part of its commitment to integrate office functionality into its networked, distributed database architecture. Some Tandem personnel have been quoted as stating that, ultimately, IMT aims to ensure that Tandem's own networks form the communications backbone for Fortune 1000 companies.

Certainly these are the correct steps for Tandem to take to position itself as a force in the department computer marketplace, at least as far as the technological infrastructure is concerned. While it is not yet evident whether or not there will be applications that distinguish this environment and in fact make it explode in growth, Tandem has now indicated that it understands the technological requirements and will comply with them. That is very important. Even though the introduction of SNAX

preceded that of DECnet/SNA, it is perceived by many that Tandem was slower than DEC in moving away from the philosophy that highly developed proprietary technologies were a sufficient platform for competition with the giants. Tandem's embrace of the "standards" may not yet be as explicit as DEC's, but it may be firm enough.

With its pursuit of these technological standards, plus retuning of its marketing, DEC has made substantial progress over the past several years in positioning itself into the departmental computer marketplace, adding that thrust to its strength in scientific applications. Tandem has a significant presence in commercial accounts, and may be able to move as quickly and effectively as DEC has. It is probably overly optimistic of Tandem to expect its networking technology to be the backbone for many of the Fortune 1000 companies. However, based on its communications infrastructure, Tandem may be better positioned technologically than anyone else to match IBM's SNA extensions and evolving network operating system.

Wang Laboratories Inc.

Wang Laboratories Inc. has historically not been as advanced as DEC and Tandem in either internal networking technologies or connectivity to IBM. In October 1983, Wang announced the Information Distribution System (IDS) family of software products, which facilitate two types of communications:

-- Communication among Wang VS systems, across existing IBM network links.

-- Communication between VS systems and IBM hosts.

IDS made available bisynchronous and SNA point-to-point and multipoint links. Its functioning depends on software that runs under CICS on the IBM host (see Figure XXI) and makes possible file-transfer and store-and-forward functions.

IDS represents an attempt by Wang to build a software infrastructure for networking. The requirement for dedicated Wang software under CICS might be problematic, since such software represents a potential point of failure and a performance bottleneck, and the requirement involves work on IBM equipment in order to achieve Wang connectivity. Tandem and DEC have not found it necessary to adopt this expedient. However, this may not be a serious concern in the long term.

Wang Systems Networking (WSN) is a framework -- not yet an architecture -- that Wang intends to evolve toward the OSI model, and is built upon the IDS transport services. WSN is now only an umbrella term describing all Wang networking products implemented according to the WSN strategy. These products include:

-- WSN Services -- file transfer, remote access, VS and 3270 terminal emulation, virtual terminal interface, store and forward, resource sharing.

-- WSN Applications -- Wang Office (calendaring, E-Mail, file sharing), gateways to DCA/DIA, etc.

-- WSN Transports -- multipoint lines, point-to-point lines, X.25, Wangnet, IDS/bisync/SNA.

Figure XXI

Wang's Information Distribution System

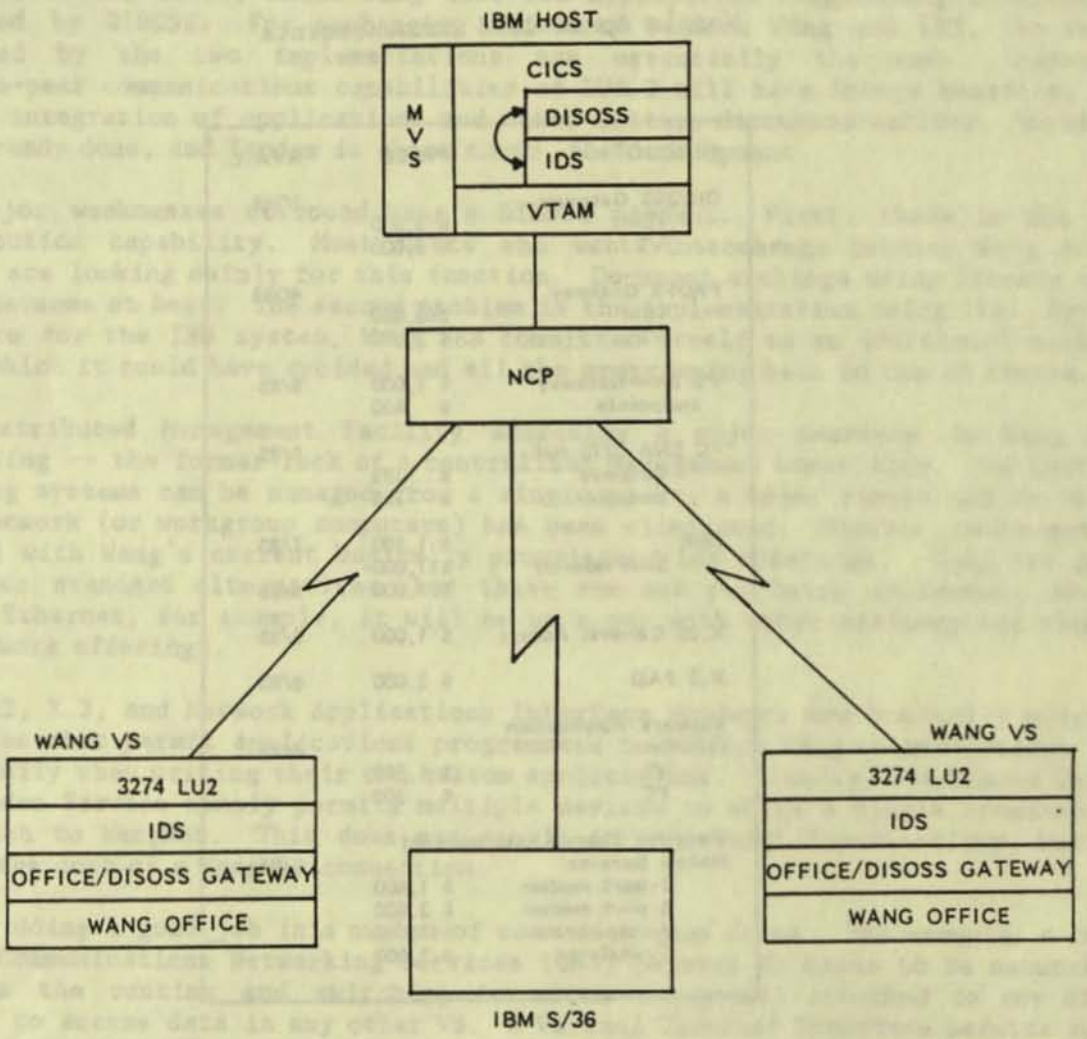

WANG OFFICE/DISOSS GATEWAY

Much of this was a repackaging of existing facilities. The third release of WSN, announced in October 1984, began shipping earlier this year. While there is not yet a least-cost routing facility, Wang is delivering such pieces as the capability of one node monitoring the performance of adjacent nodes on the network. Wang, therefore, has a lot of technology to build before it is really competitive in networking technologies.

- Wang's April Announcements

This April Wang announced nine new communications products (see Figure XXII) that fit within the WSN framework. Four of these products address connectivity to IBM, but not all the issues are handled yet.

Figure XXII

Wang's April 1985 Announcements

PRODUCT†	PRICE	AVAIL.
DISOSS Gateway		1Q86
IBM	$ 3,000	
VS	$ 3,000	
PROFS Gateway		4Q85
IBM	$40,000	
VS	$ 2,000	
VS/SNA Gateway	$ 1,000	9/85
endpoints	$ 600	
PC SNA 3270 Net		9/85
Gateway	$ 395	
endpoints	$ 795	
DMF	$ 1,500	7/85
SNA Monitor	$17,000-$24,000	9/85
X.25 General Access	$ 1,000	8/85
X.3 PAD	$ 2,000	8/85
Network Application Interface		8/85
VS	$ 1,500	
PC	$ 200	
WangNet Shared Interconnect Modem Service		6/85
2-port modem	$ 1,600	
8-port modem	$ 3,600	
Administrator	$11,000	
Translator	$ 3,500	

Formerly, Wang required that each Wang system be individually connected to an IBM host. While Wang Systems Networking (WSN) -- Wang's communications architecture -- provided fairly rich facilities for communications between Wang systems, Wang users could not specify one system in a Wang network to act as a gateway to the host. The Wang VS/SNA Gateway addresses this problem by permitting one VS in a network to be designated the IBM gateway. Any terminal attached to any VS in the Wang network can now use this single gateway to access the IBM environment under 3270 emulation. While this is a plus for Wang users, it is not unique; other vendors, notably DEC, offer similar facilities. Similarly, the PC SNA 3270 Network allows a single PC on a Wang PC network to be designed as a gateway to perform 3270 emulation for other PCs attached to the network. Obviously this is a desirable, in fact necessary, function, but it is only equivalent to what DEC and Tandem did some time ago.

Wang's DISOSS interchange product provides library services only for revisable form DCA documents. Wang users will be able to search the IBM DISOSS library for documents which have been previously stored, but the product provides no distribution

capabilities. By comparison, DEC's DDXF provides both library and distribution services, and revisable as well as final form (view and print) only. DEC's gateway also implements LU6.2, while Wang uses the Application Programming Interface (API) provided by DISOSS. For exchanging documents between Wang and IBM, the functions provided by the two implementations are essentially the same. However, the peer-to-peer communications capabilities of LU6.2 will have future benefits, such as easier integration of applications and other matters discussed earlier, for which DEC has already done, and Tandem is about to do, the development.

Two major weaknesses surround Wang's DISOSS product. First, there is the lack of distribution capability. Most users who want interchange between Wang and IBM's DISOSS are looking mainly for this function. Document exchange using library services is cumbersome at best. The second problem is the implementation using IDS. By writing software for the IBM system, Wang has committed itself to an additional maintenance chore which it could have avoided had all the programming been on the VS system.

The Distributed Management Facility addresses a major weakness in Wang Systems Networking -- the former lack of a centralized management capability. Now that multiple Wang systems can be managed from a single point, a major reason not to install a Wang network (or workgroup computers) has been eliminated. However, users must still contend with Wang's current basically proprietary LAN offerings. Wang has promised to offer standard alternatives, but these are not yet being delivered. When Wang offers Ethernet, for example, it will be on a par with other minicomputer vendors in its network offerings.

The X.25, X.3, and Network Applications Interface products are basically programming utilities that permit applications programmers to utilize Wang communications facilities easily when writing their own custom applications. Finally, the Shared Interconnect Modem Service simply permits multiple devices to share a single broadband modem to attach to WangNet. This does not result in additional functionality, but simply lowers the cost of a WangNet connection.

Wang is doing a good job in a number of communications areas. For example, a function called Communications Networking Services (CNS) permits VS hosts to be networked and provides the routing and switching to allow a terminal attached to any of these systems to access data in any other VS. A Virtual Terminal Interface permits applications programs running in one VS to access similar data or applications in another VS. In addition, Wang provides a higher degree of document transfer capabilities among its various systems than IBM, or nearly any other vendor, does.

The dilemma facing Wang is that the data processing department has a very large voice in the implementation of office computing, meaning that competitors to IBM have to be clearly superior to win. Wang may have a friendlier user interface for some applications, but has been late to recognize the importance of coexistence with IBM communication standards. Wang does not yet support any IBM workstations other than the personal computer, leaving a large installed base of IBM terminals that Wang cannot reach with its systems. While it is talking about opening up WSN to industry standards, Wang suffers from the (correct) perception that it has held on to a proprietary communications architecture and has a large task ahead to change that position. In addition, while the company can provide peer-to-peer communications at the workgroup level, it offers no centralized network management tools.

Wang's announcements address many of the shortcomings in its communications offerings, but do little more than bring the company up to par with DEC or other workgroup computer vendors. Combined with a migration to more industry-standard, physical-LAN options (Ethernet, Starlan, etc.), for which Wang has announced no timetable, this is positive for Wang. However, in matching other vendors, Wang is trying to hit a moving target, and there is no guarantee that by the time these products are actually delivered (given Wang's history of software delays), another vendor will not have raised the stakes. The direction Wang is taking is generally a good one (although it is questionable whether writing code to actually run on IBM systems is a good idea for a workgroup vendor), but whether the announced delivery schedules are early enough to keep up with the competition remains a major question.

The irony of the situation is that Wang is truly ahead of IBM in offering office functions with a high degree of transferability among a wide variety of workstations and systems. However, Wang must accelerate its development of communications facilities, both connectivity with IBM, especially LU6.2, and within WSN, in order to compete effectively in the department computer arena, particularly with IBM.

Summary

A summary tabulation of OEM capabilities for connectivity to IBM including Data General, Hewlett-Packard and AT&T, is presented in Figure XXIII.

Figure XXIII

Summary of OEM Connectivity

	LU2	LU6.2	DIA	DCA	SNADS
ATTIS (3B)	Now	10/85	-	-	-
DEC	Now	Now	-	Now	-
Tandem	Now	SOI	-	3Q85	-
Wang	Now	SOI	-	1Q85	-
Data General	Now	1Q86	-	11/85	-
Hewlett-Packard	Now	SOI	-	SOI	-

SOI means statement of intent

Tandem has the strongest base SNA technology of these vendors, and its gateway interface is a smoothly defined extension of the core Guardian/Expand software. This strength reflects Tandem's historical markets -- from day one it has sold to accounts that already had equipment processing mainstream applications, most of which were IBM's. Tandem's most significant presentation is in the banking and finance industry, in which, as of March 1985, it claimed a total of more than 2050 processors installed at more than 240 customer sites, including 21 out of the top 25 U.S. bank holding companies. That represents approximately 30 percent of Tandem's installed base, and is strong IBM territory. Tandem has also done well in communications, manufacturing and the federal government, each of which represents about 10 percent of Tandem's installed base.

While DEC also has a well-designed SNA capability, its gateway structure may be not quite so well integrated into DECnet as Tandem's is (into Guardian/Expand). DEC's realization of these communications requirements came relatively late. Of course DEC's installed base (there are over 40,000 VAX systems installed, compared to around 6,800 Tandem processors) provides it with a larger base for Research and Development (R&D) -- DEC spent $630.7 million in 1984 on R&D, compared to Tandem's $52.5 million, so there is good reason to expect its development efforts to match Tandem's.

Tandem announced its involvement in the "SNA extensions" -- DCA, DIA, LU6.2 -- three to eight months later than DEC, depending on the function. DEC announced its DCA capability in October 1984 and its LU6.2 capability in April 1985, while Tandem made statements of direction in June. Tandem's shipment of these technologies will probably occur in mid-1986, a year after DEC's.

To the extent that this gap is significant, it may represent a lack of explicit requirements by Tandem's customer base. LU6.2 is a technology that hasn't been apparent to end users. While DCA is apparent, it may be the case that Tandem's customers did not contemplate the intense use of these machines for text applications to the extent that DEC's customers did for their VAX machine. It is <u>not</u> our perception that this gap indicates any less of a technological capability on Tandem's part. The gap is probably a function of positioning uncertainty.

Wang Laboratories announcements, in particular the design of its connectivity infrastructure, do raise the issue of base technological capabilities. In addition to the reliance of IDS upon a component running under CICS on the IBM mainframe, the interface to DISOSS uses the DISOSS Application Programming Interface rather than LU6.2. Additionally, the DCA conversion technology (licensed from Soft-Switch Inc.) will apparently operate on the IBM host -- Wang is not migrating that technology to its VS equipment. It would appear that Wang's technology is the least integrated of the three. However, the purity of the base technology is certainly not the most important issue here.

For the record, to complete the tabulation of R&D expenditures for 1984, Wang's were $160.5 million and IBM's $3.148 billion.

While the technological soundness of the communications architecture is important, it is certainly overshadowed, in terms of importance as success factors with customers, by friendly presentation, usability, timely delivery and quality of field service. However, those issues are not independent of the soundness of the base technology.

Departmental computer functionality today is usually expressed in terms of office applications and connectivity to IBM mainframe SNA environments. While it may be the case that, in the long term, distribution of data and applications out of the "glass houses" will have the dominant influence on departmental computer products, the necessary technologies have not been developed to the point where they are a serious factor.

We have seen how IBM is quietly evolving its own communications infrastructure in order to accommodate its customers' growth, reliability and (last!) usability requirements. Unlike the situation in the 1975-1980 time frame, during which SNA was being introduced, the vendors of departmental computer equipment now realize that participation in that technology is essential to their success in this marketplace. Tandem and DEC have taken the lead in that technological development.

A factor, or dilemma, that all these vendors must address is that the data processing department has a very large voice in the implementation of office and departmental computing. That can translate into the constraint that competitors to IBM have to be clearly superior to win. The OEMs -- DEC, Wang, Tandem, H-P -- may have user interfaces that are more friendly for many applications, but often have been late to recognize the importance of embracing IBM's communications architecture, and the fact that mere coexistence with those technologies is insufficient. These OEMs have suffered from the often correct perception that they talk adherence to industry standards but persist with proprietary communications technologies; they therefore have a large task ahead of them to alter that perception. In addition, while most of these vendors have indicated that they will provide peer-to-peer communications at the workgroup level, there is little in the way of rich network management facilities to manage all the workgroup systems from a central site, or for that matter from multiple control centers. Most of the OEM-supplied network management facilities are little more than data collection tools. Even if the functionality and usability of the OEM product have advantages, it requires an aggressive, perhaps even brave, data processing manager to discard IBM's centralized network management tools in favor of a communications architecture that does not have similar facilities.

We see that it is a fundamental issue for the vendors of departmental computers whether they can extend their advantages over IBM in functionality and usability while narrowing the gap in base communications technologies, at a time when IBM is very much a moving target.

Bibliography

JP Gray, RJ Hansen, P. Holman, MA Lerner and M Pozevsky: "Advanced Program-to-Program Communication in SNA." IBM Systems Journal Vol. 22 #4, pp 248-318 1983

BC Housel and CJ Scopinich: "SNA Distribution Services," IBM Systems Journal, Vol. 22, #4, pp. 319-343 1983

Mr. De Sousa: "Electronic Information Interchange In An Office Environment:" IBM Systems Journal Vol. 20 #1, pp. 4-22 1981

T Schick and RF Brockish: "The Document Interchange Architecture: A Member of a Family of Architectures in the SNA Environment." IBM Systems Journal Vol. 21 #2, pp. 220-244 1982.

RJ Sundstrom and GD Schultz: "SNA's First Six Years: 1974-1980." Proc. Fifth ICCC, Atlanta, Ga. pp. 578-585 Oct. 1980.

Systems Network Architecture Format and Protocol Reference Manual: Architecture Logic for LU Type 6.2. SC30-3269-2 (Available at IBM Branch Offices)

Systems Network Architecture Concepts and Products. GC30-3072 (Available at IBM Branch Offices)

Introduction to APPC: GG24-1584 (Available at IBM Branch Offices)

CICS/VS 1.5 ISC Design Guide -- ZZ10-9897 (Available at IBM Branch Offices)

GLOSSARY

ACF

Advanced Communications Function is IBM's term for the SNA facilities for multisystems networking. These enhancements were introduced in 1977/78. While some IBM publications refer to this SNA release as SNA-2, this report considers it "Phase II" of SNA, since that makes sense in terms of connectivity.

APPC

Advanced Program to Program communications is a synonym for LU6.2, please see that definition in this Glossary.

CICS

The Customer Information Control System is IBM's strategic general purpose subsystem for the implementation of transaction-processing applications. CICS invokes customer-written application programs in response to transactions entered at teleprocessing terminals, and provides the services needed by those applications to retrieve and update data in files and return a response to the terminal which invoked them.

CNS

Communications Networking Services is a Wang software facility that permits Wang VS hosts to be networked, and provides the routing and switching to allow a terminal attached to any of these systems to access data in any other Wang VS.

Communication Solutions Inc.

CSI is an ISV that specializes in selling to OEMs the technology to interface with SNA networks. Its Access/SNA product line, written in C language, now includes LU2, Synchronous Data Link Control (SDLC) and bisynchronous capabilities. Capabilities for DIA and LU6.2 have been promised for later this year.

DCA

Document Content Architecture is a machine-independent data stream that presently describes the form and meaning of the information in a document. However, it was designed to be extended to define other types of information, such as facsimile, graphics or digitized voice. This will allow files or documents merging these forms of information to be transmitted electronically via the technologies analyzed in this report.

Currently, IBM support two levels of DCA. Revisable form text DCA specifies format controls such as those for type style, indentation, and line spacing. Final form text DCA handles format controls such as those for margin placement, page size dimensions and type style definition. One should view revisable form text DCA as defining the logical structure of a document, while final form text DCA defines a primitive format that is completely device-independent.

DDXF

The DISOSS Document Exchange Facility is DEC's software package that enables users of DEC's All-In-One office system to access both the distribution and library functions of DISOSS. DDXF has capabilities for final form data streams. The revisable capability is packaged separately (see External Document Exchange). DDXF became available in June, 1985; it involves software operating within VAX and IBM environments.

DECnet-DOS

This DEC software package provides IBM PCs with fairly sophisticated capabilities for cooperative processing in a VAX environment. In particular, the PC can access VAX files and peripherals. DECnet-DOS became available in June, 1985.

DECnet/SNA VMS APPC Programming Interface

Available in June, 1985, this DEC software facility provides VMS programs with the ability to interact via the conversational verb interface with IBM's LU6.2 data stream.

DECnet/SNA VMS 3270 Data Stream Programming Interface

This DEC software package gives VME programs access, via the 3270 data stream, to programs operating on IBM hosts. This "window" became available in June, 1985.

DHFC

The Distributed Host Command Facility is a DEC software package that provides support within the (VAX) VME operating system for IBM's 3270 data stream. In principle, DECnet can now serve as a backbone network connecting IBM 327X terminals to a remote IBM host. DHFC became available in April, 1985.

DIA

Document Interface Architecture is defined by IBM as a set of services performed by peer communications processes. A more functional view of DIA is as a set of communication programs -- i.e., software that uses communication facilities such as SNA, or others, to carry out functions requested by users. These users can be people, application programs or devices; the key thing is that a user represents the source or recipient of information.

For reasons of compatibility across the product line, it is IBM's intent that DIA use LU6.2 facilities, perhaps mediated by SNADS. However, DIA function is independent of the LU type, and can and does use LU2 (3270 data stream) or asynchronous facilities.

DISOSS

The Distributed Office Support System is a family of IBM program products that provide library and document distribution functions, as well as document translation services. This software operates only under the MVS operating system, and uses CICS for access by terminal users.

DISOSS has been slow to gain acceptance because of functional lacks and nonusability. IBM is filling in the capabilities with the Personal Services packages, which really should be thought of as "DISOSS/36," "DISOSS/38" etc.

Displaywrite

The Displaywrite products are a set of program products for word processing, based on the Displaywriter implementation, and operating in various environments. This appears to be IBM's strategic direction for word processing, and it helps to present uniform user interface across multiple architectures. Announced so far are Displaywrite/1, 2 and 3 for the PC and Displaywrite/36 for the System/36.

Expand

A Tandem software facility, operating together with Guardian, that provides various communications functions. Expand is roughly comparable to a combination of IBM's VTAM and NCP.

External Document Exchange

External Document Exchange is DEC's facility to allow users of All-In-One to both access and edit DISOSS library documents. DEC chose to package this product separately from DDXF, and made it available in July, 1985.

Guardian

Guardian is Tandem's operating system for its lines of fault-tolerant computers -- Nonstop I, Nonstop II and TXP.

IDS

The Information Distribution System is a family of Wang software products which facilitate communications of two types: among Wang VS systems (across IBM links), and between Wang VS systems and IBM hosts.

IMS

Information Management System is IBM's strategic subsystem for large DB/DC applications. IMS invokes customer-written application programs in response to transactions entered at teleprocessing terminals and provides the services needed by those applications to retrieve and update data in DL/1 databases and return a response to the terminal which invoked them.

IMS applications run in what are called message-processing regions. As needed, they communicate with the IMC Control Region which is located in a separate MVS address space. This splitting of the program product into many components (which is called granularity) is a large reason why, at this time, IMS exploits the 308X architecture more fully than CICS does, and also has somewhat superior integrity.

ISC

Inter-System Communications is a facility that permits an application in a subsystem to communicate with another application in another subsystem. The two subsystems may be of the same type (e.g., CICS/VS to CICS/VS) or of different types (CICS/VS to IMS/VS). ISC is built up on an LU Type 6 data stream. It was first introduced in CICS/VS Release 1.4, and is used to increase availability and provide Virtual Storage Constraint Relief.

IMT

Information Management Technology is a group of software products, the first of which were announced by Tandem in June, 1985, intended to provide connectivity across different architectures. The initial products concentrate on connectivity at the workstation level.

ISV

An Independent Software Vendor is just that -- a firm whose business is selling software, and which is not a subsidiary of IBM or any OEM.

LU

A Logical Unit (LU) is a term within the SNA technology that is usually defined as an end user's "port" into the network. The term is also often taken to mean a particular terminal or application program. The third meaning is that LUs are types of data streams.

This terminology is confusing, and the ambiguity of the notation contributes to the difficulty of comprehending SNA. Within this report, we will usually use the term LU to denote data streams.

LU6.2

Logical Unit Type 6.2 is a relatively new SNA protocol that is the basis for a revamping of SNA's infrastructure. It is device independent, and provides the facilities for peer to peer communications between two programs. Among other things, this means that programs and small computer systems will use LU6.2 (along with Physical Unit Type 2.1) to establish and control sessions with each other, without necessarily any participation by a host mainframe. In this manner IBM will evolve SNA away from its current host based centralized organization.

When LU6.2 is installed in devices such as "PCs" and "cluster controllers," such necessities as multiple sessions among those devices and S/370 hosts will be available.

LU6.2 will also be the foundation for a "network operating system," meaning that all of the devices on the network will use that protocol to talk intelligently to manage resources and perform network recovery.

IBM will package LU6.2 as a base component plus option sets, so that LU6.2 can be installed in small machines like Scanmaster, Displaywriter and the PC at a low price point.

The terms "LU6.2" and "APPC" (advanced Program to Program Communications) have been used synonomously by IBM, and in this report we will not distinguish between them.

MSNF

The Multi-System Networking Facility is an SNA feature, implemented in ACF/VTAM or ACF/TCAM, that supports communication among multiple host processors. A prime benefit of MSNF is that a terminal (user) can now access (programs operating on) multiple computers. That capability was lacking in the initial release of SNA.

MVS

Multiple Virtual Storage is IBM's strategic operating system for large-scale database/data communications environments. While MVS is more expensive, in terms of charges and resources used, than other IBM-supplied operating systems, it provides more reliability and device support.

NCP

The Network Control Program is the operating software for the 3705 and 3725 communications controllers. The executable code is generated in a S/370 host and downloaded to the 37X5.

OEM

Original Equipment Manufacturer is the generic term used in this report for manufacturers of departmental computers.

Personal Services

Personal Services is a family of program products, one for each office system machine -- S/36, S/38, PC -- that extends the IBM office network architecture to that environment. Personal Services/X (where X is 36, 38, etc.) provides document exchange capabilities with the other family members as well as calendar management and other office administrative functions.

We view Personal Services as just a packaging and renaming of functions originally presented or promised in DISOSS -- in other words, Personal Services/36 essentially could have been named DISOSS/36; IBM chose to rename these facilities for delivery on office systems.

Primary LU

The Primary Logical Unit controls the session between two LUs. The primary performs session initiation and termination, and receives log-on requests from the Secondary LU. In an application program -- terminal session, the program is the primary LU.

PU

Within the SNA architecture, a Physical Unit is the component that manages and controls the physical resources of a node. These resources include attached links (communication lines) adjacent link stations (terminals) and the processor. It makes sense to think of a PU as just a hardware/microcode/software implementation that has the capability to interpret SNA data streams. There are four basic PU types: PU Type 1 is a Terminal (3478); PU Type 2 is a Cluster Controller (3274, 3276); PU Type 4 is a Communications Controller (3705, 3725; actually, the NCP) and PU Type 5 is a Host Computer (actually, it is termed System Services Control Point and realized in VTAM or TCAM).

IBM has implemented an enhanced Cluster Controller Physical Unit Type, namely PU Type 2.1. Its two key new capabilities are to have multiple "upstream" links, and the so-called "Single Node Control Point" facility, whereby the device itself can initiate sessions and tear them down, without involving a PU Type V. This new PU Type is one of

the two basic building blocks (the other is LU6.2) for the current revamping and evolution of SNA by IBM.

Rabbit Software Corp.

Rabbit Software Corp. is an ISV that has announced a software package called "SNA-plus," written in C language and targeted for vendors offering Unix environments. Scheduled for shipment later this year, SNA-plus is claimed to include LU6.2 and some DIA capabilities.

Session

An SNA session is a logical connection between two network addressable units. The activation and deactivation of sessions incur significant overhead in network traffic. Sessions compete for network resources such as bandwidth. It is IBM's (or Raleigh's) intent that LU6.2 facilitate more efficient management of these processes.

SNA

Systems Network Architecture is a set of rules for designing and implementing communications programs. More formally, it is the description of the logical structure, formats and protocols for transmitting information through networks and controlling the configuration and operation of networks. These rules and descriptions are set forth in the various Format and Protocol Logic Manuals.

The original purpose of the layered structure of SNA was to allow the application programs and terminal users to be independent of and unaffected by the way specific SNA services and facilities are provided. That purpose has been largely achieved, and IBM deserves accolades for its pioneering achievement. Its next step is to build more technology independence within SNA, and LU6.2 is the keystone of that strategy.

SNADS

SNA Distribution Services is an architecture for interchanging data through an SNA network in a store-and-forward fashion. One benefit of SNADS is that a user need not be aware of the availability of the system to which he wishes to distribute information. SNADS is an SNA extension that provides an asynchronous interface.

While SNADS may be suitable in certain situations for applications that require subsecond response time, it is mostly appropriate for applications without strong timeliness requirements, such as document distribution, file transfer and electronic mail.

SNADS requires that LU6.2, and the software that implements SNADS, reside in the same processor as the LU6.2 support.

Therefore, there is nothing in the SNADS architecture to prevent a network of small processors from implementing store-and-forward networking using SNADS -- without involving a large S/370 device with DISOSS.

SNAX

SNA Communication System is Tandem's SNA gateway. SNAX both makes a Tandem system act as an SNA host and enables Tandem devices (terminals) to access applications running on IBM hosts.

Soft-Switch

The name of a software package, as well as of the ISV supplying it, that performs translation of document across the formats of many hardware vendors.

Systems Center

An ISV specializing in software packages implementing SNA technologies. Its current product line consists of the Network Data Mover products, that perform information transfer within SNA environments.

Systems Strategies Inc.

SSI is an ISV wholly owned by AGS Computer. Its product line is written in C language and targeted at vendors of Unix environments. The products provide these environments with software links to SNA networks. The currently available products offer remote job entry and LU2 facilities. An LU6.2 capability is scheduled for the near future.

TCAM

The Telecommunications Access Method is an IBM program product that resides in a host processor and directs message traffic in a network, also controlling other communication among terminals and application programs running under an OS/VS operating system. TCAM conforms to SNA, but also allows the use of binary synchronous communication and start/stop devices in the network TCAM; it therefore has certain facilities lacking in VTAM, in particular message handling and queuing and non-SNA device support. IBM is actively trying to persuade TCAM users to migrate to VTAM.

Transfer

Transfer is a Tandem software facility that provides transport and delivery services, and is roughly comparable to IBM's DIA and SNADS operating together.

VTAM

The Virtual Telecommunications Access Method is an IBM program product that controls communication between terminals and application programs running under DOS/VS, DOS/VSE, OS/VSI and MVS. VTAM is IBM's strategic telecommunications access method, and has the following responsibilities:

- Controls resource allocation in the network;
- Establishes, controls and terminates sessions between Logical Units; and
- Provides operator and programmed interfaces to SNA facilities.

WSN

Wang System Networking is a framework, or umbrella, including all Wang networking products, built upon the IDS transport services. It is Wang's intent to build WSN into an architecture that will evolve toward the OSI model.

XRF

The Extended Recovery Facility is a reliability enhancement for an IMS/VS database/data communications environment. It provides the capability of creating a "mirror" IMS system and synchronizing it to the primary subsystem. XRF improved availability by causing the mirror IMS/VS system to take over the workload from the primary -- in case of some, but not all, disruptive events.

MARTIN SIMPSON & COMPANY, INC.
150 BROADWAY
NEW YORK, N.Y. 10038

MEMBER
NEW YORK STOCK EXCHANGE, INC.
MIDWEST STOCK EXCHANGE, INC.

TEL: (212) 406-5200
TELEX: 66544
FAX: 406-0088

MARTIN B. C. SIMPSON
President
ROBERT S. ANDERSON
Vice President
CECELIA M. BRANCATO
Vice President
J. T. T. BRUNE
Vice President & NYSE Member
CARMINE CORBO
Vice President & Controller
ROBERT C. DUNNE, JR.
Vice President
FRANK HALPERN
Vice President & Secretary
STEPHEN M. KOBAN
Vice President
MICHAEL M. NESBITT
Vice President
RAYMOND REED
Vice President
LE-ELLEN SPELMAN
Vice President

November 19, 1984

see page 4

TECHVIEWS

FROM: Martin B. C. Simpson

SUBJECT: Is European Technology Foundering - or Are There Signs of Renewed Potential?

We have recently returned from visits with the top management of more than twenty leading technology companies in Europe. Some of these are European in their origins and culture, such as Ericsson, Nixdorf, Olivetti and Philips. Others are American companies, although closely involved with European business activities for up to a century in some cases. A number of the companies are household names, such as IBM, Burroughs and NCR, with billions of dollars in revenues each year. We also interviewed several smaller companies, American, Israeli and European, that are growing at annual rates of 25-50%, with excellent profit prospects, in our analysis, and relatively modest price/earnings ratios. Companies falling into this category include Norsk Data and Scanvest Ring, two rapidly growing Norwegian firms, and Scitex, based in Israel.

In this brief review we touch upon some of the highlights and trends that we believe will be important for investors during the next twelve months. Our detailed comments and analyses will soon be published in a more comprehensive report.

Macroeconomic and Political Trends

European economic activity is in a recovery mode, fueled by the recent surge in U.S. growth and further aided by the strength of the American dollar. This has caused a number of European economies to improve from the low point in 1982, with certain technology companies enhancing their relative position through local manufacturing. Overall, we believe Europe's real economic growth is running at 2-3%, and that this is likely to be sustained in 1985. As usual, the countries have different factors affecting them:

Britain: The present coal miners' strike has reached a critical stage. It looks as though Arthur Scargill, the president for life of the National Union of Mineworkers, with strong Communist leanings and financial backing from Libya and the Soviet Union, may be beaten if the drift back to work continues. Picket violence has escalated and intimidation is widespread. Symbolically, the failure of the union to achieve what appears to be increasingly capricious goals in the context of 13% unemployment in the United Kingdom (3-4 million out of work, by most counts) could provide the sort of watershed for British labor relations that was effected by the showdown early in the first Reagan term with the air controllers' PATCO union. We believe that Mrs. Thatcher is likely to prevail, given her current increased popular support as a result of the abortive attempt to kill her in Brighton last month.

France: It is difficult for many investors in the United States to comprehend the extent to which French socialism has impacted the business environment. The dramatic decline of the French franc, which has fallen approximately 40% vis a vis the U.S. dollar since 1981, is only the outward and visible sign of the rigors of doing business in France. Well managed and effective companies, such as IBM, have escaped many of the vicissitudes heaped upon their less fortunate U.S. and Japanese counterparts by the French government. However, other foreign technology companies, that are not subject to preferential French treatment as a result of export considerations are often held ransom by archaic labor practices, an anti-business corporate environment, and a level of bureaucratic mismanagement that rivals that evident in many lesser-developed countries. Due to the gradual economic integration of the EEC countries and the nature of rapidly advancing technologies in computing and communications, we suspect the French government's penchant for trying to swim upstream against the more conservative global tide may provoke a sharp political reaction at some point in the future. This will hopefully lead to less counterproductive attitudes relating to entrepreneurial activity and real economic growth.

Germany: The German economy is also being affected by political problems. The present scandal embroiling local politicians relating to the Flick affair may lead to less stability in Germany, with traditional German economic strength challenged by a more turbulent political outlook. The Flick holding company is reported to have paid 25 million Deutsche marks (more than $8 million U.S.) to politicians and political parties in the late 1970's, at a time when it successfully lobbied for an 800 million Deutsche mark ($266 million) tax exemption on capital gains.

In general, these political, social and economic factors have helped reinforce the appeal of the ostensibly overvalued U.S. currency. They may continue to provide a basic demand for American dollars that would otherwise be less attractive in purely economic terms, especially if the present pause in U.S. economic growth proves short-lived. Given the backdrop, the financial implications may be as follows:

1. The U.S. dollar may continue to be "overvalued".

2. The contraction of international liquidity, that has occurred with some consistency now for several years, may be continued.

3. Deflationary tendencies may cause the prices of commodities (gold, oil, etc.) to give further ground, with "real" interest rates declining in the U.S. and a substantial upward revision in price/earnings multiples for well managed technology equities whose growth prospects are maintained. We are struck by the bargains that appear to exist among the technology stocks that we follow on both sides of the Atlantic. Many of the equity prices of technology companies have been beaten down by fears following the well-publicized over-expansion of the venture capital field in recent years, combined with the demise of a number of highly visible corporations that, in our view, had been heading for disaster for some time. This is therefore an excellent time to be adding to positions in selected technology equities, in spite of (or perhaps because of) the widespread public pessimism that now is apparent relating to the outlook. In our opinion, many smaller technology companies are selling at only 20% to 30% of their true value at present. If we are correct in this appraisal, more venturesome investors are apt to see selected technology stocks double or even quadruple in value over the next two to three years.

In Table 1 below we review the current stock prices for favored U.S. technology companies that we believe may provide above average capital appreciation over the years ahead:

Table 1

			Earnings Per Share				Price/Earnings Ratio				Long-Term Trend In Est. EPS Growth
Company	Symbol	Recent Price	1983	1984E	1985E	1986E	1983	1984E	1985E	1986E	Rate
Burroughs	BGH	$53 1/2	$4.60	5.50	6.70	7.60	11.6	9.7	8.0	7.0	11%
Data General	DGN	53 1/4	0.97	2.60A	4.30	5.65	54.9	20.5	12.4	9.4	18%
Digital Equip.	DEC	101 1/4	5.00	5.73A	8.25	10.65	20.3	17.7	12.3	9.5	20%
IBM	IBM	120	9.04	10.70	12.30	14.10	13.3	11.2	9.8	8.5	15%
NCR	NCR	25 5/8	2.64	2.95	3.25	3.65	9.7	8.7	7.9	7.0	11%
Tandem Computers	TNDM	17 1/8	0.76	0.81	1.25	1.60	22.5	21.1	14.3	10.7	25%
Ungermann Bass	UNGR	14	0.10	0.35	0.60	1.00	140.0	40.0	23.3	14.0	40%
Wang Labs	WANB	27 1/4	1.16	1.52A	1.90	2.50	23.5	17.9	14.3	10.9	30%

A - Actual
E - Estimates of Martin Simpson & Company, Inc.

Reviewing Table 1, we believe that substantial capital appreciation potential exists for patient portfolio managers who invest in the equities presented. We are especially enthusiastic about companies such as IBM, Digital Equipment, Tandem Computers, Ungermann Bass and Wang Labs, in view of the above-average long-term growth prospects and the market positions achieved in industry areas that are likely to grow at rates well in excess of GNP. We do not buy the argument that the technology arena is now so fraught with competition that few companies deserve a premium multiple. In our analysis, the advent of the microprocessor will lead to dramatic growth for at least the next decade. Well managed participants in the computer, minicomputer, office automation and telecommunications fields should be able to take advantage of this technological revolution in a highly profitable way. While new technologies can hardly be characterized as "risk free", we believe the present environment of gloom and uncertainty has resulted in a minimal price/earnings premium for top quality growth equities. The investment opportunity is therefore significant, in our analysis.

In Table 2 we review the current stock prices and growth prospects for a number of indigenous technology companies in Europe, converted to U.S. dollars at current rates.

Table 2

Company	Fiscal Year	Price	Earnings Per Share				Price/Earnings Ratio				Long-Term Trend in Growth Rate
			1983	1984E	1985E	1986E	1983	1984E	1985E	1986E	
British Telecom (a)	PF 3/31	$1.60	.23	.22	.17	.19	7.0	7.3	9.4	8.4	10%
L M Ericsson (b)	12/31	$30 1/4	2.57	2.80	3.25	3.50	11.8	10.8	9.3	8.6	15%
Logica (a)	6/30	$5.06	0.11	0.11	0.16	0.20	46.0	46.0	31.6	25.3	30%
Nixdorf (c)	12/31	$177	5.32	6.65	8.00	9.30	33.3	26.6	22.1	19.0	20%
Norsk Data (d)	12/31	$39 1/4	$1.90	2.15	2.75	3.50	20.7	18.3	14.3	11.2	25%
Olivetti (e)	12/31	$3.20	.46	.32	.37	.42	7.0	10.0	8.6	7.6	12%
Philips (f)	12/31	$15 3/4	$1.42	1.85	2.20	2.40	11.1	8.5	7.2	6.6	15%
Scanvest Ring (d)	12/31	$ 8.35	.30	.67	.90	1.25	27.8	12.5	9.3	6.7	40%
Scitex	12/31	$15 1/2	$1.08	1.30	1.70	2.40	14.4	11.9	9.1	6.5	30%

E - Estimates of Martin Simpson & Company, Inc.

PF - Pro Forma, assuming the new capital structure had been in place, with taxes paid in fiscal 1985 and 1986 and common shares issued and outstanding.

(a) £ = $1.238 throughout

(b) $1 U.S. = 8.595 Swedish Krone throughout

(c) $1 U.S. = 3.007 German Marks throughout

(d) $1 U.S. = 8.715 Norwegin Krone throughout

(e) $1 U.S. = 1,863 Italian Lira throughout

(f) $1 U.S. = 3.395 Dutch Guilders throughout

There are several reasons why current purchases of European technology equities could work out rather well, in our judgment:

1) The high appreciation in the U.S. dollar relative to European currencies may be reversed at some point, which could cause the earnings for European equities to accelerate in U.S. terms.

2) Many foreign technology companies are growing quite rapidly at present, with management skills that are comparable in quality with those at U.S. corporations.

3) International markets served may be less mature, leading to superior growth prospects over the next several years for European companies with a strong product position and effective marketing and management capabilities.

4) The high U.S. dollar has caused some technology companies, such as IBM, to buy more of their component needs abroad, thereby stimulating the technology transfer already in process in several European countries.

5) There is an increasing awareness among international investors that local technology companies can provide above-average capital appreciation possibilities. This should heighten the appeal of well-managed high technology companies abroad.

6) As communications advance and distances shrink around the world, the appeal of capitalism will generate additional investors in growth equities in the technology field, both in the U.S. and abroad, in our judgment.

7) New markets that are opening up, including China, India, Korea and other countries in the Far East, should further expand the impact and opportunities for electronics products and technology.

Therefore, while we do not minimize the problems posed by currency fluctuations, foreign cultural and accounting differences and a host of other factors that may mitigate against foreign investment, we remain convinced that investment opportunities with above average growth potential will continue to appear in Europe.

The companies covered in Table 2 all appear to have interesting long-term investment merit. Brief comments on each appear below:

British Telecom:

This company should participate in the growth foreseen for international telecommunications in the United Kingdom. The terms of the present stock offering provide more than 13% in dividend yield on the initial flotation and we believe the stock is likely to go to a premium when issued.

L M Ericsson:

This leading international telecommunications company currently faces an increasingly competitive industry environment, in our opinion, with recent acquisitions in the process of being digested. IBM and Rolm, ITT and Rovsing, AT&T and Olivetti, STC and ICL, and a host of other companies entering the telecommunications arena cause us to be less enthusiastic about Ericsson's shares at present.

Logica — This well-managed software company has had an excellent growth record. The premium price of the stock causes us some concern on a short-term basis, but we like the long-term fundamentals.

Nixdorf — This German computer company has come public only recently and is one of the most aggressively managed competitors, in our view. The recent pricing of IBM products at the low end of the market in Germany, combined with the likelihood that Burroughs will become more effective in penetrating the German market over the years ahead, suggests that Nixdorf stock may be somewhat fully priced currently. Long-term, we believe Nixdorf will continue to achieve above average revenue and earnings growth.

Norsk Data — This Norwegian minicomputer company is benefitting from buoyant demand for the company's products in Scandinavia, Britain, and, to a lesser extent, West Germany, where the Dietz acquisition is posing some unexpected challenges. Long-term, Norsk is likely to grow at a rate of at least 20-30% and the shares are attractive for capital appreciation, in our view.

Olivetti — This Italian office automation giant has a balance sheet that is still less strong than that of other large technology companies. The 25% equity purchase by AT&T provides interesting possibilities for mutual collaboration. Short-term, the lower Italian government research subsidies and more shares outstanding could cause the stock to be a less interesting vehicle, in our analysis.

Philips — Continues to benefit from more stringent cost controls and better management direction. The shares still look undervalued, in spite of their having increased more than 30% in price in the last twelve months, with earnings per share in U.S. dollar terms (using GAAP accounting) up 70% in the first 9 months of 1984.

Scanvest — This aggressive Norwegian vendor of information systems was founded in 1973 and is highly profit-oriented. The company is generating about $100 million in revenues currently. The business has relied to some extent on acquisitions and internal growth is perhaps closer to 30% than the 70% rate reported for the 1984 first half. As the price/earnings multiple is still relatively modest, we believe the stock could show significant capital appreciation over the years ahead. Competitors in the U.S. include AT&T, ITT, Rolm and Collins' Radio, among others, so that we regard the shares as more speculative than many of the other companies reviewed in this report.

Scitex — Earnings are growing at 20-30%, with some acceleration possible in 1985 due to the introduction of attractive new products. The company is located in Israel, but is rapidly expanding its presence in the U.S., Europe and Japan. The shares appear to be cheap and we believe they should perform well over the years ahead in those portfolios that can tolerate the risks of owning a venture still heavily dependent on political developments in the volatile Middle East.

Further information is available on request

Cray Research

Much like Apollo's situation, Cray's outlook in 1985 is very strong. This pioneer in supercomputers has dominated its market niche and faced limited competition from Control Data. While it is much more difficult to enter this niche than the engineering workstation segment, it is now facing the challenge from three Japanese vendors for the first time, two of which offer a S/370 compatible approach. Since some 40% of all Fortran programs are run on IBM mainframes, a S/370 compatible approach is a logical one because it will be able to share existing mainframe peripherals and will automatically vectorize existing Fortran programs without conversion. In part to protect its customer base and keep up technologically with the Japanese vendors, which it sees as its principal competitors, IBM is making a major effort in this area. A Cray I class machine should be forthcoming in the 1986/1987 time frame. The implications are that the low and medium ends of the supercomputing market will be quite competitive involving industrial customers and universities.

Cray is likely to migrate its products upward and concentrate on the high end in order to maintain its current high margins. It is not clear to us how fast this segment will grow and hence, our best estimate for Cray's growth rate past 1985 is 20%. Longer term, the growing Japanese leadership in high-performance VLSI circuits threatens all U.S. computer vendors of high-end equipment according to the Gartner Group and is of course, of particular concern to U.S. supercomputer companies such as Cray. This Japanese effort is particularly evident in their emphasis in galium arscenide, which is expected to be the future high-speed logic and memory technology.

Again we believe the stock's P/E will decline in 1985 to allow for slower future growth of around 20% versus the approximately 30% experienced during the past few years. We would sell the stock on strength.

[redacted]

IBM is Tandem's most frequently encountered competitor, involved in approximately 90% of the bids. As Tandem increases the performance range of its processors, it is positioning itself as the new distributed mainframe company and as an alternative to IBM in core data processing applications among the largest industrial corporations and financial institutions in the world. What it lacks in customer relations and depth of support relative to IBM, it makes up for by the inherent superiority of its system architecture, which was originally designed for the on-line transaction environment. Speculation about an IBM response has centered around its development of a fault tolerant architecture based on its System 38 and the possibility that such a system will be introduced around 1987. In view of IBM's emphasis on TPF-2 as its strategic product for on-line transaction processing, and its current five incompatible families of intermediate processors, it is unlikely that IBM will introduce such a revolutionary multiprocessor architecture anytime soon. According to the Gartner Group, IBM's System 370 architecture is likely to evolve to become more like its System 38. It will probably require 10 years to migrate to machine-independent SNA/LAN architecture and truly relational data base management. Hence, for the forseeable

...ure Tandem will be able to continue to enjoy system architectural superiority with its ...ulti-processor approach.

The reason behind IBM's emphasis on the TPF-2 as a strategic product, which is difficult to install and support, is because of increasing competitive pressure from Tandem. Current large IBM users find that DB2 (IBM's new relational data base system) tops out at 25-30 transactions per second, and its IMS system (which is not relational) tops out at 125 transactions per second. Since Tandem can offer a throughput of 50-100 transactions per second with a large TXP processor configuration and can also offer relational data base and superior peer-to-peer networking capability, IBM is forced to emphasize the high throughput of its TPF-2 system, an adaptation of the ACP system used for airline reservations, which is difficult to use. As technology advances, we believe Tandem should be able to maintain its technological edge over IBM for at least the rest of this decade.

We are recommending the stock for purchase at current prices because of: (1) the profit margin recovery we envision for the company over the next two years; (2) the imminent series of new product introductions; and (3) the likelihood of limited direct system competition from the major computer firms over the next five years.

David Wu, CFA

Goldman Sachs

Investment Research

Tandem Computers

John C. Levinson

September 28, 1984

Investment Research

Goldman Sachs

September 28, 1984

TANDEM COMPUTERS

	Current Price	Earnings Per Share* 1983	1984E	1985E	P/E Ratio 1984	1985	Indicated Dividend	52-Week Price Range
TNDM	16	$0.76	$0.75	$1.20	21.3X	13.3X	Nil	41 - 13

DJIA 1217
S&P 400 188

*Fiscal year ends September 30; earnings are based on 40.8 million shares outstanding in 1983, 41.4 million in 1984, and 42.0 million in 1985.

INVESTMENT SUMMARY

It has been very difficult to make money in Tandem shares over the past six months; twice, we have seen quarterly earnings disappointments peel more than 35% off the stock price. In fact, investor dismay goes well beyond the past six months; nearly four years ago, when Tandem last raised equity, the stock was at $22, a multiple of 63 times trailing 12-month earnings. Today, after only an 18% cumulative earnings gain in the past four years, the stock is some 27% below that level, at a multiple of 21 times (see Figure 1).

We have become increasingly positive on Tandem Computers. Tandem has been a very controversial stock over the past several years, particularly during the two most recent quarters, and we have clearly been on the negative side (see our May 8th Report). So far this year, the stock is down 57% while the average minicomputer company's stock is down about 10%. We know of no major firm currently recommending the shares, and in fact, many major accounts appear unwilling to listen to anything other than a negative story on the stock or company. We are clearly in the vast minority in being somewhat constructive. Given the stock's sharp correction and some discernible changes at the company, we now believe Tandem shares have some speculative appeal at current levels. Our opinion is based on the following four points.

FIGURE 1

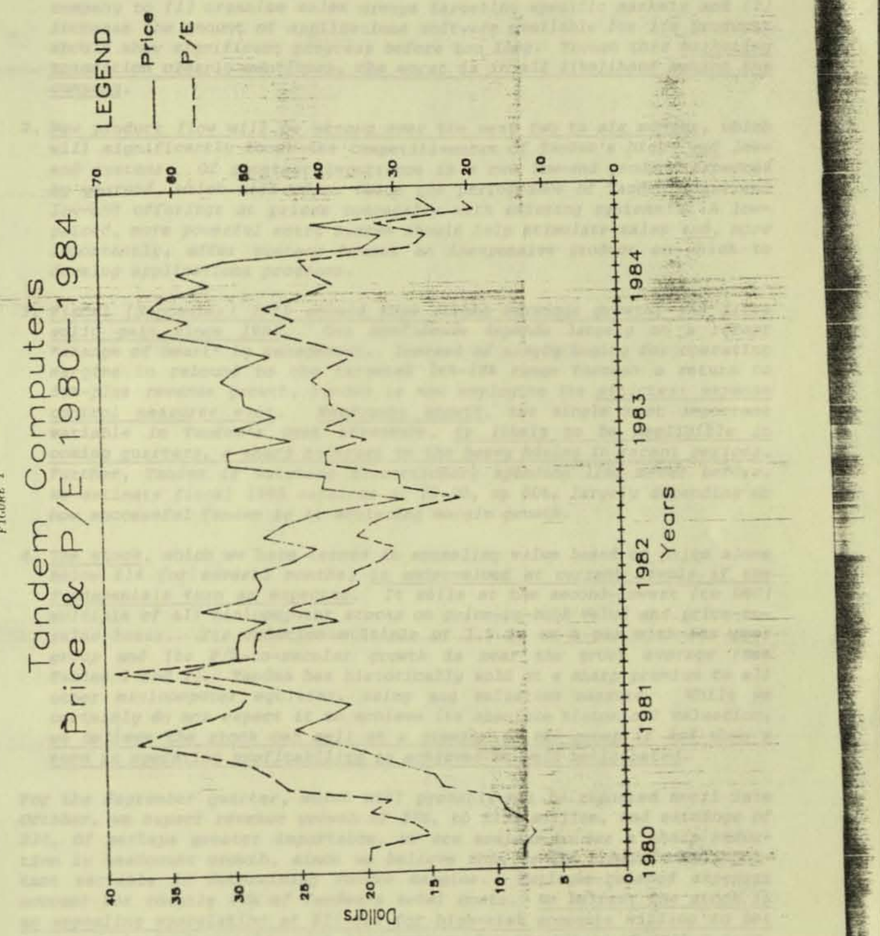

1. *The marketing transition that has plagued the company for the past two and a half to three years will probably end sometime in the next several quarters. It is impossible to pinpoint when Tandem's transition to a more market-focused company will gel, but recent all-out efforts by the company to (1) organize sales groups targeting specific markets and (2) increase the amount of applications software available for its products should show significant progress before too long. Though this marketing transition clearly continues, the worst is in all likelihood behind the company.*

2. *New product flow will be strong over the next two to six months, which will significantly boost the competitiveness of Tandem's high- and low-end systems. Of greatest importance is a new low-end product expected by yearend, which will offer twice the performance of Tandem's current low-end offerings at prices comparable with existing systems'. A low-priced, more powerful entry system should help stimulate sales and, more importantly, offer systems houses an inexpensive product on which to develop applications programs.*

3. *Fiscal (September) 1985 should show strong earnings growth, the first solid gain since 1981. Our confidence depends largely on a recent "change of heart" by management. Instead of simply hoping for operating margins to rebound to the targeted 16%-18% range through a return to 50%-plus revenue growth, Tandem is now employing its strictest expense control measures ever. Headcount growth, the single most important variable in Tandem's cost structure, is likely to be negligible in coming quarters, a sharp contrast to the heavy hiring in recent periods. Further, Tandem is watching discretionary spending like never before. We estimate fiscal 1985 earnings at $1.20, up 60%, largely depending on how successful Tandem is at achieving margin growth.*

4. *The stock, which we have termed an appealing value based on price alone below $14 for several months, is undervalued at current levels if the fundamentals turn as expected. It sells at the second-lowest (to DEC) multiple of all minicomputer stocks on price-to-book value and price-to-sales bases. Its relative multiple of 1.5 is on a par with its peer group and its P/E-to-secular growth is near the group average (see Tables 1 and 2). Tandem has historically sold at a sharp premium to all other minicomputer equities, using any valuation measure. While we certainly do not expect it to achieve its absolute historical valuation, we believe the stock can sell at a premium to the group if and when a turn in operating profitability is achieved or well anticipated.*

For the September quarter, which will probably not be reported until late October, we expect revenue growth of 27%, to $150 million, and earnings of 23¢. Of perhaps greater importance, we are anxious to see a sharp reduction in headcount growth, since we believe this is the single most important variable in determining future margins. Employee-related expenses account for roughly 40% of Tandem's total costs. We believe the stock is an appealing speculation at $15-$20 for high-risk accounts willing to bet on a turn in fundamentals, with less concern for the timing of such a turn.

Table 1

Earnings Estimates, Fiscal 1983-1985E

Company	Fiscal Year	Sept. 27 Price	Earnings 1983	Earnings 1984E(a)	Earnings 1985E(a)	P/E Ratio Absolute 1984	P/E Ratio Absolute 1985	P/E Ratio Relative 1984	P/E Ratio Relative 1985
Data General	Sept.	$51.75	$0.97	$1.50	$1.40	20.7 x	11.8 x	2.0 x	1.3 x
Digital	June	97.25	1.00	2.15	3.00	14.4	11.7	1.4	1.1
Hewlett Packard	Oct.	30.00	1.68	2.15	3.00	17.1	12.7	1.7	1.4
Prime	Dec.	16.25	0.68	1.10	1.55	14.8	10.3	1.4	1.2
Tandem Computers	Sept.	16.25	0.76	0.75	1.20	21.7	13.5	2.1	1.5
CPT Corp.	June	8.50	1.07	1.05 E	1.10	8.1	7.7	0.8	0.9
IBM	June	19.50	0.45	1.40 E	2.00	13.9	9.8	1.4	1.1
Wang Labs	June	26.50	1.16	1.51 E	1.80	17.5	14.7	1.7	1.7
Burlow PLC	Sept.	8.88	0.64	1.20	1.00	7.4	8.9	0.7	1.0
Dasgate	June	6.25	0.33	0.95 E	0.70	6.6	8.9	0.6	1.0
Tandon	Sept.	7.75	0.48	0.85	1.10	9.1	7.0	0.9	0.8
Average						14.0	10.6	1.4	1.2
IBM	Dec.	125.88	9.04	10.60	12.30	11.9	10.2	1.2	1.2
S&P 400	Dec.	188.62	14.73	16.50	21.50	10.2	8.9	1.0	1.0

Table 2

Calendar P/E and Valuation Statistics, Calendar 1984-1985

	Sept 27 Price	Cal '84 EPS(a)	Cal '85 EPS(a)	P/E Ratio Absolute 1984	P/E Ratio Absolute 1985	P/E Ratio Relative 1984	P/E Ratio Relative 1985	P/E-to-Secular Growth (c) Secular Growth (b)	P/E-to-Secular Growth (c) 1984	P/E-to-Secular Growth (c) 1985	Return on Equity(d)	Price-to-1984 Book	Price-to-1984 Revenues	Book Value Per Share	Revenues Per Share
Data General	$51.75	$2.90	$4.50	17.8 x	11.5 x	1.8 x	1.3 x	25 %	0.71 x	0.46 x	11.4 %	1.92 x	1.08 x	$24.08	$47.95
Digital	97.25	7.65	10.00	12.7	9.7	1.2	1.1	22	0.58	0.44	9.9	1.20	0.89	73.38	108.72
Hewlett Packard	30.00	2.40	3.05	15.8	12.5	1.6	1.4	22	0.72	0.57	15.2	2.24	1.46	14.56	26.09
Prime	16.25	1.10	1.55	14.8	10.5	1.4	1.2	20	0.74	0.52	15.1	2.08	1.21	6.73	13.43
Tandem Computers	16.25	0.80	1.30	20.3	12.5	2.0	1.4	27	0.75	0.46	8.5	1.65	1.30	9.03	13.59
CPT Corp.	8.50	1.10	1.15	7.7	7.4	0.8	0.8	18	0.43	0.41	12.5	0.91	0.67	8.25	12.77
IBM	19.50	1.25	1.15	11.1	9.1	1.1	1.0	20	0.56	0.45	14.7	1.53	0.94	11.02	20.74
Wang Labs	26.50	1.65	2.05	16.1	12.9	1.6	1.5	27	0.59	0.48	15.6	2.32	1.49	9.25	17.81
Burlow PLC	8.88	1.25	1.15	7.1	7.7	0.7	0.9	25	0.28	0.31	15.0	0.99	0.81	7.73	10.97
Dasgate	6.25	0.75	0.85	8.3	7.4	0.8	0.8	25	0.33	0.29	17.1	1.33	0.84	3.95	7.47
Tandon	7.75	0.90	1.15	8.6	6.7	0.8	0.8	20	0.43	0.34	15.8	1.26	0.68	5.26	8.76
Average				12.8	9.8	1.3	1.1	23	0.61	0.43	13.7	1.58	1.04	15.79	26.21
IBM	125.88	10.60	12.30	11.9	10.2	1.2	1.2	15	0.79	0.68	19.7	2.13	1.63	48.63	77.00
S&P 400	188.62	18.50	21.30	10.2	8.9	1.0	1.0	9	1.13	1.00	14.0	1.69	NA	154.00	NA

(a) Excludes one time extraordinary gain from reversal of DISC benefits.
(b) Annual earnings growth rate through 1987.
(c) Estimated absolute P/E divided by secular growth.

Recent Results: Very Disappointing to Most

Tandem's third (June) fiscal quarter earnings were 23¢ versus 21¢ on revenues of $142 million, up 29%. This was better than our expectations of 20¢ and $132 million, respectively, yet well below general Street expectations of 28¢-30¢ and $150 million. Sequential revenue growth for the June period was a very respectable 28%, the best gain in six years. Though some investors had expected a revenue spillover to aid June's level, our conservatism was based largely on a view that any spillover effect would be modest, at best. We believe that the spillover impact has perhaps $5-$10 million.

Given the revenue levels achieved in the quarter, Tandem's operating margins were soft. Gross margins were 59.3%, a shade lower than expected because of a previously announced (September) Non-Stop II upgrade program; the TXPs shipped under this program carried lower margins than "normal" TXP products. While R&D spending was in line with our expectations, SG&A levels were some $5 million (or 10%) greater than anticipated, a direct result of the past two quarters' aggressive hiring programs. In fact, Tandem's headcount was up 5% sequentially in the June quarter, on the heels of March's 9% growth. As shown in Table 3, Tandem's hiring has been far more aggressive than any other minicomputer company's over the past year or so. Since an estimated 40% of Tandem's costs are employee-related (below the 60% average for the minicomputer group because of Tandem's lack of vertical integration), the high SG&A spending levels were clearly influenced by this robust headcount growth. Beyond this, however, several extraordinary SG&A costs were absorbed this quarter: a higher reserve for doubtful accounts, a higher-than-average proportion of salary increases, and some variable marketing costs budgeted using a higher revenue assumption. Despite these extraordinary costs, most of the higher SG&A levels are believed to be employee-related, and hence, control of headcount growth is clearly one of the most important determinants of future spending levels. Interest income and the quarter's tax rate were both in line with our expectations.

Table 3

Quarterly Employment Growth, 1983-1984

Period	DEC	DGN	HWP	PRM	TNDM	WANB	Average
1983							
Mar	3%	(2)%	2%	5%	1%	6%	3%
Jun	4	0	1	2	4	9	3
Sept	3	2	3	3	6	6	4
Dec	5	4	4	2	5	7	5
1984							
Mar	3	4	3	4	9	6	5
Jun	6	5	3	5	5	4	5
Current Year-Over-Year Growth							
Employment	18%	14%	13%	14%	27%	25%	19%
Revenues	34	47	35	32	29	52	38

Table 4

Customer and Unit Shipment Trends, Calander 1981-1984

Number of Processors Shipped

Year	Mar	Jun	Sept	Dec
1981	279	325	367	405
1982	390	373	374	396
1983	370	386	621	573
1984	463	569		

Number of Customers Shipped to

Year	Mar	Jun	Sept	Dec
1981	101	108	131	120
1982	105	121	135	116
1983	119	115	159	158
1984	133	171		

Number of First Time Users

Year	Mar	Jun	Sept	Dec
1981	40	37	53	44
1982	26	31	38	35
1983	25	19	47	33
1984	25	39		

First Time Users as a % of Total Customers

Year	Mar	Jun	Sept	Dec
1981	0.40	0.34	0.40	0.37
1982	0.25	0.26	0.28	0.30
1983	0.21	0.17	0.30	0.21
1984	0.19	0.23		

In the June quarter, Tandem shipped 569 processors to 171 customers, both healthy numbers. Many of these represent shipments of TXP upgrades to existing accounts, which adds an upward bias to both figures. An important barometer of Tandem's future prospects is its ability to penetrate new accounts. Tandem shipped to 39 new customers in the June quarter, its best showing in the past three quarters and second best in two years. New customers represented a respectable, but not overwhelmingly strong, 23% of total customers. Trends in customer and unit shipments are shown in Table 4.

What's going on at Tandem?

Many are perplexed by Tandem's relatively lackluster year-over-year revenue growth given (1) the company's historical premium growth relative to almost all medium and large computer vendors, (2) the strong economy that has substantially lifted sales of nearly all competitors, and (3) Tandem's strong, TXP-related current product cycle. We believe the slowdown is largely a continuation of a transition phase Tandem has been in for two years. To gain perspective, it is valuable to take a brief look at where the company came from.

From its founding in 1975 through March 1982, Tandem had one of the (if not the) industry's most spectacular, and predictable, revenue and earnings growth records. During that six-year period of "hyper-growth," its revenues grew to an annuallized rate of $345 million and, perhaps even more impressive, its employment approached 2,000 people, nearly all trained and dedicated to the direct sale and support of Tandem's products worldwide. Things changed abruptly in June 1982 (the first significant disappointment at Tandem). As seen in Table 5 and perhaps more easily in Figure 2, since that June quarter two years ago, Tandem's quarterly sequential growth has averaged 7%, a compound annual rate of 31%. This is in sharp contrast to the 17% (87% annualized) average sequential growth of the prior two-year period. Note also the volatility of quarterly growth of late, relative to its consistent previous record.

Our point is this; Tandem's period of hyper-growth is gone forever. We believe that three key factors caused this prolonged slowdown and are largely responsible for Tandem's recent quarterly disappointments. We review them below and explain what actions Tandem has begun to meet these challenges.

1. Loss of Price/Performance Edge at Low End. Years ago, Tandem's computers almost sold themselves, based on superior price/performance characteristics and, perhaps more important, two unique qualities — fault-tolerance (its most recognized trait) and modular expandability (an extremely easy, efficient, and inexpensive way to add power). As time passed and Tandem grew larger, its commitment to its installed base became more pronounced. In fact, every product announcement since the initial Tandem Nonstop of 1976 was more powerful than the previous model, a strategy that obviously greatly benefited existing major

Table 5

Revenue Growth, Fiscal 1980-1984E

(millions)

Period		Revenues	Gain: Sequential-Quarter	Gain: Year-to-Year
1980	1Q	$20.8	14.9 %	100.0 %
	2Q	24.9	19.7	99.2
	3Q	29.2	17.3	94.7
	4Q	34.1	16.8	88.4
	Year	$109.0	—	94.6
1981	1Q	$40.6	19.1	95.2
	2Q	47.4	16.7	90.4
	3Q	55.9	17.9	91.4
	4Q	64.5	15.4	89.1
	Year	$208.4		91.2
1982	1Q	$71.0	10.1	74.9
	2Q	74.1	4.4	56.3
	3Q	79.8	7.7	42.8
	4Q	87.2	9.3	35.2
	Year	$312.1		49.8
1983	1Q	$94.1	7.9	32.5
	2Q	96.0	2.0	29.6
	3Q	110.4	15.0	38.3
	4Q	117.9	6.8	35.2
	Year	$418.4		34.1
1984	1Q	$126.4	7.2	34.3
	2Q	111.2	(12.0)	15.8
	3Q	141.9	27.6	28.5
	4QE	151.5	6.8	28.5
	YearE	$531.0		26.9

FIGURE 2

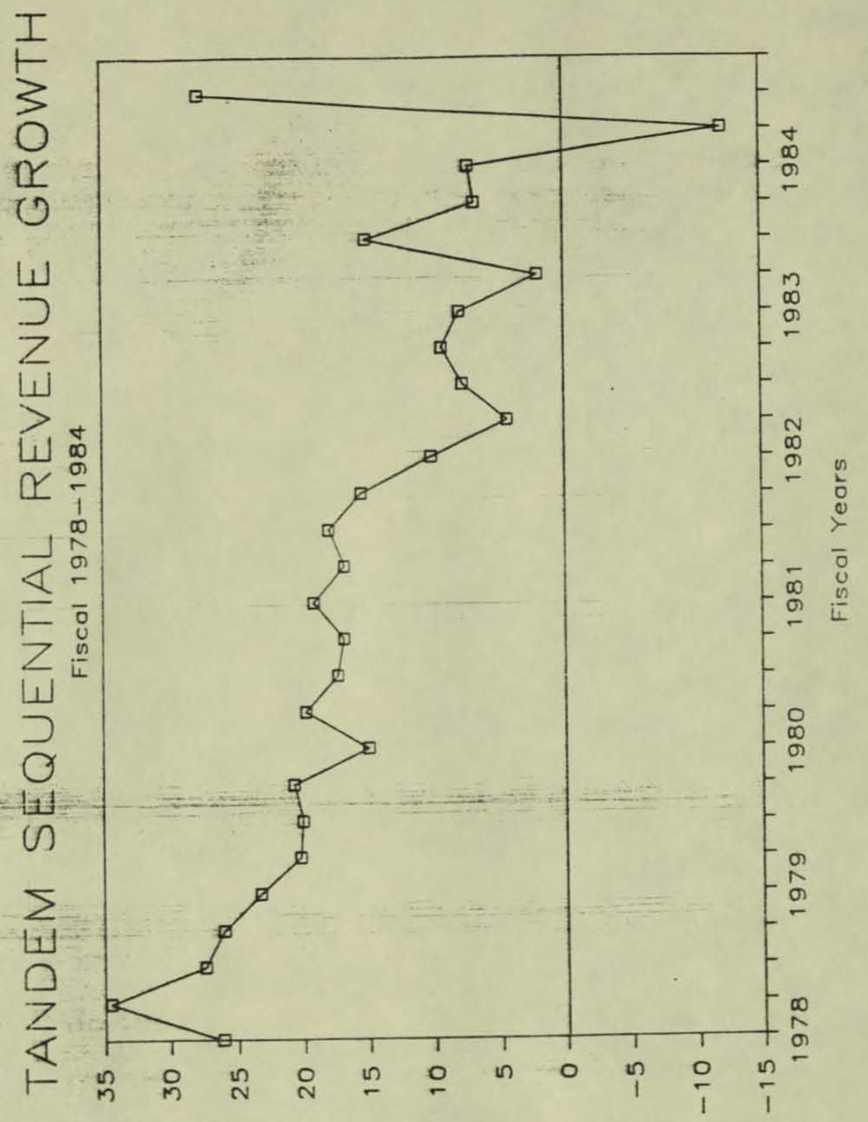

customers that had standardized on Tandem's proprietary languages and required expanded processing power. Furthermore, since the late 1970s, Tandem reduced processor prices at its aging low end only modestly. Concurrently, fault-tolerant competition aimed at Tandem's neglected low end market sprang up, the most notable being Stratus, and other mid-range minicomputers from traditional competitors (DEC, IBM, and so on) became inherently more reliable through technical advances. The lack of effort in the low end market is perhaps best shown by Tandem's product mix shift since the introduction of the TXP high-end product at the beginning of fiscal 1984. It immediately captured 40% of Tandem's shipments, and now accounts for more than 70%. In the process, the dollar volume from lower-end products has been cut in half (see Table 6).

Table 6

Estimated TXP Sales Contribution, Fiscal 1984
(millions)

	1st Qtr	2nd Qtr	3rd Qtr
TXP	$ 40	$50	$ 85
Other (Non Stop I and II)	69	41	34
Total shipments	$109	$91	$119
	As a % of Total		
TXP	40%	55%	71%
Other (Non Stop I and II)	60	45	29
Total shipments	100%	100%	100%

Several weeks ago, Tandem took its first major step in years to bolster its competitive position at the low end. It cut prices on NonStop I and II products 13%-30%, depending on configuration and volume. We expect further significant moves at the low end, including a new release of the Guardian operating system, which could boost performance 50%-75%, and a new low end system (see page 12). These products will be announced within the next three to four months.

2. Focus on Large-System Marketing. As Tandem's competitive thrust waned at the low end, success in selling higher-end products proved frustratingly slow to emerge. Over the years, Tandem has developed strong system level software (DBMS, languages, networking, and so on) to complement its relatively strong high-end price/performance, yet it failed to fully recognize the major differences in marketing large systems to Fortune 500-sized accounts. The longer sales cycle associated with these larger orders also proved to be a surprise to Tandem, as the company has repeatedly misjudged the amount of time required to close such deals. This has added significant volatility to its quarterly revenues and has caused some frustration in Tandem's salesforce, prompting a greater-than-normal turnover during the last six to nine months.

During the past year, Tandem has established seven industry-oriented marketing groups -- airlines, distribution, federal government, financial, retail, telecommunications, and manufacturing. Each group has about six headquarters employees plus regional support personnel to better target specific accounts and markets. Further, better-quality and focused applications software is expected; Tandem now has more than 30 software houses writing applications programs for its NonStop systems, compared with less than 10 just 9 months ago. A transition to this type of marketing organization with more applications software is necessary to improve Tandem's batting average at large accounts, but takes time to bear fruit. Also, successful evolution into a better marketing organization should improve recent turnover trends.

3. Increasing IBM Competition. The more a company sells high-end products, particularly those with mainframe power, the more it comes into direct marketing competition with IBM. Tandem's TXP competes in many (if not most) cases with IBM's 308X line, IBM's most important product from revenue, profit, and customer control standpoints. Successful product competition against IBM is not impossible; Tandem's products are superior from a price/performance standpoint, and Tandem's on-line transaction processing system is tough to match, even for IBM. On the other hand, when it comes to marketing, we believe IBM outshines Tandem more often than not. By focusing on increasingly larger sales, Tandem has "backed into" a position where it is competing with IBM in the market the computer giant holds dearest, the market for large mainframe applications that are central to an account's data processing operation.

   Increasing marketing competition from IBM is probably the most difficult hurdle Tandem has to negotiate to get back on the fast-growth track; it is also perhaps the most difficult for investors to evaluate. We believe there are two consequences of stiffer IBM competition: loss of sales bids and lengthened selling cycles when IBM throws its top marketing talent at a particular sales situation. Tandem claims to have lost only one major deal to IBM in the June quarter; on the other hand, while difficult to quantify, sales cycles are often longer when selling large TXP systems against IBM. To help with this transition, Tandem has recently hired a number of former IBM marketing and sales people. One of the most important is Lawrence McGraw, an 18-year IBM veteran who is now the vice president responsible for all of Tandem's marketing and field support. Mr. McGraw has been instrumental in forging a more marketing-oriented company, with vertical market groups supported by headquarters.

Revenue Outlook: New Product Flow will Aid Competitive Posture

Tandem's near-term revenue growth is somewhat uncertain because of the continuing marketing transition and recent price cuts. However, there are three major positives -- the strong economic environment, Tandem's good (and soon to be improved) product cycle, and the rapid underlying growth of the on-line transaction processing market (estimated at 30%). We have discussed our positive assessment of the economic environment at length in past reports (see our May Quarterly). In short, our Economic Research

Group expects a strong capital spending environment to continue throughout calendar 1985.

Tandem is in the midst of a favorable high-end (TXP) product cycle. We now believe that the high performance TXP accounts for more than 70% of sales revenue. Also apparent is the declining contribution from Tandem's older, less powerful NonStop I+ and II products (see Table 6). Within the next three to four months, we expect to see Checkmate, a new low-end product priced at the same level as the NonStop I+ ($75,000-$100,000), offering nearly twice its performance, or about 2.5 MIPS for an entry system. Checkmate will be a small "under the desk" system suitable for an office environment. It will probably make obsolete both the NonStop I+ and II, and give Tandem a far better low-end offering to more effectively compete with Stratus and others, and allow current customers to add inexpensive systems on an existing or new Tandem network. In addition to Checkmate, by early next year, Tandem will announce several new workstations, disk subsystems, and a major new release of its operating system.

Further out, we anticipate that by the end of calendar 1985, Tandem will announce a high end follow-on to its TXP product, dubbed Rainbow, offering four MIPS of performance per processor (or eight MIPS for an entry level, two-processor system), about double the capacity of the TXP. This product will, in all likelihood, not be fully compatible with its existing product line. While able to use the same communications network and share all data, we doubt if Rainbow will be fully applications software compatible.

In short, we believe that from an economic and product standpoint, Tandem is positioned to achieve even stronger revenue growth through calendar 1985 than its average of the past several years. The ongoing marketing transition at Tandem, however, adds a degree of uncertainty to our forecast. We have assumed that despite the above positive forces, Tandem's revenue growth will be about 5%-7% sequentially (in line with growth for the past two years) for the foreseeable future, with the notable exception of its seasonally weak second (March) fiscal quarter. In all, we expect fiscal 1985 revenues of $660 million, up 25%. Our quarterly forecast is shown in Table 7.

Emphasizing Margins: A Change of Heart at Tandem

During meetings with Tandem several times this summer, it became clear that the company had begun to turn over a new leaf on expense control. In most, if not all meetings of the past two difficult years, Tandem emphasized how it would achieve margin expansion through higher revenue growth, usually mentioning growth targets of 35%-50%. During this period, such revenue growth and margin expansion simply failed to take hold; operating margins have been largely bound by the 10%-13% range, far below Tandem's 16%-18% proclaimed target. We believe Tandem is now going through a change of heart in its approach to margins. The company openly talks about strong expense control measures and is currently (for the first time ever) implementing a selective lay-off program. Consequently, we expect headcount

Table 7

Tandem Computers, Income Statements, 1983-1985E

(millions)

	Fiscal 1984 Q1	Q2	Q3	Q4E	Fiscal 1985 Q1E	Q2E	Q3E	Q4E	Fiscal Years 1983	1984E	1985E
Sales	$108.9	$91.2	$119.1	$125.0	133.8	120.4	138.4	149.5	$360.1	$443.8	$542.1
Service	17.9	20.0	22.9	25.0	27.0	29.0	31.0	34.1	58.2	85.8	$121.1
Total Revenues	$126.4	$111.2	$141.9	$150.0	$160.8	$149.4	$169.4	$183.6	$418.3	$529.5	663.2
Cost of Revenues	$50.4	$47.2	57.8	$60.0	64.3	$62.0	66.1	$70.7	$168.7	$215.5	$263.1
SG&A	48.2	49.1	56.3	60.0	62.0	58.0	64.4	69.8	160.6	213.6	$254.2
R&D	10.8	12.9	13.5	14.5	15.0	15.5	16.1	17.4	39.2	51.7	864.0
Operating Profit	$16.9	$2.0	$14.3	$15.5	$19.5	$13.9	$22.9	$25.7	$49.8	$48.7	$81.9
Operating Margin	13.4%	1.8%	10.1%	10.3%	12.1%	9.3%	13.5%	14.0%	11.9%	9.2%	12.4%
Net Interest	$1.1	$1.1	$1.2	$1.0	$1.0	$1.0	$1.0	$1.0	$0.7	4.5	$4.0
Pretax Profit	$18.0	$3.1	$15.6	$16.5	$20.5	$14.9	$23.9	$26.7	$50.5	$53.2	$85.9
Pretax Margin	14.2%	2.8%	11.0%	11.0%	12.7%	10.0%	14.1%	14.5%	12.1%	10.0%	13.0%
Tax	$7.9	$1.2	$6.3	$6.9	$8.6	$6.3	$10.0	$11.2	$19.7	$22.3	$36.1
Tax Rate	44.0%	37.3%	40.6%	42.0%	42.0%	42.0%	42.0%	42.0%	39.0%	42.0%	42.0%
Net Income	$10.1	$2.0	$9.2	$9.6	$11.9	$8.6	$13.8	$15.5	$30.8	$30.8	$49.8
EPS	$0.24	$0.05	$0.23	$0.23	$0.28	$0.20	$0.33	$0.36	$0.76	$0.74	$1.18
Average Shares	41.8	41.8	41.0	41.0	42.0	42.2	42.4	42.6	40.8	41.4	42.3
As a % of Revenues											
Sales	85.9%	82.0%	83.9%	83.3%	83.2%	80.6%	81.7%	81.4%	86.1%	83.8%	81.7%
Service	14.2	18.0	16.1	16.7	16.8	19.4	18.3	18.6	13.9	16.2	18.3
Cost of Revenues	39.9	42.5	40.7	40.0	40.0	41.5	39.0	38.5	40.3	40.7	39.7
SG&A	38.1	44.2	39.7	40.0	38.6	38.8	38.0	38.0	38.4	40.3	38.3
R&D	8.6	11.6	9.5	9.7	9.3	10.4	9.5	9.5	9.4	9.8	9.7
Net Interest	0.9	1.0	0.9	0.7	0.6	0.7	0.6	0.5	0.2	0.8	0.6
Net Income	8.0	1.8	6.5	6.4	7.4	5.8	8.2	8.4	7.4	5.8	7.5
Year-Over-Year Growth											
Sales	33%	11%	26%	23%	23%	32%	16%	20%	32%	23%	22%
Service	45	46	45	53	51	45	36	36	47	47	41
Total Revenues	34	16	29	27	27	34	19	22	34	27	25
Cost of Revenues	33	25	28	26	27	31	14	18	54	28	22
SG&A	36	29	39	32%	29	18	17	16	25	33	19
R&D	21	31	36	39	38	21	19	20	17	32	24
Operating Profit	45	(81)	5	10	15	592	59	66	22	(2)	68
Pretax Profit	54	(69)	12	12	14	373	53	62	8	5	62
Net Income	41	(69)	10	9	18	337	50	62	3	0	62
EPS	36	(70)	10	9	18	333	45	56	0	(1)	58
Average Shares	4	3	0	(0)	0	1	3	4	4	2	2
Sequential Quarter-to-Quarter Growth											
Sales	6.9%	-15.5%	30.5%	5.0%	7.0%	-10.0%	15.0%	8.0%			
Service	9.5	11.8	14.2	9.4	8.0	7.4	6.9	10.0			
Revenues	7.2	(12.0)	27.6	5.7	7.2	(7.1)	13.4	8.4			
SG&A	5.8	1.9	17.6	3.8	3.3	(6.5)	13.9	5.6			
R&D	4.3	18.5	5.1	7.3	3.4	3.3	3.8	8.4			

Table 8

Tandem Computers, Annual Income Statements, 1980-1985E

(millions)

	1980	1981	1982	1983	1984E	1985E
Sales	$98.0	$186.9	$272.6	$360.1	$443.8	$542.1
Service	11.0	21.5	39.6	58.2	85.8	$121.1
Total Revenues	$109.0	$208.4	$312.1	$418.3	$529.5	$663.2
Cost of Revenues	$40.8	$75.5	$109.3	$168.7	$215.5	$263.1
SG&A	40.0	74.6	128.5	160.6	212.6	$254.2
R&D	8.8	17.8	33.6	39.2	51.7	$64.0
Operating Profit	$19.3	$40.4	$40.7	$49.8	$49.7	$81.9
Operating Margin	17.7%	19.4%	13.0%	11.9%	9.4%	12.4%
Net Interest	$1.8	$10.7	$6.0	$0.7	4.5	$4.0
Pretax Profit	$21.1	$51.1	$46.7	$50.5	$54.2	$85.9
Pretax Margin	19.3%	24.5%	15.0%	12.1%	10.2%	13.0%
Tax	$10.4	$24.5	$16.9	$19.7	$22.8	$36.1
Tax Rate	49.3%	48.0%	36.1%	39.0%	42.0%	42.0%
Net Income	$10.7	$26.5	$29.9	$30.8	$31.4	$49.8
EPS	$0.35	$0.72	$0.76	$0.76	$0.76	$1.18
Average Shares	30.2	37.0	39.2	40.8	41.4	42.3
As a Percent of Revenues						
Sales	89.9%	89.7%	87.3%	86.1%	83.8%	81.7%
Service	10.1	10.3	12.7	13.9	16.2	18.3
Cost of Revenues	37.5	36.3	35.0	40.3	40.7	39.7
SG&A	36.7	35.8	41.2	38.4	40.2	38.3
R&D	8.1	8.6	10.8	9.4	9.8	9.7
Net Interest	1.6	5.1	1.9	0.2	0.8	0.6
Net Income	9.8	12.7	9.6	7.4	5.9	7.5
Year-Over-Year Growth						
Sales	98%	91%	46%	32%	23%	22%
Service	96	96	84	47	47	41
Total Revenues	95	91	50	34	27	25
Cost of Revenues	98	85	45	54	28	22
SG&A	92	86	72	25	32	20
R&D	89	103	89	17	32	24
Operating Profit	99	109	1	22	(0)	65
Pretax Profit	109	142	(9)	8	7	59
Net Income	117	148	12	3	2	59
EPS	77	103	6	(0)	(0)	55
Average Shares	21	22	6	4	2	2

growth this quarter and through next year to remain very modest by historical standards -- 2%-3% sequentially versus the 6% average for the prior four quarters. In fact, the company has a goal of zero headcount growth in the September and December quarters, a goal we think will be difficult, but achievable.

More importantly, Tandem is now budgeting (for the first time, we believe) expenses for reasonably conservative revenue growth of 25%-30% in fiscal 1985. In addition, many other discretionary expenses are being curbed with the clear objective of having operating margins in the 16%-18% range four quarters from now. While we have certainly not jumped on the 16% margin bandwagon, we are optimistic that Tandem will see some meaningful expansion next year. Specifically, we beieve gross margins will improve 1.0-1.5 percentage points, to 61% or more, by late next year, the chief positive influence being a further strengthening of Tandem's product cycle. Near term, we expect only a modest improvement, to 60%, resulting from an expected smaller number of Non-Stop II upgrade sales (which hurt June margins).

Given reasonable progress in slowing headcount growth, SG&A and R&D spending increases should moderate substantially. We expect that in the September quarter, Tandem will spend modestly more in both categories than in the June quarter, then maintain those levels (as a percent of revenues) through most of fiscal 1985.

As shown in Table 8, Tandem's gross margins have dropped considerably since fiscal 1982. This is consistent with the gross margin pressure seen by many other companies, a result of both aggressive IBM pricing tactics and, to a lesser degree, competition from startup vendors who offer better hardware price/performance. It is important to note that we do not expect Tandem's operating margins to improve by dramatically reversing the gross margin pressure of the past two years. Most of the margin gain should come from tighter discretionary spending controls and slower employment growth, evidenced by lower SG&A as a percent of revenues.

Improving Accounts Receivable and Inventory Turnover

In analyzing Tandem's June balance sheet (see Table 9), we note two important factors. First, its receivable turnover of 86 days is better than any in the past six years, and improved by 13 days from March's level. In part, this reflects a smoother shipment schedule throughout the quarter, courtesy of some revenue spillover from March. Also, we believe that the high turnover indicates that Tandem did not push an inordinate amount of product out the door during the last week in June. Second, Tandem's finished goods inventory was down nearly $6 million (12%) from March, and its inventory turnover improved to 1.96 times, indicating that Tandem was not totally surprised by June's revenue levels. Finished goods inventory was down sharply from March levels, and importantly, raw materials were up, indicating that production plans remain at a healthy level.

Table 9

Balance Sheet and Cash Flow Data, 1981-1984

Receivable Turnover (a)

Year	Mar	Jun	Sept	Dec
1981	110	117	99	112
1982	112	104	102	110
1983	104	112	91	96
1984	99	86		

Inventory Turnover (b)

Year	Mar	Jun	Sept	Dec
1981	1.60	1.37	1.37	1.20
1982	1.14	1.06	1.08	1.20
1983	1.38	1.72	1.96	1.94
1984	1.76	1.96		

Operating Cash Flow (c)

(millions)

Year	Mar	Jun	Sept	Dec
1981	($13.4)	($21.0)	$10.0	($32.4)
1982	(32.3)	(10.7)	(26.3)	(24.6)
1983	10.7	17.3	24.5	(10.0)
1984	11.3	(7.7)		

(a) Quarter-end receivables divided by quarterly revenues times 90 days.
(b) Trailing 12-month cost of goods sold divided by quarter-end inventories.
(c) Excludes external sources of capital.

Tandem's operating cash flow in the quarter was a modest negative $7.7 million. We view its cash balance of $86 million at June 30 as more than sufficient to fund growth for at least the next one to two years, depending on the company's growth rate.

John C. Levinson
(212) 902-6775

SHEARSON
AMERICAN EXPRESS

Tech Trends Seminar

Computer Systems
Fault Tolerant Alternatives

Donald H. Brown
December 6, 1983

An American Express Company

Table of Contents

Introduction & Summary

John Levy

Design innovations to capitalize on technological advances have been applied to fault-tolerant, high reliability systems– a product class formerly served only by TANDEM. Some of these efforts have been conducted by small start-up firms such as Auragen and Synapse, as discussed below, while Stratus has already passed the organizational hurdles of becoming an exciting, publicly traded, small company.

John Levy will lead our panel on fail-safe, transaction systems. He worked at Digital Equipment, where he did the backplane for the VAX/780, and he later designed for Tandem. Subsequently, he was the sixth member of the LISA development team at Apple before he cashed in his stock options and started his own company, Computer Structures. He knows something about systems from a designer's standpoint.

To a computer architect, the three companies on this panel present a lovely variety of hardware approaches. In fact, I don't think three more different examples for case studies could be found. It is a real pleasure to see them. It is this type of innovation that will break the performance bottlenecks in high reliability systems.

As an introduction to high reliability systems, there are three pure strategies for dealing with failures in computer systems. A pure strategy is one which has a simple, consistent rule. One pure strategy for dealing with a failure is to ignore it. You hope that the failure causes the machine to stop, and you start over. This is the usual strategy for home computers. Of course, if the problem was caused by a software error, the same thing will probably happen again.

A second pure strategy is to protect the circuits from failures, either by finding components that will not fail, or by building combinations of circuits that won't fail as a unit. This is the sort of strategy used by NASA in things that go up in space. There isn't much chance to repair it once it's up there.

The third pure strategy is to defend against the consequences of a failure–to be immune.

No one of these strategies is, by itself, good enough for real practical systems. That's the way it is with most pure strategies. The first one–ignore the failure–is not adequate because failures do occur, and not all application programs can be started over. The second–by using high-reliability components–is susceptible to unforeseen failures. There is an old saying among programmers who write diagnostic test programs that "diagnostics are good at finding known failures." It is the unforeseen ones that are hard to handle. The third strategy–defend against the consequences–is a good idea, but not enough by itself. Programs and circuits, which are added to a system to make it immune, can themselves fail and add to the failure rate.

Beyond these strategies, processors and memories aren't the only parts of a computer that can run amuck. Lots of other things can go wrong, some of them beyond the control of the system designer. A disk built by outside suppliers must still be supported by the systems manufacturer. Telephone lines introduce static as the data goes over the line. Storage media can fail at my house because of my two year old spreading peanut butter on a diskette. Also, not to be ignored is operator error. The "oops phenomenon" implies that "it takes a human to really screw things up." An otherwise immune system can be brought down by a person doing the wrong thing at the operator console.

Finally, a couple of comments about transaction processing, since that seems to be a major area in which people want high availability. Transaction processing is a matter of doing relatively small things over and over. Quite often the demand, such as in an airline reservation system, is to do lots of transactions at the same time. The challenge is to get high throughput, a pipeline in which lots of things are going at once. In the pipeline, none of the tasks are complete, while still more are pouring in through the communications links. Response time is not the big issue; the challenge is to complete tasks in the system, to cram transactions through the pipe and still keep the system going.

Transaction processing is Input/Output intensive, moving bits back and forth between a disk or some other secondary storage medium and the computer. Little actual computing is needed in shuffling the data. The system has to spin those disks and move those read/write heads. Transaction processing requires a well thought-out data base access scheme. Data in a large storage disk farm must be well organized to be found efficiently. Transaction processing systems are not built by taking a simple data base system and an average applications programmer and putting them in a room and saying, "Make me a transaction processing system." It doesn't work that way. In particular, the task requires rollback recovery mechanisms, to insure that transactions are complete and to maintain the integrity of the data base.

Failure rates can be traded off with time to recover. Assuming failures, the system will have to recover. So there is a choice. The designer either makes the failure rates as low as possible, and then is willing to spend a fair amount of time recovering on the rare occasions that the failure occurs. Or, the designer creates a quick recovery while accepting a higher failure rate.

Differences in Design Approaches

I have been trying to sum up for myself the similarities and differences between these systems. Stratus basically has excellent detection and has what we would call on the

technical side, "hot backup". In other words, there is absolutely no delay in implementing the recovery from a hardware failure. That is expensive in hardware, on the face of it, with duplication. But, it achieves virtually no hardware failures seen by the software as a result. A neutral comment is that of course Stratus still has to deal with all the other recovery problems that any software or communications failure would create.

Synapse has good detection, and backup using module substitution. Synapse achieves a high utilization of hardware because of common memory, and because all processors are running at once without dedicating a processor to "backup." Synapse has to deliver a lot of software in order to make the applications easy. Synapse does have the detection problems, which I have not reviewed, and it will have to deal with restarts upon hardware failure.

Auragen may have difficulty with the detection of errors, and is willing to do lots of reprocessing in case of failure while counting on a low failure rate. Auragen has a high utilization of hardware, by not dedicating much processing time on a "backup." The system requires the company to create lots of software, with the objective of making application software "easy" for the user, partially by incorporating UNIX.

Finally, the architecture isn't what is going to sell these systems. The massive amount of software required to make something useful is what is so apparent when you look at these three vendors. To me, it is hard to imagine that the success of any one of these vendors would depend solely on the hardware design. Obviously, it has to be done well enough and it has to work well.But, that is not what is going to make it a big success.

Users are now smart enough to demand that the software be structured in a way that parallels the actual work. I call it a maturing of the marketplace. It will require lots of processing, lots of processors, and lots of software that works well and which the user never sees. All these parts have to be adequately designed, implemented and maintained–hidden away down there in the "automatic transmission," while the user uses the shift lever. I think that is why the market for computers can continue to expand.

Stratus Computer

Bob Freiburghouse

My responsibility at Stratus is Vice President of Software Engineering. I was one of the three founders of Stratus. I have a long career in the computer industry going back to the early '60's. Much of my time was spent at Honeywell and General Electric in various technical and management positions. In 1974, I left Honeywell to form my own company, a small software company which supplied compilers principally to the minicomputer industry. In late 1979, I got together with Bill Foster, President of Stratus, and we put together a business plan and launched the company.

I will give you a very brief history of Stratus and then concentrate most of my slides and most of my time explaining how we achieve fault tolerance in our system. We are unique in our approach, and the concepts are both interesting and simple. Non-technical people can easily understand how we do it.

This is the history of Stratus. We began in May 1980 with our first round of financing. We launched the company, at a time when venture money was scarcer than it is today, with $1.7 million. Throughout the development phase, we stayed on schedule and within budget. We announced a product in

Summary Product Comparison

Dimensions	Auragen	Synapse	Stratus	Tandem
FCS (First Customer Shipment)	10/83	8/83	1982; roughly 120 shipped 11/83	1976
Minimum System Price	$138,000	$350,000+	$150,000 plus, typically $200,000	$150,000+
Target Market	Low-end and mid-range OEM.	High-end, performance End-user initial orientation	Low-end to mid-range	Mid-range
Approach HW	Clusters of modular processors and memory.	Shared memory; distributed cache.	"Quadded" redundancy & fault detection by comparators	Duplication, Medium power mini's.
Approach SW	"Environment preservation" Backup processes in virtual memory; Message system	High-throughput task queuing/ integral database and transaction processing functions.	Straightforward, assuming no hardware faults.	Sophisticated " backup process" method; Message system.
Emphasis in Sales/Positioning	"IBM style": sell the functions on bottomline return; communications capabilities; UNIX	"Mainframe" market; "transaction/sec." or transaction/ dollar.	Fault tolerance, service, Full range software. No emphasis on "how we do it" at low level.	Fault tolerance, modular expansion, service, distributed systems & communi-cations capabilities.
Major Risks	Complex product development task for backup scheme and incorporating fault tolerance in the UNIX kernel. Executive processor must detect and manage for significant cluster failures.	Large software development task! May have underestimated continuing support and development costs? May encounter IBM response?	Will have to push aggressively to establish stable market share.	Installed base could excessively influence conversion to a new generation; large organization not quick to respond to challenges.
Special Strengths	Can substitute hardware modules at will with evolution of technology; may get better performance per dollar than Tandem. Nixdorf agreement.	Good technical staff, seem to understand market requirements.	High software expertise in corp. Easiest to deal with "reliability" features- because they are low level. Technically led management.	Good (and field proven) fault detection; Good sense of user needs; large range of tools needed by knowledgeable user (or OEM). Still strong technical organization.
Comments	Some checkpointing still required System-defined default values may not optimize operation for some particular installations. They are definable at time of installation.	Not in the same price class! Still need European distribution.	Technically led management. Olivetti agreement.	Still the one to beat; Recovering from p.r. deficit due to sales curve, etc.

November 1981; we shipped in February 1982. In 1982, which was not a full year of shipments, we did $5.5 million in revenue. As of October, we have shipped 120 systems to about 40 customers. So we have been shipping systems for seven quarters and have been profitable for the last four.

At announcement, our system was a very complete product from the standpoint of how much software was available. The popular programming languages in software packages were announced at the time the system was originally introduced, and subsequently released. Initial software enhancements included transaction processing software, screen forms management software, communications software – most importantly, support for IBM 3270 terminals. Fortran and Pascal languages bring us to a total of five major programming languages. More recent hardware and software products to date include SDLC communications support, additional disks, and remote service networking.

The list of software is extensive – our proprietary operating system, transaction processing facilities, lots of communications capability, our own global networking capability, (StrataNET), 3270 terminal support, X.25 and X.29. The list goes on. While most people focus on our unique hardware approach to fault tolerance, the software represents a major accomplishment. We are software oriented. Most of our R&D money is spent on software.

People ask us who uses our machines. The financial community is a major part – 29%; on-line services account for 24% of our installed base. But, there are a wide variety of other areas not normally associated with buying fault-tolerant computers. We think this is the beginning of a trend. As people realize they can buy a fault-tolerant computer which can be programmed as easily as any other computer and at no real penalty in price or performance, they will demand fault tolerance. It just hasn't been available before.

Stratus is unique in the way in which we achieve fault tolerance. We set out from the beginning to build a computer that was programmable in the same way as everybody else's computer was programmable. You do not have to write programs for Stratus using any specialized rules or any transaction processing package. You can run ordinary programs on our machine, and they run in a fault-tolerant mode. We achieve that by using redundant hardware – hardware that is redundant at a very low level. From the buyer's perspective, continuous processing entails no extra software overhead and, most important, no specialized application development. Existing application codes from existing computers convert to ours as easily as moving to anyone else's machine, while still obtaining the benefits of continuous processing. And the user can do all that at a price/performance level that is competitive.

Two fundamental concepts contribute to fault-tolerant operation. The central processor is one printed circuit board, about 15 inches by 18 inches. Every field replaceable component in the system self-checks its own failures. How do we do it? Simply by duplicating the logic with two sets of processor logic on each processor board. For example, the addition of two and three, producing five, is done in parallel by separate sets of circuitry. The output from each set is then compared with the output from the other set. If they are equal, the output is accepted. If, for some reason, the output were not equal, it would not be allowed to leave the board. If that were all, we would simply have a computer that came to a screeching stop with an error. But, the boards themselves are duplicated in turn, with each board self-checking with two sets of logic. That makes four sets of logic which perform the operation simultaneously on separate field replaceable boards. The left half of each board compares results with the right half of the board; separate boards compare the results in turn.

There are two principles involved here. The boards, or any field replaceable component, are self-checking. They detect their own failures and do not output bad data under any circumstances. The second principle achieves continuous operation by having duplicating components and building blocks, in this case processor boards. When there is failure on one board, the output is stopped from that board and shifts to the alternate board.

Considering the entire computer system, all the components – memory disk controllers, communications controllers, link controllers – are self-checking and in almost all cases redundant in the same manner as the processor boards. It is not sufficient to simply have extra processors because everything can break. The controllers and the devices are redundant, and self-checking. The bus that serves as the interconnection or data path between all these components is also redundant and self-checking. So is the power supply, as are the fans that cool the system. So, essentially, everything inside of the box is redundant and, most important, self-checking. Self-checking means that the failure is detected at such a low level such a primitive level in the hardware, that the failure is invisible to the software. That is the key. It is invisible for the most part to the operating system software, and it is absolutely invisible to the customer, to the application program. That is why you can run ordinary programs on a Stratus computer.

Basic units can be expanded in local area networks. Up to 32 processing modules can be linked together to operate as one computer system. The modules can be physically placed all over an office part or an office building. Beyond local area networking, our proprietary StrataNET networking software is based on the X.25 industry standard.

A local area network, anywhere in the U.S.A., with a bunch of processing modules hooked together and talking to each other, can communicate over a packet network to other clusters. They can also communicate by more primitive means, X.25 or 3270, to other computer systems, in particular IBM systems. The systems are very simple, with few controls to minimize the possibilty of operator errors. There are almost no controls. There is no operator console.

One of the benefits of the design is that customers write ordinary programs for the system. They can bring existing programs from other computers and transfer them to Stratus.

One of the other unique benefits about Stratus is that the machine continues to be fully functional even though it has one or more broken components. Additional capabilities make the system serviceable from a central location. For example, the system detects a fault, where the comparator on a board discovers a discrepancy. It shuts the board off. The board does not communicate with the system any more; no data goes on the bus. But an interrupt signal is sent. Operating system software picks up the interrupt signal and says "oops, there was a component that just went out." Now the job continues to process perfectly normally, but the software then exercises the board, running diagnostics. The computer system is still functioning fully from the user's perspective at the same speed. As diagnostics are run, the unit is automatically placed back into service and synchronized with its partner if the unit appears to check out normally. The fault is recorded in the log as a transient failure that would have brought down a normal computer system and resulted in a service call.

If the component proves to be defective by not passing the diagnostic test, then it is a true failure. The unit is removed from service, a red light appears on the outside of the machine, and a red light appears on the board. Most importantly, the computer now dials Stratus in Natick, Massachusetts, and reports its own failure automatically. It says, "I have a failure of a CPU in slot number six." The system transmits the CPU serial number and the repair and fault history automatically. The broken computer has called for its own service and reported which part needs to be replaced. In most cases, the customer is not even aware of the failure, although he can monitor the activity. He can just put the computer in a room somewhere and forget about it.

This is probably the first computer system that can truly be serviced from a central location in this way. In spite of failure, the system is still fully functional. We go beyond that to provide – in the service center – responses to customer questions. A customer who wants to order manuals or extra parts calls the same number. One central location services the customer, monitors the status of the customer's system, and application or system software problems.

It is easy for us to separate hardware and software problems. Hardware problems result in little red lights. Everything else is software. One of the great difficulties in the industry typically has been fixing the blame for what is wrong. Diagnosing the problem is difficult. Repairing the problem is usually the easiest. Stratus has gone a long way toward simplifying the most difficult part of the problem.

We have a large enough customer base now to establish real data. Our customers are very happy with the system. Customers love the central service, and even overuse it. We get enormous call lists because people find it so convenient to call up and say, "How do I do this, how do I do that? What does this mean?" The mechanism really works. Our service costs are way below the industry average. The costs of servicing these machines have been addressed in the design of the computer.

Q: What about recovery mechanisms?

A: Well, first of all, we have to accept that faults will occur. Unless they are detected at a low enough level, the result of that fault will be incorrectly processed data, and recovery then becomes a crucial aspect of the problem. The problem then becomes one of defining backup for a restart. That whole issue is eliminated by our approach. The system detects the failure so early that it can't have any consequence and redundant components produce the right answer at the right time. As a result, our own software people do not have to worry about recovery, neither does the user. Our competitors provide transaction processing software which requires programming according to the rule of the TP software package to obtain recovery and restarting capability. But that is a different kettle of fish than being able to run any old program you choose in a fault-tolerant manner. Lots and lots of applications don't fit very nicely into the TP mold which are difficult to write. We have lots of customers who have taken batch jobs off IBM mainframes and programs off Wang computers, all sorts of things that were converted in a straight-forward manner and run on Stratus equipment. The painless conversion is only possible if this type of approach is used – that is, detect the fault before it can have any consequences and therefore eliminate the need to recover.

Q: How fast are the buses?

A: The connection between the modules is 1.4 megabytes/second. There are actually two connections and, because of the protocols that are used on them, the full band with the both connections can be used, so it is 2.8 megabytes/second.

Q: Do you have statistics showing what the incidence of failure is with redundant components and self-checking service?

A: The system is designed to withstand a single failure. However, the nature of this design permits lots of combinations of multiple failures. But if you have a broken CPU, you can't take a failure on your other CPU. You could take a failure on memory or disk or whatever. How many times it actually went down because there were two failures? To date, I only know of one instance when both disks went out at the same time.

Q: What about communications?

A: The X.25 comm can run at 56KB/second. We don't have very many X.25 networks; there are perhaps two or three. The great majority of the communications is to IBM mainframes in which we look like 3270 controllers. IBM is the world in which we all live, in particular our machines. We support IBM 3270 terminals completely. In fact, our customers, if they want, can just use 3270 terminals.

Synapse Computer Corporation

Mark Leslie

My name is Mark Leslie, from Synapse Computer Corporation. I am one of the two principal founders of Synapse, which was started in early 1979. We got our first funding in November 1980, also in the days when there was a little less venture capital than there is today. We have opened sales offices around the country, with eleven opened by July. Our first customer shipment was December 30, 1982.

Synapse focuses on high performance, flexible expansion, and simple application implementation. We have built much of what we call solutions-oriented machines for transaction processing applications. We are simplifying the implementation complexity of transaction systems. We are at the beginning of a trend. Over the next five to ten years, fault tolerance will be required in all computer systems in all segments of the marketplace. In 1983, high-technology computer systems don't provide even the same reliability and MTBF (mean time between failure) as the automobile. If a car broke down an hour a week, you would trade it in for a new one. I think there will be a shift in this industry to a high degree of fault tolerance of various different kinds.

Traditional computing systems started out doing batch work which converted business transactions from paper to machine-readable form, which ran them against master files and generated reports. Most of today's development work by users concentrates on on-line access to data, to enter information at the point of origin and to avoid intermediate steps of going from paper to cards or tape. Transactions occur while a person standing in front of a terminal updates the data base in real time. The 1960's represented the generation of batch processing. The high cost of computers created an artificial environment. The price/performance gains of the 1970's enabled us to put computers closer to people and introduced the generation of distributed processing. The next step, the 1980's, represents on-line transaction processing which brings the interaction with the computer closer to the event changing the database in time. It is a very fundamental issue. At the time the person walks away from the computer, the computer has completed the transaction. Going on-line constitutes as fundamental a move as distributed processing. On-line equals the significance in terms of time that distributed processing represented in terms of location.

Significant artifical barriers remain in cost and risk. Needs of the transaction processing user are significant – fault tolerance and application availability, integrity, throughput, expandability, and flexibility for changes.

Transaction processing includes a wide range of users who tend to be in businesses willing to take more risk and pay more money to go on-line. For instance, a teller machine application enables a bank to be a leader and gain market share, or to protect market share as a follower. In both cases, banks spend more money and take more risk because the application is in the mainstream of their business. As we drive the risk and cost of transaction processing down, the market widens to include more applications in an on-line fashion. The American Bankers' Association estimates that the number of electronic transactions will rise from 1.8 billion now to 16 billion in 1986 – a 55% growth rate. Other studies suggest that the transaction processing market will grow to $20 billion in 1985.

Several traditions in computer architecture have developed to meet the needs of transaction processing. One is the compatible-family architecture from IBM and DEC. A family of similarly designed computer systems requires a new system when the user runs out of capacity. The replacement process generally encounters problems in compatibility; upgrades are expensive and disruptive. These computers are generally not fault tolerant. On the positive side, the systems offer the advantage of a centralized data base accessible by all terminals. That single advantage explains why the IBM world doesn't upgrade networks of smaller 4300's, but turns to large scale 3033's or 3081's.

Multiple-computer architectures were popularized in the 1970's by Tandem, Wang and Datapoint. The philosophy concentrates on modules with a certain processing power, terminal capability and data capacity. Cloning the modules expands capacity through an external bus. This architecture is outstanding for work station activity, when the terminal primarily processes data locally. Also, upgrades are less disruptive. Examples are graphic work stations, document preparation and word processing. The use of the system by

other terminals is casual and incidental. In fact, a graphics or engineering work station still functions as a stand-alone unit without considering other work stations in the network. On the other hand, the approach loses efficiency in transaction processing, because of data distribution. For example, an inquiry on a Mr. Jones may reside in a data base outside the boundaries of an individual computer system. The request searches for where the data might be, acquires the data, transmits it across a bus, operates on it and finally sends it back to its original location for safe storage. As the number of systems in a network increases, the probability that the data resides someplace else goes to 50% with two processors in a network and to 67% in the case of three systems.

These multiple computer architectures may be fault tolerent. Upgrades are less disruptive. On the negative side, the approach requires a centralized data base capability, constant tuning and load balancing. Interprocessor overheads are high. Control constitutes a significant issue.

Synapse has a concept of an expansion architecture, which allows us to grow from a small system to a large system in an orderly and extremely modular fashion. When we do grow to a larger system, all terminals have access to the computer power and access to all the data and the entire data bases. That is truly a unique capability in the industry. Shared main memory and central control provide for upgrades that are small and non-disruptive. The system is self-tuning and includes fault tolerance.

An elegant and pure analogy simply and intuitively illustrates the difference between a multiple system environment for transaction processing and a Synapse architecture. Consider a bank lobby with two people in line. I am from New York originally, an aggressive "Type A" personality, and I always get into the shortest line. Unfortunately, the guy in front of me is in the vending machine business from Atari, and he's got two suitcases full of quarters. We're both waiting to make our deposits. Although I thought I was on the shortest line, I am actually in the longest line. That has probably never happened to you, but it happens to me all the time. This problem, called multiple queue/multiple server, illustrates the basic architectural weakness of existing multiple-computer approaches.

In our analogy, the teller is a processor, and each person in line is a terminal connected to that processor. It is not uncommon, in a multiple-computer system, to have one processor swamped with activity, while another is idle.

Synapse architecture allows for a single line of patrons in your bank lobby. Each teller says, "Next, please." As a result, everybody gets more even-handed access to the system and a better utilization of the basic resources.

The analogy also carries into the distribution of the data base. Assume, in the non-Synapse system, that each of four tellers has a 25% chance of having my bank records. That means that, three out of four times, he has to leave his own system to find the records, go get the records, make the transaction and then return to his computer system. The more computer systems there are, the more likely that the data resides elsewhere. The more processors required, the more the data is divided and the more inefficient it gets. The Synapse system has a single set of files; all the processors have access to them with no interprocessor communication.

Consider modular expansion, as it is illustrated by the analogy: When we expand our system, we don't have to be concerned with which data to put where, which people to put where, and how to structure our terminals (called load-balancing and tuning). We basically make our file system bigger, add tellers to increase service and make our line longer. We do not have to redivide the database. With multiple-computer architectures, database efficiency erodes as the number of fault-tolerant processors rises. Each of the multiple processors is backed up by its adjacent neighbor, and can therefore be configured at only half its capacity (if you want to assume 100% throughput failure). For a two-processor job, Synapse simply requires a third processor, which is available to increase transaction performance in a healthy system. That is the concept of N + 1 for fault tolerance. Thus, Synapse would require eleven processors to do a ten-processor job, rather than twenty or thirty highly-tuned, balanced, loosely-coupled processors required in a multiple-computer system.

As a summary, there is no interprocessor communications activity. Excess capacity is equal to N + 1. There is no distribution of files, and the growth is automated within the architecture. The internal expansion bus is 64-megabytes per second (about the same performance characteristics as an IBM 3081). A multiple computer architecture includes a disk file, a CPU, a memory and some terminals attached at each of the computer nodes – those who need more specific power must order another system. In Synapse, I/O subsystems, processors and memories are separate components. Those who use more specific capability order only the additional component. For example, put in another I/O controller with another disk. Within the boundary of one system, users have the flexibility of buying the specific capacity required, configured exactly to their needs. Some users will want many IOPs (Input/Output processors) with many terminals; others will want many GPPs (general purpose processors) to grind numbers; a third group may want huge disk storage. The system is uniquely adaptable. We can choose which we want to add, and we can build a very balanced system, an extremely adaptable system.

And, of course, a single processor can be externally bussed to be part of a larger network. Synapse provides a unique performance range from the class of a very large supermini (VAX/780) up to the largest mainframe (IBM 3081). We can increment that power on-line, without interrupting the application by sliding a processor in while the system is running.

In software, a comprehensive and sophisticated set of tools automates on-line transaction processing. In 1974, DEC, Data General and HP provided a bag full of tools to build a hot-backup minicomputer. Very smart end-users assumed responsibility for the success of the systems. Tandem standardized the product. The essence of commercial innovation standardizes what everybody is currently doing in a non-standard way. In on-line transaction processing today, users buy a bag of tools and take responsibility for a wide range of sophisticated and intellectually demanding challenges. We at Synapse have "productized" those tasks.

The issues of data integrity and recovery and fault tolerance are all transparent to the application designer and implementer. A broad menu of software includes all the tools required. The last time we took a rough count of our software, we had about a million lines of code implemented on Synapse software. The operating system and the relational DBMS (data base management system) provides unique functionality. On top of the relational DBMS is the next layer of services. On top of that is the transaction processing manager. On top of that is the user. So he interfaces and layers to a system that basically buffers him from all these issues.

Beta testing was completed at the end of June. One installed customer is in communications, another one is a division of a Fortune 500 company with government OEM business. We have recently installed a system at Bank of Ohio in Columbus, Ohio, for a centralized batch system with satellites for transaction processing. We are very delighted to penetrate an account with IBM and Tandem already in place.

In summary, we believe that on-line transaction processing is the fastest growing segment of the data processing market. Next to personal computers, it represents the fastest growing segment in the high technology business today. We have brought to the market both an innovative hardware and innovative software architecture. We have "productized" the approach, and made a contribution in commercial innovation to on-line transaction processing. We address low-cost implementation for very high performance machines. High barriers to entry in this marketplace enhance long-term visibility. We have an experienced management team and strong financial backers.

Q: *With beta tests completed, what do the benchmarks show about performance?*

A: Although we don't have complete information yet because it is still quite early, with multiprocessing, the trick is to put processors in place without having memory contention in the connecting buses. Measuring bus contention and utilization, a ten processor system running all of our processors at 100% flat out, and shutting off one of our two buses to simulate twice the number of processors, we obtained a linear increment in performance for every additional processor. Ten processors still used only 20% of the capacity of one bus. Major technological advances will enable multiprocessors to break current performance bottlenecks. For comparison, the IBM 308x product line obtains 1.5 times the performance of one processor from a dual system, 1.9 times for a triple, 2.1 for a quad and back down to 1.9 times the performance of one processor when five multiprocessors are linked together. At five, the system crosses the edge of the curve. Synapse went to ten, and then shut off one of the buses, and still obtained ten times the performance characteristics on the bus. The result suggests that the system can expand much further without suffering performance degradation.

Q: *What about contention for disk accessing?*

A: All tasks in the system reside in a single list in a main shared memory that every processor accesses equally. A disk access occurs independent of any other processor. When a disk device completes processing at the I/O level, it re-enters the single list and may not be picked up by the original processor on the next access. It may be picked up by a different processor.

Q: *What about fault tolerance and recoverability?*

A: Throughout, the system has all the traditional fault-detection mechanisms. In addition to ECC memories, fault detection is distributed throughout the system. When a fault is detected, a board is taken out of service. If there has been an impact on the data base, software rolls back to a consistent state and rolls forward from there. Our worst case for recovery is probably minutes and our best case is very few seconds. Our worst case of data base recovery should be contrasted with the outright loss of a data base entailing multiple hours, or perhaps days, of effort in other systems. In this context, Synapse has built the ultimate in recoverability.

Auragen Systems

Rick Martin

My name is Rick Martin. I am president of Auragen Systems. My previous experience was all with IBM as the 370/138-148 product line manager back in the mid-70's, and the 4331 product manager just prior to joining Auragen in late 1981. Our doors opened for business in April 1981. A

prototype of our system ran at the NCC in May, 1982. Our first customer shipment, a beta site, was October 11th to Nixdorf; the first U.S. beta site received their system in November. Our original round of financing totaled $2.3 million back in 1981. A $3 million R&D partnership closed in late 1982, and $10 million was added in June of this year, bringing the total to $15 million. A major OEM contract has infused roughly $1.5 million in the form of advanced royalties. That is primarily a marketing partnership targeting Europe.

Everyone agrees the marketplace for fault tolerance is very large – no argument. One primary inhibitor to fulfilling the potential growth is the limited number of vendors in the market. One study shows that the transaction market portion of the fault tolerant market in 1986 is projected to be about $17 billion of the total $22 billion market. Perhaps only $3 billion will be realized because there simply will not be enough vendors. So there is plenty of room out there for competitors, including the three of us today as well as others. Our marketing philosophy focuses on banking and securities. We also focus on communications – both AT&T and ITT systems, as well as communication-based systems, and the Federal government.

In distribution, the OEM value-added re-marketers will represent 30-40% of our business.

Today, I would like to discuss our distributed architecture, fault-tolerance price performance, and programmer productivity. Our distributed architecture makes our system flexible, very high performance and fault tolerant. An executive processor, based on a Motorola 68000, has 128K onboard a RAM. It provides the fault tolerance automatically. There is absolutely nothing required on the part of the programmer to achieve fault tolerance in the system, which is accomplished by a separate processor. Another board, the work processor, does customer's applications. It is two 68010's running in parallel. This is not the same concept as Stratus where they actually run the same job and check each other; ours runs two different jobs. Associated with the work processor is memory, from one to eight megabytes which is demand paged. Communication processors are also 68000 based with memory of 128K. They handle terminal lines and protocols. The disk tape processor, a very high-speed proprietary design at three megabytes per second, is not based on the 68000. Frankly, the 68000 today is not fast enough, requiring a custom-designed board. Overall, our system includes very specialized boards that have specialized hardware as well as specialized software. Tied together for fault tolerance, at least two clusters share peripherals, disks, tapes, printers, terminals, communication lines, and so forth. We can interconnect across what we call the system buses, up to 32 processing clusters. If someone configured the total machine, it would have almost 200 Motorola 68000's in a system. It is very, very modular. The pieces are very small and you can increment them as you need to. If you need more processing power, you add processors. If you need communication power, you add communication processors. If you need more data base, you add more tape/disk processors.

There are three advantages to this approach:

1. *Flexibility*, which comes in several pieces, the first of which is technology. Each of those boards in fact is fully independent from each other. They only talk over the buses. The disk/tape and the work processors are changing now based on new technology. By mid 1984, 68020-based work processors will be available. Tremendous flexibility. All the changes are independent of each other. The same flexibility is available for the customer to configure exactly the system needed for his application. For example, one customer system will have five or six work processors and attach up to 5,000 terminals. So, it is very heavily biased toward the communication side with about 50 communication processors.

2. Because of the distribution of tasks across the various processors, we can *functionally specialize* in the processors. The fault-tolerant work in the executive processor does not interfere with the application programmer, or communications, or data base. For performance reasons, code can be shifted from one processor to another to optimize throughput for various types of systems. For the particular customer with a tremendous number of terminals, code is moved out of the executive processor to the communications processor. What we can do for price performance is to focus the resources on the customer's needs.

3. The other advantage from a customer point of view is *modular growth*. Inexpensive, single boards provide for expansion of the system. Adding additional work, and executive processors with a megabyte of memory, costs about $30,000. Communication processors cost $2,000-3,000. Tremendous growth capability allows for over 200 Motorola 68000's in a fully configured system.

Let's talk about the different approaches to fault tolerance, what is common and different among all of the fault-tolerant people. Any transaction system consists of three pieces: 1) A front-end process which handles terminals and collects data from the users, 2) a back-end piece that deals with the data base or file manipulation, and 3) in the middle, an application written by a customer in COBOL, PASCAL, or whatever, runs. These tasks define transaction processing regardless of fault tolerance.

December 12, 1983

Computer Systems
Donald H. Brown

EAGLE CORPORATION (OTC-EGLC)+

12/6/83 Price	52-Week Range	Earnings Per Share* 1983	1984E	1985E	P/E Ratio 1984E	1985E	Dividend	Yield
$8 3/4	$25-7	$0.02	$0.25	$0.65	35x	13.5x	---	---

*Fiscal year ends June. There are 14.6 million common shares outstanding, of which 764,900 traded in October.

+Shearson/American Express maintains an investment banking relationship with the company and has underwritten public offerings of securities of this firm within three years. Shearson/American Express makes a market in this issue.

Investment Opinion

One of the deep fears of any investor in small high tech stocks relates to the thinness of management normally running the operations. When the company president dies in a tragic and untimely death, which was the case with Eagle, nerves start to become a little frayed. Times get even tougher when the industry subsector shortly thereafter enters a consolidation phase with major players - such as Computer Devices, Osborne, and Victor - falling onto hard times, and others such as Fortune Systems and Vector Graphic getting new management. With this as a backdrop, Eagle has managed to roll right along. In fact, the management under Ron Mickwee has increasingly come out fighting -both at the October financial industry conference on technology in Monterey and then again at the November tech industry Comdex show in Las Vegas.

Eagle has recently accomplished two key objectives. First , the company has revamped the product line to include full compatibility to the IBM PC/DOS system. Although the company's PC product had previously come "close" in the compatibility race, close isn't really good enough; the shortcoming is now fixed in the new releases. Further, the company has come in with three new product versions - the Spirit series of " transportables" at 33 pounds with a 10 megabyte hard disk, a desktop PC PLUS with floppies, and a desktop PC PLUS-XL with a hard disk. Second, the company appears to be well along in upgrading its presence in the retail distribution channel. BusinessLand was the first to sign up in June, but Computer Factory, CompuCraft, Simtec, and Gateway all followed. Most recently, Eagle signed ComputerLand for deliveries starting in March, 1984. Even though each of the ComputerLand franchises has considerable freedom in ordering independently, the backing of the central staff represents a coup. As a result, Eagle is positioned to score in the current race to capitalize on the short-term window of "PC-DOS" compatibility to improve the company's penetration of the markets. At least for the next six to nine months, the trend is likely to be upward.

On the negative front, Eagle has yet to prove that it can deliver the level of profitability required to assure long-term survival in a sub-sector where competition will sharply intensify. The latest

Additional information available upon request.

EAGLE COMPUTER (cont'd)

quarter came in at a low 29.7% gross margin and miniscule .6% operating margin. Comparing the results with other small companies, the gross margin compares reasonably well with COMPAQ and Columbia; all three trail TeleVideo's leading 50% record. Still, Eagle has taken steps in pricing the 8-bit older systems up, and in shooting for a 40% level of gross margin on the newer products that offer the potential for a significant improvement in profitability in the coming year. With a little bit of luck, and a continued high energy level on the part of the management and the employees, Eagle could turn in a reasonable and possibly dramatic upswing. By no means, however, is the potential for 25¢ in per share for the June '84 year, and 65¢ for the June '85 year, assured. The shares are not at all attractive to the shy and retiring, given the high volatility and hectic pace of the sector. We consider the current price to reasonably reflect current competitive conditions in PC's and the low level of profitability of the company currently.

EGLC.FCST
TABLE I
EAGLE COMPUTER
12/1/83

($ millions)	1984FY					
	Sept.83-A	Dec. 83	Mar. 84	June 84	FY1984	FY1985
Revenues	$14.3	$17.5	$20.0	$24.0	$75.8	$115.0
% Chge Y/Y	519.1%	266.7%	100.0%	106.6%	164.1%	51.7%
% Chge Q/Q	23.1%	22.4%	14.3%	20.0%	NA	NA
Operating Income	$0.1	$0.9	$1.9	$3.2	$6.1	$15.5
% Revenues	0.6%	5.0%	9.5%	13.5%	8.0%	13.5%
Interest Inc.(Exp)	-$0.1	$0.0	$0.0	$0.0	-$0.1	$5.3
Pretax	$0.0	$0.9	$1.9	$3.2	$6.0	$20.8
Margin	0.2%	5.0%	9.5%	13.5%	8.0%	18.1%
Taxes	-$0.7	$0.4	$0.8	$1.3	$1.8	$8.6
Rate	46.1%	41.5%	41.5%	41.5%	29.9%	41.5%
Net Income	$0.7	$0.5	$1.1	$1.9	$4.2	$12.2
% Chge	-200.1%	2911.0%	44.9%	1688.1%	2348.9%	187.5%
EPS (Diluted)	$0.01	$0.03	$0.08	$0.13	$0.25	$0.65
% Chge	-110.0%	NA	25.2%	1180.7%	-925.9%	161.5%

EAGLE COMPUTER, INC. QUARTERLY INCOME SUMMARY
11/29/83

	1983FY						
	10/2/82	1/1/83	4/2/83	7/2/83	10/2/83	1982FY	1983FY
Total Revenues	$2,309	$4,772	$10,002	$11,615	$14,295	$5,963	$28,698
% Chg. Y/Y	NA	NA	NA	NA	NA	NA	381.3%
% Chg. Q/Q		106.7%	109.6%	16.1%	23.1%	NA	NA
Costs & Expenses							
Cost of Sales	$1,744	$3,418	$6,910	$8,355	$10,040	$4,431	$20,427
% of Revs.	75.5%	71.6%	69.1%	71.9%	70.2%	74.3%	71.2%
R&D	$379	$362	$551	$609	$890	$919	$1,901
% of Revs.	16.4%	7.6%	5.5%	5.2%	6.2%	15.4%	6.6%
Marketing, G & A	$871	$950	$1,651	$2,417	$3,283	$1,917	$5,889
% of Revs.	37.7%	19.9%	16.5%	20.8%	23.0%	32.1%	20.5%
Total Costs & Expens	$2,994	$4,730	$9,112	$11,381	$14,213	$7,267	$28,217
% of Revs.	129.7%	99.1%	91.1%	98.0%	99.4%	121.9%	98.3%
Operating Income	-$685	$42	$890	$234	$82	-$1,304	$481
% of Revs.	-29.7%	0.9%	8.9%	2.0%	0.6%	-21.9%	1.7%
Interest Expense	$32	$25	$74	$177	-$58	$112	$308
Pre-Tax	-$717	$17	$816	$57	$140	-$1,416	$173
Margin	-31.05%	0.4%	8.2%	0.5%	1.0%	-23.7%	0.6%
Taxes	$0	$0	$49	-$49	$59	-$690	$0
Rate	0.0%	0.0%	6.0%	-86.0%	42.1%	48.7%	0.0%
Net Earnings	-$717	$17	$767	$106	$81	-$726	$173
% Chg. Y/Y	NA	NA	NA	NA	NA	NA	-123.8%
% Chg. Q/Q			4411.8%	-86.2%	-23.6%		
EPS	-$0.10	$0.00	$0.06	$0.01	$0.01	-0.10	-$0.03
% Chg. Y/Y	NA	NA	NA	NA	NA	NA	-68.0%

December 12, 1983

Microcomputer Peripherals
Harvey C. Allison

TANDON CORP. (OTC-TCOR)+

12/6/83 Price	52-Week Range	Earnings Per Share* 1983	1984E	1985E	P/E Ratio 1984E	1985E	Dividend	Yield
$18 3/4	$35-15	$0.47	$1.10	$1.70	17.0x	11.0x	---	---

*Fiscal year ends in September. Of the 51 million shares outstanding, 11,972,200 shares traded in October.

+Shearson/American Express makes a market in this issue.

Investment Conclusion

We are adding Tandon to the Shearson/American Express Recommended List. Tandon's stock may not reflect the extent to which margins will improve over the next six months or the substantial cost advantage Tandon enjoys in the Winchester disk drive product line. Tandon is sharply off its earlier highs this year on news of large receivables write-offs from Victor Technologies, rumors of vertical integration by systems manufacturers, and non-recurring charges that destroyed fourth quarter earnings. Much of the decline has been an over-reaction, and presents a buying opportunity. Revenues can double next year, and profit margins will improve. We are using $1.10 for fiscal 1984 and $1.70 for fiscal 1985.

Tandon has a 5% cost advantage over competitors in the Winchester disk drive industry based on Tandon's internal media production alone. The success of this media effort has surprised everyone, including Tandon, and its merit does not rely on the often disputed higher performance of plated media, but on the cost. Tandon produced about 50,000 5.25" disks in September and about 70,000 in October, in contrast to the 80,000 produced by the entire industry during all of 1982. Tandon's simple three-step process will be automated beginning in the second quarter, and production rates will continue to climb to 200,000 per month by the end of Tandon's year next September. Yield by then will be over 90%, compared to current yield of about 80%.

Offshore manufacturing and vertical integration add to this cost advantage in Winchester and in floppy product lines. Winchester disk drives and 1/2 height floppy disk drives will be moved offshore where the costs are lower for materials and where significant tax advantages exist. Tandon will also be doing its own casting and injection molding next year, and will expand production of magnetic heads.

By product line, the greatest percentage growth will occur in the Winchester disk drive area which accounted for about 25% of fiscal 1983's revenues. Tandon will triple its unit production rate in Winchesters over the next year from 17,000-20,000 units per month to 50,000-60,000. The floppy business will also show strong growth. As the IBM PC and its clones continue to rout other

Additional information available upon request.

Tandon Corp. (Cont'd)

microcomputer manufacturers, the market for PC-standard floppy media manufactured by Tandon explodes. Even firms not in the IBM-compatible market are turning to this drive standard for media compatibility at the expense of 8" and other 5 1/4" drive products.

Product and Market Background

Tandon's primary product markets are for Winchester disk drives which currently account for about 30% of revenues, and floppy disk drives which account for about 60% of revenues. The market for Winchester disk drives is estimated to grow 95% in 1984, and the market for floppy disks is estimated to grow about 85%. In this market environment, we expect Tandon to increase its Winchester disk drive production by 200%, expanding its market share about 50%. Floppy disk sales and increased sales in other product areas are expected to grow at least 50%. Tandon is the primary supplier of floppy disk drives to IBM.

Tandon's two primary product lines are floppy, or flexible, disk drives and fixed, or Winchester, disk drives. Both products are used in microcomputers. Every serious microcomputer has at least one floppy disk drive. Floppy disk drives permit the diskette containing the data to be removed from the drive, allowing the user to keep data on an unlimited number of diskettes, but to read from only one at a time. Fixed, or Winchester, disk drives do not allow the media to be removed, but a fixed disk drive can store many times as much data as a floppy disk drive and the information can be read into the computer or written out to the disk much faster than when using a floppy disk drive.

Fiscal 1983

Tandon recently reported earnings for the fourth quarter of 1983 ending September 30, 1983. As expected, most of the Victor Technologies' debt was written off in this quarter which reduced earnings to $0.03 from $0.12 last year. What was not expected were the additional charges to Research and Development, and to Marketing and Administration. R&D was up 2% as a percentage of sales over the previous quarter, and M&A was up over 4%. Sales were up 93% year to year and 17% sequentially.

Without the Victor charge, earnings would have been about $0.17 per share. There were three other unusual charges during this period: Currency losses for the entire year were taken all at once in this period; an account for the accrual of employee benefits was retroactively adjusted to reflect an increase in the period of eligibility; and substantial research expenses were incurred during the rapid development of the 10MB subsystem product introduced at Comdex. While none of these three charges were extraordinary, two charged the fourth quarter with expenses incurred in previous periods.

Tandon Corp. (Cont'd)

Gross margins improved slightly during the fourth quarter, and are expected to improve more dramatically during the following year. The two week shutdown during the first quarter is expected to mask sequential operating improvements until the second quarter, when both sales and margins will show sharp increases. Our quarterly estimates for next year are:

		1983A	1984E
Dec	Q1	$0.12	$0.18
Mar	Q2	$0.15	$0.25
Jun	Q3	$0.17	$0.31
Sep	Q4	$0.03	$0.34

Tandon

Income Statement	1981	Q1	Q2	Q3	Q4	1982	Q1	Q2	Q3	Q4	1983	Q1E	Q2E	Q3E	Q4E	1984E
Net Sales	$54,152	$25,976	$30,014	$44,245	$50,254	$150,490	$52,539	$70,664	$83,173	$96,992	$303,368	$103,408	$132,363	$169,424	$194,838	$600,033
% Sequential Incr.			16%	47%	14%		5%	34%	18%	17%		7%	28%	28%	15%	
% Year To Year Incr.	138%					178%	102%	135%	88%	93%	102%	97%	87%	104%	101%	98%
Cost of Sales	$37,567	$18,381	$20,725	$30,955	$35,201	$105,262	$36,662	$51,465	$61,200	$70,493	$219,821	$75,488	$93,978	$118,597	$136,387	$424,449
% of Sales	69.4%	70.8%	69.0%	70.0%	70.0%	69.9%	69.8%	72.8%	73.6%	72.7%	72.5%	73.0%	71.0%	70.0%	70.0%	70.7%
Gross Profit	$16,586	$7,596	$9,290	$13,290	$15,053	$45,228	$15,877	$19,199	$21,973	$26,499	$83,547	$27,920	$38,385	$50,827	$58,451	$175,584
% Sales	30.6%	29.2%	31.0%	30.0%	30.0%	30.1%	30.2%	27.2%	26.4%	27.3%	27.5%	27.0%	29.0%	30.0%	30.0%	29.3%
Operating Expenses																
Research & Development	$2,870	$970	$1,265	$1,886	$1,921	$6,042	$2,438	$2,766	$3,842	$5,495	$13,741	$4,653	$6,221	$7,963	$9,157	$27,995
% Sec. Incr.			30.4%	49.1%	1.8%		26.9%	13.5%	10.0%	80.7%		-66.1%	33.7%	28.0%	15.0%	
% of Sales	5.3%	3.7%	4.2%	4.3%	3.8%	4.0%	4.6%	3.9%	3.7%	5.7%	4.5%	4.5%	4.7%	4.7%	4.7%	4.7%
Marketing & Admin.	$5,736	$2,351	$3,028	$3,897	$4,522	$13,798	$4,339	$4,694	$5,353	$10,444	$24,829	$8,273	$10,589	$13,554	$15,587	$48,003
% of Sales	10.6%	9.1%	10.1%	8.8%	9.0%	9.2%	8.3%	6.6%	6.4%	10.8%	8.2%	8.0%	8.0%	8.0%	8.0%	8.0%
Reserve for Acc. Rcvbl.										$12,000	$12,000					
Operating Income	$7,980	$4,275	$4,997	$7,507	$8,610	$25,388	$9,100	$11,739	$13,578	($1,440)	$32,977	$14,994	$21,575	$29,310	$33,707	$99,587
% of Sales	14.7%	16.5%	16.6%	17.0%	17.1%	16.9%	17.3%	16.6%	16.3%	-1.5%	10.9%	14.5%	16.3%	17.3%	17.3%	16.6%
% of Pretax	98.6%	95.3%	80.5%	90.1%	92.2%	89.5%	90.6%	89.6%	91.3%	86.9%	90.7%	95.6%	91.1%	95.0%	95.9%	94.5%
% Seq. Incr.			16.9%	50.2%	14.7%		5.7%	29.0%	15.7%	-110.6%			43.9%	35.9%	15.0%	
Interest Income	$817	$395	$1,233	$1,058	$582	$3,268	$647	$1,759	$1,431	NA	NA	NA	NA	NA	NA	NA
Interest Expense	($707)	($182)	($19)	($232)	$144	($290)	$301	($394)	($145)	NA	NA	NA	NA	NA	NA	NA
Total Other Income	$110	$213	$1,213	$826	$726	$2,979	$948	$1,364	$1,286	($217)	$3,381	$685	$2,118	$1,547	$1,427	$5,776
% of Pretax	1.4%	4.7%	19.5%	9.9%	7.8%	10.5%	9.4%	10.4%	8.7%	13.1%	9.3%	4.4%	8.9%	5.0%	4.1%	5.5%
Pretax Income	$8,090	$4,488	$6,210	$8,333	$9,336	$28,367	$10,048	$13,103	$14,864	($1,657)	$36,359	$15,679	$23,693	$30,858	$35,134	$105,363
Income Taxes	$3,585	$2,020	$2,908	$3,727	$3,977	$12,632	$4,317	$5,318	$6,317	($3,252)	$12,700	$6,585	$9,714	$12,343	$13,702	$42,344
Tax Rate	44%	45%	47%	45%	43%	45%	43%	41%	42%	-196%	35%	42%	41%	40%	39%	40%
Net Income	$4,505	$2,468	$3,302	$4,606	$5,359	$15,735	$5,731	$7,785	$8,547	$1,595	$23,659	$9,094	$13,979	$18,515	$21,432	$63,019
% Seq. Incr			33.8%	39.5%	16.3%		7.0%	35.8%	9.8%	-81.3%		470.2%	53.7%	32.4%	15.8%	
% Y-Y Incr.	198.9%	208.5%	228.7%	240.0%		249.2%	132.2%	135.8%	85.6%	-78.2%	50.4%	58.7%	79.6%	116.6%	1243.7%	166.4%
% of Sales	8.3%	9.5%	11.0%	10.4%	10.7%	10.5%	10.9%	11.0%	10.3%	1.6%	7.8%	8.8%	10.6%	10.9%	11.0%	10.5%
Earnings per Share																
EPS Primary	$0.13	$0.06	$0.07	$0.10	$0.12	$0.36	$0.12	$0.15	$0.17	$0.03	$0.48	$0.18	$0.25	$0.31	$0.34	$1.10
EPS Fully Diluted	$0.13	$0.06	$0.07	$0.10	$0.12 (est)	$0.36	$0.12	$0.15	$0.17	$0.03	$0.47	$0.18	$0.25	$0.31	$0.34	$1.10
Primary Shares	34,548	39,535	44,510	44,532	43,910	43,288	46,142	50,837	51,060	51,130	49,792E	51,200	55,149	59,279	63,719	57,337
Fully Diluted Shares	34,688	39,535	44,510	44,532	43,911	43,290	46,990	50,839	51,144	51,130	50,826E	51,200	55,149	59,279	63,719	57,337

SHEARSON AMERICAN EXPRESS

TECHNOLOGY TRENDS

December 12, 1983

A QUICK READ ON TECH STOCKS

COMPUTER SYSTEMS

IBM's Sugar Plum For Christmas

IBM's meeting with analysts yesterday planted joy in the hearts of the bull, and sent Scrooge running. Dancing with dreams of Wall Street sugar plums in the form of future stock splits, rising stock dividends, and buttressed by the reality of upward earnings revisions and IBM's own stock purchases of as many as eight million of its own shares for internal purposes, one highly respected Wall Street buy-side analyst even predicts a $200 upside price target through 1985. Not bad, and not unreasonable either. The meeting is all the more surprising against an historical tradition of plodding, frustrating analysts' meetings serving as a lesson in evasion as meaningful as a "space wars" arcade game. Yesterday, IBM spoke directly to the investment concerns impacting their future in a highly informative and unusual meeting. The session was all the more important because IBM is just now in the process of finalizing its 1984/1985 plans.

Overall, IBM suggested that the rate of revenue gain for the coming year is likely to fall in the 15% to 20% range. The company alleviated concerns on two critical issues affecting expectations for the coming twelve to eighteen months. First, an acceleration in new equipment purchases is anticipated—to the heady level of 50% in year-to-year gains for 1984. The trend will more than offset the vulnerability of difficult comparisons against recent 1983 results of installed rental equipment. Second, the company suggested that it could manage the product cycle transition to the new 308X computers (Sierra Series) widely expected to be announced between March of 1984 and March of 1985. (The company stated that an 18% increase in modules for the 308X series is anticipated in the coming year.)

Long-term, IBM expects growth to exceed the 14.2% average rate of the latest seven years.

We are revising our earnings estimate to $8.95 from $8.90 for FY1983, and to $10.60 from $10.25 for FY1984. We will be using a preliminary $12.25 estimate for 1985. The company reported $7.39 in 1982.

On balance, improved confidence in the outlook for the coming year is likely to shift attention to the potential for 1985. Our own view is that the shares should easily hit new highs and hold those highs, with the $135-$145 range as a target. A 50% premium to the market on $12.25 in 1985 per share does indeed suggest that an upside target price of $190 represents a distinct possibility. Even conservative investors can buy this tech stock now.

Other highlights:

1. Investments in property plant and equipment will rise by over 20% for 1984, versus a 10% gain for 1983.
2. Marketing will decline as a percentage of revenues, offsetting lower gross margins. The adverse trend in gross margins derives from a higher percentage of activity through third-party marketing firms, and a greater proportion of low-end gear being shipped. Operating margins in the fourth quarter could be under moderate pressure, but are expected to be flat on balance for the coming year.
3. Rental conversion sales are leveling this quarter, and the company does not appear to think there are vulnerabilities from a possible decline in rental asset conversion sales in Q1 and Q2 of 1984. (Specifically, 40% of rental asset conversion sales in 1982 appeared in the fourth quarter; 25% of rental asset conversion sales in 1983 are anticipated in the fourth quarter. There is no particular pressure on Q1 and Q2.)
4. The PC unit sales will rise by 3X's, including all versions such as the PCjr.
5. Orders recorded the highest year-to-year gain year-to-date in November.
6. Other income for the third quarter included payments from Fujitsu, and not Hitachi. The fourth quarter will include, in other income, some of the proceeds from the settlement with Hitachi from the resolution of their recent legal dispute over IBM design work. IBM is legally restricted in its settlement, from discussing the results. The trade press suggests that the settlement with Hitachi represents $300 million of payments. Only a portion will be included in the fourth quarter in connection with recent litigaion while the balance will be included in revenues in the future.
7. The tax rate is now expected to average 43.0% for 1983, compared to 44.4% from 1982. A restatement, however, which has no impact on reported earnings, will shift State taxes out of selling, general, and administrative expenses, and into the tax account. On an adjusted basis, the tax rate will come in at 45.0% for 1983, against 46.4% for 1982. Currently, no change is anticipated in the tax rate for 1984.
8. IBM will purchase shares for its internal requirements in the open market beginning January 1. This year such purchases are roughly equal to eight and a half million shares.

Eagle Computer

Eagle Computer has recently accomplished two key objectives. First, the company has revamped the product line to include full compatibility to the IBM PC/DOS system. In addition, the company appears well along in upgrading its presence in the retail distribution channel. As a result, Eagle is positioned to score in the current race to capitalize on the short-term window of "PC/DOS" compatibility that has been opened by the temporary gap between the demand for the IBM PC and IBM's limited production. For the next six to nine

Communications . . . Sandra Roth
Corporate Finance . . . S.F. Accardo
Digital Systems . . . Donald H. Brown
Market Analysis . . . Frank Korth
Micro Peripherals . . . Harvey C. Allison
Special Situations . . . Neal Goldman
Editor . . . Alisa Kline

MARKET PROFILE FOR WEEK ENDING 12/8/83

CLOSE 12/8/83	COMPANY	% CHANGE 1 WEEK	% CHANGE 4 WEEKS	% CHANGE 8 WEEKS	RELATIVE STRENGTH* 3 WEEK	RELATIVE STRENGTH* 7 WEEK/CHANGE	RELATIVE STRENGTH* ABS. DEV.
$ 9.25	Altos Computer Sys.	8.8%	23.3%	8.8%	NA	NA/ NA	NA
29.50	Datapoint	7.7	16.2	15.1	4	21/-19	2.2
16.87	Prime Computer	6.3	10.6	0.7	23	75/-35	2.7
21.25	Apple Computer	5.6	8.9	- 7.0	70	79/ 21	13.7
43.12	Tandy Corp.	4.8	17.3	12.3	14	57/ 32	8.8
24.50	Bolt Beranek & Newman	4.2	3.7	- 5.7	46	91/-82	6.5
40.87	Hewlett-Packard	3.8	17.6	- 3.8	17	81/-14	2.8
35.87	Tandem Computer	2.5	5.5	4.0	28	31/-26	2.6
12.00	Ramtek	2.1	26.3	-11.1	NA	NA/ NA	NA
9.12	Eagle Computer	1.3	17.6	- 7.6	NA	NA/ NA	NA
31.87	General Instrument	0.3	- 0.7	- 8.9	52	71/ 24	8.6
25.25	Giga-tronics	0.0	5.2	- 3.8	NA	NA/ NA	NA
18.75	Electromagnetic Sci.	0.0	11.9	- 8.5	19	71/-16	2.1
7.00	Computone Systems	0.0	33.3	-7.6	NA	NA/ NA	NA
6.25	K-Tron	0.0	4.1	-16.6	NA	NA/ NA	NA
118.62	IBM	- 0.1	- 5.6	-11.1	87	68/-42	5.5
25.00	Timeplex	- 0.4	9.8	19.7	2	8/ 48	2.3
33.37	Wang Labs	- 0.7	- 2.5	- 8.8	59	61/-42	3.6
17.75	Lee Data	- 1.3	0.0	- 4.0	NA	NA/ NA	NA
28.50	NBI, Inc.	- 1.7	11.7	13.4	5	17/ 77	3.6
18.37	Alpha Industries	- 2.6	- 4.5	-18.7	71	79/-67	2.3
19.37	California Microwave	- 3.1	8.3	13.1	8	12/ 78	2.8
40.37	Analog Devices	- 3.3	9.4	- 0.6	60	57/ 14	3.8
14.50	PAR Technology	- 3.3	0.0	-39.5	NA	NA/ NA	NA
28.50	Management Science Amer.	- 3.3	- 1.7	- 7.3	NA	NA/ NA	NA
39.00	Harris Corp.	- 3.3	6.1	3.6	9	35/ 24	3.8
36.75	Commodore Intl.	- 3.5	7.7	-18.1	8	95/-30	5.6
6.00	Comserv	- 4.0	2.2	-25.0	NA	NA/ NA	NA
68.37	Digital Equipment	- 4.0	0.0	-30.9	48	100/-20	13.7
33.75	Data General	- 4.5	- 7.9	- 9.0	81	59/-35	6.7
20.75	Analogic	- 4.6	- 4.6	- 5.6	60	57/ 14	3.1
28.12	Policy Mgmt. Sys.	- 4.6	5.1	0.4	20	32/ 14	1.2
18.37	TeleVideo	- 5.7	18.5	18.5	NA	NA/ NA	NA
24.25	Avantek	- 6.2	- 2.4	4.3	31	43/ 36	2.6
20.00	Liebert	- 6.9	- 3.6	14.2	32	2/ 86	1.9
41.00	Micom Systems	- 7.3	- 1.2	-12.3	35	55/-49	3.0
15.62	Scientific-Atlanta	- 7.4	3.3	- 0.8	16	61/ 34	1.9
21.87	Convergent Tech.	- 7.9	- 5.9	7.3	NA	NA/ NA	NA
53.00	ROLM	- 9.4	- 8.0	-6.1	42	71/ 10	5.5
20.37	M/A-COM	- 9.4	- 7.4	-15.1	76	86/ 9	3.5
20.00	Quantum Corp.	-10.1	0.0	5.2	NA	NA/ NA	NA
2.12	Vector Graphics	-10.5	-22.9	-41.4	88	100/- 1	1.5
18.50	Tandon	-11.9	-15.9	-26.7	73	93/- 7	NA
10.62	Miniscribe	-12.3	-17.4	NA	NA	NA/ NA	NA
12.87	National Micronetics	-12.7	5.0	-13.4	NA	NA/ NA	NA
20.50	Cipher Data Products	-12.7	3.8	- 5.7	3	39/-31	1.7
13.50	Seagate	-13.5	-15.6	-26.0	50	75/-34	NA
16.50	Verbatim	-17.9	-23.2	-30.1	NA	NA/ NA	NA
	Average	- 3.4%	2.9%	- 5.9%	38	58/ 0	4.3

INDICES

S&P 500	0.7%	0.5%	- 2.6%
DJIA	- 1.0	2.1	0.0

GROUP PERFORMANCE AVERAGES

Communications	- 4.0%	1.2%	1.4%	36	52/- 9
Digital Systems	1.0	8.3	0.1	39	59/- 9
Peripherals	-11.7	-10.0	-15.1	42	69/-24

* Three-week trend and seven-week trend performance ratings give a percentile ranking against a 1500-stock universe. A rank of 1 is the highest and 100 the lowest. Change is an absolute number representing (the seven-week trend percentile today) minus (the seven-week trend percentile seven weeks ago.) Abs. Dev. (Absolute Deviation) is a volitility indicator.

COMPANY PROFILES

COMPANY	EPS 83E	EPS 84E	EPS 85E	PE 83PE	PE 84PE	PE 85PE	MKT CAP $ MIL	REC	ANALYST
Alpha Industries (3)	$0.85	$1.05	$NA	21	17	NA	$128	Buy	SR
†* Altos Computer Sys. (6)	0.58A	0.41	0.60	15	22	15	135		DB
Analog Devices (10)	0.95	1.35	NA	42	29	NA	691	Buy	SR
†* Analogic (7)	0.85A	1.10	1.40	24	18	14	346	Buy	SR
* Apple Computer (9)	1.28	0.60	1.00	16	35	21	1,213		DB
* Avantek	0.63	0.95	NA	38	25	NA	445	Buy	SR
Bolt Beranek & Newman (6)	0.46A	0.75	NA	53	32	NA	152	Buy	SR
* California Microwave (6)	0.65A	0.85	1.05	29	22	18	153	Buy	SR
†* Cipher Data Products (6)	0.50A	0.80	NA	41	25	NA	263		HA
Commodore Intl. (6)	2.86A	4.25	NA	12	8	NA	1,117	Buy	NG
†* Computone Systems (5)	NA	NA	NA	NA	NA	NA	16		SFA
†* Comserv	NA	NA	NA	NA	NA	NA	20		SFA
* Convergent Tech.	0.41	0.90	1.50	53	24	14	772		DB
Data General (9)	0.96	2.25	3.50	35	15	9	731		DB
Datapoint (7)	0.29	0.92	1.00	101	32	29	589		DB
Digital Equipment (6)	5.01A	4.75	8.00	14	13	8	3,784	Buy	DB
†* Eagle Computer	0.02	0.25	0.65	456	30	14	132	Buy	DB
* Electromagnetic Sci.	0.55	NA	NA	34	NA	NA	59		SR
General Instrument (2)	1.50	3.50	NA	21	9	NA	998	Buy	NG
* Giga-tronics	0.65	1.10	NA	38	22	NA	95	Buy	SR
Harris Corp. (6)	1.60A	2.25	2.75	24	17	14	1,266	Buy	SR
Hewlett-Packard (10)	1.55	2.00	NA	26	20	NA	10,192		DB
IBM	8.95	10.60	12.25	13	11	10	71,226	Buy	DB
†* K-Tron	0.15	NA	NA	41	NA	NA	22		SFA
†* Lee Data (3)	1.20	1.50	NA	14	11	NA	218	Buy	SFA
Liebert (9)	0.80A	1.20	1.50	25	16	13	265	Buy	SR
M/A-COM (9)	0.76A	1.05	NA	26	16	12	799	Buy	SR
* Management Science Amer.	0.88	1.20	NA	32	23	NA	481		HA
† Micom Systems (3)	1.30	1.80	NA	31	22	NA	587	Buy	SR
Miniscribe	R	R	R	R	R	R	27		HA
†* National Micronetics (6)	0.41A	0.65	NA	31	19	NA	83		SFA
NBI, Inc. (6)	0.47A	1.50	1.90	60	19	15	281	Buy	DB
† PAR Technology	0.75	0.90	NA	19	16	NA	109		SFA
†* Policy Mgmt. Sys.	0.65	0.85	NA	43	33	NA	421		SFA
Prime Computer	0.69	1.05	1.40	24	16	12	793		DB
* Quantum Corp	0.96	1.10	1.50	20	18	13	190	Buy	HA
†* Ramtek (6)	0.29A	NA	NA	41	NA	NA	40		SFA
ROLM (6)	1.80A	1.75	NA	29	30	NA	939	Buy	SR
* Scientific-Atlanta (6)	0.02A	0.90	NA	781	17	NA	368		SR
Seagate	0.33	0.65	1.00	40	20	13	588		HA
* Tandem Computer (9)	0.76	1.20	1.50	47	29	23	1,343		DB
* Tandon	NA	NA	NA	NA	NA	NA	939	Buy	HA
Tandy Corp. (6)	2.67A	3.10	3.55	16	13	12	4,494	Buy	DB
†* TeleVideo (10)	0.57	0.90	1.25	32	20	14	761	Buy	DB
Timeplex (6)	0.57A	0.85	1.10	43	29	22	159	Buy	SR
†* Vector Graphic	–0.43A	NA	NA	NA	NA	NA	12		SFA
* Verbatim (6)	0.64A	0.90	NA	25	18	NA	373		HA
Wang Labs (6)	1.16A	1.60	2.05	28	01	16	4,100		DB

ECONOMIC FORECAST

	This Month Current Q	This Month Next Q	Last Month Current Q	Last Month Next Q
3-Mo. T-Bills				
S/AE	8.5%	8.25%		
Consensus	8.7	8.5	8.8	8.6
AA Utility Bonds				
S/AE	12.1%	11.9%		
Consensus	12.3	12.1	12.3	12.0

	This Month Current Q	This Month Next Q	Last Month Current Q	Last Month Next Q
Indus. Prod				
S/AE	13.0%	4.42%		
Consensus	10.9	7.0	8.7	7.0
Cap. Spend. (FY1984 Est)				
S/AE	12.2%		11.3%	
Consensus	12.3		11.3	

Consensus figures from Blue Chip Economic Indicators

* Shearson/American Express makes a market in these stocks.
† Shearson/American Express maintains an investment banking relationship with these companies.

months, Eagle's trend is likely to be upward.

On the negative front, Eagle has yet to prove that it can deliver the level of profitability required to assure long-term survival in a sub-sector where competition will sharply intensify. While the shares are not for the faint of heart, they do represent a reasonable speculation.

For additional information on Eagle Computer, please see the report released by Shearson/American Express.

Eagle Computer is an investment banking client of Shearson/American Express.

Tandy November Sales Up 18%

Tandy reported very strong retail sales for the month of November, up 18% for consolidated results, and up 17% for U.S. Radio Shack. After the 10% in October, what a relief! We'd average the two months at an effective 14% since the shortfall in production for October probably slipped some sales into November. We continue to hope for 12% to 15% increases through February as discussed in last week's notes.

COMMUNICATIONS

Harris/Lanier—A Definite Plus

Lanier hosted a small group of analysts last Friday at the company's corporate headquarters in Atlanta. The meeting was an open one with the key objective being to convince the group that Lanier can sell not only stand-alone products, but also systems. In their Southern style, they went through their recruiting, evaluation, and performance-oriented marketing programs, noted "network-type" sales already made, and highlighted the ⅓ of the sales force which is systems-oriented and compensated on a salary-plus-bonus basis.

Many question arose regarding the cost of the Harris/Lanier LAN, the timing of new product introductions and the possible GSA problem resolution. We will likely see a prototype network in February and a more definitive answer to several major issues.

Clearly, dictation and 3M products, service and the finance company are very positive contributors. The System 1000 pc is now being shipped. It, too, should contribute in the second quarter, and it fills a real gap in the Harris product line, offering pc add-ons to Harris' 9000, 800, 1000 and MIND 9240 Systems.

Lanier's management insists the company will meet its objective of not diluting Harris results in any quarter. Achieving this task in the second quarter is easier than originally thought, since our estimate is $0.40-$0.45, with the lower end most likely. Bookings for Harris continue strong particularly in Semiconductor, and the third quarter and second half should reflect this. The ramp-up at Lanier should be similar to Harris' as the year progresses.

M/A-COM—When The Going Gets Tough

M/A-COM stock was off 9% last week as estimates for fiscal 1984 started to come down. We grudgingly recognized the worst-case, though we still believe low-probability, EPS estimate of $1.00 and are using a more realistic potential of $1.15 on the high side. The truth probably lies somewhere in between. The two primary factors for the downgrade were still substantial up-front costs associated with DBS and related-product development, and a slower-than-expected operating turn in Integrated Digital Products, though the quality of the backlog improves monthly.

The Prodelin move will be completed in the second quarter, and unless GI or another DBS player gets off dead center and antenna volume increases, it will continue to negatively impact Cable/Home profitability. Development costs for Video Cypher II and III continue, but just one DBS customer—like a TIME, Inc./Turner—could make all the R&D visibly worthwhile. Video Cypher II, now scheduled for fall 1984 introduction, is an interim DBS product which will offer soft video and hard audio scrambling, with the full-blown VLSI version (III) offering hard video as well, scheduled for mid-late 1985. HBO, which originally bought Video Cypher I and will begin to install those units on the East Coast in January, is now negotiating for Video Cypher II when it is available. Time Inc. will expand its scrambling program to include Cinemax in the fall as well. While a DBS service must eventually be scrambled, it does not have to include scrambling at the outset. We therefore believe Time, Inc./Turner or someone else may announce at any point the programming and service questions are resolved. MAI has clearly positioned itself on the equipment side, at a substantial cost, which will hurt 1984, but if they're right, it will be an extremely large source of revenues in the future.

Our estimate for the first quarter is $0.17-0.19 vs. 0.17 with modest improvement in the second quarter and a strong second half. MAI stock, in our view, discounts the worst case, and does not take into account the potential inherent in their leading-edge technologies. The stock remains on our Recommended List.

PERIPHERALS

Tandon as a Play on PC Demand

With the PC production trebling in the coming year, it's no wonder that IBM is getting into floppy and disk manufacturing—Tandon can't match the demand and others aren't willing to become a second source to Tandon's low cost manufacturing advantage. At the very outside limit, IBM might be able to supply 50% of its internal requirements some time in 1985 or 1986, in the face of a demand curve calling for a trebling in production of the IBM PC family in the coming twelve months.

Assuming IBM's intentions materialize, what would the effect on Tandon be? If IBM tried to make 50% of their own needs, Tandon's sales to IBM would still increase by 50% next year because of IBM's estimated tripling of production. Not a bad increase for Tandon, and quite in line with our expectations, but does anyone really believe that IBM is going to be able to keep up their end of this deal? Starting from scratch? Forget it. So why is everyone hitting on Tandon?

Long-term, Tandon will remain the primary independent supplier of the floppy drives used in the PC products and PC media-compatible products that others produce, and since everyone already has their floppy diskettes, and they always will have to be read by the new equipment, the product line is safe and can't be altered radically. There is even the possibility of a license to IBM from Tandon to allow IBM to use the Tandon design. That would be great.

Winchesters, however, are another story. A Winchester disk drive "talks" to the computer it is built into, and to no other. While there have always been standards for the form factors, power and communications protocol used by the disk drives, there is no reason that a large systems manufacturer could not define a different (read proprietary) disk drive and communications protocol for their products. It would give them a better lead on the plug-compatible competitors and it could be designed with significant cost economies. That is one problem. Tandon looks to be in good shape here since they have no vulnerable business with IBM in Winchesters, they have good customers, a small revenue base, and they have a cost advantage over the competition in the low-cost plated media they produce.

We have added Tandon to the Shearson/American Express Recommended List and expect them to earn $1.10 in 1984 and $1.70, 1985. While not all seventy-odd disk drive firms may make it, we favor firms with low costs, a proven ability to manufacture, proprietary design and we continue to recommend Tandon and Quantum.

Fault tolerance adds two requirements. Some backup hardware someplace must be available if something fails. In addition, either the tasks are running concurrently or they can be restarted to recover in case of failure. So fault-tolerant systems have redundant hardware and the ability to have the programs running concurrently in a second location.

How the systems recover defines the difference. The hardware approach, as exemplified in Stratus, runs everything parallel all the way through the system. A fault any place along the line is caught very early with control passed to another set of hardware. Recovery is instantaneous. There is no special programming required. Unfortunately, there is no additional performance available from the backup hardware. From a pure hardware point of view, productive work is going on in half the system, and the other half is indeed fault-tolerant overhead. This is hardware overhead as opposed to software overhead.

Another approach today is Tandem's checkpoint method. Every time there is some substantial activity, a checkpoint is taken. The checkpoint is to a backup piece of hardware. In case a primary process fails, the backup is up-to-date relative to where the application was at the checkpoint, so the backup can restart and continue. One advantage is that the backup processors are available to do additional work, although recovery may take a little longer. It takes a second or two on the average for a Tandem system to recover. Some of the disadvantages include frequent checkpoints. Either the programmer does them himself, or the system does them. But, a lot of overhead is involved. But the real key is that it takes special fault-tolerant programming to make this kind of system work, which is what we have seen in the marketplace since the mid-1970's.

Our approach is similar to both these alternatives. In our message-based system, the executive processor keeps track of a queue and count buffer. A message coming in is saved in case there is a failure. We don't process it. We don't take checkpoints. We simply save the input. We save the data to recreate the processing if there is a failure – as opposed to processing the data with a redundant set of hardware, or sending massive amounts of checkpoint data back and forth. In case of a failure, the input messages to the backup process are in the buffer. These messages are processed by the backup only when recovery is needed. On the average, in a very large system, it is going to take two to five seconds. So the recovery times that we quote are measured in seconds, and they are not instantaneous. In the meantime, whatever else was running continues. As new transactions come in, rather than being routed to a failed piece of hardware, they are routed elsewhere. The approach is totally automatic. There are absolutely no programming changes required. We take code off our two VAX/780's for development and run it on our hardware without changing a single line. There is no requirement for the programmer to know anything about fault tolerance.

There is also minimal synchronization. Occasionally, the input queue could build up and become very large over time. Periodically, we do synchronize the various processors in a process which is similar to Tandem's. The result is low overhead for fault tolerance localized in the one 68000 in the executive processor. From our early measurement, fault tolerance consumes 10-20% of the system, so 80-90% of the entire system is available for the customer's workload which is what the customer really buys.

The real challenge in fault tolerance is figuring out where the fault is. And once you have it, of course, you have to replace the failing piece of hardware. About 30% of our system hardware in fact covers fault detection and recovery. It is a combination of ECC (error correction code) and parity. A CMOS microprocessor on every board periodically checks for problems. If there is an intermittent failure, this little processor logs it into special, battery-powered memory. When that board comes back to us for repair, we know precisely what failed on the board.We also test the clusters by checking the parameters as they pass back and forth in the operating system to tighten up on software. Also, Stratus has a great idea, and we do have our system dialing out for testing and maintenance to identify which board may be failing.

Other issues relate to programmer productivity and ease of use. About half of our software expenditures concentrates on these challenges, while the other half concentrates on operating systems. Our relational data base is totally compatible with IBM's relational data base called SQL/DS. With our screen manager and menu development system, the programmer needs only to write his COBOL, C, PASCAL, or FORTRAN program without worrying about many basic requirements. On-line help, an English-like query language, and LISA-like user interface also serve the user. A bit map display is not currently available nor is a mouse.

Historically, programmers really had to be cognizant of all parts of the system. They had to worry about the data base management, about dealings with terminals and about the transaction manager. Menus took a lot of programming. These requirements are met as a basic part of the system. The programmers simply write as if it were a batch environment. They have no concept of the transaction processing. It is no more difficult to program our system than a conventional computer. The programmer simply makes calls into the data base manager or screen manager through dictionaries.

Over the years, no customer has ever told me a system was too fast, or too cheap or too easy to use. Mostly, they thought it was too difficult and too slow. We have really focused on price/performance. We do believe that function through performance is the way of the 1980's. And that is why we have put in all these microprocessors doing their specialized functions. Our entry system price is about $140,000 for two clusters which provide 1.7 MIPs (something in excess of a VAX/780), two megabytes of memory and two 80 megabyte disks. All of the productivity tools in the basic system are bundled, without charging a separate price for the productivity tools. Everything is there to encourage their use up-front.

Q: What language did you use to write your code?

A: Our operating system is written primarily in C. But the same portability will hold for COBOL, FORTRAN, BASIC, and PASCAL. We have already tested it to date between the VAX and the Altos systems. Obviously, the object code on the VAX is DEC's object code, and ours is for the 68000. We simply recompile and run. We have started the COBOL testing. We see no reason why the process won't be working in all languages.

LC 12-11
Additional information available upon request

Tandem, MKt Res. Reports

MARTIN SIMPSON

& Company, Inc.
115 Broadway, New York, N. Y. 10006

Tel: (212) 349-7450
Telex 66544

TECHNOLOGY STOCK SERVICE
Update #2
July 25, 1983

Tandem Computers End-User Survey

After several years of explosive growth, Tandem, the leading independent vendor of fault-tolerant systems for the on-line transactions processing market, experienced a major setback nearly one year ago as order rates fell and revenues and earnings were restated downward to reflect over-aggressive booking policies. Accounting practices have now been tightened and recent orders have shown some modest improvement, leading to our expectation that 1984 will be an excellent year for the company.

We recently conducted a survey of 200 Tandem users (30% of its customer base). Overall, the results were quite positive, although several problem areas were identified. Of the greatest significance, fully 97% of those polled had either definite or probable plans to expand their Tandem installations during the coming year. These users expected to add nearly 900 CPU's, which represents more than two quarter's of Tandem shipments. On the negative side, 22% of the users were unsure whether they would choose Tandem if the decision were made again. The reasons were basically twofold: high cost and support. Tandem's price/performance has deteriorated as the operating system has expanded and as competitive system prices have come down. Software support was felt to be an area in need of improvement. A new 16-bit system should be introduced in the fall and should answer some of the overhead/performance problems.

On the 32-bit issue (and Tandem's lack thereof), 62% felt that such a system was either very important (21%) or somewhat important (41%). Tandem's entry into this market is not likely until 1984. While current customers were willing to wait, prospective ones may not.

While we have some concerns regarding the current quarter now that the trade-in program has ended and the new products have not yet been introduced, we expect a very strong revenue and earnings rebound next year aided by a buoyant economic environment and favorable product cycle. We recommend purchase of the stock for substantial potential capital appreciation.

Tandem Computers, Inc.
(OTC - TNDM)

Recent Price:	$29
1982-1983 Price Range:	$14-34
1985 Price Objective:	$60
Annual Compound Appreciation:	44%
Estimated 12 Month Downside Risk:	25%

Year (9/30)	Per Share	Price/ Earnings
1980	$0.35	82.9x
1981	$0.72	40.3x
1982	$0.76	38.2x
1983E	$0.80	36.2x
1984E	$1.40	20.7x
1985E	$1.90	15.3x

Dividend: None Yield: Nil

Capitalization ($000's)

Long-term Debt and Leases	$ 21,102	8%
Common Equity	250,988	92
Total	$272,090	100%

The company has 41 million shares outstanding. Tandem's equity at the current market price is valued at approximately $1.19 billion.

(E) Estimates of Martin Simpson & Co., Inc.

years of experience with the company. New customers represented a small percentage of the total. This may possibly be explained by two factors (1) newer users may be less likely to be members of the International Tandem User's Group and thus were excluded from the population (2) shipments to new customers as a percentage of total placements have been declining for the company as a whole.

TABLE 3

Tandem: Average Years Of Usage

No. Of Years As A Tandem Customer	% Of Users
1	16%
2-3	47
4-5	32
6-7	5
	100%

Average of All Users: 3.0 Years

Source: Martin Simpson & Company, Inc.

Geogrpahically, the U.S. represented a disproportionately large percentage, 85%, of the sample. U.S. users were farily evenly split between the eastern and western regions. Europe accounted for only 5% of the sample, yet comprises 25% of Tandem's customer base. Participation from Canadian, Asian and Middle Eastern users was also underrepresented. This is attributable to two factors (1) no telephone interviews were conducted outside the U.S. (2) foreign users were under significantly more time pressure to return the questionnaires.

Table 4

Geographic Distribution of Tandem Users

Eastern U.S.	44%
Western U.S.	41
Canada	4
Europe	5
Asia, Middle East	6
	100%

Source: Martin Simpson & Company, Inc.

Survey Methodology And Sample Size

We recently completed a major survey of Tandem users. Nearly 200 users were contacted, representing approximately 30% of Tandem's user base, a very significant sample size. 170 responses were received by mail through questionnaires sent out by the International Tandem Users Group **(See Appendix)** while the remainder were surveyed in phone interviews conducted by Martin Simpson & Company, Inc. The survey was conducted in the April - June 1983 timeframe.

The 193 users surveyed had 1,454 Tandem CPU's installed. The sample included both NonStop I and NonStop II models, with an edge to the latter as can be seen in Table 1.

TABLE 1

Number of Tandem CPU's Surveyed By Model

NonStop I	552	38%
NonStop II	902	62
	1,454	100%

Source: Martin Simpson & Company, Inc.

The average installation size was 7.5 processors per user. The typical NonStop II installation was more than one-third larger than the average NonStop I installation.

TABLE 2

Average Size of Installation By Model

	No. of Processors
NonStop I	5.3
NonStop II	7.2
Total	7.5

Source: Martin Simpson & Company, Inc.

Tandem User Profile

The typical Tandem user had been a customer of the company for approximately 3 years. The distribution, as set forth in Table 3, is heavily weighted towards users with 2-3

Distribution of the users by industry was highly correlated with the breakdown for the company as a whole, as can be seen in Tables 5 and 6.

TABLE 5

Breakdown By Industry of Tandem Users Surveyed

Banking	19%
Insurance	2
Other Financial	6
Services	17
Manfacturing	14
Government	10
Medical	7
Software House/OEM	5
Education	4
Communications	4
Retail	3
Utilities	3
Transporation	1
Distribution	1
Entertainment	1
Computer Programming	1
Agriculture	1
Religion	.5
No Response	.5
Scientific	0
	100%

Source: Martin Simpson & Company, Inc.

TABLE 6

Breakdown By Industry Of Tandem Users Corporate

Banking	19%
Other Financial	10
Manufacturing	11
Communications	10
Government	9
Medical	7
Distribution	3
Transportation	3
Other	28
	100%

Source: Martin Simpson & Company, Inc.

Banking, along with other financial sectors such as insurance and brokerage firms, is Tandem's most important industry group, accounting for nearly 30% of sales. This has risen several percentage points over the past two years reflecting both increased penetration in this marketplace as well as a recession-related decline in the manufacturing sector. Tandem has systems installed at 70% of the 25 largest banks in the U.S., and has captured many leading U.K. banks as well. The trend towards the widespread acceptance of automated teller machines as well as the possible emergence of home banking, presents an enormous marketing opportunity for Tandem. These types of real-time, transaction processing applications demand the fault-tolerance and data integrity that Tandem systems provide.

Manufacturing usage of Tandem systems as a percentage of the total has declined due tc the recent recession. Certain geographical areas, such as Germany, Canada, and the MidWest U.S., were particularly hard hit. As the economy recovers, we would expect a resurgence in this sector.

Service industries (consulting, legal, etc.) have grown in importance to Tandem as has the communications industry. While users in the communications field were underrepresented in our survey, we expect dramatic growth in this area for Tandem, particularly as such services as Telextex grow in importance.

The government and medical markets are both quite significant to the company, accounting for about 16% of the total. Software houses and original equipment manufacturers (OEMs) comprise a fairly small sector, while the scientific market is of virtually no importance to the company.

In terms of system acquisition, the overwhelming majority had purchased their machines (78%). Of those who had leased, many were among Tandem's newer customers. The third category comprises those organizations who use Tandem systems but do not own or lease them.

TABLE 7

Tandem Computers: System Acquisition Method

Own	79%
Lease	14
Use, but does not own or lease	7
	100%

Source: Martin Simpson & Company, Inc.

System Selection

As would be expected, Tandem's fail-safe capability was the primary factor users cited in the selection process. Nearly 50% of those surveyed mentioned redundancy as the key reason for installing Tandem. Ease of system expansion (particularly the fact that no reprogramming is necessary) was cited as the most important factor in the system selection by 24% of the users. Networking and database management software were cited by 12% and 11% of those surveyed, respectively. Price/performance was mentioned by a mere 1%, which correlates with the general feeling expressed throughout the survey that Tandem's systems were too expensive.

TABLE 8

Primary Reason Given For Selecting Tandem

Fail-Safe Capability	48%
Expandability	24
Networking Software	12
Data Base Management Software	11
Tandem's Reputation	1
Price/Performance	1
Other (1)	3
	100%

(1) Software, IBM Communications, support, ease of developing new applications.

Source: Martin Simpson & Company, Inc.

Other factors which users mentioned as key reasons in the selection process were (single responses):

1. **Software**
2. **IBM Communications**
3. **Support**
4. **Ease of Developing New Applications**
5. **Desired Software Packages Ran on Tandem**
6. **Block Structured Language**
7. **Easy Operating System Maintenance**

Users were asked to rate the features listed in Table 9 in order of their importance in the selection process. The higher the number, the more important the feature. While this shows the same data as in Table 8, its presentation in Table 9 shows a tighter range in the value placed on the various

criteria. Again, redundancy and expandability ranked highest, while data base management software, networking software, price/performance and the company's reputation were all rated in the same general range.

TABLE 9

Reasons For Selecting Tandem

Fail-Safe Capability	5.5
Expandability	5.3
Data Base Management Software	3.7
Networking Software	3.6
Price/Performance	3.5
Tandem's Reputation	3.3

Scale 1 = Not Important
7 = Very Important

Source: Martin Simpson & Company, Inc.

In terms of seriously evaluating competitive systems at the time the decision to install Tandem was made, Digital Equipment was most often mentioned (by 52% of those polled) followed closely by IBM (49%). The other leading minicomputer vendors, Hewlett-Packard, Data General and Prime, were mentioned nearly one-third of the time each. The BUNCH companies (Burroughs, Univac, NCR, Control Data, and Honeywell) were evaluated much less frequently, while foreign vendors were rarely looked at. At the time that most customers made their equipment evaluations, the new fail-safe start-up companies either were not in existence or did not represent a viable alternative. While 3% of those surveyed had seriously considered Stratus (and rejected it because of the company's lack of experience and/or less powerful product), less than 1% had looked at Synapse and no one had evaluated Computer Consoles, Parallel Computing, August Computers or Sequoia. Fully 15% of the users had not evaluated any competitive equipment, stating that Tandem offered the only hardware/software solution that fit their needs.

TABLE 10

% Of Tandem Users Evaluating Competitive Equipment

None	15%
Digital Equipment	52
IBM	49
Hewlett-Packard	31
Data General	27
Prime	19
Burroughs	14
Honeywell	10
Sperry Univac	5
NCR	5
Perkin-Elmer	5
Wang	3
Stratus	3
Datapoint	2
General Automation	1
Control Data	1
Modcomp	1
ICL	1
GEAC	1
Siemens	.5
Fujitsu	.5
Gould SEL	.5
Qantel	.5
MAI - Basic Four	.5
Synapse	.5
Harris	.5
Texas Instruments	.5

Source: Martin Simpson & Company, Inc.

Table 11 sets forth the ratings results on thirteen characteristics of Tandem, ranked on a scale of 1 (poor) to 5 (excellent). No category was ranked lower than 2.56 (fair-good) and the majority faired in the good - very good category. Again, as would be expected, CPU reliability ranked best at 4.33, while peripheral reliability was more than a full point lower. Disk drive and printer failures were commonly mentioned as problems. While service and support of hardware ranked high, software support was given a low rating (this issue will be discussed later) as was applications software (customers wanted to see more from Tandem.) Service responsiveness recieved the lowest overall ranking, with users commenting that while major problems were quickly corrected, minor, but troublesome issues, took Tandem a long time to resolve. Systems software products were all highly regarded, with rankings from 3.0 to 3.95. Overall satisfaction with Tandem was quite high, averaging 3.9.

TABLE 11

Tandem: System Ratings

		% Of Users Responding
CPU Reliability	4.33	100%
Peripheral Reliability	3.23	95
Service Quality	3.61	100
Service Responsiveness	2.56	100
Hardware Support	3.61	100
Software Support	2.88	100
Applications Software	2.70	78
Operating Systems Software	3.65	100
Data Base Management Software	3.70	87
Networking Software	3.95	47
Applications Languages	3.50	93
Program Development Tools	3.00	87
Overall Satisfaction with Tandem	3.90	100

Scale: **1 = Poor**
2 = Fair
3 = Good
4 = Very Good
5 = Excellent

Source: Martin Simpson & Company, Inc.

Users were asked what they liked most about Tandem, the results of which are presented in Table 12. Once again, reliability was far and away the most highly praised feature. Users often commented that Tandem's claim to be truly fault-tolerant was correct. Few customers mentioned any downtime problems other than system crashes which occured when Tandem maintenance personnel were working on the system **(See Table 13)**. Other complaints were (1) there is no redundancy capability in line controllers (2) Tandem's NonStop philosophy can lead to a false sense of security if back-up functions are lost without being brought to one's attention and thus certain failures need to be highlighted or require operator acknowledgment and (3) NonStop is difficult to implement in Fortran and TAL.

Many users complimented Tandem's treatment of its customers, with the most frequent comment being a "feeling that the company wanted the user to be successful." Following on the same theme, its overall customer support was also well-regarded.

The system's performance and design (both hardware and software) were frequently cited as Tandem's greatest strengths, as was its ease of expansion. Most competitive vendors require users who need more computer power to upgrade to a higher-level

model, which is not only expensive and disruptive, but which often requires some degree of reprogramming or software modification. With a Tandem system, new CPU's can be added as desired without interupting system performance. This has been an enormous strength for Tandem and one which several of the newer computer vendors (eg Convergent Technologies) have also adopted.

Several users praised the excellent quality of Tandem's management and staff, with many stating that the level of professionalism exhibited was unsurpassed in the industry. Tandem's ease of use and hardware service were also considered to be strong points.

Additional positive features mentioned infrequently and thus not listed in Table 12 included: data integrity, ability to recover from system failures, education centers, and NonStop I/NonStop II compatibility.

TABLE 12

Tandem Features Liked Best

	No. Of Times Mentioned
Reliability	65
Tandem Attitude/Support	27
System Architecture/Performance/Quality	26
Expandability	20
Quality of Employees	18
Ease of Use	15
Hardware Service	14
System Software	9
Data Base Management System	8
Software Development Tools	6
Fast Response To Major Problems	4
Networking Software	4
Enscribe Software	3

Source: Martin Simpson & Company, Inc.

There were many more separate categories given in response to what users liked least about Tandem. Leading the list was cost. Many users felt that Tandem systems were significantly overpriced. While Tandem has, from the start, commanded a premium for its fail-safe capability, this premium has increased over time due to several factors: (1) competitive systems (not necessarily fault tolerant) have undergone dramatic downward price adjustments while Tandem's prices have stayed about the same, thus widening the price gap (2) new releases of the operating system require increased overhead to maintain, thus raising system cost. Many customers were sensitive to this issue, complaining that they were forced to maintain an increasingly costly operating system which has grown to support

additional software products which these users did not require. Several usrs commented that it now took four processors to run the Cobol Compiler, where previously it required two, making it not only larger, but slower as well. Clearly, the issue of growing operating system overhead, which has a direct bearing on price/performance, will have to be addressed in new products from Tandem. Other cost issues related to maintenance, memory and peripherals, all of which were felt to be priced out of line with Tandem's competitors.

A major source of contention with Tandem was the issue of software support/quality control. Many users felt that software support was badly in need of attention. Customers were concerned that it took too long to get responses to questions from the company, in some cases, months. Software documentation, in particular, was felt to be poor. Quality control has deteriorated, according to some users, resulting in software with so many bugs it could not run. Specific issues raised with respect to software support and software products were:

1. **Poor language support**
2. **Inadequate development tools - need more than an editor and TAL compiler**
3. **Lack of interactive query language**
4. **Poor user orientation ("it's better on a P.C.")**
5. **Would like to see a manufacturing control package and job accounting package for batch programs**
6. **Unavailibility of applications software**
7. **Complex structure of TAL: need language between TAL and Fortran/Cobol**
8. **Software is untested**
9. **NonStop software is cumbersome**
10. **Difficulty in predicting system utilization of an application.**

Problems with peripherals related to several factors, besides cost: (1) difficulty in interfacing non-Tandem terminals (2) poor sequential I/O performance and lack of tape-handling procedures (3) 600 l.p.m. printer failures (4) lack of disk compression utility (5) difficulty in diagnosing sporadic disk errors (6) lack of a one command utility to remove disk fragmentation and (7) lack of word processing support on the 6530 terminal.

As with most end-user surveys there is a certain degree of inconsistency. For instance, while staff quality was mentioned as a strength in Table 12, it was also felt to be a weakness. Regional differences can best explain this, as most of the complaints concerning employees were related to field service offices in particular locations. In terms of Tandem's management, five customers expressed concern about its unorthodox style which made it harder to sell the company to the user's top

management. A few users also felt that customers were getting involved, against their wishes, in internal company politics.

Other areas of concern were service, performance (largely related to the operating system problems discussed earlier) and documentation. Several users were upset about parts replacement (reluctance on the part of Tandem to replace, rather than fix, an obviously defective part) and its microcode policy (requiring it when not needed) which was felt to be a "marketing ploy".

A number of customers felt that while Tandem was great for interactive tasks, its ability to handle batch processing applications was weak and in need of greater support. Additionally, a few users expressed concern that Tandem's pace of new product development, both hardware and software, was too slow.

Other drawbacks that were singled out included: (1) sloppy hardware upgrades (missing parts) (2) pressure to buy before discount period runs out and (3) lack of leasing provisions.

TABLE 13

Tandem Features Liked Least

	No. Of Times Mentioned
High Cost	26
Software Support	25
Software Quality Control	20
Operating System Overhead	16
Staff Quality	14
Peripherals	14
System Performance	12
Field Service	11
After - Sales Support	9
Maintenance Downtime	7
Software Development Tools	7
Documentation	6
Batch Processing Capability	6
Tandem Management Philosophy	5
Microcode Policy	4
Parts Replacement	4
Not Enough New Products	4

Source: Martin Simpson & Company, Inc.

Users were then asked whether they would still choose Tandem if the decision were being made today. 78% answered "yes", without reservation. The central theme appeared to be

that Tandem still had the best hardware/software solution for their needs, despite some very real problem areas. 7% were unsure, stating that the competitive environment had changed sufficiently since the original decision was made, presenting other alternatives which would have to be evaluated. Fully 15% stated that they would not choose Tandem again, with the most common reason being its high cost. Other reasons were (1) would prefer IBM (2) would prefer Stratus (3) poor high-speed communications support and (4) disappointment with Tandem's system performance and support. It should be noted that the majority of the 15% responding "no", were planning to add to their Tandem installations within the coming year, underscoring the importance of an installed base.

TABLE 14

% Of Users Who Would Select Tandem If Decision Were Made Today

Yes	78%
Maybe	7
No	15
	100%

Source: Martin Simpson & Company, Inc.

32-bit systems have rapidly become the standard architecture in the minicomputer industry. These systems provide four principal advantages over 16-bit minicomputers. First, 32-bit word length allows more direct memory addressing capability, permitting direct access of up to 4 billion words or 16 billion bytes versus 64,000 words or 128,000 bytes for 16-bit processors. As a result, 32-bit systems have higher throughput and are therefore more productive. Third, computations are more precise as more significant digits can be carried through calculations. Finally, program development is made easier by the longer word length.

Tandem's product line is still 16-bit. Customers were asked how they felt about a Tandem 32-bit processor. 62% felt it was either somewhat (41%) or very (21%) important, while 35% considered the issue irrelevant. 32-bit systems are generally most appropriate in scientific applications, particularly those with heavy number crunching tasks. As can be seen in Table 5, none of the Tandem users surveyed fell into this group. While many users felt that the type of processor was not important as long as the system did the job, some felt that it would improve system throughput. Others felt that since the industry was adopting 32-bit architecture, Tandem should too. A 32-bit Tandem system is not likely to be introduced before 1984 and will probably <u>not</u> be fully compatible with Tandem's current 16-bit

systems. While most users can wait until next year (though perhaps not new customers), the compatibility issue will prove to be a thornier issue, in our opinion.

TABLE 15

Importance Of Tandem 32-Bit Processor

Very Important	21%
Somewhat Important	41
Not Important	35
Undecided	3
	100%

Source: Martin Simpson & Company, Inc.

"ENCOMPASS", Tandem's relational data base management system, has acheived excellent acceptance. As can be seen in Table 16, 56% have installed or plan to install the product.

TABLE 16

Tandem: ENCOMPASS Usage Plans

Have Installed or Plan to Install	56%
Do Not Plan to Install	40
Undecided	4
	100%

Source: Martin Simpson & Company, Inc.

Far fewer users, 18%, were planning to add "Transfer." This is a new software product, which based on the "Expand" network, will integrate electronic mail, facsimile and communications. It will tie together all types of equipment, regardless of the location or application.

TABLE 17

Tandem: TRANSFER Usage Plans

Have Installed or Plan to Install	18%
Do Not Plan to Install	70
Undecided	12
	100%

Source: Martin Simpson & Company, Inc.

Table 18 sets forth the breakdown of installations polled by model. NonStop II's have a decided edge over NonStop I's, while 18% of the users had configurations consisting of both models.

Table 18

Breakdown of Tandem Installations By Model

NonStop I	36%
NonStop II	46
Both Models	18
	100%

Source: Martin Simpson & Company, Inc.

Tandem introduced the NonStop II in 1981. A trade-in policy was instituted whereby customers received larger and larger credits for the NonStop I the longer they waited to upgrade. This was done in to provide an orderly transition to the new product line. At the final stages of the trade-in program, 100% credit was given for NonStop I CPU's traded in for a NonStop II. (CPU's account for approximately 50% of total system value). The trade-in program ended on June 30, 1983. As can be seen in Table 19, 38% of those polled had already upgraded while an additional 15% planned to do so (and we assume they already have since the poll was taken before the program ended). 13% had decided not to upgrade. Reasons given included (1) not enough additional power to justify the additional cost (2) paying for what you don't need (3) too costly, even with the trade-in credit and (4) would prefer to front-end the system with microcomputers and (5) are waiting for the NonStop III. Only 8% had started out with NonStop II's, reflecting the diminshing contribution from new customers, versus old.

TABLE 19

NonStop II Expansion Plans

Already Upgraded to NonStop II	38%
Plan to Upgrade to NonStop II	15
Started With NonStop II	8
Have Decided Not to Upgrade	13
Haven't Made a Decision Yet	18
No Response	8
	100%

Source: Martin Simpson & Company, Inc.

Table 20 sheds some light on the pricing issue at Tandem. Average revenues per CPU have nearly doubled in six years, despite declining prices for the industry as a whole. Tandem has effectively maintained its prices, but is now running into customer resistance. It should be noted, of course, that newer CPU's are also much more powerful.

TABLE 20

Tandem Revenue Per CPU

	CPU's	Revenue Per CPU (1) ($000)
1977	69	$111.6
1978	176	138.1
1979	389	144.0
1980	653	166.9
1981	1,210	172.2
1982 (2)	1,542	202.4
1983E (2)	1,950	220.0

(1) Includes Service
(2) Increase in Revenues Per CPU are somewhat distorted on the high side as shipments represent net numbers and do not take into account returned NonStop I's.
(E) Estimate

Source: Martin Simpson & Company, Inc.

Perhaps the most impressive part of this survey is the fact that 97% will either definitely (58%) or probably (39%) expand their installations during the next twelve months, while only 3% will not. Not only does this demonstrate the value of an installed base and how difficult it is to convert software to another vendor, but also underscores customer loyalty to Tandem. Many of the users who complained the loudest about certain issues were the ones who were expanding their installations the most! However, combined with the results shown in Table 14 where 22% would possibly not choose Tandem again, it highlights the fact that while existing customers may be loyal, new customers may be hard to get unless some of the price/performance and support issues are resolved.

TABLE 21

User Expansion Plans Within Next Twelve Months

Will Definitely Expand	58%
May Expand	39
Will Not Expand	3
	100%

Source: Martin Simpson & Company, Inc.

Table 22 sets forth the percentage of users planning to expand within the eight categories listed. 69% of those polled were planning to add more Tandem processors. (This does not include NonStop II upgrades.) A total of 860 CPU's are expected to be added over the next twelve months, averaging 4.46 per user. Most were adding one or two CPU's while one customer planned to add 30 processors. The 860 CPU's represent more than two quarters' worth of Tandem shipments, indicating that if customers follow through with their plans, 1984 should be a very strong shipment year for the company.

TABLE 22

Tandem Expansion Plans By Product Category: CPU's, Memory, Disk Drives, Terminals and Printers

	% Of Users Expanding
CPU'S	69%
Tandem Memory	61
Non - Tandem Memory	5
Tandem Disk Drives	76
Tandem Terminals	37
Non - Tandem Terminals	49
Tandem Printers	21
Non - Tandem Printers	52

Source: Martin Simpson & Company, Inc.

TABLE 23

Tandem CPU Expansion Plans Within Next Twelve Months

No. Of CPU's Planned (1)	% Of Users
1	37%
2	34
3-10	22
11-20	5
21+	2
	100%

Total CPU's Planned: 860
Average Planned CPU's User: 4.46
(1) Does not Include Upgrades of NonStop I's to NonStop II's

Source: Martin Simpson & Company, Inc.

61% of those polled planned to add Tandem memory. The average addition was 6.8 megabytes per user and a total of 1,320 Mb. Only 5% were planning to buy foreign memory (a total of only 34 Mb). However, several users felt that non-Tandem memory would become more prevalent unless Tandem lowered its memory prices. To date, this has certainlly not been the case.

TABLE 24

Tandem Memory Expansion Plans Within Next Twelve Months

Amount of Memory Planned	% Of Users
Less than 1 Mb	26%
1.5 - 4 Mb	46
5 - 10 Mb	13
Greater than 10 Mb	15
	100%

Total Tandem Memory Planned: 1.32 Gb
Average Tandem Memory Planned Per User: 6.84 Mb

Source: Martin Simpson & Company, Inc.

TABLE 25

Non-Tandem Memory Expansion Plans Within Next Twelve Months

Amount Of Non-Tandem Memory Planned	% Of Users
Less Than 1 Mb	11%
1.5 - 4 Mb	33
5 - 10 Mb	22
Greater Than 10 Mb	0
Undecided	34
	100%

Total Amount of Non-Tandem Memory Planned: 34 Mb
Average Planned Amount of Non-Tandem Memory Per User: 180K

Source: Martin Simpson & Company, Inc.

76% of the users (more than any other category) planned to buy more disk drives, for a total of 1,525 units, or an average of nearly 8 drives per user. The most popular sizes were 128 Mb and 240 Mb capacities.

TABLE 26

Tandem Disk Drive Expansion Plans Within Next Twelve Months

No. Of Tandem Disk Drives Planned	% Of Users
1-2	59%
3-6	28
7-100	10
101+	3
	100%

Total Disk Drives Planned: 1,526
Average Disk Drives Planned Per User: 7.91

Source: Martin Simpson & Company, Inc.

While 37% of those surveyed intended to purchase Tandem terminals, not too far below the 49% planning to add third-party terminals, the difference is enormous in terms of the absolute number of units to be added. Approximately 1,300 Tandem Model 6530 terminals will be bought, averaging 6.8 per user, with the majority purchasin very small quantities. However, these same users plan to buy nearly 17,000 non-Tandem terminals, for an average per user of 86 units. As would be expected, cost is the

major factor and several customers complained that Tandem priced its own terminals at approximately three times the third-party competition.

TABLE 27

Tandem Terminals Expansion Plans

No. Of Tandem Terminals Planned	% Of Users
1-2	19%
3-5	25
6-10	17
11-40	29
41-60	4
61-100	2
101+	4
	100%

Total Number of Tandem Terminals Planned: 1,311
Average Number of Tandem Terminals Planned Per User: 6.8 Units

Source: Martin Simpson & Company, Inc.

TABLE 28

Non-Tandem Terminals Expansion Plans

No. Of Non-Tandem Terminals Planned	% Of Users
1-2	10%
3-5	15
6-10	22
11-40	24
41-60	13
61-100	5
101-500	5
500-1000	4
1000+	2
	100%

Total Number of Non-Tandem Terminals Planned: 16,676
Average Number of Non-Tandem Terminals Planned Per User: 86 Units

Source: Martin Simpson & Company, Inc.

The same situation prevailed for printers. More than half of the users polled planned to add third-party printers, for a total of 1,053 units, or an average of 5.5 units per customer. Only 21% expected to buy Tandem printers, for a total of 153 units.

TABLE 29

Non-Tandem Printer Expansion Plans

No. Of Non-Tandem Printers Planned	% Of Users
1	21%
2-5	27
6-50	44
51+	8
	100%

Total Number of Non-Tandem Printers Planned: 1,053
Average Number of Non-Tandem Printers Planned Per User: 5.5 Units

Source: Martin Simpson & Company, Inc.

TABLE 30

Tandem Printer Expansion Plans

No. Of Tandem Printers Planned	% Of Users
1	53%
2-5	33
6+	14
	100%

Total Number of Tandem Printers Planned: 153
Average Number of Tandem Printers Planned Per User: .8 Units.

Source: Martin Simpson & Company, Inc.

Conclusion

The survey results were generally quite positive but did identify some major problems that Tandem must cope with. Overall, most users felt that Tandem still offered the best total solution for their needs. Redundancy and easy expandability were the most important reasons in the selection process, followed by software, particularly data base management systems and networking. While competition in the fault-tolerant marketplace has intensified, Tandem still has a very significant lead, in our analysis. The newer entries (Stratus, Synapse, August, Parallel, No Halt) are largely still untested, often lack software, and most are less powerful than Tandem. Several of the leading computer vendors (Digital Equipment, IBM, and Hewlett-Packard) have introduced systmes with some degree of redundancy but none offer truly fault-tolerant systems. Designing a fail-safe system which is compatible with existing software is an enormous engineering challenge.

Clearly, the survey results indicate a high degree of customer loyalty as evidenced by the 97% who will either definitely or probably expand their installations in the coming year. However, the price/performance question is a thorny issue, particularly in terms of operating system overhead. The new products, due out later this year, should address this problem. Until these products are announced, it may prove difficult to add significant numbers of new customers.

The on-line transaction processing market is one of the fastest growing segments of the computer industry. Tandem has an excellent foothold in this market, particularly the banking segment which is growing at roughly twice the rate as the industry as a whole. The company's potential continues to be enormous and $1 billion in revenues by 1986 appears to be a very attainable goal.

Le-ellen Spelman

Additional Information Available on Request

ITUG CONFIDENTIAL TANDEM SURVEY

Dear Tandem User,

This brief questionnaire has been designed to provide the ITUG Board of Directors and the survey sponsor, Martin Simpson & Co., with objective, impartial, and unbiased information about Tandem users. Please fill this out immediately and return it in the attached envelope. The envelope has been pre-addressed for your convenience. The postage has been prepaid for all respondents returning the questionnaire from a U.S. location. International postal regulations prohibit us from prepaying mail sent from outside the U.S., so please attach whatever postage is necessary to return this.

All of your answers will be kept STRICTLY CONFIDENTIAL. No individual respondents will be identified. By guaranteeing you complete confidentiality, we hope you will respond freely and accurately.

Thank you very much for your understanding and consideration. Please mail this on or before May 20, 1983 so that your opinion is counted. We need to hear from you and appreciate your cooperation.

INSTRUCTIONS: Please circle or mark the response(s) which are most appropriate. If you would like to make additional comments or explain any of your answers further, please use Page 2 (the back) of this sheet. Thanks!

A. Do you or your organization own, lease, or use a Tandem system? Please circle the answer which best describes whether or not you are an official Tandem customer. Please circle only one answer.) <4>

The Tandem computer system used by me and/or my organization . . .

[1] . . . is owned by me or my organization.

[2] . . . is leased by me or my organization.

[3] . . . is used, but not owned or leased, by me or my organization.

B. How many years have you been a Tandem customer or a Tandem user? (If you have been a customer or user for less than 1.5 years, please put down "1" year.) <5>

(a) I have been a Tandem customer for _____ year(s).

(b) I have been a Tandem user for _____ year(s).

C. What is your *principal* application on or use of your Tandem? <6-7>

D. Which *single* category below best describes your organization? (Please circle only one answer. Thank you.) <8-9>

[1] Banking
[2] Finance, excluding banking or insurance
[3] Insurance
[4] Education and/or academic
[5] Medical
[6] Manufacturing
[7] Government
[8] Transportation
[9] Services
[10] Scientific
[11] Other (please specify) _____

E. What is the geographic location of your primary Tandem system? (Please circle only one answer. Thank you.) <10>

[1] Eastern U.S. (East of the Mississippi)
[2] Western U.S. (West of the Mississippi)
[3] Europe (including U.K. and Scandanavia)
[4] Canada
[5] Other (please specify the country) _____

Please rank your reasons for selecting Tandem. Use a scale of 1 thru 7, where "1" = the least important reason and "7" = the most important reason for selecting Tandem. Thank you.

[____] F. NonStop capability <11>
[____] G. Tandem's reputation <12>
[____] H. Expandability <13>
[____] I. Data Base management software <14>
[____] J. Price/Performance <15>
[____] K. Networking capability <16>
[____] L. Other (please specify) _____ <17>

M. What other vendors were *seriously* evaluated before you chose Tandem? (Please circle as many as apply. Thank you.) <18-35>

[01] NCR
[02] Prime
[03] Data General
[04] Hewlett-Packard
[05] Synapse
[06] August Computers
[07] Burroughs
[08] Honeywell
[09] Stratus
[10] Paragon Systems
[11] IBM
[12] Parallel Computing
[13] DEC
[14] Sequoia
[15] Prime
[16] Computer Consoles
[17] Other (please list) _____

N. If the decision were being made today, would you still select Tandem? <36>

[1] Yes
[2] No

O. Please briefly explain your answer to the previous question. Thank you. <37-38>

Please rate Tandem on each of the items, below, using the following scale:
1 = Poor, 2 = Fair, 3 = Good, 4 = Very Good, 5 = Excellent, 6 = Don't Know/Not Applicable

P. [____] CPU reliability <39>
Q. [____] Peripheral reliability <40>
R. [____] Service quality <41>
S. [____] Service responsiveness <42>
T. [____] Hardware support <43>
U. [____] Software support <44>
V. [____] Applications software <45>
W. [____] Operating Systems software <46>
X. [____] DBMS software <47>
Y. [____] Networking software <48>
Z. [____] Applications languages <49>
AA. [____] Program development tools <50>
BB. [____] Overall satisfaction with Tandem <51>

Do you use or plan to use the following Tandem products?

CC. TRANSFER [1] Yes [2] No <52>

DD. ENCOMPASS [1] Yes [2] No <53>

EE. How important is a 32-bit Tandem processor to you? (Please circle one answer.)

[1] Very important [2] Somewhat important [3] Not important

How many Tandem system CPU's do you have or use?

FF. NonStop I CPU's: _____ <55-56> GG. NonStop II CPU's: _____ <57-58>

HH. Is your system or the one you use composed of TNS I's or TNS II's? The system I use consists of <59>

[1] only NonStop I's [2] only NonStop II's [3] both NonStop I's and NonStop II's

II. If you have or had NonStop I's, did you or do you plan to upgrade to NonStop II's?

[1] Already upgraded to NS II [2] Plan to upgrade to NS II

[3] Started with NS II—never had NonStop <60>

[4] Have decided not to upgrade [5] Haven't made a decison yet

JJ. If you answered "4" or "5" in the last question, please briefly explain why you are not or may not upgrade. Thank you. <61-62>

KK. Will you be expanding your Tandem system in the next 12 months?

[1] Yes, definitely [2] Maybe [3] No, definitely not <63>

LL. If you answered, "Yes" or "Maybe", ("1" or "2" in the preceding question), please describe how you plan to expand your system or installation(s) and give the approximate number of additional pieces of equipment, as listed below, which you plan to acquire (remember, your organization's name is not identified with this response so all information will be kept strictly confidential). "Our general plans are to

_____ <64-65>

MM. Number of additional CPU's to be acquired: _____ <66-67>

NN. Amount of additional Tandem memory: _____ <68-69>

OO. Amount of additional non-Tandem memory _____ <70-71>

PP. Number of additional Tandem disk drives: _____ <71-73>

QQ. Number of additional Tandem terminals: _____ <74-75>

RR. Number of additional non-Tandem terminals: _____ <76-77>

SS. Number of additional Tandem printers: _____ <78-79>

TT. Number of additional non-Tandem printers: _____ <80-81>

(In the event you do not have a specific number, please give your best estimate for the preceding questions. Thank you.)

1. _____ 2. _____

VV. What is your *single* biggest *compliment*, if any, that you have for or about your Tandem or about Tandem Computers Incorporated? (That is, what do you like most?) <86-87>

_____

WW. What is your *single* biggest complaint, if any, that you have about your Tandem or about Tandem Computers Incorporated? (That is, what bothers you the most?) <88-89>

_____

XX. Name the *single* most important suggestion you can make to improve the International Tandem User's Group: <90-91>

_____

***** If you have additional comments, suggestions, or would like to explain an answer in detail, please use the following space. Please feel free to call Aaron Harber at (303) 443-2620 for clarification of any question.

_____

_____

_____

_____

_____

_____

_____

_____

_____

_____

_____

_____

_____

Please return this questionnaire (no later than May 20, 1983) to:

ITUG CONFIDENTIAL TANDEM SURVEY
P. O. Box 2184
Boulder, CO 80306
U.S.A.

Thank you very much for your cooperation. The results will be published for everyone in the ITUG NEWSLETTER. Again, return this questionnaire on or before May 20th. Thanks!

INDUSTRY PERIODICAL
March 1980 RI 417/20
L. Duane Kirkpatrick, First Vice President
San Francisco, (415) 392-7200

RANDOM-ACCESS MONTHLY

This service of the Technology Group is a monthly review of events and trends affecting companies in the technology-related industries. Comments and observations in this publication are oriented toward clients wishing to maintain general awareness of the conditions affecting investment decisions. *Random-Access Monthly* supplements our more detailed analyses of individual companies and industry groups.

The IBM 4331, the smaller of the two 4300 systems introduced in January 1979, addresses a very critical segment of the computer marketplace. It falls into the price/performance range where the overlap between the larger minicomputers and the smaller mainframes is most visible. This analysis compares the IBM 4331 with several of the most important competitive products now on the market. The general conclusions are:

* The 4331 is primarily designed to protect IBM's installed customer base.
* The system is best used in a batch processing environment.
* The 4331 alters total system pricing, decreasing the importance of the central processor.
* Most mainframe competitors have responded with comparable price/performance offerings, but the traditional 20-30% advantage has been erased, which focuses their attention on protecting their own customer bases rather than attacking that of IBM.
* The minicomputer companies continue to offer superior price/performance products, which preserves their advantage in traditional markets. Altered priceing strategies may make it more difficult for several of these companies to prosper in the commercial data processing marketplace in the longer run.

CONTENTS

ADDITIONAL INFORMATION ON COMPANIES MENTIONED IN THIS REPORT IS AVAILABLE ON REQUEST.

DEAN WITTER REYNOLDS INC.

TECHNOLOGY STOCK SUMMARY

Company	3/18/80 Price	52-Week Price Range	Earnings per Share 1978	1979E	1980E	12 Months 12/79E	12/80E	Percent Change	P/E 12/80E	Relative P/E 12/80E	Latest 12 Mos. Revs. ($Mil.)	Shares ($Mil.)	Dividend	Est. 5-Year EPS Growth Rate	Rec.	Analyst
Adv. Micro Devices (3/31)	$38	$46—31	$1.48	$2.90	$3.00	$2.57A	$2.95	15%	12.9	162	187	7.5	Nil	25%	B/H	MCD
Amdahl Corp.	20	45—16	2.81	1.02A	1.50	1.02A	1.50	47	13.3	168	300	16.8	$0.40	20	B/H	LDK
Auto. Data Processing (6/30)	33	41—29	1.84	2.21A	2.55	2.37	2.83	19	11.7	147	409	15.2	0.64	20	B/H	LLB
Beckman Instruments (6/30)	25	33—20	1.35	1.66A	1.85	1.81A	2.30	27	10.9	137	487	19.8	0.32	20	B	AFC
Burroughs Corp. (o)	66	88—65	6.21	7.45A	8.25	7.45A	8.25	11	8.0	101	2,831	41.0	2.60	15	B/H	SCD
Computer Sciences (3/31)	19	24—11	0.98	1.04A	1.55	1.33	1.72	29	11.0	139	416	15.2	Nil	18	B/H	LLB
Control Data (o)	49	63—32	4.88	6.86A	8.20	6.86A	8.20	20	6.0	75	2,273	17.0	0.60	15	B	LLB
Data General (9/30)	60	75—46	4.00	4.82A	5.60	4.82A	5.95	23	10.1	127	540	10.4	Nil	25	B	SCD
Datapoint Corp. (7/31)	102	123—68	4.10	5.84A	7.50	6.80	8.30	22	12.3	155	247	4.4	Nil	25	B/H	MCD
Digital Equipment (6/30)(o)	65	83—51	3.40	4.10A	5.35	4.70A	5.95	27	10.9	138	2,031	44.9	Nil	25	B	SCD
Eastman Kodak (o)	46	67—43	5.59	6.20A	6.00	6.20A	5.85	(6)	7.7	97	8,028	161.4	2.90	12	B/H	EGG
Electronic Data Sys. (6/30)	21	28—19	1.54	1.82A	1.95	1.93	2.08	8	10.9	137	312	13.5	1.00	15	H	LLB
Floating Point Systems (10/31)(m)	18Ask	18—9	0.78	0.53A	1.00	0.53A	1.00	89	18.0	227	33	3.9	Nil	25	H	LLB
Fluke, John, Mfg. (9/30)	24	27—16	1.20	1.59A	1.95	1.64A	2.00	22	12.0	151	110	5.9	Nil	17	H	AFC
Four-Phase Systems	32	49—30	2.58	3.25A	3.25	3.25A	3.25	0	9.8	124	178	5.2	Nil	25	B	MCD
Harris Corp. (6/30)	30	40—25	1.95	2.25A	2.55	2.49A	2.75	10	10.9	138	1,049	30.2	0.72	15	H	AFC
Hewlett-Packard (10/31)(o)	60	71—46	2.63	3.43A	4.50	3.43A	4.50	31	13.3	168	2,520	59.4	0.40	21	B	AFC
Honeywell, Inc. (o)	82	101—64	8.48	10.95A	11.20	10.95A	11.20	2	7.3	92	4,210	21.9	2.60	12	H	SCD
Intel Corp. (m)	64Ask	75—42	2.16	3.59A	4.20	3.59A	4.20	17	15.2	192	663	22.1	Nil	25	B/H	MCD
International Business Machines (o)	59	80—59	5.32	5.16A	5.85	5.16A	5.85	13	10.1	127	22,863	583.4	3.44	14	B/H	LDK
Intersil Inc. (9/30) (m)	22Ask	33—12	1.27	1.65A	1.85	1.70	1.90	12	11.6	146	130	6.1	0.24	25	B/H	MCD
Lanier Business Products (5/31)	23	34—23	1.43	1.88A	2.35	2.05	2.50	22	9.2	116	215	6.1	0.44	17	B/H	EEG
M/A-Com (9/30)	33	42—20	1.10	1.48A	1.85	1.57A	1.95	24	16.9	213	129	5.7	0.24	20	B/H	AFC
Minnesota Mining & Mfg.	48	60—47	4.83	5.59A	6.00	5.59A	6.00	7	8.0	101	5,400	117.2	2.80	13	B	SCD
Motorola Inc. (o)	54	64—38	4.04	4.96A	4.50	4.96A	4.50	(9)	12.0	151	2,714	31.0	1.40	11	H	MCD
National Semiconductor (5/31)(o)	32	41—20	1.72	2.57A	3.50	3.15	3.50	11	9.1	115	719	13.3	Nil	17	H	MCD
NCR Corp. (o)	69	82—57	6.90	8.78A	9.65	8.78A	9.65	10	7.2	90	3,003	26.7	1.60	15	B/H	SCD
Northern Telecom	36	48—34	2.92	3.16A	3.45	3.16A	3.45	9	10.4	132	1,625	33.6	0.69	15	H	AFC
Perkin-Elmer (7/31)(o)	39	47—26	1.77	2.53A	3.35	2.79	3.50	25	11.1	140	869	20.1	0.72	15	B	AFC
Pitney-Bowes (o)	30	37—23	3.32	4.02A	4.50	4.02A	4.50	12	6.7	84	1,025	13.3	1.40	12	B	EGG
Polaroid Corp. (o)	20	42—20	3.60	1.10A	2.30	1.10A	2.30	109	8.7	110	1,362	32.9	1.00	12	B/H	EGG
Prime Computer	28	33—14	0.73	1.44A	1.80	1.44A	1.80	25	15.6	196	153	11.8	Nil	30	B/H	SCD
Reynolds & Reynolds (9/30)(m)	29Ask	34—25	2.56	3.15A	3.50	3.22	3.63	13	8.0	101	190	4.7	1.08	17	B	LLB
ROLM Corp. (6/30)	41	49—29	0.51	1.50A	2.15	1.89A	2.45	30	16.7	211	134	7.6	Nil	30	H	AFC
Scientific-Atlanta (6/30)	37	46—25	1.39	1.76A	2.40	2.02A	2.65	31	14.0	176	156	4.8	0.18	20	B/H	AFC
Shared Medical Systems (n)	17Ask	19—16	0.72	0.90A	1.10	0.90	1.10	22	15.5	195	82	6.0	0.28	23	B/H	LLB
Sperry Corp. (3/31)(o)	50	60—42	5.73	6.75	7.50	6.53A	7.20	10	6.9	88	4,586	40.5	1.56	12	B/H	SCD
Storage Technology (o)	16	23—13	1.14	1.58A	1.90	1.58A	1.90	20	8.4	106	479	24.6	Nil	20	B/H	LLB
Tektronix, Inc. (5/31)	45	64—47	3.19	4.28A	4.85	4.66A	5.20	12	8.7	109	878	18.2	0.84	17	B	AFC
Teradyne, Inc.	36	45—16	2.03	2.32A	3.05	2.32A	3.05	31	11.8	149	122	3.5	Nil	20	B/H	AFC
Texas Instruments (o)	86	108—78	6.15	7.58A	7.40	7.58A	7.40	(2)	11.6	147	3,224	22.8	2.00	15	H	MCD
Tymshare, Inc.	47	58—34	2.31	3.06A	3.70	3.06	3.70	21	12.7	160	183	4.9	Nil	22	B/H	LLB
Xerox Corp. (o)	54	69—54	5.67	6.69A	7.30	6.69A	7.30	9	7.4	93	6,604	84.1	2.80	12	B	EGG
S&P 400	116	134—107	13.17	15.45	14.60	15.45	14.60	(6)	8.3	100	—	—	6.38	11	—	—

Notes: Fiscal years ending on dates other than 12/31 shown in parentheses after company names; for fiscal years ending 3/31, date under prior years. (m) Dean Witter Reynolds Inc. makes a market in this over-the-counter security. (n) Dean Witter Reynolds Inc. and/or its stockholders may in the normal course of business have a position in this over-the-counter security. (o) Listed options are available on this issue. ¹Stock dividend, estimated cash value on ex-distribution date. (s) Adjusted for recent stock dividend or split.

Highlighted Changes from 3/10/80

Eastman Kodak: 1980 estimate raised from $5.85.
National Semiconductor: 1980 estimate raised from $3.40.
Rank Organisation: Dropped from coverage.
Teradyne: 1980 estimate reduced from $3.15.

1. SUMMARY

IBM Defensive Response

The IBM 4331 is viewed as a defensive system designed to protect IBM's customer base; it is less an offensive machine introduced to capture new market share. The impact of the 4331 will be felt most strongly by those competitors which have been most aggressive in attacking the IBM base in the recent past.

The IBM 4331 is significant in the IBM product line because:

- It places IBM's batched-oriented "370-type" system in the low-price market to compete head-on with the high-end minicomputers and competitive small mainframes.
- It will significantly alter the role of the vertically integrated hardware manufacturer, system houses, and the process of distributing small systems to end-users because the relative pricing of hardware and software has changed, and because the computer processing unit (CPU) and memory are only a very small percentage of total system price.
- It overlaps in price with two other new IBM systems, the 8100 and System/38, yielding competition within IBM and alternatives for the customer.
- It confirms new plateaus of technology and price performance that others will have to attain during 1980.
- It will significantly enlarge the highly elastic small-scale, general-purpose computer market, because of lower-price computation, memory, and disk storage.

The competitive and market objectives of the 4331 in decreasing order of importance to IBM are:

- to prevent further erosion of the $5 billion of installed IBM System/370 models 115 through 138 (as well as older System/360 models 20 through 40) by providing a compatible and competitive growth system, primarily for the batch processing market.
- To provide a compatible satellite small batch processor for current large-system users.
- To provide a product for Data Processing Division salesmen to sell to large distributed data processing prospects.
- To provide an alternative upward growth path for System/3 and System/34 users to prevent erosion of this base.
- To provide a low-cost system that utilizes IBM's large operating systems, which changes the selling effort from the raw hardware competitive arena to a "total system" competition.

4331 Market is Batch Processing

The major data processing market, when measured by dollar value of computer usage, has been, and for a while will continue to be, batch data processing.

Large distributed data processing installations, while on the increase, are typically shut down at night and turned into batch processing systems for one-half to two-thirds of their operating time.

Many systems that appear to be on line are actually performing on-line data collection and batch processing functions.

It is in this market, batch data processing, that the 4331 will be the most effective, and where it is primarily aimed by IBM's marketing efforts.

Although the 4331's on-line capabilities have been improved relative to the 370, they are primarily accommodations for the mixed environment of batch and on-line processing, and the 4331 will have negligible impact on such traditional minicomputer markets as:

- Process control
- Scientific computation
- Time sharing
- Interactive small business systems
- Transaction processing.

Vertical Integration is Required

The 4331 is the first major step by IBM toward changing the distribution pattern of computers because it substantially changes the relative values of the system components. IBM has always set the tone for computer industry marketing practices, and the 4331 represents a significant shift in that tone.

IBM is in the systems business: Selling a total problem solution to the end-user. In the past, the solution was primarily a piece of hardware sold on its value rather than on a multiple of its cost. With the rapid decrease in hardware costs, this old formula is not sufficient to provide acceptable revenue growth.

The basic instruction execution of the 4331 is more powerful than a 360/50, the commercial computer workhorse of the 1960's. This is an improvement in price/performance over the 360/50 by a factor of 10 in 15 years. With the 360, software was free, terminals were rare, and the CPU and memory represented about one-half of the total system price.

With the advent of the 4331, the economics have changed, and a typical total systems price is broken down as follows:

CPU Memory	20%
Peripherals	38
Terminals	12
Maintenance[1]	12
Software[1]	18
	100%

[1]Rental and monthly charges are capitalized as if they were a five-year lease, using 2.3% of purchase price per month as the capitalizing factor. This is equivalent to 12.5% interest.

Market success under these conditions will require a vertically integrated manufacturer which has good systems software and controls its maintenance force.

IBM Products Compete for Market

IBM now has three new products in approximately the same purchase price range: $100,000-500,000. These are the 4331, System/38, and the 8100. Each of these offers different features for different markets, but they compete by offering alternatives to the users and, by doing so, tend to block competition's move into this price range. These differences are discussed more fully in Section 3.

4300 Technology Is State of Art

IBM is using new state-of-the-art proprietary technology which promises to produce logic components at a cost of at least one-fourth that of currently available parts. IBM's financial strength allows it to price these components based on their future cost rather than their current cost, which could be as much as ten times higher than expected at full run rate.

The lower manufacturing costs could take at least two years to achieve as IBM builds up production capacity and obtains manufacturing experience in the new technology. In the interim, the higher manufacturing costs will depress IBM profits for the 4331. However, at the end of the product cycle, the reverse condition should prevail.

The net result is that IBM has produced a system that gives substantially better price/performance than other mainframes which have similar cost structures. But the 4331 is still not as cost effective as the major minicomputer vendors.

Price/Performance

Table 1 illustrates the price/performance of the major competitive systems in the 4331 class. The systems chosen were those most likely to be involved in a competitive battle with the 4331, typically in a batch environment. If a directly competitive machine was not available, then a machine in the comparable price range was chosen.

Table 1

RELATIVE SYSTEM RANKINGS
(Price Per MIPS)

Vendor	System	Performance (MIPS)	CPU Price ($K)	System Price ($K)	System Price/MIP ($K)
Tandem	(2-CPU)	1.3	50	185-600	142
Perkin-Elmer	3220	.7	45	135-560	193
Digital Equipment	VAX 11/780	.9	100	210-660	233
Univac	V77-800	.6	40	150-620	250
Data General	C/350	.4	54	150-520	375
Prime	550	.4	70	190-660	475
Prime	450	.3	50	170-620	567
Honeywell	6/53	.2	43	145-500	725
NCR	8555	.3	90	220-600	733
Honeywell	64/350	.3	155	250-650	833
Honeywell	64/320	.2	81	170-550	850
IBM	4331	.2	68	180-575	900
Burroughs	2930	.3	140	290-800	967
Digital Equipment	2020	.3	130	300-900	1000
NCR	8455	.17	50	170-550	1000
Hewlett-Packard	3000 III	.12	110	190-550	1583
Hewlett-Packard	33	.07	55	120-500	1714
Univac	90/30	.15	140	280-799	1867

The performance is measured in millions of instructions per second (MIPS). The price is based on typical systems or a minimum Central Processing Unit (CPU) as defined by the manufacturer. These figures should only be used as a very rough guide to the relative value of the systems.

Table 1 ranks the systems on the relative system price per MIPS. The lower the number in the last column of the table, the more effective the system appears to be. The 4331 ranks number 12 on the list and is over six times as expensive as the best, the Tandem dual-processor system. Perkin-Elmer's 3220 shows in second place in this ranking. DEC's older VAX 11/780 is in third position.

4300 Closes Current PCM Market

Prior to the 4300 announcements, the plug-compatible mainframes (PCM's) were the most serious threat to IBM because the user could switch vendors at negligible cost, without reprogramming. The 4331 and 4341, with their advanced price/performance and new technology, have made it more difficult for PCM's to compete.

The cost of entering the PCM market was relatively small. A vendor could design a cost-competitive machine for an engineering expense of under $1 million. At least one PCM designed a processor for an engineering cost of $250,000. The engineering cost rises dramatically with increasing machine size, especially in the 370/158 or 3031 class, just above the 4341.

Most of the PCM's were also offering a 30-50% price advantage over IBM and now offer no advantage to users over the 4300 series, except that of early delivery. In order to compete, the PCM's have had to introduce new models.

We expect that one or two manufacturers may compete successfully with newly designed systems, the most likely competitors being National Semiconductor and Magnuson. Both these vendors may seek specialized narrow market segments to compete directly against the IBM 4300 systems.

Mainframers Not Heavily Impacted

Mainframers such as Burroughs, Honeywell, and Univac must penetrate the IBM user base if they hope to grow significantly faster than the market. However, at the low end of the market, their growth has been due as much to finding new users as to switching IBM users.

With the announcement of the 4331, these vendors must now look to protect their own installed base, which could become attracted to IBM's better price/performance. However, the cost of reprogramming will keep most users from making a rapid switch to a noncompatible machine.

With the 4331 announcement, all mainframers except Univac were quick to counter with their own new announcements to show their users a path for growth and, therefore, reduce the temptation to switch to IBM. Generally, these new systems are about equal to the 4331 and do not offer the traditional 30% better price/performance offered by non-IBM machines.

Short-term, the mainframers will not be significantly impacted by IBM. But long-term, the impact could be more severe as they are forced to become more vertically integrated, from components to peripherals, in order to maintain their margins. Their growth will also be slowed to the extent that past revenue growth has come from switching IBM users to their own equipment.

Minicomputer Impact Will Be on Future Market Potential

Unlike the mainframe companies, most of the minicomputer companies, Hewlett-Packard being a notable exception, have a better price/performance than the 4331. There is no need for them to rush new announcements. Some of the current mini offerings—such as the Tandem system, Digital Equipment's VAX, or Perkin-Elmer's 3220—offer three to six times better price/performance than the 4331. Normal new-product offerings will improve this ratio even more, so that, on a price/performance basis, the minis can easily compete with IBM.

Over the longer term, the outlook for the minicomputer companies must be viewed in a different context. In order to maintain their fast growth and grow faster than the old electronic data processing (EDP) market, the minis will have to add to their traditional markets by entering the commercial data processing marketplace, where they will compete head-to-head with IBM.

The raw price/performance is only one factor in making a sale in this market. With the competition's performance margin reduced because of IBM's improved product offerings, other factors such as peripherals, software, and services, where IBM has more entrenched strength, will come more fully into play.

The other new IBM products—the System/38, aimed at first-time and on-line interactive users; the 8100, aimed at the distributed processing markets; and the Series/I, aimed at the large end-user and process-control markets—will compete in the province of the minicomputer manufacturers.

IBM has less inherent advantage in these markets but, through sheer force of size and reputation, may capture a sizeable share of an expanded market.

2. THE IBM 4331

An Improved 370/125

The 4331 is primarily an improved version of the 370/115 and 125 models, using new technology. The basic internal machine organization is very similar. However, the data path width on the 4331 is significantly improved to 32 bits, providing much higher performance.

The 4331 is constructed of completely new technology, with 700 circuits on a chip versus two circuits for a 370, and 64k bits of memory on a chip versus 2k bits on the 370. As the cost of a chip is relatively constant, once full production is reached, the 4331, even with more circuits, will be considerably less expensive to build than the replaced 370 models.

4331 Differences

The net result is that the 4331 has an improved price/performance. It offers slightly less performance than a 370/148, at a price less than the 370/115, for a five to six times overall improvement in price/performance.

The 4331 has a MIPS rating of 0.2 and a price of $180,000 for a minimum system, leading to a price of $900,000 per MIPS. The system price includes a capitalized value for monthly maintenance and software charges.

The major new features of the 4331 are:

* New architecture (i.e., instruction set)
* Improved input/output capabilities
* New peripherals
* Reliability, availability, and serviceability
* New software.

Architecture

Architecture, as defined by IBM, is the instruction set and other system functions visible by and available to the user. It is not the internal data flow design of the computer.

Because the 4331 is microprogrammed, the architecture could be easily changed, if desired. However, IBM changes architecture very slowly, because of the large cost of developing system software to support the new architecture and the need for compatibility to preserve the user's software investment.

Initially, IBM's large-system architecture was defined by the System/360 in 1964. Approximately a dozen new instructions were added for the System/370 architecture in the early 1970's. The 4331 introduces another new architecture with about a dozen new instructions.

The new architecture is called Extended Control Program Support (ECPS). The user instruction set is equivalent to the 370, to preserve compatibility for user programs. The privileged instruction set used by the operating system is completely changed, which will protect the new software from unlicensed usage.

The 4331 user has two choices of architecture: The older 370 or the newer 4331 ECPS mode. The decision as to which mode to use is determined at the time the system power is turned on and can be changed for different job environments by changing microcode. It is feasible for IBM to alter the microcode to emulate competitive machines, just as was done for the 360.

Using the new architecture mode in conjunction with an improved version of the old operating system that puts more system functions into microcode, the operating-system overhead is substantially reduced and could be as much as one-third that of a standard 370, considerably enhancing actual system performance.

The primary difference between 370 architecture and the 4331 ECPS architecture is in the use of single-level addressing for main storage and disk I/O. Previously, to perform disk I/O, the operating system had to know the physical track, sector, and byte address of the data. With ECPS, the system only has to know the page number. The disk controller's microcode maps the page address into the correct physical address.

The system programmers and users are no longer concerned about the physical compatibility or layout of the storage media. The I/O controller takes care of the problem. This will allow IBM to change I/O media and develop new compatible disk drives.

Input/Output

The 4331 I/O is considerably improved over the 370, not only in the addressing previously described but also in speed. The total I/O rate of a 4331 is 2.3 million bytes per second, compared to the 1 million bytes per second for the average competitors in this price class.

In a batch environment, tapes and disks are read sequentially, and I/O speed rather than MIPS could be the primary measure of total system performance.

Another aspect of the 4331 I/O, similar to the 370/125, is the use of integrated I/O adapters. Functionally, the adapters include additional microcode used to perform the channel and I/O functions. The adapters degrade system performance because either the adapter microcode or the central processing microcode can be running at one time. However, the adapters are much lower in cost than the alternative of separate controller boxes.

Peripherals

Coupled with the new 4300 announcements were several new peripherals. The most important announcement was the 3770 disk drive with 570 million bytes of storage. Its incremental price of $23,000 per unit, or $40 per million bytes, is the lowest in the industry.

The new drive also has a read/write speed of 1.8 million bytes per second, almost half again as fast as other disk drives. This is a definite improvement for a batch environment where the disk is read sequentially. The 3770 attaches to the 4331 through an integrated disk adapter. Average access time is 20 ms.

As an alternative to the 3770, IBM also offers the 3310 disk drive, the eight-inch rigid disk called "Piccolo," which has spawned a large number of imitators. The 3310 stores 64 million bytes, with an average access time of 27 ms. and a data transfer rate of one megabyte per second.

Both of the new IBM disk drives introduced the notion of fixed-block architecture. All data are recorded in preformatted 512-byte blocks which are directly addressable under the single-level addressing scheme mentioned earlier.

A new printer, less expensive than the older models, complements the peripherals lineup. The 3203 model 5 allows printing at 1,200 lines per minute, further enhancing the batch processing orientation of the 4331.

Reliability, Availability, Serviceability

A major facet of the IBM system is the increased emphasis on reliability and maintainability. This is reflected in the maintenance costs, which are 0.25% of the purchase price per month compared to 0.5-1% for most other manufacturers.

The increase in system reliability is due to:

* A decrease in component count, because of the high-density chips
* Error check and correcting (ECC) bits in memory
* Extensive internal checking
* Extra heavy-duty power supply.

The increase in maintainability is due to:

* Improved microdiagnostics that will pinpoint the failing field replaceable unit (FRU) 80% of the time and reduce the problem to one of three FRU's most of the remaining times.
* Maintenance of the systems by swapping good FRU's for bad, which are repaired at a central site and not at the customer site.
* A separate maintenance console with microdiagnostics that allows the 4331 to be quickly switched on-line to a remote service center for error checking and test.

Software

The software for the 4331 is basically the same system software as used on the System 360/370. The primary operating system for the 4331 will be 370 DOS or the enhanced version, Extended Control Program Support: Virtual Storage Extended (ECPS:VSE), which uses the new 4300 architecture and puts certain common routines into microcode.

Another popular operating system for the 4331 will be the Virtual Machine (VM) operating system. The VM system supports a simple interactive time sharing system and will be the primary operating system for interactive environments.

With the 4331's relatively slow speed and small memory, it cannot adequately handle the larger MVS operating system. Additional software is needed to make the 4331 a high performance interactive system, competitive with the minis. For example, an interactive VAX will support two to three times the number of users as a 4331.

The software prices have been revised, not only for the 4331, but for all IBM systems. If a user elects to use 370 architecture and 370 software, then the 370 prices apply. If the user elects to upgrade the soft-

ware to take advantage of the new 4331 architecture, then there are additional software charges.

This improved software costs over $1000 per month, or 20% of a fully loaded system price. In a more typical system, the extra software cost will be from $300-$800 per month, or 15-20% of the total system price.

Table 2 summarizes the strengths and weaknesses of the 4331.

Table 2

IBM 4331 STRENGTHS AND WEAKNESSES

STRENGTHS

-370 Compatible
-Low price
-Low manufacturing cost
-Made by IBM
-Batch software
-High reliability, availability, serviceability
-Fast I/O
-Low electrical power
-New disk drives

WEAKNESSES

-Poor performance in an interactive environment
-Slow performance compared to minis
-Maximum main memory size of 1 million bytes
-Requires a system programmer

3. RELATION TO THE OTHER IBM PRODUCTS

In price, the 4331 competes with two other IBM products:

* The 8100—Data Processing Division (DPD)
* The System/38—General Systems Division (GSD).

Each of these systems has certain strengths and weaknesses which aim them at certain markets:

* The 8100 is a low-cost system and has two models. One model is priced lower than the 4331, and the higher-priced 8100 model is just about equal in price/performance to the 4331.
* The System/38 also has two models. The lower-priced is about equal in price/performance to the 4331, and the higher is midway in price between the 4331 and 4341.

The characteristics used by IBM to segment the market among these systems are:

* Is the application system design and development controlled and implemented at the local user site or at the central computer site?
* Is a systems programmer required?
* Are the applications batch- or transaction-oriented?

The managerial control of the system is a strong issue to IBM because of the sales territory for the DPD and GSD sales forces. If a company is centrally controlled and is not a first-time user, it is DPD sales territory. If the control is at the local site or is a first-time user, then it is GSD sales territory.

The use of systems programmers is related to the degree of customer sophistication. The smaller and less sophisticated the customer, the less desire or need for an expensive systems programmer. This definition also coincides with the GSD sales territory.

As the systems become larger and more complex, the talents of a systems programmer are required.

All three of the low-cost IBM systems can support a communications environment, such as remote job entry or satellite processors, but the market distinctions are in the managerial control and first-time user issues. Table 3 summarizes the market positions of the three IBM systems.

Table 3

MARKET POSITIONS

	4331	8100	System/38
Batch/Transaction Oriented	Batch	Transaction	Transaction
System Programmer Required	Yes	Yes	Usually No
System Design & Control	Local	Central	Local
Sold By	DPD	DPD	GSD
User	Sophisticated	Sophisticated	Unsophisticated

Omitted from these definitions and, therefore, a weakness in IBM's low-end computer line, are some of the traditional mini markets:

* Interactive computations
* Scientific computations
* Process control
* Transaction processing.

The Series/1 addresses some of these markets but with a much lower price/performance than the 4331.

4. BURROUGHS CORP.

Summary

Burroughs should feel little impact from the 4331 and stands to maintain its position in the small-business-system marketplace.

Burroughs announced two machines to counter the 4331. In its small stand-alone small-business-system line, it announced a dual-processor 1885 configuration. However, this is sufficiently below the 4331 in both price and performance (about $30,000 in price and 20-30% less in performance) that it was not analyzed further.

The most important machine in the 4331 price/performance range is the 2930, which is available as either a uniprocessor or in a multiprocessor configuration with up to four central processors.

Although the 2930 protects Burroughs from IBM on a price/performance comparison, Burroughs is still exposed without a mainframe-type system in the $150,000-200,000 range for a complete system, well under the 2930's minimum system price of $290,000.

The strength of the 2930 will be in the transaction- and communication-oriented markets, where the combined communication front-end processors and the network software will make this a strong contender in the low-price market.

Price/Performance

The single CPU 2930 system offers a MIPS rating of 0.3, or 50% faster than the 4331's MIPS rating of 0.2, and at a price for a minimum complete system of $290,000 leading to a price per MIPS about 7% higher than the 4331. A multiprocessor system lowers the price per MIPS by about 30% per processor, bringing Burroughs to the same price/performance as IBM, but at about twice the price—$340,000 for a minimum system.

The 2930's strength is in a communication environment, where the 4331 is relatively weak. The 2930 uses programmable front-end processors to handle the communications, and their effective MIPS processing rate must be added to the main processor to provide a better comparison. The 2930 also uses separate integrated microprocessors for the disk controllers, further increasing the effective MIPS rate. The 4331 has integrated disk controllers, but they time share the central processor functions. Therefore, the Burroughs 2930 has a much lower price per effective MIPS than the 4331.

The basic CPU and memory of the Burroughs is over twice the price of the 4331—$140,000 versus $68,000 for the latter. Part of this price difference is the disk controller included in the Burroughs price.

Burroughs' main memory, at an incremental price of $15,000 per million bytes, is tied with IBM and lower than any other surveyed machines. The memory uses 16k-bit chips and will probably result in lower profit margins for Burroughs than other parts of the product line.

The logic circuitry of the 2930 is quite advanced, using a chip with 800 gates, compared to IBM's new circuitry containing 700 gates per chip.

Peripherals

Burroughs' disk drives are based on the old removable 3330 philosophy and are priced at $140 per million bytes versus $40 for IBM's new 3770. Because the drives use an older technology, they are slower, with access times on the order of 30 ms. and data rates of 1 million bytes per second versus 20 ms. and 1.8 million bytes per second for the IBM 3770.

Reliability and Serviceability

The Burroughs maintenance and reliability philosophy is similar to IBM's in that the 2930 has a separate maintenance processor. While the IBM maintenance processor can only operate when the system is stopped, the Burroughs maintenance processor operates continuously and simultaneously with the main processor to detect faults.

Part of the monitoring process includes monitoring external power fluctuations and temperature to assure proper system operation. Only the 4331, among the surveyed systems, also offers this feature.

The maintenance philosophy is reflected in the maintenance prices, which are relatively low, 0.7% of the purchase price per month. This is the lowest, next to IBM, and a standoff with NCR among the mainframers.

The high-reliability philosophy, coupled with the multiprocessor configuration, gives Burroughs a very high availability, which is needed for the transaction market.

The transaction-processing market is discussed more completely in the section on Tandem.

Software

The Burroughs software, although more than adequate, is not as extensive as some of the others, and its data base management system, priced at $17,000, is among the most expensive.

On a total system basis, Burroughs has one of the lowest software prices, at 6-7% of the total system price compared to 15-20% for IBM.

5. DATA GENERAL

Summary

Data General is one of the companies that will be somewhat impacted by the 4331, but in an indirect fashion.

As a major minicomputer manufacturer moving from the traditional markets of OEM, scientific, real-time response, etc., into the end-user markets, especially the IBM customer base, Data General now faces a new competitor in IBM. With this new alternative solution available, many users will elect to remain with IBM, especially those who are batch-oriented.

Price/Performance

Data General has been among the most aggressive leaders in technology, and the price and performance Data General offers bear this out.

The Data General system that comes closest to the IBM 4331 in price is the Eclipse 350 (C350). On a total system basis, a C350 runs from $150,000 to $500,000 compared to the IBM 4331 price of $180,000 to $575,000. On a CPU-only basis, the price is comparable to the C350 priced at $54,000 versus the 4331 at $68,000.

The C350 is a new machine, announced in the fall of 1978, and fills what had been a gap in the complete line of business-oriented processors offered by Data General, ranging from the C150 at $80,000 for a complete system to the M600, which can cost $1 million for a complete system.

On a performance basis, the C350 ranks about the middle of the surveyed systems, with a MIPS rating of

0.4, after adjustment for the fact that the C350 is a 16-bit machine. Benchmark testing confirms this rating but, depending upon which benchmarks are used, the C350's performance could range from 0.2 to 0.4.

Buyers of the C350 can adjust the system performance by the type of memory they purchase. A noninterleaved memory is standard, but a two- or four-way interleaved memory can be purchased, raising the performance by 15-20% for a nominal $2,000-6,000 price, depending upon memory size and degree of interleaving. With interleaving, the C350 becomes equivalent to a 32-bit wide memory in speed.

On a price per MIPS basis, the C350 is priced at $375 to $1,300, depending upon system size, less than one-half IBM's $900 to $2,800.

The high end of the Eclipse family is capped by the M600 model which, on a system basis, is about $50,000 more than the C350.

Although the main processors of the C350 and M600 are identical, the M600 produces significantly faster performance through an additional high-speed arithmetic unit, up to eight-way memory interleaving, and a front-end/back-end multiprocessor architecture that increases the effective MIPS rate by adding the speed of the two processors.

Input/Output

The standard I/O speed for the Eclipse series is relatively slow, with a speed of 1.5 million bytes per second. Burst I/O rates are higher.

Peripherals

Data General's peripherals appear adequate for the market segments they currently occupy. However, IBM's 3370 is clearly superior. The largest disk drive offered is 277 million bytes, with an incremental cost of $136 per million bytes versus IBM's $40 per million bytes. For a large, billion-byte data base-oriented system, this would result in a price difference of $100,000 per system.

Data General was one of the last minicomputer companies to react to IBM's new memory pricing stategy by cutting the C350 price from almost $70,000 per megabyte to $27,000. While the latter is significantly higher than IBM's $15,000, it is less than a 2% difference on the total system price.

Reliability and Serviceability

The maintenance and reliability of the Eclipse is standard for a minicomputer with few extra features to aid in maintenance except for microdiagnostics. Data General, however, has given some extra attention to power supply design, a major part of system unreliability, allowing the sytem to operate in a 20% under-voltage environment.

The maintenance prices of the Eclipse are among the highest of the surveyed machines, running about 22-23% of the total system price, or about 0.8-0.9% per month of the hardware price.

Software

On an overall basis, the Data General software is one of the weakest of the surveyed systems, although Data General has strengthened it with several recent improvements.

Data General offers no data base management system, although one may be bought from a third party.

The software offering of Data General is relatively strong in the remote job entry (RJE) environment, allowing the system to act as a remote processor to an IBM system. IBM 3270 emulation and X.25 support are also available.

The COBOL offered compared favorably with other minicomputers but is substantially lacking in features to be able to compare it to a mainframe COBOL.

Data General's new operating system, AOS, first offered in 1977, has many good features and is a substantial improvement over the earlier RDOS.

6. DIGITAL EQUIPMENT CORP.

Summary

Digital Equipment Corp. (DEC), along with Prime, may be one of the two companies most strongly impacted by the IBM 4331. This impact will be evident in a possible lessening of the growth rate rather than in an outright decline in sales, but it will be difficult to measure since DEC and Prime may find it harder than before to penetrate IBM's user base and may put their efforts into other markets, thus continuing at the same growth rate as in the past.

There are three DEC products that compete against the 4331 on a price basis. The first is the DEC 2020, a 36-bit system which DEC considers a mainframe. The second is the VAX 11/780, DEC's 32-bit minicomputer. The third is the older PDP 11/70, a 16-bit minicomputer. While the latter is the most popular of the three DEC systems, it will not be the main future competitor in the commercial marketplace and therefore was not further analyzed.

The VAX was first announced in 1977, and although it is upwards compatible from the older PDP/11 family, in its own native mode it is a completely different system offering full 32-bit instructions and virtual memory capabilities. The VAX 11/780 is the main future thrust of DEC. The single model in this family is expected to be extended in the near future with additional members.

Markets

The DEC System 20 is primarily aimed at the interactive computing market, and about 600 units of all models have been sold. The model 2020, the smallest and most comparable to the 4331, has had about 200 units installed for a total value of about $100 million.

The major market tapped by DEC for the 2020 has been the in-house time sharing market. Many companies have switched from time sharing service bureaus to an in-house 2020 and achieved a reduction in expenses plus an increase in capacity.

Because of its timesharing orientation, it is unlikely that the 2020 will be impacted by the 4331 in this submarket. However, about half the 2020's have been sold as stand-alone batch processors to replace existing systems, including 360/30's and other older IBM equipment. This portion of the 2020 market will be severely impacted by the 4331, if not dried up completely. This is a relatively small dollar volume for DEC, at $50 million per year, and will have negligible impact on DEC's total corporate picture.

The initial market for the VAX is the traditional minicomputer market of scientific and interactive work. Some early users have commented that, with the interactive programming development capabilities of the VAX, software development costs have been one-half to one-third that of other systems. As the VAX software improves, this system will be extended more into the commercial data processing marketplace.

The 4331, while not as fast as the VAX, will slow DEC's entrance into the batch market by providing the IBM user with a fully compatible machine for upward growth.

Price/Performance

The DEC 2020 is a 36-bit computer, aimed at the mainframe market. Its MIPS rating is 0.3, 50% faster than the 4331's 0.2, and its price is also 50% higher than the 4331, running from $300,000 to $900,000 for a complete system.

On a price per MIPS basis, the 2020 is slightly more expensive than the 4331, at $1,000 to $2,600 per MIPS versus the 4331's $900 to $2,800, depending upon system size and configuration.

Recent user benchmark on the 2020 found that in the interactive environment it was considerably more productive than either the 4331 or the VAX 11/780.

The VAX picture is different. The VAX 11/780 has a MIPS rating of 0.9, the highest of any of the systems analyzed except the Tandem dual-processor system, and a system price about 20% higher than the 4331, at $210,000 to $660,000.

The price per MIPS of the VAX was among the lowest of any of the systems analyzed, ranging from $200,000 to $700,000 per MIPS, one-fourth that of the 4331 and ranking third behind the Tandem system and the Perkin-Elmer 3220.

DEC's main memory was originally priced at $40,000 per megabyte, but after the 4331 announcement, DEC reduced this price down to a very competitive $18,000 per megabyte.

Peripherals

As a batch processor, the 2020 has relatively good peripherals, with 176-megabyte-capacity disk drives and a 1.2-million-byte-per-second tape drive. These were the equal of the IBM equipment until the announcement of the new 3770 disk on the 4331.

The VAX has a very high I/O rate for its price, at 8 million bytes per second, three times that of the 4331.

The peripherals for the VAX are DEC's 176-million-byte RP06 disk drive, the same as used on the 2020, and a very slow 45-inch-per-second tape drive. This tape is unsuitable for batch processing and can only be used for disk backup, further underlining the fact that DEC is not seeking to sell the VAX in the batch market.

The RP06 disk drive has the highest price of any of the systems analyzed, with an incremental price of $193 per million bytes, almost five times that of IBM's $40 per million bytes.

DEC's CRT, the VT 100, is an extremely competitive device, and next to IBM's 3270 has proven so popular that there is now up to a two-year wait for a terminal. It is priced at $1,900 for a plain terminal, or $2,200 for a 132-character-per-line version with extra features. This compares to $3,000 for a plain 3270 and more for more sophisticated versions.

Reliability and Serviceability

DEC is one of the few companies that places strong emphasis on reliability and maintainability. The VAX has a separate maintenance computer (an LSI 11) for monitoring the behavior of the system and performing diagnostic work. This emphasis is paid back to the user with low maintenance prices that run about 0.5-0.6% of the purchase price per month, well below the "normal" mini price of 0.9% and the lowest percentage of any of the companies except IBM. On a large system, IBM's maintenance cost is 15% of the total price versus the VAX's 16%, a negligible difference. The next lowest is Tandem, at 18%, while most systems are in the 20-22% range.

The DEC 2020 does not place as much emphasis on reliability as the VAX, and its maintenance prices are about 20% higher, but still under those of a typical minicomputer.

Software

DEC bundles most of it software into the system price, reducing its apparent software costs and raising its apparent hardware costs.

The VAX was designed from the ground up as a completely integrated system; that is, the software and hardware were designed together, not separately. However, like most new systems, it has not yet reached its full potential. In order to bring the machine to the market as quickly as possible, DEC used much of the older PDP/11 software and is gradually phasing in the native mode software to achieve VAX's full potential.

An example is the COBOL compiler. The original PDP/11 version is very slow and, on benchmark tests, was one-half the speed of the DEC 2020 COBOL, a machine one-third the MIPS rating of the VAX. The new COBOL is being phased in by DEC and is approaching the native speed in its capabilities.

The current line-up of native mode software is still relatively weak and includes only an operating system, DECNET network program, COBOL, and FORTRAN compilers. No database management systems exist, and RPG, PL/1, and APL languages are not supported.

In the interim , the full range of PDP 11/70 software is available for use on the VAX, and this is the widest range of software available for any minicomputer.

The DEC 2020 has a relatively full complement of software, with a good COBOL compiler, DECNET networking software, and a database management system.

7. HEWLETT-PACKARD

Summary

Hewlett-Packard (HP) is in an ambiguous position relative to the 4331. On a straight price/performance chart, the HP 3000 is the lowest ranked except for the Univac 90/30, which shouldn't be compared as it is completely noncompetitive. However, from a sales standpoint, the 4331 has had a negligible impact on HP, and the future impact will also be relatively light.

Market

The HP systems are aimed at a completely different market than the 4331. They are primarily aimed at first-time, unsophisticated users for interactive- and transaction-oriented applications. The database software on the HP 3000 is the most sophisticated of any of the minicomputer systems; this greatly simplifies software development, making it easy for new users to program.

HP's user orientation plus its strong IBM-compatible RJE communications software has allowed HP to make a significant penetration into the IBM user base as a distributed processor vendor. Because the HP 3000 is designed to be easy to program by end-users at remote sites, HP will be impacted more by the System/38 than by the 4331 or even the 8100 series.

Price/Performance

HP has two systems that are price competitive with the 4331. They are both relatively new. The larger 3000 Series III was first delivered in late 1978; the smaller model 33 was first delivered in early 1979. The latter uses HP's advanced Silicon-on-Sapphire (SOS) technology, which reduces the CPU logic to three chips.

These systems bracket the 4331 in price: The model 33 is 30% lower than the 4331, and the Series III is about 10% higher.

On the MIPS scale, both HP systems have less performance than the 4331, by at least 40% for the Series III and at least 60% for the Model 33. In absolute speeds, the MIPS ratings are 0.07 and 0.12, respectively, compared to 0.2 MIPS for the 4331.

On a price/performance basis, both HP systems are nearly twice as expensive as the 4331, with price per

MIPS running from $1,500 to $7,000 versus the 4331's $900 to $2,900, depending upon system size and configuration.

It is difficult to make a direct comparison based only on a MIPS analysis because of major internal structure differences between the HP and IBM computers, such as the fact that the 3000 is a 16-bit minicomputer and uses stack architecture for intermediate storage, while the 4331 is a 32-bit machine and uses registers for storage.

The number of terminals supported is a better measure of performance than MIPS for this environment. Because HP's systems are designed as interactive systems, they will support more terminals than a 4331. The model 33 will support 10 to 15 terminals, compared to IBM's 8 to 10, while the Series III should support 20 to 30 terminals.

HP has reacted to IBM's announcement by cutting in half the incremental cost of main memory from over $30,000 per million bytes to $15,000 per million bytes. This has a very small effect on a total HP system price, reducing the total price less than 5% for a typical system. Its primary impact—the ability to say that HP offers the same price as IBM—is psychological.

HP also reacted to the IBM 4331 by dropping CPU prices on both the Model 33 and Model III and bundling some of the software into the hardware price. The net result of these three changes was a total system price reduction on the order of 10%.

Input/Output

For a batch environment where the I/O speed is important, the 3000 may be slower than the 4331. The Model 33 has an I/O speed of one million bytes per second; the Series III has a speed of 2.8 million bytes per second, serving a single device. This compares to the 4331's aggregate data rate of 2.3 million bytes per second and the ability to serve two 1-million-byte devices simultaneously.

Peripherals

Compared to the 4331, HP peripherals are very limited. The largest disk drive offered by HP has a capacity of 120 MB and a restriction of 8 drives, or 1 billion bytes per system. The Model 33 has an even smaller capacity. The 4331 has a capacity of nine billion bytes of on-line storage.

The tape drive offered by HP, at 1600 BPI and 45 IPS speed, is too slow to be used in a batch environment but is adequate for backing up the small HP disk drive. The speed and capacity are inadequate to back up a larger drive, and if HP announces a larger disk, a new tape drive will be needed.

The HP CRT's are very easy to use and very sophisticated. HP was one of the first to offer a movable screen to improve operator productivity and the first to put function keys directly on the CRT, making them easier to use.

Reliability and Serviceability

HP's systems designs are strongly oriented toward high reliability and low maintenance costs, and HP is one of the few vendors that offers remote debugging capabilities for its system. A separate maintenance support processor is used as a diagnostic processor.

HP maintenance prices are less than those of most other manufacturers, accounting for about 20% of the total system price, two to five percentage points lower than most others. When comparing just the hardware, the maintenance costs run about 0.5% of the purchase price per month, well below the typical minicomputer cost of 1% per month and about 20% above IBM's prices for a large system.

Software

HP software is among the strongest offered on any of the minis. It is a completely integrated and proprietary system. The major thrust of the software is interactive- and transaction-oriented. The database and file inquiry languages are among the best.

Neither the HP software nor hardware is oriented toward the batch processing market. The COBOL is among the weakest, supporting only the 1968 version, not the 1974 like most other companies.

HP is also going counter to the industry trend of unbundling software. Its latest announcements bundled in previously separated software. Because most of the software is bundled in, HP's apparent software price is the cheapest, 5% of the total system price.

Bundling reduces the number of options and makes it easier for an unsophisticated user to buy from HP. All of this keeps pointing HP to its ultimate end-user market, the first-time user.

8. HONEYWELL, INC.

Summary

Honeywell, like Univac, can compete against the 4331 from two different levels—minicomputer and mainframe. At the minicomputer level, Honeywell competes with the Level/6 Model 53 and at the mainframe level with the Level/64 Model DPS 320 (Distributed Processing System).

Both the Level/6 and Level/64 now have nearly identical MIPS ratings and system prices, but they differ vastly in applications approach. Neither should be strongly impacted by the 4331 in their typical markets.

Markets

The Level/6 is a typcial 16-bit minicomputer, with good software for a mini but poor for a mainframe. A strong data communication capability and supporting software make it a good distributed data processing product. The software support is not aimed at the first-time user market; it is more complex and harder to use than many other minicomputer offerings. The Level/6 does not have the raw performance of other 16-bit minis in its price class.

The Level/64 Model DPS 320 is also a strong communication-oriented system. Its primary purpose is to protect the Honeywell installed base of H200 class systems. There are over 800 of these systems installed, with an aggregate purchase value of about $1 billion. The Level 64/DPS is expected to extend this base into the Distributed Processing System market.

Both Honeywell systems are strong contenders in the communication environment. The Level/6 is stronger for small transaction-oriented systems; the Level/6 will be more competitive to the 8100 than the 4331.

Price/Performance

On a price per MIPS basis, the 64/320 is lower than the 4331, but does not provide the usual 20-30% price/performance advantage users would expect for switching from IBM. As a result, Honeywell salesmen will be constrained in attempts to capture new customers from the IBM base.

Honeywell also offers a larger version, the 64/350. It is 50% faster, with a MIPS rating of 0.3, at a price increase of 50% for a small system, resulting in a price per MIPS slightly higher than the 64/320.

The 64/350 aims at a gap in the IBM line by offering a system at a price between 4331 and 4341. However, this gap should be plugged by IBM with the expected upgrade announcement to the 4331. (Subsequent to

the detailed comparisons used in this analysis, Honeywell announced the Level 64/DPS-330, which supersedes both the 320 and the 350. This new machine offers approximately 10% greater performance than the DPS-350 and should be a strong contender in this segment of the marketplace.)

The Level/6 minicomputer system model most directly price competitive with the 4331 is the Model 53, announced in early 1978. The basic instruction rate of the Model 53 is slightly faster than the 4331; however, the Model 53 is only a 16-bit machine. To some extent, the 16-bit deficiency is overcome by a two-way interleaved memory which has many of the speed characteristics of a 32-bit wide memory.

On an overall basis, the adjusted MIPS rate is 0.2, equivalent to the 4331. The price per MIPS, on a system basis, is substantially lower than the 4331, at $725,000 per MIPS, because the CPU is less expensive.

Peripherals

On an incremental basis, Honeywell's prices are among the best. The fixed media disk drive is priced at $50 per million bytes versus IBM's $40 per million bytes. Main memory has an incremental price of $20,000 per million bytes versus IBM's $15,000. Both differences, while sizeable in percentage, are negligible in terms of total dollars for the system sizes analyzed.

Reliability and Maintainability

Neither of the Honeywell systems is particularly noteworthy in the maintainability or reliability areas. Only single processor systems are offered, and the maintenance prices are in the middle ground, running about 20% of system prices.

Software

Honeywell has relatively mature software, with both Level 6 and Level 64 software having migrated through several generations.

The Level/6 software is not as first-time-user-oriented as other minicomputer software, and more systems knowledge is required to use it.

Both Level/6 and Level/64 software have strong communication capabilities. The Level/6 offers the TOTAL data base management software, from CINCOM. The COBOL is quite weak when measured against IBM or even against Honeywell's own Level/64 COBOL system.

The Level/64 offers strong data base management systems support languages.

9. NCR CORP.

Summary

NCR responded quickly to the 4331 with a brace of products bracketing the 4331. The new 8455 is 10% under the 4331 in performance, while the new 8555 is 50% faster.

These machines provide a powerful base for NCR to continue its aggressive expansion program. The impact NCR will feel from the 4331 will be relatively mild, as the new systems are not only strong enough to protect the current NCR base but also provide a reasonable alternative to the 4331, especially to those customers who cannot wait for the long lead times on delivery of the 4300's.

From a hardware and software standpoint, NCR is in the strongest position among IBM's mainframe competitors for the following reasons:

- NCR offers a full range of five compatible systems ranging in price from $100,000 to several million dollars.
- NCR's new fixed media disks compete with IBM in price at $50/MB, compared to IBM's $40/MB, a negligible difference even for a large system.
- NCR offers a full multi-processor virtual memory operating system.
- NCR's strong emphasis on reliability and maintainability give it one of the lowest maintenance prices of any mainframe vendor except IBM, at 0.7% of the system price per month versus IBM's 0.5% per month.
- NCR's strong communication emphasis is combined with low-price integrated communication controllers on the smaller machines.
- NCR's main memory price of $20,000/MB is among the lowest of any vendor and is inconsequential in difference on a system basis versus IBM's $15,000/MB.

Markets

Markets for the 8455 or 8555 will be found among NCR's traditional customers, primarily in banking, retailing, and manufacturing.

No announcements were made concurrently with the new processors that would suggest any substantive change in NCR's marketing strategies.

System Description

All the new systems use NCR's new Virtual Resource Operating System (VRX) to provide a full virtual memory capability. The larger of the machines, the 8555, is also available in a multi-processor configuration to provide a high availability system with up to four processors. The 8555 dual processor system would be in the same price and performance class as the 4341.

NCR offers a further advantage in that its low-end Criterion series, which starts in price at $50,000 for a complete system, is upwards compatible with the 8455 and 8555 systems. This provides even the smallest user complete upwards growth visibility and offers the large user the ability to configure distributed processing systems, which are program compatible, with a very wide variety of processors.

Price/Performance

The 8455 hardware is less expensive than the 4331 by over 25%. However, this is not the total price seen by the user. NCR software is among the most expensive of all the companies examined, at approximately 20% of the total system price. When the extra software expense is added to the slightly higher maintenance costs, the total system price of the NCR 8455 is a wash with the IBM 4331.

The 8455 is 15% under the 4331 in performance, with a MIPS rating of 0.17, and about equal in price; the 8555 is 50% faster, with a MIPS rating of 0.3, and only 20% more expensive, with a system price per MIPS of about $733. For multi-CPU configuration, the 8555 would have an even lower MIPS price.

This gives the NCR 8555 the best price/performance ratio of the mainframe products analyzed.

Input/Output

The 8555 has the highest internal bus speed, 36 megabytes per second, of any of the systems analyzed. Despite the high internal bus speed, the I/O speed of the 8455 and 8555 is only two megabytes per second, about equal to the speed of the 4331.

Peripherals

Although NCR is competitive with IBM on disk drive prices, at $50 per million bytes, the operating capabilities are inferior. The average access time of the newest NCR drive announced in 1979 is 35 ms. versus IBM's 20 ms., or 50% slower. This time is most critical in a transaction environment. In a batch environment, the read/write speed could be even more critical, and this is 1.2 million bytes per second for NCR versus 1.8 million bytes per second for IBM's 3770—also a 50% difference.

Software

The software on the NCR is a reasonable configuration of languages and facilities, including a network control language. CINCOM'S TOTAL is used as the data base management system.

10. PERKIN-ELMER

Summary

Perkin-Elmer could stand to benefit the most from the 4331 announcement. Far from being adversely impacted by the announcement, Perkin-Elmer could seize the opportunity offered by IBM's slow delivery schedule to enhance its position significantly.

Perkin-Elmer's new 32-bit system, the 3220, first delivered in the summer of 1979, is similar to the 370 architecture and, therefore, is more directly comparable to the 4331 than most other systems studied.

Perkin-Elmer can deliver a system with 4341 capability at less than a 4331 price.

Perkin-Elmer's even newer 3240 system, announced in October 1979, could fill the price gap between the 4331 and 4341, while offering performance substantially better than that of the 4341.

The ability of Perkin-Elmer to take advantage of the 4331 opportunity will depend upon the speed at which it can react to the marketplace.

Markets

Perkin-Elmer's traditional marketplace has been in the scientific, engineering, simulation and related markets, and the 4331 opportunity would represent a major change in market direction. Management statements have identified a strategic desire to enter the commercial market. However, the specific marketing efforts to date seem tentative.

Perkin-Elmer's thrust into the commercial marketplace could be heavily aided by an IBM-compatible version of the 3220. A minor modification to the hardware and software would allow the 3220 to run any 360 and 370 user-written code without modification. This would include all of the IBM compilers and any code generated by the compilers.

The 3220 is not a PCM in that it cannot run IBM operating systems, nor can it plug into IBM peripherals, but it can run the user object code.

Price/Performance

Typical system prices for the Perkin-Elmer 3220 are $135,000 to $560,000, compared to IBM's $180,000 to $575,000.

On a system price per MIPS, Perkin-Elmer's 3220 is less than $200 per MIPS, well below the 4331 at $900 per MIPS.

Input/Output

The I/O speed of the 3220 matches its internal speed with a rate of eight million bytes per record, or about three times faster than the 4331.

Peripherals

Perkin-Elmer uses the CDC storage module, as do many other minicomputers, for disk drives and has a relatively high price of $164 per million bytes compared to IBM's $40—a very significant price difference. However, CDC is just beginning delivery of its Winchester-style drives that are interface-compatible to the storage module disk drives, which could improve Perkin-Elmer's relative disk drive prices.

Perkin-Elmer also offers 6250 BPI tape drives if required, the same speed as IBM's drives.

Perkin-Elmer makes a very low priced CRT that sells for about $500 in quantity. The use of this CRT substantially reduces the price of a large system.

Reliability and Maintenance

From a reliability and maintainability viewpoint, the 3220 is a middle-of-the-road machine. The main console can be used to initiate diagnostics that can be run to determine gross machine faults, and an error-logging facility keeps track of current and potential hardware problems. On a price basis, the PE 3220 is right in line with other minis, with maintenance prices at 0.9% of the purchase price per month, well above IBM's 0.5%.

Software

Perkin-Elmer offers RJE support to enable the use of the 3220 as a remote job entry system to a 370.

The FORTRAN compiler is one of the most sophisticated available, having an excellent optimizing capability.

The rest of the Perkin-Elmer software is relatively weak. There is no data base management software or networking capability. The COBOL is comparable to other minis but lacks features that would make it acceptable in the mainframe world.

There is no support for PL/1 or the APL languages.

An interactive operating system provides a reasonable level of support for interactive processing, and a transaction oriented monitor will support 30-50 terminals for transaction processing.

11. PRIME COMPUTER

Summary

Prime could be among those most seriously impacted by the 4331 announcement, but in an indirect fashion, by reducing the opportunity to continue to invade the IBM base.

Markets

Prime's initial market thrust was in the scientific and interactive processing market, where Prime's high speed and virtual memory capability provided a significant advantage.

Prime had this market niche to itself until DEC's VAX 11/780 provided a superior product on the higher end and WANG provided virtual memory systems both above and below Prime in price.

- High system availability
- High system reliability
- Relatively low data transmission per transaction
- High data base integrity.

On-line data entry, where data are captured but not processed, constitutes the bulk of the current on-line distributed data processing market. This is not the same as transaction processing. Time sharing and interactive computations as found in engineering or planning type applications are also different from the transaction market.

High availability and high reliability, although related, are not the same system design characteristics. A system with high reliability could have low availabilty or vice versa. A system with hundreds of operators that is unavailable for even short periods, idling the operators, would be very expensive in lost business and operators' time. A financial transaction system that processes a single million-dollar transaction incorrectly would be unacceptable, even though the system had high availability. These applications put a high premium on data base integrity.

Performance Characteristics

The MIPS rating has less value in measuring systems performance in a transaction environment. The important measurement is number of transactions per hour. This measurement is related to the total system throughput, which includes:

- Hardware speed
- Applications software
- Systems software
- Data base software
- Disk and terminal I/O speeds.

It is often impossible to determine how many transactions a system will process until the system is designed and built. It may also be difficult to estimate the number of transactions to be processed because volumes may grow drastically as users become familiar with and learn to use the system.

Tandem, with its unique multi-processor approach, is able to satisfy these multiple requirements:

- Modular growth
- Availability
- Reliability.

If a user misjudges the transaction processing requirements of his/her system, he/she may plug in another central processor at a small incremental cost to provide the additional computing capability. Adding the processors is typically a one- or two-hour job and can be done while the system is running.

The 4331 is available only in a single processor configuration, thus not solving the user's growth problem in nearly as satisfactory a manner as Tandem.

The largest installed Tandem system currently uses 14 processors. It is possible to use 16 processors in one system. Up to 255 systems may be connected in a distributed data processing network. The multiprocessor configuration also provides high availability by automatically sharing the work load among the other processors if one processor stops working.

High reliability is obtained by extensive self-diagnosis within the Tandem system.

Price/Performance

On a MIPS basis, a single Tandem processor is three times faster than a 4331 processor. However, Tandem systems are sold with a minimum of two processors, resulting in over six times the performance of a 4331 at approximately the same price.

On a transaction processing application, the Tandem system should support five to ten times the number of transactions that a 4331 does.

Each program in a Tandem environment can address 128k bytes of memory for program space and an additional 128k bytes for data. Each processor can support 256 programs running at the same time. A single processor can have 2 million bytes of memory, and a Tandem system can have 4 MB. of memory.

On a total system price, the Tandem system and the 4331 are comparable, differing by less than $5,000 for a small system and less than 15% for the large system, but Tandem's greater performance gives it a price per MIPS under $150,000 versus IBM's $900,000 per MIPS.

Peripherals

The IBM 3770 disk drives have considerably faster access time—20 ms. versus 28 ms.—than the CDC storage module used by Tandem. This could be a mitigating influence on the transaction rate.

However, the Tandem "mirrored Database" software can write two copies of the data on different disks. Each disk is assigned to perform half the reads, which can cut average access time to about 14 ms. In many on-line applications, the system performs four to six reads for every write. Therefore, the average access time in the Tandem system is considerably better than the raw hardware would suggest.

The IBM 3770 disk drive, incrementally, is about one-third the price of the Tandem disk, at $40 per million bytes versus $130 for Tandem. For a large system, this is a $100,000 difference. However, if Tandem upgrades to the CDC Winchester, which is interface compatible, the difference would be reduced.

Tandem uses the Lear Siegler CRT, one of the lowest-priced CRT terminals available. It is priced at one-half the IBM 3270 CRT, reducing the cost of a large Tandem system.

The I/O channel on the Tandem system is capable of transferring data at 4 megabytes per second. That is nearly twice as fast as the 4331 and gives further advantage to Tandem's total throughput.

Reliability and Maintainability

The 4331 design stresses reliable components, so that the single processor configuration will provide high availability. The design philosophy includes extensive testing and diagnostics to analyze and locate faults and facilitate quick repair.

The Tandem approach is completely different. Tandem has designed in extensive checking to make sure the processor is operating reliably, but once a processor fails, the workload is automatically switched to another processor. This approach reduces the component count needed for a single processor. Maintenance is performed by swapping a good processor for a defective processor, while the system is running.

The 4331 architecture, as currently defined, precludes a multiprocessor design, but even if the archi-

tecture expands to allow multiprocessing, major new software would be needed to allow more than two processors to be tied together. IBM's most sophisticated software provides for load sharing, but does not handle the dynamic case of taking the workload from a failing processor and giving it to another processor to continue where the first one left off.

The maintenance price for an IBM 4331 is half that of a Tandem system. The dollar difference, at $500 to $1,000 per month, depending on system size, translates to $20,000 to $40,000 in equivalent purchase price value in favor of IBM.

Software

A major factor in total system cost is the cost of developintg user programs. Most Tandem systems will be installed for new applications, often where the user is converting from a batch to a transaction processing environment.

Although the Tandem software is not as extensive as IBM's, it is sufficient for the applications most often approached.

Tandem offers all the major languages, such as COBOL, FORTRAN, and Data Base Management, inquiry programs, report generators, screen formatters, and networking software.

More importantly, the system is oriented toward on-line program development and transaction processing. Many facilities that would have to be specially coded for a 4331 are already available in Tandem's system. Tandem users report that program development, using Tandem's proprietary language, TAL, often takes one-third or less the time necessary on an IBM system. This could be a very significant dollar saving for a Tandem system user.

The Tandem software has considerably less overhead than the IBM software and provides substantially greater throughput in the transaction environment.

13. UNIVAC

Summary

Univac could be heavily impacted by IBM's 4331 unless it introduces new equipment within the next six to twelve months.

Like Honeywell, Univac competes in two areas—the mainframe market and the minicomputer market.

The V77-800, the largest of the Univac minicomputers, is the most likely competitor in the minicomputer market. Overall, the V77-800 is a highly competitive machine for a terminal-oriented environment. It will have less relevance in the batch market, which is the 4331 stronghold.

The 90/30 is the most direct competitor in the mainframe market.

The 90/30

The 90/30 is Univac's main small-end mainframe. Although it was competitive with the IBM/370, it is not competitive with the 4331, as it was introduced in 1974. The 90/30 is at least 30% slower and almost twice as expensive, for a price performance ratio two-and-one-half times worse than the 4331.

The main appeal of the 90/30 is to the approximately 200 users of the smaller 90/25 looking for a compatible growth machine and to the 2,000 customers of the much older 9200 and 9300. More importantly, there are approximately 2,000 current 90/30 users, with installed equipment valued at a purchase price of about $1 billion. These users could consider moving either to the 4331, for an increase in performance at one-

half the monthly price of the 90/30, or to the 4341, which would give a five- to six-times increase in performance for the same price.

To date, probably few, if any, Univac users have ordered the 4331. Most orders taken by IBM have been for customer upgrades. But the second round of competitive marketing could begin within the next year, and it is in that environment that the Univac sales force will definitely need better price/performance products in this area.

It is important to note that a substantial portion of the 90/30 customer base formerly consisted of IBM customers using the older 360/20's and 360/30's. Without a competitive offering, Univac might find some customers returning to the IBM fold.

The V77-800

Poor as Univac's mainframe competitive position is, the minicomputer picture is much better.

Univac purchased Varian Data Machines, a pioneer and one-time leader in the microprogrammed minicomputer market, in June of 1977. Varian was merged into Univac and is now managed under the same part of the corporate structure as the BC/7 small business computer.

In 1976 Varian announced the V77 family. Univac's first announcement since acquiring Varian occurred in 1978, with the V77-Model 800, which predated the 4331 announcement by a month.

The market for the V77-800 is the typical minicomputer market, with a heavy emphasis on communication-oriented applications. More recent announcements, such as an extensive manufacturing software package, signal Univac's intention to use the V77 in the end-user market.

The V77-800 should feel a negligible impact from the 4331, except for the typical minicomputer problem of the IBM customer who has no place to go and might have chosen Univac before the 4331 announcement.

Price/Performance

The V77-800 stands fourth in price/performance of all the surveyed systems, three-and-one-half times better than the 4331. The hardware is about $30,000 less for the small system; it is somewhat more expensive on the large system because of the costly Univac disk.

The raw MIPS performance of the V77-800 is almost three times that of the 4331, after adjustment for the 16-bit word of the Univac. Direct MIPS comparisons must be weighted by the fact that, like most minis, the V77-800 is a 16-bit machine, while the 4331 is a 32-bit machine, requiring twice as many memory fetches for the same amount of data.

The main memory speed for the V77-800, of 600 ns. for a 16-bit word, is considerably enhanced by 2k bytes of 150 ns. cache memory. With an expected hit ratio of 93-95%, the effective overall memory speed is 200 ns. Even after adjustment for the 16-bit wide memory, the V77-800's effective memory speed is two-and-one-half times as fast as that of the 4331.

The writeable control store used to microprogram the V77-800 is 64 bits wide and has a cycle time of 165 ns.

A major advantage of the Univac system is the ability to attach up to four systems together to provide a high-availability system. The software support for this configuration is relatively limited at the present time.

The CPU and main memory price of the V77-800 is almost one-half that of the 4331 and is the major reason for its price advantage. Main memory is among the cheapest of the surveyed systems, at $24,000 per megabyte. This is a negligible price difference in a one-million-byte system that sells for $400,000.

Input/Output

The I/O speed of the V77-800 is fast for a minicomputer, at two million bytes per second, close to the speed of the 4331's 2.3 million bytes per second.

Peripherals

Univac disks are relatively expensive at $125/MB. compared to the $40/MB. for the 4331. For the large system, with one billion bytes of disk, this could cause a $100,000 price difference.

Reliability and Maintainability

Univac's maintenance prices are in line with minicomputer standards of 1% of the purchase price per month, compared to the 0.25% on the 4331. On a total system basis, this is $5,000 to 10,000 per year, or $20,000 to 40,000 capitalized price to the user on an equivalent purchased system.

Software

The V77-800 software is relatively inexpensive, ranging from $7,500 for its small system to about $50,000 for all the features available on a large system.

The software coverage of the V77-800 is among the best of the minicomputers. It includes all the standard programming languages, RJE facilities, use of CINCOM's TOTAL, rated one of the best for data base management, and a networking facility.

The COBOL is competitive in the minicomputer arena but weak in an IBM competition, because it lacks features typically used in a batch processing environment.

Two different operating systems, VORTEX and SUMMIT, provide for tailoring the system to an interactive or transaction processing environment.

This report was prepared with assistance from a computer industry consultant.

Table 4

CURRENCY VALUES
(Average United States Dollar Equivalents and the EDP Weighted United States Dollar)

	Pound (United Kingdom)	Franc (France)	Mark (West Germany)	Lira (Italy)	Yen (Japan)	Dollar (Canada)	EDP Weighted Dollar* (United States)
1972	2.5008	.1983	.3136	.0017	.0033	1.0094	.902
1973	2.4510	.2254	.3776	.0017	.0037	0.9998	.822
1974	2.3403	.2081	.3872	.0015	.0034	1.0226	.853
1975	2.2216	.2335	.4073	.0015	.0034	0.9830	.833
1976	1.8048	.2094	.3874	.0012	.0034	1.0141	.904
1977	1.7509	.2039	.4317	.0011	.0036	0.9412	.905
1978	1.9337	.2230	.5050	.0012	.0049	0.8777	.817
First Quarter	1.9203	.2130	.4883	.0012	.0043	0.8942	.842
Second Quarter	1.8428	.2192	.4816	.0012	.0046	0.8886	.838
Third Quarter	1.9502	.2210	.5032	.0012	.0053	0.8757	.805
Fourth Quarter	2.0213	.2389	.5470	.0012	.0053	0.8523	.782
1979							
First Quarter	2.0280	.2334	.5369	.0012	.0049	0.8449	.799
Second Quarter	2.1011	.2295	.5311	.0012	.0045	0.8646	.803
Third Quarter	2.2353	.2380	.5554	.0012	.0045	0.8583	.785
Fourth Quarter	2.1645	.2444	.5718	.0012	.0041	0.8517	.791
1980							
January	2.2668	.2454	.5741	.0012	.0042	0.8637	.781
February	2.2625	.2401	.5624	.0012	.0040	0.8737	.792
Year-to-Year Change	11.7%	2.4%	3.9%	--	(18.4)%	4.3%	(0.8)%

*Index value weighted by six-country 1977 IBM revenues (May 1970 = 1.00).
Sources: *Federal Reserve Bulletin*, *Wall Street Journal*, Estimates of Dean Witter Reynolds Inc.

Table 5

OFFICE EQUIPMENT SHIPMENTS
(Dollar Amounts in Millions)

	Shipments	Percent Change	New Orders	Percent Change	Ending Backlogs	Percent Change	Ending Inventories	Percent Change	Orders/ Shipments	Backlog/ Inventories
1963	$ 3,503	8%	$ 3,869	9%	$1,492	33%	$ 833	3%	1.10	1.79
1964	3,859	10	4,395	14	2,028	86	1,036	24	1.14	1.96
1965	4,265	11	4,936	12	2,699	33	1,192	15	1.16	2.26
1966	5,963	40	6,754	37	3,490	29	1,335	12	1.13	2.61
1967	5,732	(4)	6,195	(8)	3,953	13	1,480	11	1.08	2.67
1968	6,215	8	6,608	7	4,346	10	1,661	12	1.06	2.62
1969	7,420	19	8,203	24	5,129	18	2,086	26	1.11	2.46
1970	7,693	4	7,258	(12)	4,694	(8)	2,288	10	0.94	2.05
1971	6,909	(10)	6,732	(7)	4,517	(4)	2,110	(8)	0.97	2.14
1972	8,605	25	9,161	36	5,073	12	2,084	(1)	1.06	2.43
1973	10,054	17	10,397	13	5,416	7	2,647	27	1.03	2.05
1974	12,179	21	12,979	25	6,216	15	3,325	26	1.07	1.87
1975	11,528	(5)	11,607	(11)	6,295	1	3,035	(9)	1.01	2.07
1976	13,724	19	13,191	14	5,762	(8)	3,088	2	0.96	1.87
1977	15,791	15	16,556	26	6,527	13	3,547	15	1.05	1.84
1978	18,571	18	19,382	17	7,431	14	4,449	25	1.04	1.67
1978										
First Quarter	3,733	6	3,913	5	6,707	12	3,786	19	1.05	1.77
Second Quarter	4,418	15	4,932	14	7,221	12	4,093	25	1.12	1.76
Third Quarter	4,640	19	4,791	20	7,374	13	4,254	29	1.03	1.73
Fourth Quarter	5,780	26	5,746	28	7,341	12	4,449	25	0.99	1.65
1979										
First Quarter	5,037	35	6,758	73	9,062	35	4,802	27	1.34	1.89
Second Quarter	5,249	19	5,513	12	9,148	27	5,270	29	1.05	1.74
Third Quarter	5,352	15	4,828	1	9,000	22	5,433	28	0.90	1.66
October	1,990	18	1,813	0	7,942	9	5,862	35	0.91	1.35
November	2,007	15	2,401	31	8,337	13	5,996	38	1.20	1.39
December	2,753	22	2,561	28	8,145	14	5,981	34	0.93	1.36
January (Prelim)	1,580	23	1,617	12	8,182	12	6,159	37	1.02	1.33

Source: U.S. Department of Commerce; compiled by Dean Witter Reynolds Inc. (These unpublished data are furnished by the Bureau of the Census as reported by companies in the industrial categories covered by SIC code 357, Office Computing and Accounting Machines.)

Table 6

UNITED STATES IMPORTS AND EXPORTS*
OF COMPUTERS AND RELATED EQUIPMENT

	United Kingdom		France		West Germany		Italy		Japan		Total World	
	Imports	Exports	Imports	Exports	Imports	Exports	Imports	Exports	Imports	Exports	Imports	Exports
Trade Value (Millions)												
1974	$14	$281	$37	$234	$42	$260	$11	$ 82	$52	$238	$381	$2,152
1975	22	266	41	222	36	269	21	80	60	183	404	2,183
1976	20	302	60	271	51	290	35	85	149	233	630	2,512
1977	34	420	93	315	47	350	45	102	150	271	730	3,181
1978	52	567	95	395	73	515	43	134	245	332	961	4,214
1979												
First Quarter	15	184	30	121	26	153	14	37	86	100	293	1,259
Second Quarter	22	178	35	118	25	172	6	43	81	98	345	1,285
Third Quarter	24	176	32	131	22	195	7	45	68	111	296	1,367
October	14	59	7	42	5	66	3	13	22	41	110	458
November	14	72	8	49	5	70	3	20	24	34	108	504
December	11	79	6	56	8	79	2	18	17	51	113	580
Year to date	100	748	118	537	91	735	35	176	298	435	1,265	5,453
Import/Export Ratio												
1974	0.05		0.16		0.16		0.13		0.22		0.18	
1975	0.08		0.18		0.13		0.26		0.33		0.19	
1976	0.07		0.22		0.18		0.41		0.64		0.25	
1977	0.06		0.30		0.13		0.33		0.55		0.30	
1978	0.09		0.24		0.14		0.32		0.74		0.23	
1979 (Year to date)	0.13		0.22		0.12		0.20		0.69		0.23	
Imports (Percent Increase)												
1974	0		54%		40%		10%		0		11%	
1975	57%		11		(14)		91		15%		6	
1976	(9)		46		42		67		148		56	
1977	70		55		(08)		29		0.6		16	
1978	53		2		55		(4)		63		32	
1979 (Year to date)	92		24		25		(19)		22		32	
Exports (Percent Increase)												
1974		33%		21%		24%		55%		17%		28%
1975		(5)		(5)		3		(2)		(23)		1
1976		14		22		8		6		27		15
1977		39		16		21		20		16		27
1978		35		25		47		31		23		32
1979 (Year to date)		32		36		43		31		31		29

*Imports - shipments to U.S.; exports - shipments from U.S.
Sources: U.S. Department of Commerce. Compiled by Dean Witter Reynolds Inc.

DEAN WITTER REYNOLDS INC.

RESEARCH DEPARTMENT

RESEARCH MANAGEMENT

Lee H. Idleman, CFA, Senior Vice President	Director of Research	(212) 437-5011
Richard D. Rippe, Ph.D., First Vice President	Assistant Director of Research	(212) 437-5092

ECONOMIC/STRATEGY GROUP

Arnold X. Moskowitz, Vice President	Economics	(212) 437-5025
Martin S. Feldstein, Ph.D., President, NBER	Economics Advisor	

BASIC INDUSTRIES GROUP

Craig W. Fanning, Vice President	Electrical Equipment	(212) 437-5075
Timothy L. Jones, Vice President	Building and Related	(212) 437-5068
James R. McCann, Vice President	Machinery	(212) 437-5087
J. Clarence Morrison, Vice President	Metals	(212) 437-5055

CONSUMER GROUP

Lawrence Adelman, CFA, Vice President	Packaged Foods, Commodity Processors, Beverages	(212) 437-5090
Jeffrey B. Edelman, Vice President	Retail Trade, Textiles and Apparel	(212) 437-5114
Emma W. Hill	Cosmetics	(212) 437-5141
Frederick W. Kopf, CFA, Vice President	Retail Trade	(212) 437-5048
Joseph S. Phillippi, Vice President	Automotive Parts, Tire and Rubber	(212) 437-5073
Victor J. Raskin, CFA, First Vice President	Lodging, Restaurants, Mobile Homes	(212) 437-5015
Michael A. Ward, CFA, First Vice President	Auto Manufacturers	(212) 437-5110
John W. Weiss, CFA	Packaged Foods, Commodity Processors, Beverages	(212) 437-5034

ENERGY GROUP

Eugene L. Nowak, First Vice President	Domestic Petroleum	(212) 437-5086
Joel Price, CFA, First Vice President	Coal	(212) 437-5165

FINANCIAL GROUP

Lawrence W. Cohn, CFA, Vice President	Banking	(212) 437-5640
Jerry Gitt, CPA, First Vice President	Savings and Loans, Western Banks	(415) 392-7200

HEALTH-CARE GROUP

David B. Lippman, CFA, Vice President	Drugs	(212) 437-5052

TECHNOLOGY GROUP

L. Duane Kirkpatrick, First Vice President	Technology Group Leader	(415) 392-7200
Linden L. Berkheimer	Computer Services and Peripherals	(415) 392-7200
Adam F. Cuhney, Vice President	Instruments, Telecommunications	(415) 392-7200
Stephen C. Dube, CFA, Vice President	Computers	(212) 437-5007
Mark C. Dyniewicz, Vice President	Semiconductors, Distributed Data Processing	(415) 392-7200
Eugene G. Glazer, CFA, Vice President	Office Systems, Photography	(212) 437-5112
Aharon Orlansky, Vice President	Technology – Europe	Paris (011331) 261-8450

UTILITY GROUP

Kenneth Hollister, CFA, First Vice President	Electric, Natural-Gas, Telephone Utilities	(212) 437-5013
Joseph Garcia, CFA	Electric, Natural-Gas Utilities	(212) 437-5096
Mary Ellen Guzewicz	Junior Analyst	(212) 437-5499

GENERALISTS

Fred Anschel	Special Situations	(212) 437-5100
William G. Becker	Corporate Client Information Service	(212) 437-5046
Leonard G. Collins, CFA, Vice President	Special Situations	(415) 392-7200
E. Coe Kerr	Corporate Client Information Service	(212) 437-5094
David E. J. Maxwell, CFA	Special Situations	(415) 392-7200
Charles J. Nicolosi, Vice President	Special Situations	(212) 437-5017
Charles W. Shaeffer	Special Situations	(212) 437-5050
William D. Tichy	Special Situations	(415) 392-7200

TECHNICAL ANALYSIS/COMPARE

Frank D. Korth, CFA, Vice President	Director of COMPARE	(212) 437-5118
John F. Connallon	COMPARE Analyst	(212) 437-5082
Donald W. Kimsey, Assistant Vice President	Technical Analyst	(212) 437-5163
Ross R. Ritchey	COMPARE Analyst	(212) 437-5080

PROFESSIONAL SUPPORT SERVICES

Daniel N. Reissig, Assistant Vice President	Editor-in-Chief	(212) 437-5077
Julia B. Shibla, Assistant Vice President	Chief Librarian	(212) 437-3732
Beth Morris	Research Department Librarian	(212) 437-5106

August 22, 1984

To: Staff
FYI

To: Distribution

From: Steve Schmidt

Subject: IDC report on our low-end market emphasis

****

Attached is a report prepared by International Data Corporation about today's announcement of our new pricing.

This is a reuslt of a pre-announcement meeting with Aaron Goldberg of IDC, myself and Gina Burr. Other pre-announcement meetings Gina arranged included Peter Lowber of the Yankee Group, Jeff Beeler of Computerworld, Skip Bushee of InfoCorp, and Omri Serlin of the Fault Tolerant newsletter.

In addition, Information Systems News picked up rumors about this announcement early last week. We arranged an interview with their reporter; the resulting story will break this week. It will be interesting to see whether the story is positive, considering the reporter's attitude last week. We'll keep you posted.

IDC is a leading market research and information management consulting organization. IDC will distribute it to its customers, typically large end-user organizations.

Although this report won't guarantee that press coverage will be this positive, it will help.

--- END ---

INTERNATIONAL DATA CORPORATION

1333 Lawrence Expressway, Suite 220, Santa Clara, CA 95051 • (408) 244-5710

Tandem Broadens Product Line

Tandem has recently revamped their product line (announced August 20, 1984) to allow the large user organization to take better advantage of the inherent capabilities of their systems. Although there are no major new product introductions, the company has lowered pricing and developed packaged system solutions to better meet the needs of the user. The actual changes in pricing and configuration follow closely the move toward development of the corporate information architecture now going on at many large user sites.

The particulars of the announcement include,

- Two bundled NonStop 1+ systems, a 2 processor and 4 processor configuration, without mirrored disks and controllers that sell for 32 percent and 12 percent less than previous comparable configurations.

- A bundled NonStop II system that is base priced at $129,500, a 24 percent reduction.

- The development of a trade-in program that allows users to move from the NonStop 1+ to a NonStop II or TXP, or from a NonStop II to TXP and receive a credit from Tandem for 60 to 80 percent of the current list price.

On the face of things it may appear that these announcements are little more than price restructuring, however that is a very elementary viewpoint. To go along with these single quantity price reductions is a very aggressive discount schedule. For example, the smaller NonStop 1+ in a quantity of only 20 systems is priced at $55,000. This has important implications for system/site selection choices.

What this means to the user is slightly more difficult to explain. Transaction processing applications are one of the crucial areas of new application development in many large organizations. At the present time, many of these applications are implemented on systems that are designed to support this transaction environment as well as other systems tortured into supporting on-line environments. The problem arises when the actual geographic sites that are part of the transaction environment are not of sufficient size to justify the installation of a $100,000 and above transaction processing system.

The result has been that the transaction systems are used where they can be justified, and smaller minicomputers are used where a transaction system's price used to be excessive. Because of this, non-compatible software environments exist, line costs reach larger proportions due to inefficient concentration, and excessive system overhead happens due to all the "conversion" of data streams. The Tandem announcement helps to limit this problem for transaction applications.

The new pricing of the Tandem system allows for the installation of a totally compatible system at sites where the previous choice was a minicomputer of the Series/1 class (larger 4955 and 4956). This may sound a bit hard to believe, and in fact is at first glance. Yet looking at the quantity price of $55,000 which includes peripheral devices and software for applications and operating system, the prices are quite equal. For comparison, a 4956 model Series/1 configuration with software and similar hardware would list in quantity 1 between $50,000 and $60,000. Of course, IBM does offer volume discounts, but through more efficient application environments with single vendor solutions, the cost difference is minimal. Also, the upgrade policy announced allows users to be assured of a very high residual value.

For users, this announcement helps to prevent problems that arise in transaction applications where three environments have necessarily existed, the mainframe, transaction application system, and small mini at the geographic site. In being able to put transaction computing systems at more of the individual sites, coherent information from the transaction system as a whole is more valid, and immediacy of the data base is improved.

In lowering the cost of their entry level systems, Tandem has made it easier for users to develop a transaction application without worrying if all sites that need systems will be able to cost justify them. The product line from Tandem now stretches from price points equal to the larger Series/1 all the way to 308X class machines. This increased flexibility in the choice of specific equipment at each individual site could make transaction applications feasible that previously included too many smaller sites where transaction systems were not affordable.

Tandem NonStop Series
Computer Systems

■ PROFILE

Operating System • GUARDIAN, multiprocessor, multiprogramming, network-oriented, "fault-tolerant" • up to 255 concurrent tasks per processor

Database • ENCOMPASS; relational, distributed; includes transaction processing facilities

Communications/Networks • GUARDIAN/EXPAND operating system supports networks of multiprocessor nodes; End-to-End packet protocol • additional support includes X.25 virtual circuit handling and IBM 3270/3271 emulation

Languages • FORTRAN, COBOL, MUMPS, and TAL (Tandem Application Language)

Applications Development • substantial facilities for developing transaction processing and distributed database applications

Models • 16-bit NonStop and NonStop II use the same software and peripherals; NonStop II features increased physical/virtual memory address space, 32-bit-wide access path to main memory, loadable control store, and support for local/remote system control/diagnostics

Memory • up to 2M bytes per processor

Disk • 64M to 2 billion bytes per disk controller; 32 controllers per processor maximum

Terminals/Workstations • 1 hard copy or display terminal per port; 32 ports per asynchronous controller; 32 controllers per processor maximum

Printers • 1 per printer controller; 32 controllers per processor maximum; 200 cps to 900 lpm

First Delivery • May 1976

TYPICAL PRICE RANGES

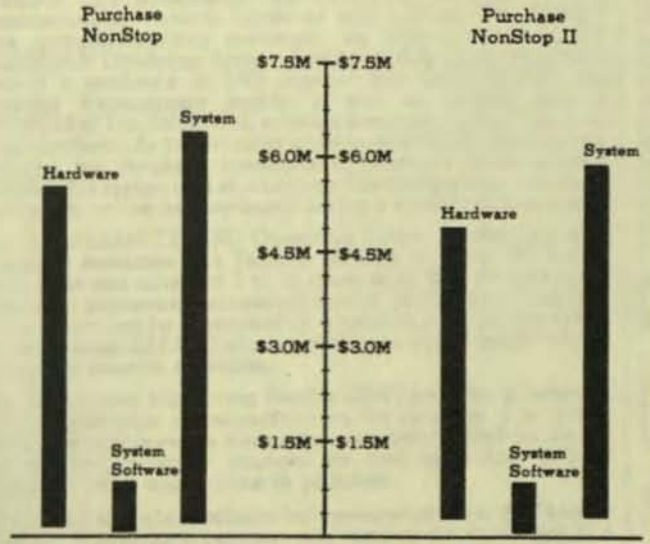

Hardware pricing includes annualized maintenance ($/mo x 60). Software pricing includes annualized service ($/mo x 60). Systems not available from Tandem on rental or lease arrangement.

Systems Delivered • 1817 processors as of May 1981; average of 2.5 processors per system

Comparable Systems • no other systems offer the same degree of internal fault tolerant operation; several vendors, such as Digital Equipment, Perkin Elmer, and SEL, offer dual ported disk controllers and shared memory for a distributed environment

Vendor • Tandem Computers, Inc.; 19333 Vallco Parkway, Cupertino, CA 95014 • 408-725-6000

Distribution • through direct sales and service offices in major metropolitan areas of USA, Canada, Europe, United Kingdom, and Japan; through distributors in Mexico, South America, Australia, and Finland

■ ANALYSIS

The Tandem NonStop and NonStop II Systems are the result of a dedicated effort to fill a need for a "no-fail" system in the computing industry. Previously, vendors used duplicate systems working in parallel with rudimentary intercommunication for switching back and forth when a failure occurred. Few applications other than telephone, airline traffic control, and airline reservation systems truly required non-stop operation. Thus, the normal computer mode has been single processor operation with dual or multiprocessor operation the exception. When a company begins to use a computer to control many aspects of its business online, the consequences of the computer being down for any length of time (even if it happens rarely) become more serious. This leads experts to predict that multiprocessor, non-stop operation will become the normal mode of operation for online systems.

Tandem Computers, Inc. is the only company that has addressed fault-tolerant transaction processing exclusively. The hardware, software, and systems organization are all designed for non-stop operation. This applies not only for continuation of processing in the case of component failure, but also for expanding processor power, doing periodic maintenance, and increasing network size. The Tandem implementation of a non-stop system concept is

Tandem NonStop Series
Computer Systems

comforting to users expanding the active role of computers in business activities. Other vendors will have a hard time catching up with Tandem. Mainframe vendors have a large hardware and software investment in single processor systems. The architecture of such hardware and software will require major changes to support non-stop operation effectively. But, change they must. Customers will demand it because Tandem has shown it can be done.

The Tandem product line currently consists of the original NonStop and the new NonStop II. The two systems are based on different 16-bit processors, but are almost totally software-compatible and use the same peripherals and ancillary equipment. Compared to the NonStop, the NonStop II offers increased physical and virtual memory address space, a 32-bit-wide access path to main memory, a loadable control store, and support for local/remote system control and diagnostics. Configurability of the systems is the same, although the NonStop II will support a larger main memory when Tandem begins to produce memory modules based on 64K-bit chips, overcoming current physical packaging limitations.

Tandem's GUARDIAN operating system can support from 2 to 16 processors of the same type and runs in each processor. At regular intervals each processor sends out status information to all other processors. If a processor misses its regular checkpoint, it is assumed to be down and a backup processor takes over. For each process running on a Tandem system, one GUARDIAN system executes it, while another GUARDIAN system maintains checkpoint data so it can run the process if the active GUARDIAN system fails.

The GUARDIAN operating system has reduced the problem of writing application programs to run in a multiprogramming, multitasking, multiprocessor environment by treating every program as a named process consisting of program code and program data. The only thing a process can do is execute its own code, access its own data, and communicate with other processors via messages. It cannot change its own code; the processor has no instructions to write into code areas. Thus, all application programs produce re-entrant code which can be used by multiple applications. Every processor in a Tandem configuration has a process directory of process names and their locations. The directory indicates the location of the backup process as well as the active process. In a multiprocessor environment, one of the tasks of GUARDIAN is to send messages to update all of the process directories in the system when a backup process becomes active.

Tandem has extended the non-stop concept to its communications network hardware and software. The ENVOY data communications manager, an integral part of the GUARDIAN Operating System, maintains data communications even if a processor or I/O channel fails. ENCOMPASS, the database management facility, is also an integral part of GUARDIAN. The ENSCRIBE software manages "mirror" volumes of the database. As for any other process, one GUARDIAN system performs the database management functions while another GUARDIAN system acts as a backup. The database can reside on one system or can be distributed across a number of processors.

The GUARDIAN/EXPAND Operating System handles non-stop processor resources in a Tandem network of up to 255 nodes. Each node can consist of 2 to 16 processors; thus the maximum number of processors per network is 4080. NonStop and NonStop II processors can be intermixed on a network level, but not within the same node. EXPAND allows nodes to be added online without disrupting network operation.

The Transaction Monitoring Facility (TMF) provides a network-wide identification of transactions on the database. If a system module failure prevents transaction completion, TMF backs out the pending database changes for that transaction without disrupting other transactions in progress.

PATHWAY software interfaces the user application to the Tandem network. A monitor supervises and controls the processes in a transaction processing system. A user writes an application or service program and defines screen formats via PATHWAY SCREEN COBOL. Terminal Control Process (TCP) interfaces the terminal to the service program.

Tandem offers a Data Definition Language (DDL) for the user to specify the system database to the ENCOMPASS Data Base Management System. ENFORM is the non-procedural language for querying ENCOMPASS and for writing reports on the database.

The body of software available for the TANDEM NonStop System is impressive in its overall support for handling transaction processing in modern network configurations. Some of it, such as the ENCOMPASS distributed database management software, and GUARDIAN/EXPAND is relatively new. So far, the Tandem software has performed as advertised.

☐ Strengths

Tandem offers fault-tolerant operation and effective distributed operations with minimum fuss by users. Expansion of processing capabilities is also accomplished without fuss. The Tandem NonStop hardware and software have been operating successfully since 1976. They have proved they can do it. The number of customers who have provided testimony to its success is unprecedented.

☐ Limitations

Tandem has been doubling its revenues from the Tandem NonStop each year for the five years since its first delivery. This tremendous growth rate, which is hard to maintain and manage, makes it difficult to train the numbers of field service personnel needed to support the installed base. Tandem is still relatively small and will remain so for the next few years, even though it is growing quite rapidly. Its marketing and sales support are thus currently limited to major population areas.

■ SOFTWARE

☐ Terms and Support

Terms • completely unbundled; monthly maintenance fees are applied to all software packages in addition to a one-time license fee; many software packages require hardware microcode options which are separately charged.

Support • local on-site for remedial maintenance online and for updates to installed packages; maintenance is performed while system is running.

☐ Operating Systems

GUARDIAN • multiprogramming, multiple-processor executive that supports interprocessor communication for transparent non-stop operations • a copy of GUARDIAN runs in each processor; up to 255 concurrent tasks per processor; each GUARDIAN copy is aware of all other processors in the system; creates user process pairs with one program in one processor active and executing and the other program in the second processor passive but receiving backup messages to update status and current data; passive program begins execution if primary process fails • supports mirrored disk volumes and standardized access to peripherals • simplifies user and program interface to system with command interpreter and system procedure calls • allows online expansion or repair • supports database and transaction processing via ENCOMPASS database management software and Transaction Monitoring Facility (TMF).

Base GUARDIAN • base operating system; includes ENSCRIBE Database Record Manager • program number is T16/9000 for NonStop, T16/9050 for NonStop II:

$11,000 lcns	$110 serv

GUARDIAN • base operating system; includes ENSCRIBE Database Record Manager, XRAY monitor, and Spooler • program number is T16/9020 for NonStop, T16/9070 for NonStop II:

13,000	130

GUARDIAN/EXPAND • extension to GUARDIAN for networking • a copy runs in each processor • can support up to

LCNS: one-time fee. SERV: monthly maintenance cost. Prices effective as of May 1981.

Tandem NonStop Series
Computer Systems

255 nodes, each composed of from 2 to 16 processors for a theoretical total of 4080 processors • End-to-End protocol for ensuring data integrity from sender to receiver • provides multiple communications paths between nodes • runs Network Control Process at each node for automatic best path routing to other nodes and for network status logging • includes Network Line Handlers; each handler manages 1 full-duplex communications line or 1 X.25 virtual circuit • includes Network Utilities to monitor the network; traces data paths through the network, tracks events, displays network statistics, and determines status of individual processors in the network • provides network-wide identification of transactions • requires GUARDIAN/EXPAND Microcode • program number is T16/9007 for NonStop, T16/9057 for NonStop II:

10,000	100

☐ Database Management

ENCOMPASS • collection of facilities, including ENSCRIBE, DDL, TMF, ENFORM, and PATHWAY • runs under GUARDIAN or GUARDIAN/EXPAND operating system; supports distributed database • composed of transaction definition and application control, terminal management, and database management; supplies procedures, programs, and application structures to write modular, single-threaded application programs; manages functions required to define terminal environment, such as screen formatting, data validation, data mapping, and screen sequences; provides facilities for data definition, database structure and relational database manager with query and report writing capabilities • guarantees data integrity against power failures, communications problems, and application errors • furnishes modules needed to control and monitor application programs at runtime • supports programs written in COBOL, FORTRAN, MUMPS, and TAL languages • program number is T16/9104 for NonStop, T16/9114 for NonStop II:

$22,000 lcns	$220 serv

ENSCRIBE Database Record Manager • included with ENCOMPASS • may be used alone or with data definition language (DDL) for centralized database administration • key-sequenced, entry-sequenced, and relative file structures; multikey access with 3 index options (exact, approximate, or generic key match); data and index compression; cache buffer, transparent to the program, holds most recently referenced data blocks in memory for easy access • user can allocate more memory to cache to increase system transaction rates; multivolume files of up to 4 billion bytes on 1 disk or partitioned across up to 16 disk volumes; a separate processor can control each partition; record and file locking to control concurrent file accesses.

Data Definition Language (DDL) • included with ENCOMPASS and available separately • single unified language for design and description of a relational database • used interactively or in batch mode to produce data dictionary with standardized data definitions that can be combined to form complex data structures • produces output for File Utility Program (FUP) under ENSCRIBE to create files and source statements for language definitions to COBOL, FORTRAN, and TAL compilers • features multiple dictionaries for efficient distributed processing • provides access to dictionary from ENFORM, the relational/query report writer • RECORD statements define file name and type, record name and type, field name and type, field groups, ENSCRIBE file access keys, and file sequence information. DEFINITION statements describe data (in terms of fields and groups) independently of any record • DDL dictionary provides a database of maintained information and supports ENFORM with mapping data to describe file types and access methods • program number is T16/9100:

3,000	30

ENFORM • component of ENCOMPASS • relational query/report writing language to retrieve records from a database compiled and defined by DDL • consists of Compiler/Report Processor and Query Processor, which can be remotely located at different nodes to reduce volume of data transmitted • features shared access to database and data dictionary, transparent access, query/report development facilities for customization to fit application, and interactive or program use; user can specify new relationships at runtime in addition to those established during database design; words in language can be changed to German, French, or other foreign language • requires ENSCRIBE software and microcode and decimal arithmetic package • program number is T16/9102:

7,000	70

TRANSACTION MONITORING FACILITY (TMF) • component of ENCOMPASS • system-level facility to maintain consistent database during concurrent transaction processing; database can be either on one system or distributed over multiple nodes • runs under GUARDIAN/EXPAND operating system, which provides network-mode identification of transactions • audit trails stored on distributed mirrored disk volumes to maintain high reliability and availability; transparent to user application program; can be copied to tape to release disk space; all audit data on records at a particular node are contained in the audit trails at that node • 2-phase process updates all audit buffers on each node before unlocking records and making changes to the distributed database; ensures uniformity across all nodes even if communications are lost or node fails; distributed transaction is backed out if a failure occurs before the changes to the various databases are made permanent with the application program call to end transaction • guides operator through recovery procedure after catastrophic failure that damages one or more files; uses online dumps of database files and audit trails to recover damaged files • program number is T16/9016 for NonStop, T16/9066 for NonStop II:

10,000	100

PATHWAY • included in ENCOMPASS and available separately • allows user to write transaction processing application in non-stop environment; consists of Terminal Control Process (TCP), Screen COBOL program, Screen Builder Utility, and PATHWAY Monitor (PATHMON); supports IBM 3270, Tandem 6520 multipage display, and Tandem 6510 terminals • user writes "server" program to implement application in COBOL, FORTRAN, MUMPS, or TAL; user designs screen formats with the interactive Screen Builder Utility which generates the Screen COBOL source statements to describe the formats; user defines PATHWAY parameters (TCP's, terminals, and server processes) to PATHMON • user defines the terminal handling in Screen COBOL, including screen formats, input and output data mapping, data validation and consistency checks, transaction routing, and overall application control; a program is associated with each terminal on the system • TCP passes a transaction record to the user-written server program, which performs the functions and replies; server program is terminal independent • PATHMON supervises and controls the transaction processing system for a single system or for a network; it balances loads, giving server programs a proportionate share of transactions, creating new copies of server programs when all current ones are busy; accepts commands to start TCP's, terminals, and server programs, to display system status, to control error logging, to shut down PATHWAY system, and to control transactions • a system can run several independent PATHWAY systems simultaneously, each identified by its PATHMON process name; each system can support multiple independent user applications • requires T16-2002 ENSCRIBE microcode and T16/2001 Decimal Arithmetic package • program number is T16/9103:

8,500	85

☐ Communications/Networks

A Tandem network is based on the GUARDIAN/EXPAND Operating System, which supports up to 255 nodes with 2 to 16 processors per node for a maximum of 4080 processors per network. NETMON, the EXPAND network control monitor, can reside in any one system in the network where it logs changes in network status, logs and displays changes in processor status in remote system, and displays network traffic. NETMON also provides two network utilities: MAPS can display status of network as seen from any remote system, and PROBE can display the current path to any remote system or to any node of the network. EXPAND supports automatic packet forwarding and best path message routing. When a line fails, EXPAND automatically selects the next best path for routing and retransmission.

Tandem NonStop Series
Computer Systems

In conjunction with the ENCOMPASS Data Base Management System, EXPAND networks can provide access from any node to network-wide database files that are distributed across remote network nodes. The location of the file is transparent to the user.

Each node in a network is connected to other nodes in the network by a dedicated private line or by a public X.25 packet carrier. To help users establish networks of Tandem Computers, the EXPAND software includes implementation tools including Network Control Process (NCP), End-to-End Protocol, and Network Line Handlers. NCP runs at each network node to establish communication paths, store routing information, calculate best path routing, and monitor and log network status changes. Tandem has defined a packet switched End-to-End protocol to guarantee the integrity of data from sender to receiver.

The Network Line Handlers implement the End-to-End protocol. The Direct Connect Line Handler, included at no extra cost, manages a full-duplex dedicated line between two systems.

X.25AM-X.25 Access Method Line Handler • manages one X.25 virtual circuit • program number is T16/9010 for NonStop, T16/9060 for NonStop II:

$2,000 lcns	$20 serv

ENVOY Data Communications Manager • runs as an integral part of the GUARDIAN Operating System to interface transaction processing applications to a data communications network • composed of a set of data communications processes that control data communications lines and fulfill requests from applications through the file management system; supports Tandem as a supervisor or tributary in a multipoint network and as a station in a point-to-point network • supports asynchronous, bisync, ADM-2, TINET, Burroughs, IBM SDLC, and ISO HDLC protocols • synchronous controllers operate with Bell System (or equivalent) 201, 208, 209, and 212 modems; asynchronous controllers operate with Bell System (or equivalent) 103, 113, 303, and 212 modems • Trace Facility records line events in a trace table; maintains line statistics as long as line is in use; provides test programs to verify that communications are functioning properly.

EXCHANGE Remote Batch Workstation • subsystem designed to support 2780 and 3780 emulation • accepts batch input from any media (terminal, mag tape, or disk); allows file transfer between Tandem Systems • operates in conversational, command file, and program modes; includes command interpreter and server programs; maintains nonstop environment • program number is T16/9004:

2,000	20

AM3270 Access Method • interfaces up to 255 IBM 3270 terminals and printers to a Tandem system • handles bisync communication lines; each line can support multiple controllers; each controller can interface multiple 3270 terminals • a Communications Utility Program (CUP) performs operations associated with lines, networks, and devices; works in complementary fashion with TR 3271 to allow pass-through access to IBM host • requires 6202 Byte Synchronous Controller for automatic polling, ASCII/EBCDIC code conversion, and check character generation and checking • program number is T16/9011 for NonStop, T16/9061 for NonStop II:

2,000	20

TR3271/Access Method • emulates one or more 3271 Cluster Controllers • simplifies connecting Tandem to an IBM-compatible host; used in complementary fashion with AM3270 access method pass-through mode; allows terminals connected to Tandem to access IBM-compatible hosts; supports 3270 devices connected to bisync, multipoint communications line • requires Tandem 6202 Byte Synchronous Controller • program number is T16/9012 for NonStop, T16/9062 for NonStop II:

4,000	40

AM6520 Access Method • supports synchronous communications for T16/6520 and T16/6524 terminals operating in multipoint mode • program number is T16/9013 for NonStop, T16/9063 for NonStop II:

2,000	20

☐ Program Development/Languages

COBOL • ANSI X3.25, 1974 level; extensions allow program to set non-stop mode; CHECKPOINT verb can pass information to backup process; automatic checkpoint on execution of OPEN and CLOSE statements • extended I/O to allow READ and WRITE verbs to be used for LOCK and UNLOCK operations; used with ENSCRIBE for files and record locking; allows concurrent COBOL programs to share common database for interprocess communication • allows direct access to operating system procedures by binding the routines to runtime system during compilations • includes utility to interface programs to SPOOLER • supports multilingual environment with FORTRAN, TAL, and MUMPS • supported by DDL for consistent data definitions; can be used with ENFORM to query a database and write reports • requires T16/2001 Decimal Arithmetic Package microcode, T16/9002 ENSCRIBE software, T16/2002 ENSCRIBE Microcode, and T16/2003 COBOL Microcode for each processor • program number is T16/9201:

$7,000 lcns	$70 serv

FORTRAN • conforms to ANSI FORTRAN X3.9, 1977; extensions provide for non-stop operation; can be used with DDL and ENSCRIBE, and interprocess communication • CHECKPOINT passes information for backup; checkpoint automatic with OPEN and CLOSE verbs • DDL can define FORTRAN record structures; FORTRAN READ and WRITE extended to use ENSCRIBE; LOCK allows concurrent access at either file or record level • READ and WRITE can use GUARDIAN facilities for interprocess communications • requires T16/2006 FORTRAN microcode, T16/2005 Floating Point Arithmetic, and T16/2001 Decimal Arithmetic Package for each processor • program number is T16/9202:

6,000	60

TAL (Tandem Application Language) • macro assembler language • block structured; used for applications and systems programming • assembler produces re-entrant object code; provides for both string and bit manipulation • included as part of standard software with GUARDIAN.

MUMPS (Massachusetts General Hospital Utility MultiProgramming System) • meets ANSI 1977 MUMPS language standard • high-level, interpretive language for interactive transaction processing; originally developed for hospital data management applications; also used for commercial applications; extended to run under GUARDIAN in non-stop environment • interfaces to COBOL, FORTRAN, TAL, and PATHWAY programs • can utilize ENSCRIBE Data Base Manager, ENFORM Query/Report writer, and EXPAND communications network • requires T16/2009 MUMPS microcode, T16/9002 ENSCRIBE software, T16/2002 ENSCRIBE microcode, and T16/2001 Decimal Arithmetic Package • program number is T16/9203:

7,000	70

☐ Applications Development Aids

SCREEN BUILDER • integral part of PATHWAY software • an interactive screen definition utility that allows the user to define screen formats directly from a terminal; output is a series of SCREEN COBOL source statements that are compiled under the PATHWAY TCP.

☐ Other Facilities

SORT • runs as a separate program • may be initiated by interactive commands, from a text file, or from a program • supports multiple file types; sort keys may be ascending or descending, string or integer, or multiple fields • included as part of standard software with GUARDIAN.

SPOOLER • stores output files for later retrieval • collectors accept applications output and store on disk; print processes print spooled output; Spool Supervisor monitors, controls, and routes each spooled output job; usually runs in multiple processors to increase throughput • requires T16/9002 ENSCRIBE software and T16/2002 ENSCRIBE microcode for each processor • program number is T16/9101:

$2,000 lcns	$20 serv

XRAY • online monitoring tool for measuring system performance • used for balancing job mix, tuning application, and managing system growth • collects and scans data to

Tandem NonStop Series
Computer Systems

measure use of various hardware components including processors, data communications lines, disks, terminals, processes, disk file opens, and so on; produces reports and plots on usage; takes less than 1% of processor workload; requires no special equipment, except T16/3900 Diagnostic Link Subsystem for first 2-processor system plus one T16/3901 Diagnostic Link Extension for each additional processor pair in the system • program number is T16/9006 for NonStop, T16/9056 for NonStop II:

2,500	25

■ HARDWARE

□ Terms and Support

Terms • outright purchase only • separate installation charge for components added in field; purchase price includes initial system installation.

Support • standard monthly maintenance is Monday to Friday, 8:00 AM to 5:00 PM with 8-hour response time; extended maintenance agreements up to on-site, full-time support are negotiable options.

□ T16 System

All systems are multiprocessor configurations, consisting of 2 to 16 processors, which are connected by a dual interprocessor link called the Dynabus. GUARDIAN operating system handles transmission of backup information over Dynabus via SEND instruction. Data is transmitted in packets of 16 words (5 data words and one checksum word) up to a maximum of 32,767 words. Each processor has two 16-word buffers to receive data from other processors over both interprocessor buses. Receiving processor is interrupted only when one of its buffers is full.

NonStop systems are based on the T16/1412-1 processor, and NonStop II systems are based on the T16/1420 processor. The two processors can be intermixed on a GUARDIAN/EXPAND network level, but not within the same node. The processors are software-compatible at object- and source-code levels. Only privileged programs that access operating system data space or the "environment" register requires modification to be transported from NonStop to NonStop II systems. The NonStop II processor features a number of enhancements as compared to the original NonStop processor. The NonStop II is capable of directly addressing up to 16M bytes of main memory, although the maximum capacity is currently 2M bytes due to physical packaging constraints. More of main memory can be allocated as I/O buffer space on the NonStop II. In addition, a 32-bit-wide data access path to main memory is used. Virtual memory address space—1G bytes—is also increased. Additional features of the NonStop II include a loadable control store, facilitating microcode updates, and an Operations & Service Processor that supports local/remote system control and diagnostics.

The NonStop and NonStop II processors contain two pipelined, microcoded processing units. One unit operates as the central processor and Dynabus controller; the other operates as an I/O processor. The NonStop II processor also includes a processing unit for control of a data transceiver for diagnostics. Each processor has its own independent power supply. All I/O is DMA; total DMA throughput is 4M or 5M bytes per second on the NonStop and NonStop II, respectively. Each I/O controller contains a buffer so that all I/O transfers are in burst mode to reduce the load on the I/O bus. Each processor can be configured with up to 2M bytes of MOS error detection/correction memory. All peripherals interface to the system by dual-access controllers, which are powered by both of the processors connected to them. Thus, loss of power in one processor does not disrupt the power in the controllers connected to it.

NonStop and NonStop II System Maximums • 16 processors; 2M bytes memory per processor for total of 32M bytes per system; 1920M bytes per disk controller; 32 controllers per processor for all peripherals and data communication lines; all controllers are dual access for 2-processor system total of 256 controllers.

□ Packaged Components

T16/244-3 System • includes two T16/1412-1 Processors plus T16/7104 System Cabinet with room for 2 additional processors and 13 I/O slots for system expansion, one T16/6604 hard copy console, one T16/6303 asynchronous controller with 2 ports, one T16/7501 patch panel, one T16/3202 Magnetic Tape Controller with dual channel connections and capacity for 2 drives, and one T16/5103 Magnetic Tape drive PE, 45 ips, 800/1600 bpi:

$94,975 prch	$726 maint

T16/1412-1 Processor • includes CPU/Dynabus-control processing unit, I/O processing unit, 384K bytes of memory, virtual memory control, memory mapping and protection of 2M bytes memory, hardware multiply/divide, power-fail/auto-restart, bootstrap loader, interval timer, control panel, Dynabus controller and interface, interfacing for 32 I/O controllers, 173 instructions, power supply, and one T16/7303 battery pack for MOS memory backup:

33,500	243

T16/250 NonStop II System • includes two T16/1420 Processors plus T16/7120 System Cabinet with room for 2 additional processors and 13 I/O slots for system expansion; one T16/3910 Operations and Service Processor (OSP); dual Dynabus Redundant Inter-processor links; two 5M-byte-per-second block multiplexer channels; one T16/3202 Magnetic Tape Controller with dual channel connections and capacity for 2 drives; one T16/5103 Magnetic Tape drive, 45 ips, 800/1600 bpi; and two T16/7303 Battery Packs for backup of MOS memory:

144,475	883

T16/1420 NonStop II Processor • includes CPU/Dynabus-control processing unit; I/O processing unit; diagnostic data-transceiver-control processing unit; 512K bytes of memory; virtual memory control supporting 1G-byte address space; 8K of 32-bit-word loadable control store; 1K of 32-bit-word read-only control store; string-manipulation and double-word-arithmetic microcode; power-fail/auto-restart; bootstrap loader; interval timer; control panel; Dynabus controller and interface; interfacing for 32 I/O controllers; power supply; and one T16/7303 Battery Pack for MOS memory backup:

50,500	237

Decimal Arithmetic Package • microcode that extends standard instruction set by 40 instructions for ASCII conversion and decimal add, subtract, multiply, divide and scale • required on each processor that runs COBOL or FORTRAN object code • product number is T16/2001 for NonStop, T16/2021 for NonStop II; maintenance charge applies only to NonStop version:

2,000	22

Floating Point Arithmetic Package • microcode that extends instruction set by 40 instructions for floating point arithmetic for single (23 bits) and extended (55 bits) precision • required for each processor that runs FORTRAN object code • product number is T16/2005 for NonStop, T16/2025 for NonStop II; maintenance charge applies only to NonStop version:

2,000	22

Base GUARDIAN Microcode • extends processor so it can run the GUARDIAN software • requires software license • product number is T16/2002 for NonStop, T16/2022 for NonStop II:

1,500	NC

GUARDIAN Microcode • extends processor so it can run the GUARDIAN software; package includes Decimal Arithmetic and Base GUARDIAN microcode; NonStop II version also includes XRAY and Spooler microcode • requires software license • product number is T16/2014 for NonStop, T16/2015 for NonStop II; maintenance charge applies only to NonStop version:

3,500	22

ENCOMPASS Microcode • extends processor so it can run the ENCOMPASS software package • requires software license • product number is T16/2012 for NonStop, T16/2013 for NonStop II:

7,000	NC

PRCH: purchase price. MAINT: monthly charge for prime shift maintenance. NC indicates no charge. Prices effective as of May 1981.

Tandem NonStop Series
Computer Systems

COBOL Microcode • extends processor so it can run the COBOL compiler; not required to run COBOL object code • product number is T16/2003 for NonStop, T16/2023 for NonStop II:

500 NC

FORTRAN Microcode • extends processor so it can run FORTRAN compiler; not required to run FORTRAN object code • product number is T16/2006 for NonStop, T16/2026 for NonStop II:

500 NC

EXPAND Microcode • extends processor so it can run the EXPAND software • requires software license • product number is T16/2007 for NonStop, T16/2027 for NonStop II:

2,000 NC

ENFORM Microcode • extends processor so it can run the ENFORM software • requires software license • product number is T16/2008 for NonStop, T16/2028 for NonStop II:

2,000 NC

MUMPS Microcode • extends processor so it can run the MUMPS software package • requires software license • product number is T16/2009 for NonStop, T16/2029 for NonStop II:

2,000 NC

PATHWAY Microcode • extends processor so it can run the PATHWAY software package • requires software license • product number is T16/2010 for NonStop, T16/2030 for NonStop II:

2,500 NC

Transaction Monitoring Facility Microcode • extends processor so it can run the Transaction Monitoring Facility software • requires software license • product number is T16/2011 for NonStop, T16/2031 for NonStop II:

2,500 NC

DDL Microcode • extends processor so it can run the DDL software • requires software license • for NonStop II only; product number is T16/2032:

500 NC

SPOOLER Microcode • extends processor so it can run the SPOOLER software • requires software license • for NonStop II only; product number is T16/2033:

500 NC

EXCHANGE Microcode • extends processor so it can run the EXCHANGE Package software • requires software license • for NonStop II only; product number is T16/2034:

500 NC

XRAY Microcode • extends processor so it can run the XRAY software • requires software license • for NonStop II only; product number is T16/2035:

500 NC

X25AM Microcode • extends processor so it can run the X.25 Access Method software • requires software license • for NonStop II only; product number is T16/2036:

500 NC

AM3270 Microcode • extends processor so it can run the 3270 Access Method software • requires software license • for NonStop II only; product number is T16/2037:

500 NC

TR3271 Microcode • extends processor so it can run the TR3271 Access Method software • requires software license • for NonStop II only; product number is T16/2038:

1,000 NC

AM6520 Microcode • extends processor so it can run the 6520 Access Method software • requires software license • for NonStop II only; product number is T16/2039:

500 NC

☐ CPU

Nonstop CPU • stack-oriented with 16-bit internal registers • 16-bit addition takes 500 nanoseconds; storing a result in memory takes 900 nanoseconds • microprogrammed instruction set • micro-instruction execution time is 100 nanoseconds; claimed performance is 0.7 MIPS (millions of instructions per second) per processor and 1.3 MIPS per 2-processor system • program access to 4 memory maps including system data, system code, user data, and user code; addressing to 64K words through each map for total of 256K words • 16-bit wide data paths and memory addressing • direct and indirect addressing with and without indexing to memory as well as top-of-stack in memory • 172 instructions including string and bit manipulation and doubleword arithmetic; instruction pre-fetch while current instruction executing; no instruction can write into system or user code area.

NonStop II CPU • stack-oriented with 16-bit internal registers; microprogrammed instruction set • micro-instruction execution time is 100 nanoseconds; performance comparisons to NonStop processor not available • direct main memory addressing to 16M bytes; virtual memory addressing to 1G bytes • 32-bit wide data paths; 16-bit addressing for application programs and extended 32-bit addressing for system programs • direct and indirect addressing with and without indexing to memory as well as top-of-stack in memory • 222 basic instructions plus 20 decimal arithmetic instructions; instruction pre-fetch while current instruction executing; no instruction can write into system or user code area.

T16/3910 Operations & Service Processor (OSP) • included in T16/250 NonStop II System and available separately • serves as operator console and supports local/remote diagnostics for NonStop II systems; can be controlled locally, from a remote T16/6520 CRT Terminal, or from another (remote) T16/3910 OSP • includes microprogrammed control processor; two 1M-byte diskette drives; 300/1200-baud, Bell 212A-compatible modem; T16/6520 CRT Terminal; power supply; desk:

$14,875 prch $163 maint

Interprocessor Busses (Dynabus) • dual interprocessor busses; 13M bytes per second transfer rate per bus • can transfer data simultaneously for 26M bytes per second transfer rate • packet-multiplexed transfers between 2 to 16 processors • direct access by one processor to other processor's memory, but no memory sharing • transfers of 1 to 32K bytes per block for NonStop; 1 to 64K bytes per block for NonStop II.

☐ Memory

Main Memory • 384K bytes of 500-nanosecond MOS included in T16/1412-1 NonStop processor; 384K-byte module available for expansion to maximum of 2M bytes • 512K bytes of 400-nanosecond MOS included in T16/1420 NonStop II processor; 512K-byte module available for expansion to maximum of 2M bytes • 22 bits per word (16 data bits and 6 error detection/correction bits); corrects all 1-bit errors and detects 2-bit errors.

T16/2412 MOS Memory Module • 384K bytes; for T16/1412-1 NonStop processor:

$9,600 prch $132 maint

T16/2420 MOS Memory Module • 512K bytes; for T16/1420 NonStop II processor:

12,800 47

☐ I/O Channels

I/O Processor • integral part of each processor • controls all I/O for each CPU; supports 32 controllers with up to 8 devices per controller; table-driven from I/O control (IOC) table in memory; 256 entries in IOC, one for each possible I/O device; entry contains buffer address and count of number of bytes to be transferred; data buffers are included in all I/O controllers • all data transfers between I/O controllers and memory in DMA mode at 4M bytes per second for NonStop, 5M bytes per second for NonStop II; all I/O devices can be operating simultaneously with interleaved data transfers over I/O channel.

T16/3401 Universal Interface • dual channel, dual power from 2 processors; either processor can sustain power • controls 2 devices with 16-line parallel interface; TTL logic for line drivers

Tandem NonStop Series
Computer Systems

and receivers of one device and differential logic for the other; transfer rate up to 950K bytes per second:

$2,800 prch $20 maint

T16/3810 Tandem-to-IBM Link (TIL) • supports high-speed communications with a local IBM mainframe by emulating an IBM 3420 Model 7 tape drive; attaches to IBM 3803 Model 1 or 2 Controller on IBM mainframe; serves as first link; additional links are provided by T16/3812:

24,500 600

T16/3812 • additional link:

12,900 375

T16/3830 Tandem Hyper Link (THL) • supports connection to computer systems using Network System Corporation Hyperchannel; includes channel interface, NSC A400 adapter interface, and cables • requires T16/7506 Patch Panel:

14,900 151

T16/7506 Hyper Link Patch Panel:

350 NC

☐ Communications

All communications are through dual channel, dual power controllers connected to two processors. Either processor can sustain power and maintain communications.

T16/6202 Byte Synchronous Controller • controls 1 to 4 point-to-point or multipoint communication lines • supports up to 80K bps transmission rate per line; supports 160K bps aggregate transmission rate for all 4 lines • performs ASCII or EBCDIC character translation, block check character generation and checking, and autopolling:

$5,800 prch $29 maint

T16/6810 Peripheral Line Adapter • supports attachment of 6, 7, or 8 T16/652X CRT Terminal strings and up to two T16/5508 Printers to T16/6202 Byte Synchronous Controller; includes 1 input port and 8 output ports:

1,975 15

T16/6203 Bit Synchronous Controller • controls 1 to 4 point-to-point or multipoint lines • supports up to 56K bps transmission rate per line; performs variable length and offset translation of EBCDIC and ASCII codes; generates and recognizes control characters; automatically inserts and deletes zero bits • requires T16/7502 Patch Panel:

5,800 39

T16/7502 Synchronous Patch Panel • provides connection between synchronous controller and lines • supports up to 3 controllers (12 lines):

775 6

T16/6303 Asynchronous Controller • controls up to 2 current loop or RS-232 local or modem-connected lines • programmable line speed: 50 bps to 19.2K bps; programmable continuous read option allows use of simplex links • supports two T16/6303 extensions of 15 async lines each • requires T16/7501 Patch Panel:

3,600 18

T16/6304 Asynchronous Extension Board • T16/6303 option provides control for 15 async lines • second T16/6304 requires additional T16/7501 Terminal Patch Panel:

4,300 22

T16/7501 Terminal Patch Panel • provides connection between asynchronous controllers and terminal connections • supports up to 17 terminal ports:

775 6

☐ Disk

T16/3106 Disk Controller • dual-channel; dual-powered from 2 processors; either processor can sustain power • supports up to 8 drives in any mix • requires 2 I/O slots and T16/1504 Patch Panel:

$10,500 prch $58 maint

T16/7504 Disk Patch Panel • connects disk drives to controller (T16/3106) • includes 16 ports for 4 disk controllers:

775 NC

T16/410X Moving-Head Disk Drives • all drives use removable disk pack • 30-millisecond average seek time; 8.33-millisecond rotational delay • each drive includes 1 pack.

T16/4104 Disk Drive • 240M bytes; 11-high pack; pedestal mount:

26,500 208

T16/4105 Disk Drive • 64M bytes; 5-high pack; pedestal mount • supports one T16/4106 add-on drive:

15,500 174

T16/4106 Disk Drive • add-on, 64M bytes; 5-high pack; drawer mounts within T16/4105 cabinet:

15,200 174

T16/4109 Fixed/Moving-Head Disk Drive • 65.45M-byte total capacity; 1.45M bytes on 1 nonremovable surface with fixed-heads; 64M bytes on 5 nonremovable surfaces with moving heads:

18,270 123

T16/7202 Disk Cabinet • houses up to two T16/4109 drives:

1,500 NC

☐ Magnetic Tape

Controllers • dual channel; dual powered from 2 processors; either processor can sustain power • supports 2 tape drives • all drives include cabinet and cable; each drive separately connected.

T16/3202 Controller • for T16/5101, 5103, and 5104 drives:

$4,800 prch $23 maint

T16/3203 Controller • for T16/5105 drives; not available for NonStop II systems:

4,800 23

T16/5103 Drive • 9-track; 800/1600 bpi; 45 ips:

8,000 57

T16/5104 Drive • 9-track; 800/1600 bpi; 125 ips:

14,500 110

T16/5105 Drive • 7-track; 200/800 bpi; 45 ips; not available for NonStop II systems:

8,000 57

☐ Terminals and Workstations

Hard Copy Terminals

T16/6603 Terminal • RS-232 interface; 132 column; 30 cps:

$2,900 prch $48 maint

T16/6604 Terminal • 20 mA current loop interface; 132 column; 30 cps:

3,200 48

CRT Terminals

T16/6401 CRT Terminal • RS-232 interface; 24 lines x 80 characters; switch selectable speed ranging from 110 bps to 19.2K bps; local or modem attachment:

$1,500 prch $25 maint

T16/6402 CRT Terminal • 20 mA current loop interface; 24 lines x 80 characters; switch selectable speed ranging from 110 bps to 19.2K bps:

1,500 25

T16/6C Upper/Lower Case • for T16/6401 or 6402:

150 NC

T16/6511 CRT Terminal (Page) • RS-232 interface; 24 lines x 80 characters; character or page mode; local editing and function keys; protected/unprotected; full/half bright:

2,400 35

Tandem NonStop Series
Computer Systems

T16/6512 CRT Terminal (Page) • 20 mA current loop interface; 24 lines x 80 characters; character or page mode; features include local editing and function keys, protected/unprotected, full or half bright; switch selectable speeds ranging from 110 bps to 9.6K bps:

2,400	35

T16/6520 CRT Terminal (Multipage) • 3/6 memory pages of 24 lines x 80/40 characters • async or sync; character or block transmission; RS-232 or 20 mA current loop interfaces; point-to-point or multipoint communication at speeds ranging from 110 bps to 19.2K bps • full complement of video and data attribute specifications; local editing, program function keys and line for status display:

2,950	22

T16/6524 CRT Terminal (Multipage) • same as T16/6520, except includes RS-232 communications port for serial printer:

3,150	24

T16/6552 CRT Terminal (Page), Polling • same as T16/6511, except uses polling protocol for multidrop communications:

2,700	33

☐ Printers

T16/3302 Controller • dual channel, dual power from 2 processors; either processor can sustain power • supports 1 line printer:

$2,800 prch	$20 maint

Line Printers • 3 models; 300/600/900/1350 lpm • all printers feature 132 columns, 64-character set, and 12-channel VFU • each requires T16/3302 Controller.

T16/5502 Line Printer • 300 lpm; drum print mechanism:

11,500	140

T16/5503 Line Printer • 600 lpm; drum print mechanism:

14,000	178

T16/5504 Line Printer • 900 lpm; drum print mechanism:

21,000	202

T16/5510 Line Printer • 1350 lpm; band print mechanism:

45,500	507

T16/5A 96-Character Drum Option • available for all line printers; includes upper/lower case:

2,000	NC

T16/5X Drum Options • available for 64-character OCR-A, 64-character OCR-B, 96-character OCR-B, English, Swedish, Spanish, and German:

2,000	NC

T16/5508 Serial Character Printer • 200 cps • async; 96-character set; switch selectable current loop or modem attachment:

4,500	55

☐ Other Peripherals

T16/3305 Card Reader/Line Printer Controller • dual channel, dual power from 2 processors; either processor can sustain power • supports 1 line printer (any model) and 1 card reader:

$2,800 prch	$20 maint

T16/5301 Card Reader • 600 cpm; 80-column cards:

5,800	46

T16/3900 Diagnostic Link Subsystem • basic subsystem designed for 2-processor system; extension required for each processor pair added • includes control panel that connects in patch panel space and a printer circuit board that mounts in I/O controller board slot; contains modem equivalent to a Bell 113B; can connect to customer-supplied external modem • included free with all Tandem service contracts:

3,400	NC

T16/3901 Diagnostic Line Extension • required for each processor pair added to basic subsystem:

2,000	NC

☐ Power Supplies and Cabinets

T16/7301 Power Module • supplies power for I/O controllers • required only on large configurations:

$3,500 prch	$40 maint

T16/7302 Transformer • 15K Va; for isolation:

4,000	16

T16/7303 Battery Pack • backup for MOS memory:

1,500	28

T16/7901 Service Panel • plug-in module:

5,000	16

T16/7104 System Cabinet • for up to 4 processors and 16 I/O controllers; supports up to two T16/7801 I/O expansion boards; 71-in H x 30-in W x 32-in D:

8,600	88

T16/7120 System Cabinet • for up to 4 processors and 24 I/O controllers; 71-in H x 30-in W x 32-in D:

15,800	166

T16/7105 Patch Panel Cabinet • for up to 10 patch panels; 71-in H x 30-in W x 32-in D; can be equipped with AC power distribution:

2,500	NC

T16/7108 Patch Panel Cabinet • for up to 10 patch panels; 71-in H x 22-in W x 32-in D; no provisions for AC power distribution:

2,500	NC

T16/7107 Expansion Cabinet • required for system with more than 8 processors; provides space for 4 power modules and 4 patch panels; 71-in H x 30-in W x 32-in D:

2,500	NC

T16/7109 I/O Cabinet • provides space for 32 I/O slots; 71-in H x 30-in W x 32-in D:

8,600	88

T16/7121 I/O Cabinet • provides space for 24 I/O slots; 71-in H x 30-in W x 32-in D:

15,800	104

T16/7110 Peripheral Cabinet • provides standard 19-inch rack space; 71-in H x 22-in W x 32-in D:

2,500	NC

• **END**

INDUSTRY PERIODICAL
March 1980 RI 417/20
L. Duane Kirkpatrick, First Vice President
San Francisco, (415) 392-7200

RANDOM-ACCESS MONTHLY

This service of the Technology Group is a monthly review of events and trends affecting companies in the technology-related industries. Comments and observations in this publication are oriented toward clients wishing to maintain general awareness of the conditions affecting investment decisions. *Random-Access Monthly* supplements our more detailed analyses of individual companies and industry groups.

The IBM 4331, the smaller of the two 4300 systems introduced in January 1979, addresses a very critical segment of the computer marketplace. It falls into the price/performance range where the overlap between the larger minicomputers and the smaller mainframes is most visible. This analysis compares the IBM 4331 with several of the most important competitive products now on the market. The general conclusions are:

- The 4331 is primarily designed to protect IBM's installed customer base.
- The system is best used in a batch processing environment.
- The 4331 alters total system pricing, decreasing the importance of the central processor.
- Most mainframe competitors have responded with comparable price/performance offerings, but the traditional 20-30% advantage has been erased, which focuses their attention on protecting their own customer bases rather than attacking that of IBM.
- The minicomputer companies continue to offer superior price/performance products, which preserves their advantage in traditional markets. Altered priceing strategies may make it more difficult for several of these companies to prosper in the commercial data processing marketplace in the longer run.

CONTENTS

ADDITIONAL INFORMATION ON COMPANIES MENTIONED IN THIS REPORT IS AVAILABLE ON REQUEST

TECHNOLOGY STOCK SUMMARY

Company	3/18/80 Price	52-Week Price Range	Earnings per Share 1978	1979E	1980E	12 Months 12/79E	12/80E	Percent Change	P/E 12/80E	Relative P/E 12/80E	Latest 12 Mos. Revs. ($Mil.)	Shares ($Mil.)	Dividend	Est. 5-Year EPS Growth Rate	Rec.	Analyst
Adv. Micro Devices (3/31)	$38	$46—31	$1.48	$2.90	$3.00	$2.57A	$2.95	15%	12.9	162	187	7.5	Nil	25%	B/H	MCD
Amdahl Corp.	20	45—16	2.81	1.02A	1.50	1.02A	1.50	47	13.3	168	300	16.8	$0.40	20	B/H	LDK
Auto. Data Processing (6/30)	33	41—29	1.84	2.21A	2.55	2.37	2.83	19	11.7	147	409	15.2	0.64	20	B/H	LLB
Beckman Instruments (6/30)	25	33—20	1.35	1.66A	1.85	1.81A	2.30	27	10.9	137	487	19.8	0.32	20	B	AFC
Burroughs Corp. (o)	66	88—65	6.21	7.45A	8.25	7.45A	8.25	11	8.0	101	2,831	41.0	2.60	15	B/H	SCD
Computer Sciences (3/31)	19	24—11	0.98	1.04A	1.55	1.33	1.72	29	11.0	139	416	15.2	Nil	18	B/H	LLB
Control Data (o)	49	63—32	4.88	6.86A	8.20	6.86A	8.20	20	6.0	75	2,273	17.0	0.60	15	B	LLB
Data General (9/30)	60	75—46	4.00	4.82A	5.60	4.82A	5.95	23	10.1	127	540	10.4	Nil	25	B	SCD
Datapoint Corp. (7/31)	102	123—68	4.10	5.84A	7.50	6.80	8.30	22	12.3	155	247	4.4	Nil	25	B/H	MCD
Digital Equipment (6/30)(o)	65	83—51	3.40	4.10A	5.35	4.70A	5.95	27	10.9	138	2,031	44.9	Nil	25	B	SCD
Eastman Kodak (o)	46	67—43	5.59	6.20A	6.00	6.20A	5.85	(6)	7.7	97	8,028	161.4	2.90	12	B/H	EGG
Electronic Data Sys. (6/30)	21	28—19	1.54	1.82A	1.95	1.93	2.08	8	10.9	137	312	13.5	1.00	15	H	LLB
Floating Point Systems (10/31)(m)	18Ask	18—9	0.78	0.53A	1.00	0.53A	1.00	89	18.0	227	33	3.9	Nil	25	H	LLB
Fluke, John, Mfg. (9/30)	24	27—16	1.20	1.59A	1.95	1.64A	2.00	22	12.0	151	110	5.9	Nil	17	H	AFC
Four-Phase Systems	32	49—30	2.58	3.25A	3.25	3.25A	3.25	0	9.8	124	178	5.2	Nil	25	B	MCD
Harris Corp. (6/30)	30	40—25	1.95	2.25A	2.55	2.49A	2.75	10	10.9	138	1,049	30.2	0.72	15	H	AFC
Hewlett-Packard (10/31)(o)	60	71—46	2.63	3.43A	4.50	3.43A	4.50	31	13.3	168	2,520	59.4	0.40	21	B	AFC
Honeywell, Inc. (o)	82	101—64	8.48	10.95A	11.20	10.95A	11.20	2	7.3	92	4,210	21.9	2.60	12	H	SCD
Intel Corp. (m)	64Ask	75—42	2.16	3.59A	4.20	3.59A	4.20	17	15.2	192	663	22.1	Nil	25	B/H	MCD
International Business Machines (o)	59	80—59	5.32	5.16A	5.85	5.16A	5.85	13	10.1	127	22,863	583.4	3.44	14	B/H	LDK
Intersil Inc. (9/30) (m)	22Ask	33—12	1.27	1.65A	1.85	1.70	1.90	12	11.6	146	130	6.1	0.24	25	B/H	MCD
Lanier Business Products (5/31)	23	34—23	1.43	1.88A	2.35	2.05	2.50	22	9.2	116	215	6.1	0.44	17	B/H	EEG
M/A-Com (9/30)	33	42—20	1.10	1.48A	1.85	1.57A	1.95	24	16.9	213	129	5.7	0.24	20	B/H	AFC
Minnesota Mining & Mfg.	48	60—47	4.83	5.59A	6.00	5.59A	6.00	7	8.0	101	5,400	117.2	2.80	13	B	SCD
Motorola Inc. (o)	54	64—38	4.04	4.96A	4.50	4.96A	4.50	(9)	12.0	151	2,714	31.0	1.40	11	H	MCD
National Semiconductor (5/31)(o)	32	41—20	1.72	2.57A	3.50	3.15	3.50	11	9.1	115	719	13.3	Nil	17	H	MCD
NCR Corp. (o)	69	82—57	6.90	8.78A	9.65	8.78A	9.65	10	7.2	90	3,003	26.7	1.60	15	B/H	SCD
Northern Telecom	36	48—34	2.92	3.16A	3.45	3.16A	3.45	9	10.4	132	1,625	33.6	0.69	15	H	AFC
Perkin-Elmer (7/31)(o)	39	47—26	1.77	2.53A	3.35	2.79	3.50	25	11.1	140	869	20.1	0.72	15	B	AFC
Pitney-Bowes (o)	30	37—23	3.32	4.02A	4.50	4.02A	4.50	12	6.7	84	1,025	13.3	1.40	12	B	EGG
Polaroid Corp. (o)	20	42—20	3.60	1.10A	2.30	1.10A	2.30	109	8.7	110	1,362	32.9	1.00	12	B/H	EGG
Prime Computer	28	33—14	0.73	1.44A	1.80	1.44A	1.80	25	15.6	196	153	11.8	Nil	30	B/H	SCD
Reynolds & Reynolds (9/30)(m)	29Ask	34—25	2.56	3.15A	3.50	3.22	3.63	13	8.0	101	190	4.7	1.08	17	B	LLB
ROLM Corp. (6/30)	41	49—29	0.51	1.50A	2.15	1.89A	2.45	30	16.7	211	134	7.6	Nil	30	H	AFC
Scientific-Atlanta (6/30)	37	46—25	1.39	1.76A	2.40	2.02A	2.65	31	14.0	176	156	4.8	0.18	20	B/H	AFC
Shared Medical Systems (n)	17Ask	19—16	0.72	0.90A	1.10	0.90	1.10	22	15.5	195	82	6.0	0.28	23	B/H	LLB
Sperry Corp. (3/31)(o)	50	60—42	5.73	6.75	7.50	6.53A	7.20	10	6.9	88	4,586	40.5	1.56	12	B/H	SCD
Storage Technology (o)	16	23—13	1.14	1.58A	1.90	1.58A	1.90	20	8.4	106	479	24.6	Nil	20	B/H	LLB
Tektronix, Inc. (5/31)	45	64—47	3.19	4.28A	4.85	4.66A	5.20	12	8.7	109	878	18.2	0.84	17	B	AFC
Teradyne, Inc.	36	45—16	2.03	2.32A	3.05	2.32A	3.05	31	11.8	149	122	3.5	Nil	20	B/H	AFC
Texas Instruments (o)	86	108—78	6.15	7.58A	7.40	7.58A	7.40	(2)	11.6	147	3,224	22.8	2.00	15	H	MCD
Tymshare, Inc.	47	58—34	2.31	3.06A	3.70	3.06	3.70	21	12.7	160	183	4.9	Nil	22	B/H	LLB
Xerox Corp. (o)	54	69—54	5.67	6.69A	7.30	6.69A	7.30	9	7.4	93	6,604	84.1	2.80	12	B	EGG
S&P 400	116	134—107	13.17	15.45	14.60	15.45	14.60	(6)	8.3	100	—	—	6.38	11	—	—

Notes: Fiscal years ending on dates other than 12/31 shown in parentheses after company names; for fiscal years ending 3/31, date under prior years. (m) Dean Witter Reynolds Inc. makes a market in this over-the-counter security. (n) Dean Witter Reynolds Inc. and/or its stockholders may in the normal course of business have a position in this over-the-counter security. (o) Listed options are available on this issue. 'Stock dividend, estimated cash value on ex-distribution date. (s) Adjusted for recent stock dividend or split.

Highlighted Changes from 3/10/80

Eastman Kodak: 1980 estimate raised from $5.85.
National Semiconductor: 1980 estimate raised from $3.40.
Rank Organisation: Dropped from coverage.
Teradyne: 1980 estimate reduced from $3.15.

1. SUMMARY

IBM Defensive Response

The IBM 4331 is viewed as a defensive system designed to protect IBM's customer base; it is less an offensive machine introduced to capture new market share. The impact of the 4331 will be felt most strongly by those competitors which have been most aggressive in attacking the IBM base in the recent past.

The IBM 4331 is significant in the IBM product line because:

- It places IBM's batched-oriented "370-type" system in the low-price market to compete head-on with the high-end minicomputers and competitive small mainframes.
- It will significantly alter the role of the vertically integrated hardware manufacturer, system houses, and the process of distributing small systems to end-users because the relative pricing of hardware and software has changed, and because the computer processing unit (CPU) and memory are only a very small percentage of total system price.
- It overlaps in price with two other new IBM systems, the 8100 and System/38, yielding competition within IBM and alternatives for the customer.
- It confirms new plateaus of technology and price performance that others will have to attain during 1980.
- It will significantly enlarge the highly elastic small-scale, general-purpose computer market, because of lower-price computation, memory, and disk storage.

The competitive and market objectives of the 4331 in decreasing order of importance to IBM are:

- to prevent further erosion of the $5 billion of installed IBM System/370 models 115 through 138 (as well as older System/360 models 20 through 40) by providing a compatible and competitive growth system, primarily for the batch processing market.
- To provide a compatible satellite small batch processor for current large-system users.
- To provide a product for Data Processing Division salesmen to sell to large distributed data processing prospects.
- To provide an alternative upward growth path for System/3 and System/34 users to prevent erosion of this base.
- To provide a low-cost system that utilizes IBM's large operating systems, which changes the selling effort from the raw hardware competitive arena to a "total system" competition.

4331 Market is Batch Processing

The major data processing market, when measured by dollar value of computer usage, has been, and for a while will continue to be, batch data processing.

Large distributed data processing installations, while on the increase, are typically shut down at night and turned into batch processing systems for one-half to two-thirds of their operating time.

Many systems that appear to be on line are actually performing on-line data collection and batch processing functions.

It is in this market, batch data processing, that the 4331 will be the most effective, and where it is primarily aimed by IBM's marketing efforts.

Although the 4331's on-line capabilities have been improved relative to the 370, they are primarily accommodations for the mixed environment of batch and on-line processing, and the 4331 will have negligible impact on such traditional minicomputer markets as:

- Process control
- Scientific computation
- Time sharing
- Interactive small business systems
- Transaction processing.

Vertical Integration is Required

The 4331 is the first major step by IBM toward changing the distribution pattern of computers because it substantially changes the relative values of the system components. IBM has always set the tone for computer industry marketing practices, and the 4331 represents a significant shift in that tone.

IBM is in the systems business: Selling a total problem solution to the end-user. In the past, the solution was primarily a piece of hardware sold on its value rather than on a multiple of its cost. With the rapid decrease in hardware costs, this old formula is not sufficient to provide acceptable revenue growth.

The basic instruction execution of the 4331 is more powerful than a 360/50, the commercial computer workhorse of the 1960's. This is an improvement in price/performance over the 360/50 by a factor of 10 in 15 years. With the 360, software was free, terminals were rare, and the CPU and memory represented about one-half of the total system price.

With the advent of the 4331, the economics have changed, and a typical total systems price is broken down as follows:

CPU Memory	20%
Peripherals	38
Terminals	12
Maintenance[1]	12
Software[1]	18
	100%

[1]Rental and monthly charges are capitalized as if they were a five-year lease, using 2.3% of purchase price per month as the capitalizing factor. This is equivalent to 12.5% interest.

Market success under these conditions will require a vertically integrated manufacturer which has good systems software and controls its maintenance force.

IBM Products Compete for Market

IBM now has three new products in approximately the same purchase price range: $100,000-500,000. These are the 4331, System/38, and the 8100. Each of these offers different features for different markets, but they compete by offering alternatives to the users and, by doing so, tend to block competition's move into this price range. These differences are discussed more fully in Section 3.

4300 Technology Is State of Art

IBM is using new state-of-the-art proprietary technology which promises to produce logic components at a cost of at least one-fourth that of currently available parts. IBM's financial strength allows it to price these components based on their future cost rather than their current cost, which could be as much as ten times higher than expected at full run rate.

The lower manufacturing costs could take at least two years to achieve as IBM builds up production capacity and obtains manufacturing experience in the new technology. In the interim, the higher manufacturing costs will depress IBM profits for the 4331. However, at the end of the product cycle, the reverse condition should prevail.

The net result is that IBM has produced a system that gives substantially better price/performance than other mainframes which have similar cost structures. But the 4331 is still not as cost effective as the major minicomputer vendors.

Price/Performance

Table 1 illustrates the price/performance of the major competitive systems in the 4331 class. The systems chosen were those most likely to be involved in a competitive battle with the 4331, typically in a batch environment. If a directly competitive machine was not available, then a machine in the comparable price range was chosen.

Table 1

RELATIVE SYSTEM RANKINGS
(Price Per MIPS)

Vendor	System	Performance (MIPS)	CPU Price ($K)	System Price ($K)	System Price/MIP ($K)
Tandem	(2-CPU)	1.3	50	185-600	142
Perkin-Elmer	3220	.7	45	135-560	193
Digital Equipment	VAX 11/780	.9	100	210-660	233
Univac	V77-800	.6	40	150-620	250
Data General	C/350	.4	54	150-520	375
Prime	550	.4	70	190-660	475
Prime	450	.3	50	170-620	567
Honeywell	6/53	.2	43	145-500	725
NCR	8555	.3	90	220-600	733
Honeywell	64/350	.3	155	250-650	833
Honeywell	64/320	.2	81	170-550	850
IBM	4331	.2	68	180-575	900
Burroughs	2930	.3	140	290-800	967
Digital Equipment	2020	.3	130	300-900	1000
NCR	8455	.17	50	170-550	1000
Hewlett-Packard	3000 III	.12	110	190-550	1583
Hewlett-Packard	33	.07	55	120-500	1714
Univac	90/30	.15	140	280-799	1867

The performance is measured in millions of instructions per second (MIPS). The price is based on typical systems or a minimum Central Processing Unit (CPU) as defined by the manufacturer. These figures should only be used as a very rough guide to the relative value of the systems.

Table 1 ranks the systems on the relative system price per MIPS. The lower the number in the last column of the table, the more effective the system appears to be. The 4331 ranks number 12 on the list and is over six times as expensive as the best, the Tandem dual-processor system. Perkin-Elmer's 3220 shows in second place in this ranking. DEC's older VAX 11/780 is in third position.

4300 Closes Current PCM Market

Prior to the 4300 announcements, the plug-compatible mainframes (PCM's) were the most serious threat to IBM because the user could switch vendors at negligible cost, without reprogramming. The 4331 and 4341, with their advanced price/performance and new technology, have made it more difficult for PCM's to compete.

The cost of entering the PCM market was relatively small. A vendor could design a cost-competitive machine for an engineering expense of under $1 million. At least one PCM designed a processor for an engineering cost of $250,000. The engineering cost rises dramatically with increasing machine size, especially in the 370/158 or 3031 class, just above the 4341.

Most of the PCM's were also offering a 30-50% price advantage over IBM and now offer no advantage to users over the 4300 series, except that of early delivery. In order to compete, the PCM's have had to introduce new models.

We expect that one or two manufacturers may compete successfully with newly designed systems, the most likely competitors being National Semiconductor and Magnuson. Both these vendors may seek specialized narrow market segments to compete directly against the IBM 4300 systems.

Mainframers Not Heavily Impacted

Mainframers such as Burroughs, Honeywell, and Univac must penetrate the IBM user base if they hope to grow significantly faster than the market. However, at the low end of the market, their growth has been due as much to finding new users as to switching IBM users.

With the announcement of the 4331, these vendors must now look to protect their own installed base, which could become attracted to IBM's better price/performance. However, the cost of reprogramming will keep most users from making a rapid switch to a noncompatible machine.

With the 4331 announcement, all mainframers except Univac were quick to counter with their own new announcements to show their users a path for growth and, therefore, reduce the temptation to switch to IBM. Generally, these new systems are about equal to the 4331 and do not offer the traditional 30% better price/performance offered by non-IBM machines.

Short-term, the mainframers will not be significantly impacted by IBM. But long-term, the impact could be more severe as they are forced to become more vertically integrated, from components to peripherals, in order to maintain their margins. Their growth will also be slowed to the extent that past revenue growth has come from switching IBM users to their own equipment.

Minicomputer Impact Will Be on Future Market Potential

Unlike the mainframe companies, most of the minicomputer companies, Hewlett-Packard being a notable exception, have a better price/performance than the 4331. There is no need for them to rush new announcements. Some of the current mini offerings—such as the Tandem system, Digital Equipment's VAX, or Perkin-Elmer's 3220—offer three to six times better price/performance than the 4331. Normal new-product offerings will improve this ratio even more, so that, on a price/performance basis, the minis can easily compete with IBM.

Over the longer term, the outlook for the minicomputer companies must be viewed in a different context. In order to maintain their fast growth and grow faster than the old electronic data processing (EDP) market, the minis will have to add to their traditional markets by entering the commercial data processing marketplace, where they will compete head-to-head with IBM.

The raw price/performance is only one factor in making a sale in this market. With the competition's performance margin reduced because of IBM's improved product offerings, other factors such as peripherals, software, and services, where IBM has more entrenched strength, will come more fully into play.

The other new IBM products—the System/38, aimed at first-time and on-line interactive users; the 8100, aimed at the distributed processing markets; and the Series/I, aimed at the large end-user and process-control markets—will compete in the province of the minicomputer manufacturers.

IBM has less inherent advantage in these markets but, through sheer force of size and reputation, may capture a sizeable share of an expanded market.

2. THE IBM 4331

An Improved 370/125

The 4331 is primarily an improved version of the 370/115 and 125 models, using new technology. The basic internal machine organization is very similar. However, the data path width on the 4331 is significantly improved to 32 bits, providing much higher performance.

The 4331 is constructed of completely new technology, with 700 circuits on a chip versus two circuits for a 370, and 64k bits of memory on a chip versus 2k bits on the 370. As the cost of a chip is relatively constant, once full production is reached, the 4331, even with more circuits, will be considerably less expensive to build than the replaced 370 models.

4331 Differences

The net result is that the 4331 has an improved price/performance. It offers slightly less performance than a 370/148, at a price less than the 370/115, for a five to six times overall improvement in price/performance.

The 4331 has a MIPS rating of 0.2 and a price of $180,000 for a minimum system, leading to a price of $900,000 per MIPS. The system price includes a capitalized value for monthly maintenance and software charges.

The major new features of the 4331 are:

- New architecture (i.e., instruction set)
- Improved input/output capabilities
- New peripherals
- Reliability, availability, and serviceability
- New software.

Architecture

Architecture, as defined by IBM, is the instruction set and other system functions visible by and available to the user. It is not the internal data flow design of the computer.

Because the 4331 is microprogrammed, the architecture could be easily changed, if desired. However, IBM changes architecture very slowly, because of the large cost of developing system software to support the new architecture and the need for compatibility to preserve the user's software investment.

Initially, IBM's large-system architecture was defined by the System/360 in 1964. Approximately a dozen new instructions were added for the System/370 architecture in the early 1970's. The 4331 introduces another new architecture with about a dozen new instructions.

The new architecture is called Extended Control Program Support (ECPS). The user instruction set is equivalent to the 370, to preserve compatibility for user programs. The privileged instruction set used by the operating system is completely changed, which will protect the new software from unlicensed usage.

The 4331 user has two choices of architecture: The older 370 or the newer 4331 ECPS mode. The decision as to which mode to use is determined at the time the system power is turned on and can be changed for different job environments by changing microcode. It is feasible for IBM to alter the microcode to emulate competitive machines, just as was done for the 360.

Using the new architecture mode in conjunction with an improved version of the old operating system that puts more system functions into microcode, the operating-system overhead is substantially reduced and could be as much as one-third that of a standard 370, considerably enhancing actual system performance.

The primary difference between 370 architecture and the 4331 ECPS architecture is in the use of single-level addressing for main storage and disk I/O. Previously, to perform disk I/O, the operating system had to know the physical track, sector, and byte address of the data. With ECPS, the system only has to know the page number. The disk controller's microcode maps the page address into the correct physical address.

The system programmers and users are no longer concerned about the physical compatibility or layout of the storage media. The I/O controller takes care of the problem. This will allow IBM to change I/O media and develop new compatible disk drives.

Input/Output

The 4331 I/O is considerably improved over the 370, not only in the addressing previously described but also in speed. The total I/O rate of a 4331 is 2.3 million bytes per second, compared to the 1 million bytes per second for the average competitors in this price class.

In a batch environment, tapes and disks are read sequentially, and I/O speed rather than MIPS could be the primary measure of total system performance.

Another aspect of the 4331 I/O, similar to the 370/125, is the use of integrated I/O adapters. Functionally, the adapters include additional microcode used to perform the channel and I/O functions. The adapters degrade system performance because either the adapter microcode or the central processing microcode can be running at one time. However, the adapters are much lower in cost than the alternative of separate controller boxes.

Peripherals

Coupled with the new 4300 announcements were several new peripherals. The most important announcement was the 3770 disk drive with 570 million bytes of storage. Its incremental price of $23,000 per unit, or $40 per million bytes, is the lowest in the industry.

The new drive also has a read/write speed of 1.8 million bytes per second, almost half again as fast as other disk drives. This is a definite improvement for a batch environment where the disk is read sequentially. The 3770 attaches to the 4331 through an integrated disk adapter. Average access time is 20 ms.

As an alternative to the 3770, IBM also offers the 3310 disk drive, the eight-inch rigid disk called "Piccolo," which has spawned a large number of imitators. The 3310 stores 64 million bytes, with an average access time of 27 ms. and a data transfer rate of one megabyte per second.

Both of the new IBM disk drives introduced the notion of fixed-block architecture. All data are recorded in preformatted 512-byte blocks which are directly addressable under the single-level addressing scheme mentioned earlier.

A new printer, less expensive than the older models, complements the peripherals lineup. The 3203 model 5 allows printing at 1,200 lines per minute, further enhancing the batch processing orientation of the 4331.

Reliability, Availability, Serviceability

A major facet of the IBM system is the increased emphasis on reliability and maintainability. This is reflected in the maintenance costs, which are 0.25% of the purchase price per month compared to 0.5-1% for most other manufacturers.

The increase in system reliability is due to:

- A decrease in component count, because of the high-density chips
- Error check and correcting (ECC) bits in memory
- Extensive internal checking
- Extra heavy-duty power supply.

The increase in maintainability is due to:

- Improved microdiagnostics that will pinpoint the failing field replaceable unit (FRU) 80% of the time and reduce the problem to one of three FRU's most of the remaining times.
- Maintenance of the systems by swapping good FRU's for bad, which are repaired at a central site and not at the customer site.
- A separate maintenance console with microdiagnostics that allows the 4331 to be quickly switched on-line to a remote service center for error checking and test.

Software

The software for the 4331 is basically the same system software as used on the System 360/370. The primary operating system for the 4331 will be 370 DOS or the enhanced version, Extended Control Program Support: Virtual Storage Extended (ECPS:VSE), which uses the new 4300 architecture and puts certain common routines into microcode.

Another popular operating system for the 4331 will be the Virtual Machine (VM) operating system. The VM system supports a simple interactive time sharing system and will be the primary operating system for interactive environments.

With the 4331's relatively slow speed and small memory, it cannot adequately handle the larger MVS operating system. Additional software is needed to make the 4331 a high performance interactive system, competitive with the minis. For example, an interactive VAX will support two to three times the number of users as a 4331.

The software prices have been revised, not only for the 4331, but for all IBM systems. If a user elects to use 370 architecture and 370 software, then the 370 prices apply. If the user elects to upgrade the soft-

ware to take advantage of the new 4331 architecture, then there are additional software charges.

This improved software costs over $1000 per month, or 20% of a fully loaded system price. In a more typical system, the extra software cost will be from $300-$800 per month, or 15-20% of the total system price.

Table 2 summarizes the strengths and weaknesses of the 4331.

Table 2

IBM 4331 STRENGTHS AND WEAKNESSES

STRENGTHS

-370 Compatible
-Low price
-Low manufacturing cost
-Made by IBM
-Batch software
-High reliability, availability, serviceability
-Fast I/O
-Low electrical power
-New disk drives

WEAKNESSES

-Poor performance in an interactive environment
-Slow performance compared to minis
-Maximum main memory size of 1 million bytes
-Requires a system programmer

3. RELATION TO THE OTHER IBM PRODUCTS

In price, the 4331 competes with two other IBM products:

- The 8100—Data Processing Division (DPD)
- The System/38—General Systems Division (GSD).

Each of these systems has certain strengths and weaknesses which aim them at certain markets:

- The 8100 is a low-cost system and has two models. One model is priced lower than the 4331, and the higher-priced 8100 model is just about equal in price/performance to the 4331.
- The System/38 also has two models. The lower-priced is about equal in price/performance to the 4331, and the higher is midway in price between the 4331 and 4341.

The characteristics used by IBM to segment the market among these systems are:

- Is the application system design and development controlled and implemented at the local user site or at the central computer site?
- Is a systems programmer required?
- Are the applications batch- or transaction-oriented?

The managerial control of the system is a strong issue to IBM because of the sales territory for the DPD and GSD sales forces. If a company is centrally controlled and is not a first-time user, it is DPD sales territory. If the control is at the local site or is a first-time user, then it is GSD sales territory.

The use of systems programmers is related to the degree of customer sophistication. The smaller and less sophisticated the customer, the less desire or need for an expensive systems programmer. This definition also coincides with the GSD sales territory.

As the systems become larger and more complex, the talents of a systems programmer are required.

All three of the low-cost IBM systems can support a communications environment, such as remote job entry or satellite processors, but the market distinctions are in the managerial control and first-time user issues. Table 3 summarizes the market positions of the three IBM systems.

Table 3

MARKET POSITIONS

	4331	8100	System/38
Batch/Transaction Oriented	Batch	Transaction	Transaction
System Programmer Required	Yes	Yes	Usually No
System Design & Control	Local	Central	Local
Sold By	DPD	DPD	GSD
User	Sophisticated	Sophisticated	Unsophisticated

Omitted from these definitions and, therefore, a weakness in IBM's low-end computer line, are some of the traditional mini markets:

- Interactive computations
- Scientific computations
- Process control
- Transaction processing.

The Series/1 addresses some of these markets but with a much lower price/performance than the 4331.

4. BURROUGHS CORP.

Summary

Burroughs should feel little impact from the 4331 and stands to maintain its position in the small-business-system marketplace.

Burroughs announced two machines to counter the 4331. In its small stand-alone small-business-system line, it announced a dual-processor 1885 configuration. However, this is sufficiently below the 4331 in both price and performance (about $30,000 in price and 20-30% less in performance) that it was not analyzed further.

The most important machine in the 4331 price/performance range is the 2930, which is available as either a uniprocessor or in a multiprocessor configuration with up to four central processors.

Although the 2930 protects Burroughs from IBM on a price/performance comparison, Burroughs is still exposed without a mainframe-type system in the $150,000-200,000 range for a complete system, well under the 2930's minimum system price of $290,000.

The strength of the 2930 will be in the transaction- and communication-oriented markets, where the combined communication front-end processors and the network software will make this a strong contender in the low-price market.

Price/Performance

The single CPU 2930 system offers a MIPS rating of 0.3, or 50% faster than the 4331's MIPS rating of 0.2, and at a price for a minimum complete system of $290,000 leading to a price per MIPS about 7% higher than the 4331. A multiprocessor system lowers the price per MIPS by about 30% per processor, bringing Burroughs to the same price/performance as IBM, but at about twice the price—$340,000 for a minimum system.

The 2930's strength is in a communication environment, where the 4331 is relatively weak. The 2930 uses programmable front-end processors to handle the communications, and their effective MIPS processing rate must be added to the main processor to provide a better comparison. The 2930 also uses separate integrated microprocessors for the disk controllers, further increasing the effective MIPS rate. The 4331 has integrated disk controllers, but they time share the central processor functions. Therefore, the Burroughs 2930 has a much lower price per effective MIPS than the 4331.

The basic CPU and memory of the Burroughs is over twice the price of the 4331—$140,000 versus $68,000 for the latter. Part of this price difference is the disk controller included in the Burroughs price.

Burroughs' main memory, at an incremental price of $15,000 per million bytes, is tied with IBM and lower than any other surveyed machines. The memory uses 16k-bit chips and will probably result in lower profit margins for Burroughs than other parts of the product line.

The logic circuitry of the 2930 is quite advanced, using a chip with 800 gates, compared to IBM's new circuitry containing 700 gates per chip.

Peripherals

Burroughs' disk drives are based on the old removable 3330 philosophy and are priced at $140 per million bytes versus $40 for IBM's new 3770. Because the drives use an older technology, they are slower, with access times on the order of 30 ms. and data rates of 1 million bytes per second versus 20 ms. and 1.8 million bytes per second for the IBM 3770.

Reliability and Serviceability

The Burroughs maintenance and reliability philosophy is similar to IBM's in that the 2930 has a separate maintenance processor. While the IBM maintenance processor can only operate when the system is stopped, the Burroughs maintenance processor operates continuously and simultaneously with the main processor to detect faults.

Part of the monitoring process includes monitoring external power fluctuations and temperature to assure proper system operation. Only the 4331, among the surveyed systems, also offers this feature.

The maintenance philosophy is reflected in the maintenance prices, which are relatively low, 0.7% of the purchase price per month. This is the lowest, next to IBM, and a standoff with NCR among the mainframers.

The high-reliability philosophy, coupled with the multiprocessor configuration, gives Burroughs a very high availability, which is needed for the transaction market.

The transaction-processing market is discussed more completely in the section on Tandem.

Software

The Burroughs software, although more than adequate, is not as extensive as some of the others, and its data base management system, priced at $17,000, is among the most expensive.

On a total system basis, Burroughs has one of the lowest software prices, at 6-7% of the total system price compared to 15-20% for IBM.

5. DATA GENERAL

Summary

Data General is one of the companies that will be somewhat impacted by the 4331, but in an indirect fashion.

As a major minicomputer manufacturer moving from the traditional markets of OEM, scientific, real-time response, etc., into the end-user markets, especially the IBM customer base, Data General now faces a new competitor in IBM. With this new alternative solution available, many users will elect to remain with IBM, especially those who are batch-oriented.

Price/Performance

Data General has been among the most aggressive leaders in technology, and the price and performance Data General offers bear this out.

The Data General system that comes closest to the IBM 4331 in price is the Eclipse 350 (C350). On a total system basis, a C350 runs from $150,000 to $500,000 compared to the IBM 4331 price of $180,000 to $575,000. On a CPU-only basis, the price is comparable to the C350 priced at $54,000 versus the 4331 at $68,000.

The C350 is a new machine, announced in the fall of 1978, and fills what had been a gap in the complete line of business-oriented processors offered by Data General, ranging from the C150 at $80,000 for a complete system to the M600, which can cost $1 million for a complete system.

On a performance basis, the C350 ranks about the middle of the surveyed systems, with a MIPS rating of

0.4, after adjustment for the fact that the C350 is a 16-bit machine. Benchmark testing confirms this rating but, depending upon which benchmarks are used, the C350's performance could range from 0.2 to 0.4.

Buyers of the C350 can adjust the system performance by the type of memory they purchase. A noninterleaved memory is standard, but a two- or four-way interleaved memory can be purchased, raising the performance by 15-20% for a nominal $2,000-6,000 price, depending upon memory size and degree of interleaving. With interleaving, the C350 becomes equivalent to a 32-bit wide memory in speed.

On a price per MIPS basis, the C350 is priced at $375 to $1,300, depending upon system size, less than one-half IBM's $900 to $2,800.

The high end of the Eclipse family is capped by the M600 model which, on a system basis, is about $50,000 more than the C350.

Although the main processors of the C350 and M600 are identical, the M600 produces significantly faster performance through an additional high-speed arithmetic unit, up to eight-way memory interleaving, and a front-end/back-end multiprocessor architecture that increases the effective MIPS rate by adding the speed of the two processors.

Input/Output

The standard I/O speed for the Eclipse series is relatively slow, with a speed of 1.5 million bytes per second. Burst I/O rates are higher.

Peripherals

Data General's peripherals appear adequate for the market segments they currently occupy. However, IBM's 3370 is clearly superior. The largest disk drive offered is 277 million bytes, with an incremental cost of $136 per million bytes versus IBM's $40 per million bytes. For a large, billion-byte data base-oriented system, this would result in a price difference of $100,000 per system.

Data General was one of the last minicomputer companies to react to IBM's new memory pricing stategy by cutting the C350 price from almost $70,000 per megabyte to $27,000. While the latter is significantly higher than IBM's $15,000, it is less than a 2% difference on the total system price.

Reliability and Serviceability

The maintenance and reliability of the Eclipse is standard for a minicomputer with few extra features to aid in maintenance except for microdiagnostics. Data General, however, has given some extra attention to power supply design, a major part of system unreliability, allowing the sytem to operate in a 20% undervoltage environment.

The maintenance prices of the Eclipse are among the highest of the surveyed machines, running about 22-23% of the total system price, or about 0.8-0.9% per month of the hardware price.

Software

On an overall basis, the Data General software is one of the weakest of the surveyed systems, although Data General has strengthened it with several recent improvements.

Data General offers no data base management system, although one may be bought from a third party.

The software offering of Data General is relatively strong in the remote job entry (RJE) environment, allowing the system to act as a remote processor to an IBM system. IBM 3270 emulation and X.25 support are also available.

The COBOL offered compared favorably with other minicomputers but is substantially lacking in features to be able to compare it to a mainframe COBOL.

Data General's new operating system, AOS, first offered in 1977, has many good features and is a substantial improvement over the earlier RDOS.

6. DIGITAL EQUIPMENT CORP.

Summary

Digital Equipment Corp. (DEC), along with Prime, may be one of the two companies most strongly impacted by the IBM 4331. This impact will be evident in a possible lessening of the growth rate rather than in an outright decline in sales, but it will be difficult to measure since DEC and Prime may find it harder than before to penetrate IBM's user base and may put their efforts into other markets, thus continuing at the same growth rate as in the past.

There are three DEC products that compete against the 4331 on a price basis. The first is the DEC 2020, a 36-bit system which DEC considers a mainframe. The second is the VAX 11/780, DEC's 32-bit minicomputer. The third is the older PDP 11/70, a 16-bit minicomputer. While the latter is the most popular of the three DEC systems, it will not be the main future competitor in the commercial marketplace and therefore was not further analyzed.

The VAX was first announced in 1977, and although it is upwards compatible from the older PDP/11 family, in its own native mode it is a completely different system offering full 32-bit instructions and virtual memory capabilities. The VAX 11/780 is the main future thrust of DEC. The single model in this family is expected to be extended in the near future with additional members.

Markets

The DEC System 20 is primarily aimed at the interactive computing market, and about 600 units of all models have been sold. The model 2020, the smallest and most comparable to the 4331, has had about 200 units installed for a total value of about $100 million.

The major market tapped by DEC for the 2020 has been the in-house time sharing market. Many companies have switched from time sharing service bureaus to an in-house 2020 and achieved a reduction in expenses plus an increase in capacity.

Because of its timesharing orientation, it is unlikely that the 2020 will be impacted by the 4331 in this submarket. However, about half the 2020's have been sold as stand-alone batch processors to replace existing systems, including 360/30's and other older IBM equipment. This portion of the 2020 market will be severely impacted by the 4331, if not dried up completely. This is a relatively small dollar volume for DEC, at $50 million per year, and will have negligible impact on DEC's total corporate picture.

The initial market for the VAX is the traditional minicomputer market of scientific and interactive work. Some early users have commented that, with the interactive programming development capabilities of the VAX, software development costs have been one-half to one-third that of other systems. As the VAX software improves, this system will be extended more into the commercial data processing marketplace.

The 4331, while not as fast as the VAX, will slow DEC's entrance into the batch market by providing the IBM user with a fully compatible machine for upward growth.

Price/Performance

The DEC 2020 is a 36-bit computer, aimed at the mainframe market. Its MIPS rating is 0.3, 50% faster than the 4331's 0.2, and its price is also 50% higher than the 4331, running from $300,000 to $900,000 for a complete system.

On a price per MIPS basis, the 2020 is slightly more expensive than the 4331, at $1,000 to $2,600 per MIPS versus the 4331's $900 to $2,800, depending upon system size and configuration.

Recent user benchmark on the 2020 found that in the interactive environment it was considerably more productive than either the 4331 or the VAX 11/780.

The VAX picture is different. The VAX 11/780 has a MIPS rating of 0.9, the highest of any of the systems analyzed except the Tandem dual-processor system, and a system price about 20% higher than the 4331, at $210,000 to $660,000.

The price per MIPS of the VAX was among the lowest of any of the systems analyzed, ranging from $200,000 to $700,000 per MIPS, one-fourth that of the 4331 and ranking third behind the Tandem system and the Perkin-Elmer 3220.

DEC's main memory was originally priced at $40,000 per megabyte, but after the 4331 announcement, DEC reduced this price down to a very competitive $18,000 per megabyte.

Peripherals

As a batch processor, the 2020 has relatively good peripherals, with 176-megabyte-capacity disk drives and a 1.2-million-byte-per-second tape drive. These were the equal of the IBM equipment until the announcement of the new 3770 disk on the 4331.

The VAX has a very high I/O rate for its price, at 8 million bytes per second, three times that of the 4331.

The peripherals for the VAX are DEC's 176-million-byte RP06 disk drive, the same as used on the 2020, and a very slow 45-inch-per-second tape drive. This tape is unsuitable for batch processing and can only be used for disk backup, further underlining the fact that DEC is not seeking to sell the VAX in the batch market.

The RP06 disk drive has the highest price of any of the systems analyzed, with an incremental price of $193 per million bytes, almost five times that of IBM's $40 per million bytes.

DEC's CRT, the VT 100, is an extremely competitive device, and next to IBM's 3270 has proven so popular that there is now up to a two-year wait for a terminal. It is priced at $1,900 for a plain terminal, or $2,200 for a 132-character-per-line version with extra features. This compares to $3,000 for a plain 3270 and more for more sophisticated versions.

Reliability and Serviceability

DEC is one of the few companies that places strong emphasis on reliability and maintainability. The VAX has a separate maintenance computer (an LSI 11) for monitoring the behavior of the system and performing diagnostic work. This emphasis is paid back to the user with low maintenance prices that run about 0.5-0.6% of the purchase price per month, well below the "normal" mini price of 0.9% and the lowest percentage of any of the companies except IBM. On a large system, IBM's maintenance cost is 15% of the total price versus the VAX's 16%, a negligible difference. The next lowest is Tandem, at 18%, while most systems are in the 20-22% range.

The DEC 2020 does not place as much emphasis on reliability as the VAX, and its maintenance prices are about 20% higher, but still under those of a typical minicomputer.

Software

DEC bundles most of it software into the system price, reducing its apparent software costs and raising its apparent hardware costs.

The VAX was designed from the ground up as a completely integrated system; that is, the software and hardware were designed together, not separately. However, like most new systems, it has not yet reached its full potential. In order to bring the machine to the market as quickly as possible, DEC used much of the older PDP/11 software and is gradually phasing in the native mode software to achieve VAX's full potential.

An example is the COBOL compiler. The original PDP/11 version is very slow and, on benchmark tests, was one-half the speed of the DEC 2020 COBOL, a machine one-third the MIPS rating of the VAX. The new COBOL is being phased in by DEC and is approaching the native speed in its capabilities.

The current line-up of native mode software is still relatively weak and includes only an operating system, DECNET network program, COBOL, and FORTRAN compilers. No database management systems exist, and RPG, PL/1, and APL languages are not supported.

In the interim , the full range of PDP 11/70 software is available for use on the VAX, and this is the widest range of software available for any minicomputer.

The DEC 2020 has a relatively full complement of software, with a good COBOL compiler, DECNET networking software, and a database management system.

7. HEWLETT-PACKARD

Summary

Hewlett-Packard (HP) is in an ambiguous position relative to the 4331. On a straight price/performance chart, the HP 3000 is the lowest ranked except for the Univac 90/30, which shouldn't be compared as it is completely noncompetitive. However, from a sales standpoint, the 4331 has had a negligible impact on HP, and the future impact will also be relatively light.

Market

The HP systems are aimed at a completely different market than the 4331. They are primarily aimed at first-time, unsophisticated users for interactive- and transaction-oriented applications. The database software on the HP 3000 is the most sophisticated of any of the minicomputer systems; this greatly simplifies software development, making it easy for new users to program.

HP's user orientation plus its strong IBM-compatible RJE communications software has allowed HP to make a significant penetration into the IBM user base as a distributed processor vendor. Because the HP 3000 is designed to be easy to program by end-users at remote sites, HP will be impacted more by the System/38 than by the 4331 or even the 8100 series.

Price/Performance

HP has two systems that are price competitive with the 4331. They are both relatively new. The larger 3000 Series III was first delivered in late 1978; the smaller model 33 was first delivered in early 1979. The latter uses HP's advanced Silicon-on-Sapphire (SOS) technology, which reduces the CPU logic to three chips.

These systems bracket the 4331 in price: The model 33 is 30% lower than the 4331, and the Series III is about 10% higher.

On the MIPS scale, both HP systems have less performance than the 4331, by at least 40% for the Series III and at least 60% for the Model 33. In absolute speeds, the MIPS ratings are 0.07 and 0.12, respectively, compared to 0.2 MIPS for the 4331.

On a price/performance basis, both HP systems are nearly twice as expensive as the 4331, with price per

MIPS running from $1,500 to $7,000 versus the 4331's $900 to $2,900, depending upon system size and configuration.

It is difficult to make a direct comparison based only on a MIPS analysis because of major internal structure differences between the HP and IBM computers, such as the fact that the 3000 is a 16-bit minicomputer and uses stack architecture for intermediate storage, while the 4331 is a 32-bit machine and uses registers for storage.

The number of terminals supported is a better measure of performance than MIPS for this environment. Because HP's systems are designed as interactive systems, they will support more terminals than a 4331. The model 33 will support 10 to 15 terminals, compared to IBM's 8 to 10, while the Series III should support 20 to 30 terminals.

HP has reacted to IBM's announcement by cutting in half the incremental cost of main memory from over $30,000 per million bytes to $15,000 per million bytes. This has a very small effect on a total HP system price, reducing the total price less than 5% for a typical system. Its primary impact—the ability to say that HP offers the same price as IBM—is psychological.

HP also reacted to the IBM 4331 by dropping CPU prices on both the Model 33 and Model III and bundling some of the software into the hardware price. The net result of these three changes was a total system price reduction on the order of 10%.

Input/Output

For a batch environment where the I/O speed is important, the 3000 may be slower than the 4331. The Model 33 has an I/O speed of one million bytes per second; the Series III has a speed of 2.8 million bytes per second, serving a single device. This compares to the 4331's aggregate data rate of 2.3 million bytes per second and the ability to serve two 1-million-byte devices simultaneously.

Peripherals

Compared to the 4331, HP peripherals are very limited. The largest disk drive offered by HP has a capacity of 120 MB and a restriction of 8 drives, or 1 billion bytes per system. The Model 33 has an even smaller capacity. The 4331 has a capacity of nine billion bytes of on-line storage.

The tape drive offered by HP, at 1600 BPI and 45 IPS speed, is too slow to be used in a batch environment but is adequate for backing up the small HP disk drive. The speed and capacity are inadequate to back up a larger drive, and if HP announces a larger disk, a new tape drive will be needed.

The HP CRT's are very easy to use and very sophisticated. HP was one of the first to offer a movable screen to improve operator productivity and the first to put function keys directly on the CRT, making them easier to use.

Reliability and Serviceability

HP's systems designs are strongly oriented toward high reliability and low maintenance costs, and HP is one of the few vendors that offers remote debugging capabilities for its system. A separate maintenance support processor is used as a diagnostic processor.

HP maintenance prices are less than those of most other manufacturers, accounting for about 20% of the total system price, two to five percentage points lower than most others. When comparing just the hardware, the maintenance costs run about 0.5% of the purchase price per month, well below the typical minicomputer cost of 1% per month and about 20% above IBM's prices for a large system.

Software

HP software is among the strongest offered on any of the minis. It is a completely integrated and proprietary system. The major thrust of the software is interactive- and transaction-oriented. The database and file inquiry languages are among the best.

Neither the HP software nor hardware is oriented toward the batch processing market. The COBOL is among the weakest, supporting only the 1968 version, not the 1974 like most other companies.

HP is also going counter to the industry trend of unbundling software. Its latest announcements bundled in previously separated software. Because most of the software is bundled in, HP's apparent software price is the cheapest, 5% of the total system price.

Bundling reduces the number of options and makes it easier for an unsophisticated user to buy from HP. All of this keeps pointing HP to its ultimate end-user market, the first-time user.

8. HONEYWELL, INC.

Summary

Honeywell, like Univac, can compete against the 4331 from two different levels—minicomputer and mainframe. At the minicomputer level, Honeywell competes with the Level/6 Model 53 and at the mainframe level with the Level/64 Model DPS 320 (Distributed Processing System).

Both the Level/6 and Level/64 now have nearly identical MIPS ratings and system prices, but they differ vastly in applications approach. Neither should be strongly impacted by the 4331 in their typical markets.

Markets

The Level/6 is a typcial 16-bit minicomputer, with good software for a mini but poor for a mainframe. A strong data communication capability and supporting software make it a good distributed data processing product. The software support is not aimed at the first-time user market; it is more complex and harder to use than many other minicomputer offerings. The Level/6 does not have the raw performance of other 16-bit minis in its price class.

The Level/64 Model DPS 320 is also a strong communication-oriented system. Its primary purpose is to protect the Honeywell installed base of H200 class systems. There are over 800 of these systems installed, with an aggregate purchase value of about $1 billion. The Level 64/DPS is expected to extend this base into the Distributed Processing System market.

Both Honeywell systems are strong contenders in the communication environment. The Level/6 is stronger for small transaction-oriented systems; the Level/6 will be more competitive to the 8100 than the 4331.

Price/Performance

On a price per MIPS basis, the 64/320 is lower than the 4331, but does not provide the usual 20-30% price/performance advantage users would expect for switching from IBM. As a result, Honeywell salesmen will be constrained in attempts to capture new customers from the IBM base.

Honeywell also offers a larger version, the 64/350. It is 50% faster, with a MIPS rating of 0.3, at a price increase of 50% for a small system, resulting in a price per MIPS slightly higher than the 64/320.

The 64/350 aims at a gap in the IBM line by offering a system at a price between 4331 and 4341. However, this gap should be plugged by IBM with the expected upgrade announcement to the 4331. (Subsequent to

the detailed comparisons used in this analysis, Honeywell announced the Level 64/DPS-330, which supersedes both the 320 and the 350. This new machine offers approximately 10% greater performance than the DPS-350 and should be a strong contender in this segment of the marketplace.)

The Level/6 minicomputer system model most directly price competitive with the 4331 is the Model 53, announced in early 1978. The basic instruction rate of the Model 53 is slightly faster than the 4331; however, the Model 53 is only a 16-bit machine. To some extent, the 16-bit deficiency is overcome by a two-way interleaved memory which has many of the speed characteristics of a 32-bit wide memory.

On an overall basis, the adjusted MIPS rate is 0.2, equivalent to the 4331. The price per MIPS, on a system basis, is substantially lower than the 4331, at $725,000 per MIPS, because the CPU is less expensive.

Peripherals

On an incremental basis, Honeywell's prices are among the best. The fixed media disk drive is priced at $50 per million bytes versus IBM's $40 per million bytes. Main memory has an incremental price of $20,000 per million bytes versus IBM's $15,000. Both differences, while sizeable in percentage, are negligible in terms of total dollars for the system sizes analyzed.

Reliability and Maintainability

Neither of the Honeywell systems is particularly noteworthy in the maintainability or reliability areas. Only single processor systems are offered, and the maintenance prices are in the middle ground, running about 20% of system prices.

Software

Honeywell has relatively mature software, with both Level 6 and Level 64 software having migrated through several generations.

The Level/6 software is not as first-time-user-oriented as other minicomputer software, and more systems knowledge is required to use it.

Both Level/6 and Level/64 software have strong communication capabilities. The Level/6 offers the TOTAL data base management software, from CINCOM. The COBOL is quite weak when measured against IBM or even against Honeywell's own Level/64 COBOL system.

The Level/64 offers strong data base management systems support languages.

9. NCR CORP.

Summary

NCR responded quickly to the 4331 with a brace of products bracketing the 4331. The new 8455 is 10% under the 4331 in performance, while the new 8555 is 50% faster.

These machines provide a powerful base for NCR to continue its aggressive expansion program. The impact NCR will feel from the 4331 will be relatively mild, as the new systems are not only strong enough to protect the current NCR base but also provide a reasonable alternative to the 4331, especially to those customers who cannot wait for the long lead times on delivery of the 4300's.

From a hardware and software standpoint, NCR is in the strongest position among IBM's mainframe competitors for the following reasons:

- NCR offers a full range of five compatible systems ranging in price from $100,000 to several million dollars.

- NCR's new fixed media disks compete with IBM in price at $50/MB, compared to IBM's $40/MB, a negligible difference even for a large system.

- NCR offers a full multi-processor virtual memory operating system.

- NCR's strong emphasis on reliability and maintainability give it one of the lowest maintenance prices of any mainframe vendor except IBM, at 0.7% of the system price per month versus IBM's 0.5% per month.

- NCR's strong communication emphasis is combined with low-price integrated communication controllers on the smaller machines.

- NCR's main memory price of $20,000/MB is among the lowest of any vendor and is inconsequential in difference on a system basis versus IBM's $15,000/MB.

Markets

Markets for the 8455 or 8555 will be found among NCR's traditional customers, primarily in banking, retailing, and manufacturing.

No announcements were made concurrently with the new processors that would suggest any substantive change in NCR's marketing strategies.

System Description

All the new systems use NCR's new Virtual Resource Operating System (VRX) to provide a full virtual memory capability. The larger of the machines, the 8555, is also available in a multi-processor configuration to provide a high availability system with up to four processors. The 8555 dual processor system would be in the same price and performance class as the 4341.

NCR offers a further advantage in that its low-end Criterion series, which starts in price at $50,000 for a complete system, is upwards compatible with the 8455 and 8555 systems. This provides even the smallest user complete upwards growth visibility and offers the large user the ability to configure distributed processing systems, which are program compatible, with a very wide variety of processors.

Price/Performance

The 8455 hardware is less expensive than the 4331 by over 25%. However, this is not the total price seen by the user. NCR software is among the most expensive of all the companies examined, at approximately 20% of the total system price. When the extra software expense is added to the slightly higher maintenance costs, the total system price of the NCR 8455 is a wash with the IBM 4331.

The 8455 is 15% under the 4331 in performance, with a MIPS rating of 0.17, and about equal in price; the 8555 is 50% faster, with a MIPS rating of 0.3, and only 20% more expensive, with a system price per MIPS of about $733. For multi-CPU configuration, the 8555 would have an even lower MIPS price.

This gives the NCR 8555 the best price/performance ratio of the mainframe products analyzed.

Input/Output

The 8555 has the highest internal bus speed, 36 megabytes per second, of any of the systems analyzed. Despite the high internal bus speed, the I/O speed of the 8455 and 8555 is only two megabytes per second, about equal to the speed of the 4331.

Peripherals

Although NCR is competitive with IBM on disk drive prices, at $50 per million bytes, the operating capabilities are inferior. The average access time of the newest NCR drive announced in 1979 is 35 ms. versus IBM's 20 ms., or 50% slower. This time is most critical in a transaction environment. In a batch environment, the read/write speed could be even more critical, and this is 1.2 million bytes per second for NCR versus 1.8 million bytes per second for IBM's 3770—also a 50% difference.

Software

The software on the NCR is a reasonable configuration of languages and facilities, including a network control language. CINCOM'S TOTAL is used as the data base management system.

10. PERKIN-ELMER

Summary

Perkin-Elmer could stand to benefit the most from the 4331 announcement. Far from being adversely impacted by the announcement, Perkin-Elmer could seize the opportunity offered by IBM's slow delivery schedule to enhance its position significantly.

Perkin-Elmer's new 32-bit system, the 3220, first delivered in the summer of 1979, is similar to the 370 architecture and, therefore, is more directly comparable to the 4331 than most other systems studied.

Perkin-Elmer can deliver a system with 4341 capability at less than a 4331 price.

Perkin-Elmer's even newer 3240 system, announced in October 1979, could fill the price gap between the 4331 and 4341, while offering performance substantially better than that of the 4341.

The ability of Perkin-Elmer to take advantage of the 4331 opportunity will depend upon the speed at which it can react to the marketplace.

Markets

Perkin-Elmer's traditional marketplace has been in the scientific, engineering, simulation and related markets, and the 4331 opportunity would represent a major change in market direction. Management statements have identified a strategic desire to enter the commercial market. However, the specific marketing efforts to date seem tentative.

Perkin-Elmer's thrust into the commercial marketplace could be heavily aided by an IBM-compatible version of the 3220. A minor modification to the hardware and software would allow the 3220 to run any 360 and 370 user-written code without modification. This would include all of the IBM compilers and any code generated by the compilers.

The 3220 is not a PCM in that it cannot run IBM operating systems, nor can it plug into IBM peripherals, but it can run the user object code.

Price/Performance

Typical system prices for the Perkin-Elmer 3220 are $135,000 to $560,000, compared to IBM's $180,000 to $575,000.

On a system price per MIPS, Perkin-Elmer's 3220 is less than $200 per MIPS, well below the 4331 at $900 per MIPS.

Input/Output

The I/O speed of the 3220 matches its internal speed with a rate of eight million bytes per record, or about three times faster than the 4331.

Peripherals

Perkin-Elmer uses the CDC storage module, as do many other minicomputers, for disk drives and has a relatively high price of $164 per million bytes compared to IBM's $40—a very significant price difference. However, CDC is just beginning delivery of its Winchester-style drives that are interface-compatible to the storage module disk drives, which could improve Perkin-Elmer's relative disk drive prices.

Perkin-Elmer also offers 6250 BPI tape drives if required, the same speed as IBM's drives.

Perkin-Elmer makes a very low priced CRT that sells for about $500 in quantity. The use of this CRT substantially reduces the price of a large system.

Reliability and Maintenance

From a reliability and maintainability viewpoint, the 3220 is a middle-of-the-road machine. The main console can be used to initiate diagnostics that can be run to determine gross machine faults, and an error-logging facility keeps track of current and potential hardware problems. On a price basis, the PE 3220 is right in line with other minis, with maintenance prices at 0.9% of the purchase price per month, well above IBM's 0.5%.

Software

Perkin-Elmer offers RJE support to enable the use of the 3220 as a remote job entry system to a 370.

The FORTRAN compiler is one of the most sophisticated available, having an excellent optimizing capability.

The rest of the Perkin-Elmer software is relatively weak. There is no data base management software or networking capability. The COBOL is comparable to other minis but lacks features that would make it acceptable in the mainframe world.

There is no support for PL/1 or the APL languages.

An interactive operating system provides a reasonable level of support for interactive processing, and a transaction oriented monitor will support 30-50 terminals for transaction processing.

11. PRIME COMPUTER

Summary

Prime could be among those most seriously impacted by the 4331 announcement, but in an indirect fashion, by reducing the opportunity to continue to invade the IBM base.

Markets

Prime's initial market thrust was in the scientific and interactive processing market, where Prime's high speed and virtual memory capability provided a significant advantage.

Prime had this market niche to itself until DEC's VAX 11/780 provided a superior product on the higher end and WANG provided virtual memory systems both above and below Prime in price.

The new Prime 50 series, announced in January 1979 (models 450, 550, 650, and 750), provided an answer to DEC and gave Prime a more competitive machine.

Prime's marketing emphasis has been switching to the commercial market and, even more recently, to that part of the OEM market that sells to the commercial end-user with a version of the Reality software originally developed by Microdata.

The OEM market now accounts for 10-20% of Prime's revenue. While many of the OEM's customers are first-time users and unlikely to be impacted by the 4331, others which do have IBM 360 or 370 systems installed may elect to remain with IBM's compatible 4331 and save on their programming costs.

Prime's market strength will continue to be in interactive applications where Prime's easy-to-use software will be a big advantage against IBM's more difficult-to-use 4331 software.

Prime will also compete in the new-user market without undue pressure from IBM's 4331, but may face new competition from the System/38 in this market.

Price/Performance

The 50 series model that is most price-competitive with the 4331 is the Model 450, with total system prices ranging from $170,000 to $620,000 versus IBM's $180,000 to $575,000.

The 450 has a MIPS rating of 0.3 versus the 4331's 0.2, leading to a system price per MIPS advantage for Prime of less than $600 per MIPS versus IBM's $900 per MIPS. On a CPU-only basis, Prime maintains about the same pecentage basis, with a price of $200 per MIPS versus IBM's $300.

If the slightly larger Prime 550 is used for comparison, the system price per MIPS drops below $500. The 550 provides a 30% performance advantage, 0.4 versus 0.3 MIPS, for a 10% price increase when compared to the 450.

For hardware only, the small configurations of the Prime 550 and the 4331 are identical in price at $128,000 each, but the higher Prime maintenance costs raise the total effective system price.

Both of these Prime systems' price/performance ratings compared poorly against DEC's VAX, which has a 0.9 MIPS rating and a system price slightly over $200 per MIPS.

Prime's main memory price of $40,000 per million bytes is high compared to IBM. However, on a systems price basis, this will amount to only a 5% difference, a relatively insignificant figure.

Input/Output

The I/O speed of both Prime systems is 2.5 million bytes per second versus IBM's 2.3 million bytes per second.

Peripherals

Prime uses the CDC storage module drives with a capacity of 300 million bytes per spindle. These are incrementally priced at $110 per million bytes—less than other minicomputer companies but considerably more than IBM's $40 per million bytes. This difference will not be apparent on a total system price until the data base reaches a billion bytes. The resulting $50,000 price difference for disks would result in a 10% difference for the total system. This partially accounts for Prime's higher price on larger systems.

Reliability and Maintenance

Prime's reliability and maintenance philosophy is better than average for minicomputers. Prime is one of

the few companies that offer remote system debugging capabilities. This capability is not reflected in maintenance prices, which are at the minicomputer standard of 0.9% of the purchase price per month.

The maintenance prices for a larger system differ by over $1,500 per month in favor of IBM, adding an extra $65,000 to the purchase price of a large Prime system.

Software

Prime offers all main languages, including PL/1, COBOL, RPG, and FORTRAN.

Prime's software is relatively strong for a minicomputer company, with a full complement of languages, a data base management system, networking capability, and source-level debugging package.

The COBOL compiler is weak from a language facilities standpoint, missing many features used on large systems. It is also a relatively poor operating compiler, resulting in slow execution speed of the generated code. These factors make it more difficult for Prime to compete in the standard commercial market.

Prime's communication software offers full RJE support and a networking capability. Prime is also one of the few companies to offer the X.25 communication capability which is important in Europe, and just beginning to catch on in the United States.

Prime is one of the few mini companies supporting a library of application packages. Most of these are oriented to the scientific, rather than the business, community.

The PRIMOS operating system, like other minicomputer software, is oriented toward the interactive market and provides a reasonable level of support for transaction processing, although not to the same degree as an operating system designed specifically for transaction processing. It offers minimal support for batch processing.

The software strength and the ease of program development provides the main market strength for Prime.

12. TANDEM

Summary

Of all the companies analyzed, Tandem is the least likely to feel the impact of the 4331 announcement. The primary reason for this is that Tandem has targetted the transaction processing market, and the 4331 is not a good price performer for this application.

Markets

The transaction processing market is currently relatively small, under one billion dollars per year, but it could grow to 10-20% of the total EDP market by the middle 1980's.

The user applications that characterize this market are those in which a terminal operator keys a transaction that is processed against a relatively large data file, and only a small amount of data is returned to the operator. Specific functional requirements are:

- Large data bases
- Complex decision-making on how to process the data
- Multiple non-EDP users
- Relatively light computation

- High system availability
- High system reliability
- Relatively low data transmission per transaction
- High data base integrity.

On-line data entry, where data are captured but not processed, constitutes the bulk of the current on-line distributed data processing market. This is not the same as transaction processing. Time sharing and interactive computations as found in engineering or planning type applications are also different from the transaction market.

High availability and high reliability, although related, are not the same system design characteristics. A system with high reliability could have low availabilty or vice versa. A system with hundreds of operators that is unavailable for even short periods, idling the operators, would be very expensive in lost business and operators' time. A financial transaction system that processes a single million-dollar transaction incorrectly would be unacceptable, even though the system had high availability. These applications put a high premium on data base integrity.

Performance Characteristics

The MIPS rating has less value in measuring systems performance in a transaction environment. The important measurement is number of transactions per hour. This measurement is related to the total system throughput, which includes:

- Hardware speed
- Applications software
- Systems software
- Data base software
- Disk and terminal I/O speeds.

It is often impossible to determine how many transactions a system will process until the system is designed and built. It may also be difficult to estimate the number of transactions to be processed because volumes may grow drastically as users become familiar with and learn to use the system.

Tandem, with its unique multi-processor approach, is able to satisfy these multiple requirements:

- Modular growth
- Availability
- Reliability.

If a user misjudges the transaction processing requirements of his/her system, he/she may plug in another central processor at a small incremental cost to provide the additional computing capability. Adding the processors is typically a one- or two-hour job and can be done while the system is running.

The 4331 is available only in a single processor configuration, thus not solving the user's growth problem in nearly as satisfactory a manner as Tandem.

The largest installed Tandem system currently uses 14 processors. It is possible to use 16 processors in one system. Up to 255 systems may be connected in a distributed data processing network. The multiprocessor configuration also provides high availability by automatically sharing the work load among the other processors if one processor stops working.

High reliability is obtained by extensive self-diagnosis within the Tandem system.

Price/Performance

On a MIPS basis, a single Tandem processor is three times faster than a 4331 processor. However, Tandem systems are sold with a minimum of two processors, resulting in over six times the performance of a 4331 at approximately the same price.

On a transaction processing application, the Tandem system should support five to ten times the number of transactions that a 4331 does.

Each program in a Tandem environment can address 128k bytes of memory for program space and an additional 128k bytes for data. Each processor can support 256 programs running at the same time. A single processor can have 2 million bytes of memory, and a Tandem system can have 4 MB. of memory.

On a total system price, the Tandem system and the 4331 are comparable, differing by less than $5,000 for a small system and less than 15% for the large system, but Tandem's greater performance gives it a price per MIPS under $150,000 versus IBM's $900,000 per MIPS.

Peripherals

The IBM 3770 disk drives have considerably faster access time—20 ms. versus 28 ms.—than the CDC storage module used by Tandem. This could be a mitigating influence on the transaction rate.

However, the Tandem "mirrored Database" software can write two copies of the data on different disks. Each disk is assigned to perform half the reads, which can cut average access time to about 14 ms. In many on-line applications, the system performs four to six reads for every write. Therefore, the average access time in the Tandem system is considerably better than the raw hardware would suggest.

The IBM 3770 disk drive, incrementally, is about one-third the price of the Tandem disk, at $40 per million bytes versus $130 for Tandem. For a large system, this is a $100,000 difference. However, if Tandem upgrades to the CDC Winchester, which is interface compatible, the difference would be reduced.

Tandem uses the Lear Siegler CRT, one of the lowest-priced CRT terminals available. It is priced at one-half the IBM 3270 CRT, reducing the cost of a large Tandem system.

The I/O channel on the Tandem system is capable of transferring data at 4 megabytes per second. That is nearly twice as fast as the 4331 and gives further advantage to Tandem's total throughput.

Reliability and Maintainability

The 4331 design stresses reliable components, so that the single processor configuration will provide high availability. The design philosophy includes extensive testing and diagnostics to analyze and locate faults and facilitate quick repair.

The Tandem approach is completely different. Tandem has designed in extensive checking to make sure the processor is operating reliably, but once a processor fails, the workload is automatically switched to another processor. This approach reduces the component count needed for a single processor. Maintenance is performed by swapping a good processor for a defective processor, while the system is running.

The 4331 architecture, as currently defined, precludes a multiprocessor design, but even if the archi-

tecture expands to allow multiprocessing, major new software would be needed to allow more than two processors to be tied together. IBM's most sophisticated software provides for load sharing, but does not handle the dynamic case of taking the workload from a failing processor and giving it to another processor to continue where the first one left off.

The maintenance price for an IBM 4331 is half that of a Tandem system. The dollar difference, at $500 to $1,000 per month, depending on system size, translates to $20,000 to $40,000 in equivalent purchase price value in favor of IBM.

Software

A major factor in total system cost is the cost of developintg user programs. Most Tandem systems will be installed for new applications, often where the user is converting from a batch to a transaction processing environment.

Although the Tandem software is not as extensive as IBM's, it is sufficient for the applications most often approached.

Tandem offers all the major languages, such as COBOL, FORTRAN, and Data Base Management, inquiry programs, report generators, screen formatters, and networking software.

More importantly, the system is oriented toward on-line program development and transaction processing. Many facilities that would have to be specially coded for a 4331 are already available in Tandem's system. Tandem users report that program development, using Tandem's proprietary language, TAL, often takes one-third or less the time necessary on an IBM system. This could be a very significant dollar saving for a Tandem system user.

The Tandem software has considerably less overhead than the IBM software and provides substantially greater throughput in the transaction environment.

13. UNIVAC

Summary

Univac could be heavily impacted by IBM's 4331 unless it introduces new equipment within the next six to twelve months.

Like Honeywell, Univac competes in two areas—the mainframe market and the minicomputer market.

The V77-800, the largest of the Univac minicomputers, is the most likely competitor in the minicomputer market. Overall, the V77-800 is a highly competitive machine for a terminal-oriented environment. It will have less relevance in the batch market, which is the 4331 stronghold.

The 90/30 is the most direct competitor in the mainframe market.

The 90/30

The 90/30 is Univac's main small-end mainframe. Although it was competitive with the IBM/370, it is not competitive with the 4331, as it was introduced in 1974. The 90/30 is at least 30% slower and almost twice as expensive, for a price performance ratio two-and-one-half times worse than the 4331.

The main appeal of the 90/30 is to the approximately 200 users of the smaller 90/25 looking for a compatible growth machine and to the 2,000 customers of the much older 9200 and 9300. More importantly, there are approximately 2,000 current 90/30 users, with installed equipment valued at a purchase price of about $1 billion. These users could consider moving either to the 4331, for an increase in performance at one-

half the monthly price of the 90/30, or to the 4341, which would give a five- to six-times increase in performance for the same price.

To date, probably few, if any, Univac users have ordered the 4331. Most orders taken by IBM have been for customer upgrades. But the second round of competitive marketing could begin within the next year, and it is in that environment that the Univac sales force will definitely need better price/performance products in this area.

It is important to note that a substantial portion of the 90/30 customer base formerly consisted of IBM customers using the older 360/20's and 360/30's. Without a competitive offering, Univac might find some customers returning to the IBM fold.

The V77-800

Poor as Univac's mainframe competitive position is, the minicomputer picture is much better.

Univac purchased Varian Data Machines, a pioneer and one-time leader in the microprogrammed minicomputer market, in June of 1977. Varian was merged into Univac and is now managed under the same part of the corporate structure as the BC/7 small business computer.

In 1976 Varian announced the V77 family. Univac's first announcement since acquiring Varian occurred in 1978, with the V77-Model 800, which predated the 4331 announcement by a month.

The market for the V77-800 is the typical minicomputer market, with a heavy emphasis on communication-oriented applications. More recent announcements, such as an extensive manufacturing software package, signal Univac's intention to use the V77 in the end-user market.

The V77-800 should feel a negligible impact from the 4331, except for the typical minicomputer problem of the IBM customer who has no place to go and might have chosen Univac before the 4331 announcement.

Price/Performance

The V77-800 stands fourth in price/performance of all the surveyed systems, three-and-one-half times better than the 4331. The hardware is about $30,000 less for the small system; it is somewhat more expensive on the large system because of the costly Univac disk.

The raw MIPS performance of the V77-800 is almost three times that of the 4331, after adjustment for the 16-bit word of the Univac. Direct MIPS comparisons must be weighted by the fact that, like most minis, the V77-800 is a 16-bit machine, while the 4331 is a 32-bit machine, requiring twice as many memory fetches for the same amount of data.

The main memory speed for the V77-800, of 600 ns. for a 16-bit word, is considerably enhanced by 2k bytes of 150 ns. cache memory. With an expected hit ratio of 93-95%, the effective overall memory speed is 200 ns. Even after adjustment for the 16-bit wide memory, the V77-800's effective memory speed is two-and-one-half times as fast as that of the 4331.

The writeable control store used to microprogram the V77-800 is 64 bits wide and has a cycle time of 165 ns.

A major advantage of the Univac system is the ability to attach up to four systems together to provide a high-availability system. The software support for this configuration is relatively limited at the present time.

The CPU and main memory price of the V77-800 is almost one-half that of the 4331 and is the major reason for its price advantage. Main memory is among the cheapest of the surveyed systems, at $24,000 per megabyte. This is a negligible price difference in a one-million-byte system that sells for $400,000.

Input/Output

The I/O speed of the V77-800 is fast for a minicomputer, at two million bytes per second, close to the speed of the 4331's 2.3 million bytes per second.

Peripherals

Univac disks are relatively expensive at $125/MB. compared to the $40/MB. for the 4331. For the large system, with one billion bytes of disk, this could cause a $100,000 price difference.

Reliability and Maintainability

Univac's maintenance prices are in line with minicomputer standards of 1% of the purchase price per month, compared to the 0.25% on the 4331. On a total system basis, this is $5,000 to 10,000 per year, or $20,000 to 40,000 capitalized price to the user on an equivalent purchased system.

Software

The V77-800 software is relatively inexpensive, ranging from $7,500 for its small system to about $50,000 for all the features available on a large system.

The software coverage of the V77-800 is among the best of the minicomputers. It includes all the standard programming languages, RJE facilities, use of CINCOM's TOTAL, rated one of the best for data base management, and a networking facility.

The COBOL is competitive in the minicomputer arena but weak in an IBM competition, because it lacks features typically used in a batch processing environment.

Two different operating systems, VORTEX and SUMMIT, provide for tailoring the system to an interactive or transaction processing environment.

This report was prepared with assistance from a computer industry consultant.

Table 4

CURRENCY VALUES
(Average United States Dollar Equivalents and the EDP Weighted United States Dollar)

	Pound (United Kingdom)	Franc (France)	Mark (West Germany)	Lira (Italy)	Yen (Japan)	Dollar (Canada)	EDP Weighted Dollar* (United States)
1972	2.5008	.1983	.3136	.0017	.0033	1.0094	.902
1973	2.4510	.2254	.3776	.0017	.0037	0.9998	.822
1974	2.3403	.2081	.3872	.0015	.0034	1.0226	.853
1975	2.2216	.2335	.4073	.0015	.0034	0.9830	.833
1976	1.8048	.2094	.3874	.0012	.0034	1.0141	.904
1977	1.7509	.2039	.4317	.0011	.0036	0.9412	.905
1978	1.9337	.2230	.5050	.0012	.0049	0.8777	.817
First Quarter	1.9203	.2130	.4883	.0012	.0043	0.8942	.842
Second Quarter	1.8428	.2192	.4816	.0012	.0046	0.8886	.838
Third Quarter	1.9502	.2210	.5032	.0012	.0053	0.8757	.805
Fourth Quarter	2.0213	.2389	.5470	.0012	.0053	0.8523	.782
1979							
First Quarter	2.0280	.2334	.5369	.0012	.0049	0.8449	.799
Second Quarter	2.1011	.2295	.5311	.0012	.0045	0.8646	.803
Third Quarter	2.2353	.2380	.5554	.0012	.0045	0.8583	.785
Fourth Quarter	2.1645	.2444	.5718	.0012	.0041	0.8517	.791
1980							
January	2.2668	.2454	.5741	.0012	.0042	0.8637	.781
February	2.2625	.2401	.5624	.0012	.0040	0.8737	.792
Year-to-Year Change	11.7%	2.4%	3.9%	--	(18.4)%	4.3%	(0.8)%

*Index value weighted by six-country 1977 IBM revenues (May 1970 = 1.00).
Sources: *Federal Reserve Bulletin, Wall Street Journal,* Estimates of Dean Witter Reynolds Inc.

Table 5

OFFICE EQUIPMENT SHIPMENTS
(Dollar Amounts in Millions)

	Shipments	Percent Change	New Orders	Percent Change	Ending Backlogs	Percent Change	Ending Inventories	Percent Change	Orders/ Shipments	Backlog/ Inventories
1963	$ 3,503	8%	$ 3,869	9%	$1,492	33%	$ 833	3%	1.10	1.79
1964	3,859	10	4,395	14	2,028	86	1,036	24	1.14	1.96
1965	4,265	11	4,936	12	2,699	33	1,192	15	1.16	2.26
1966	5,963	40	6,754	37	3,490	29	1,335	12	1.13	2.61
1967	5,732	(4)	6,195	(8)	3,953	13	1,480	11	1.08	2.67
1968	6,215	8	6,608	7	4,346	10	1,661	12	1.06	2.62
1969	7,420	19	8,203	24	5,129	18	2,086	26	1.11	2.46
1970	7,693	4	7,258	(12)	4,694	(8)	2,288	10	0.94	2.05
1971	6,909	(10)	6,732	(7)	4,517	(4)	2,110	(8)	0.97	2.14
1972	8,605	25	9,161	36	5,073	12	2,084	(1)	1.06	2.43
1973	10,054	17	10,397	13	5,416	7	2,647	27	1.03	2.05
1974	12,179	21	12,979	25	6,216	15	3,325	26	1.07	1.87
1975	11,528	(5)	11,607	(11)	6,295	1	3,035	(9)	1.01	2.07
1976	13,724	19	13,191	14	5,762	(8)	3,088	2	0.96	1.87
1977	15,791	15	16,556	26	6,527	13	3,547	15	1.05	1.84
1978	18,571	18	19,382	17	7,431	14	4,449	25	1.04	1.67
1978										
First Quarter	3,733	6	3,913	5	6,707	12	3,786	19	1.05	1.77
Second Quarter	4,418	15	4,932	14	7,221	12	4,093	25	1.12	1.76
Third Quarter	4,640	19	4,791	20	7,374	13	4,254	29	1.03	1.73
Fourth Quarter	5,780	26	5,746	28	7,341	12	4,449	25	0.99	1.65
1979										
First Quarter	5,037	35	6,758	73	9,062	35	4,802	27	1.34	1.89
Second Quarter	5,249	19	5,513	12	9,148	27	5,270	29	1.05	1.74
Third Quarter	5,352	15	4,828	1	9,000	22	5,433	28	0.90	1.66
October	1,990	18	1,813	0	7,942	9	5,862	35	0.91	1.35
November	2,007	15	2,401	31	8,337	13	5,996	38	1.20	1.39
December	2,753	22	2,561	28	8,145	14	5,981	34	0.93	1.36
January (Prelim)	1,580	23	1,617	12	8,182	12	6,159	37	1.02	1.33

Source: U.S. Department of Commerce; compiled by Dean Witter Reynolds Inc. (These unpublished data are furnished by the Bureau of the Census as reported by companies in the industrial categories covered by SIC code 357, Office Computing and Accounting Machines.)

Table 6

UNITED STATES IMPORTS AND EXPORTS* OF COMPUTERS AND RELATED EQUIPMENT

	United Kingdom		France		West Germany		Italy		Japan		Total World	
	Imports	Exports	Imports	Exports	Imports	Exports	Imports	Exports	Imports	Exports	Imports	Exports
Trade Value (Millions)												
1974	$14	$281	$37	$234	$42	$260	$11	$ 82	$52	$238	$381	$2,152
1975	22	266	41	222	36	269	21	80	60	183	404	2,183
1976	20	302	60	271	51	290	35	85	149	233	630	2,512
1977	34	420	93	315	47	350	45	102	150	271	730	3,181
1978	52	567	95	395	73	515	43	134	245	332	961	4,214
1979												
First Quarter	15	184	30	121	26	153	14	37	86	100	293	1,259
Second Quarter	22	178	35	118	25	172	6	43	81	98	345	1,285
Third Quarter	24	176	32	131	22	195	7	45	68	111	296	1,367
October	14	59	7	42	5	66	3	13	22	41	110	458
November	14	72	8	49	5	70	3	20	24	34	108	504
December	11	79	6	56	8	79	2	18	17	51	113	580
Year to date	100	748	118	537	91	735	35	176	298	435	1,265	5,453
Import/Export Ratio												
1974	0.05		0.16		0.16		0.13		0.22		0.18	
1975	0.08		0.18		0.13		0.26		0.33		0.19	
1976	0.07		0.22		0.18		0.41		0.64		0.25	
1977	0.06		0.30		0.13		0.33		0.55		0.30	
1978	0.09		0.24		0.14		0.32		0.74		0.23	
1979 (Year to date)	0.13		0.22		0.12		0.20		0.69		0.23	
Imports (Percent Increase)												
1974	0		54%		40%		10%		0		11%	
1975	57%		11		(14)		91		15%		6	
1976	(9)		46		42		67		148		56	
1977	70		55		(08)		29		0.6		16	
1978	53		2		55		(4)		63		32	
1979 (Year to date)	92		24		25		(19)		22		32	
Exports (Percent Increase)												
1974		33%		21%		24%		55%		17%		28%
1975		(5)		(5)		3		(2)		(23)		1
1976		14		22		8		6		27		15
1977		39		16		21		20		16		27
1978		35		25		47		31		23		32
1979 (Year to date)		32		36		43		31		31		29

*Imports - shipments to U.S.; exports - shipments from U.S.
Sources: U.S. Department of Commerce. Compiled by Dean Witter Reynolds Inc.

Gartner Group

STRATEGIES IN
SMALL COMPUTER SYSTEMS
BIWEEKLY RESEARCH NOTES

June 9, 1982

Edited by Bill Rosser

TANDEM COMPUTERS TARGETS ANOTHER OPPORTUNITY MARKET

Having defined, developed, and dominated the "non-stop," transaction-processing market for the past seven years, Tandem is now targeting a new market which may well be as significant as its first one. This time it is corporate communications, a market which already exists (just as the transaction-processing market did), but Tandem has devised its own slant and its own particular packaging and appeal. On May 12 Tandem revealed its new plans and some initial products to support them. The fundamental concept is to use one communications network, called Transfer, to handle the entire array of information-transfer requirements. Specifically, this includes traditional data processing (DP) information, ASCII text (Telex/TWX, word processing), images (dot matrix, graphics, video) and digitized voice.

Transfer is a software implementation of a higher-level network, which takes any kind of data - electronic mail, Group 1 or Group 2 facsimile (FAX), voice messages, word processing, freeze-frame video, etc. - and combines and merges it with normal transaction processing data. This information can be sent to any other part of the network and can also be transmitted via gateways to incompatible technologies such as other mainframes, word processors, etc. Also included is the concept of the extended transaction, that is, one that can be carried out over time (such as some time before tomorrow morning or at a specific time in the future, e.g., 9 a.m. Sept. 3, 1982).

Transfer consists of two applications packages - one for electronic mail throughout the network, and another for FAX, which permits the user to input, store, and retrieve documents for Group 1 and Group 2 FAX devices. The FAX package

includes a hardware interface (based on an MC68000 chip) and also permits the transmission of digitized voice. In addition to these applications, software packages allow the user to write simple applications to merge some different device or data source with the transmission of the other data.

To add another major dimension to the capability, Tandem introduced the first integrated computer satellite communications network - Infosat - to be carried out by Tandem, American Satellite, and Vitalink. This project has created controllers, RF modems, computer/satellite communication interface modules and a new earth station design. The earth station is built around a newly developed five-watt semiconductor device which permits low-cost transmission of two parallel (but not redundant) 56-kilobaud channels (or 112 Kilobaud in all), all within the architectural framework of Tandem's non-stop design. In other words, the transmission amplifiers can be worked on while the antennas remain in operation. Also, computer service engineers can maintain the device without a special FCC license. For the customer, this team will design, install and provide the communications-network management Function for an end-to-end communications network to interconnect a number of Tandem systems. And Tandem will service it all.

Tandem hopes, in this endeavor, to gain and keep account control. While in the 1970's control may well have rested with the vendor supplying the information processing capability, Tandem's view is that account control in the 1980's will belong to the firm providing the online communications network.

A Datapro Report on

TANDEM NONSTOP SYSTEMS

Reprinted From

DATAPRO REPORTS ON MINICOMPUTERS

DATAPRO REPORTS ON MINICOMPUTERS is a management-oriented information service which covers all significant minicomputers, microcomputers, and small business computers. New and updated reports are published monthly. All opinions and evaluations contained in these reports are independent and objective, and represent the proprietary work of Datapro Research Corporation.

datapro

DATAPRO RESEARCH CORPORATION
1805 Underwood Boulevard, Delran, New Jersey 08075 (609) 764-0100

Palo Alto, CA (415) 467-6007 Phoenix, AZ (602) 263-7831 Washington, DC (301) 933-4020 Chicago, IL (312) 644-2000 Dallas, TX (214) 980-4525
A McGraw-Hill Company

Tandem NonStop Systems

MANAGEMENT SUMMARY

Tandem Computers Incorporated began operations in 1974 with the goal of developing and marketing fault-tolerant computer systems for businesses which depend heavily on the continuous availability of their computers. Such businesses are characterized by high-volume on-line transaction processing or message handling requirements. Further, Tandem sought to provide high data integrity and flexible modular expansion without reprogramming. In the ensuing years, Tandem has stuck to these goals, and in 1976 the company began marketing the 16-bit NonStop system. In April 1981, Tandem introduced their newest system, the NonStop II, a more powerful version of NonStop with 32-bit addressing.

The principles that lie behind the Tandem NonStop systems concept are not very complex. The company employs multiple processors, multiple controllers, multiple data paths between the system modules, and multiple power supplies in all system configurations. The benefits, according to Tandem, are twofold. First, the multiplexed configurations offer a high probability that at least one processor and one data path will always be operable. Secondly, a Tandem system in normal use may employ all its processing modules and data paths to some degree for running daily workloads. Thus, at any given time, a processing module may have excess capacity available to run jobs from a failed module. The capability to switch from a failed module to an operational module is accomplished automatically, through a hardware/software technique.

An added strength of this type of configuration is that defective modules can be replaced without powering down the balance of the system. Likewise, additional computing power can be introduced through the addition ➢

Tandem is the only manufacturer offering standard (rather than customized) fault-tolerant systems designed to operate continuously, despite breakdowns of single components. Tandem's products are the 16-bit NonStop system and the new 16-bit, network-oriented NonStop II system. The NonStop II is designed for on-line transaction applications requiring large numbers of interactive terminals, communications lines and large geographically distributed data bases. Aimed primarily at the end-user market, the Tandem product line consists of multiple processor modules with MOS memory, some peripherals, and two packaged systems.

MAIN MEMORY: 384K to 2 million bytes
DISK CAPACITY: Minimum—64 million bytes. Maximum—no set limit
TERMINALS: No set limit
PRINTERS: 200 cps; 600 to 1350 lpm
OTHER I/O: Magnetic tape, punched card

CHARACTERISTICS

MANUFACTURER: Tandem Computers Incorporated, 19333 Vallco Parkway, Cupertino, California 95014. Telephone (408) 725-6000.

Tandem's sole business involvement is the design, development, manufacture, marketing, and servicing of fault-tolerant, multiple-processor computer systems. Approximately 20 percent of subassembly production for printed-circuit boards and cables takes place at Tandem's manufacturing facilities, with the remaining 60 percent provided by subcontractors. The company purchases substantially all of the required ➢

The NonStop II, Tandem's new multiple processor computer system supports high-volume, on-line transaction applications with many terminals, communication lines and large local or distributed data bases. Non Stop II systems are fault-tolerant and can be modularly expanded from 2 to 16 processors (and linked in networks of up to 255 system nodes), without software changes. An Operations and Service Processor (OSP) is shown in the foreground. The OSP provides detailed system status information and diagnostic facilities to assist service personnel, and includes a modem for remote diagnosis and remote system operation.

Tandem NonStop Systems

of more processing modules, memory, or peripherals without any reprogramming or other detrimental effects on the original system.

The original NonStop system's processing modules are 16-bit units which include two microcoded pipelined processing units, one for CPU and Dynabus (interprocessor bus) control and one for I/O control; 384K bytes of 500 nanosecond semiconductor (MOS) memory with EDAC; a complete DMA-only I/O system; virtual memory control; memory mapping and protection for up to two megabytes of memory; hardware multiply/divide; power fail/auto restart; bootstrap loader; interval timer; control panel; Dynabus controller and interface; interfacing for 32 I/O controllers; a 173-member instruction set; and a power supply.

The NonStop II is Tandem's second fault-tolerant network-oriented multiple processor system, and is designed specifically for on-line transaction processing applications which require large numbers of interactive terminals, communication lines (including X.25 virtual circuits), and large geographically distributed data bases. The NonStop II system has 16-bit architecture and enables each processor to address up to one billion bytes of virtual memory.

In order to provide object code compatibility, the NonStop II processor retains the 16-bit instruction set of the original NonStop system. Programs compiled on the original NonStop system are compatible with the NonStop II, although programs that contain privileged code may require modification. The NonStop II can be used in a common data communications network without software modification.

A new processor, called the Operations and Service Processor (OSP), is a major feature of the NonStop II. The OSP communicates with the Diagnostic Data Transceiver (DDT) processor included as part of each processing unit in a NonStop II system and uses a page mode terminal to provide detailed system status information and diagnostic facilities to assist service personnel. The DDT processor monitors the status of the Dynabus interface, I/O channel processor, memory, and the Instruction Processing Unit (IPU), including the internal data paths of the IPU. The OSP allows for remote diagnosis and remote system operation through a built-in modem.

Another feature of the NonStop II is processor microcode which resides in 32K bytes of loadable control storage. The loadable control store allows for the instruction set of the NonStop II processor to be modified for new features as Tandem's software is updated.

Along with the introduction of the NonStop II, Tandem has announced an exchange program that offers existing customers the option of upgrading their on-site NonStop systems to the NonStop II. This plan provides for customers to exchange credits from returned equipment against the list price of the upgrade.

components and peripherals from other manufacturers. Established in November 1974, Tandem employs 1500 persons, including 65 in engineering and software development, 109 in manufacturing, 244 in marketing and field service, and 28 in general management. Tandem systems are marketed through 36 U.S. and 16 international sales offices.

MODELS: T16/1412-1 and T16/1420 processors and T16/244-3, and T16/250 packaged systems.

DATE ANNOUNCED: NonStop—November 1974, NonStop II—April 1981.

DATE OF FIRST DELIVERY: May 1976 for NonStop and April 1981 for NonStop II.

NUMBER INSTALLED TO DATE: Over 300 customers and over 1879 processors.

DATA FORMATS

BASIC UNIT: 16-bit words; 8-bit bytes, 32-bit double words, and 64-bit quadruple words can also be accessed and manipulated.

FIXED-POINT OPERANDS: 8, 16, 32, or 64 bits in memory or register stack as designated by the instruction type. Bytes can be loaded into and stored from the register stack; moved from one memory location to another; compared to the contents of another memory location; or scanned in a block of one, two, or four words. Bytes represent unsigned values in the range of 0 to 255. Words, double words, and quadruple words can be loaded and stored from the register stack. Both logical (unsigned) and interger (signed) arithmetic can be performed on word operands. Double words and quadruple words may only have integer arithmetic performed on them. Word operands can represent signed numbers in the range of -32,768 to +32,767 and unsigned numbers in the range of 0 to 65,535. Double-word operands can represent signed numbers in the range of -2,147,583,648 to +2,147,583,647. Quadruple-word operands can represent 19-digit numbers in the range of -9,223,372,036,854,775,808 to +9,223,372,036,854,775,807. For words, double words, or quadruple words, positive values are represented in true binary notation. Negative values are represented in two's-complement notation with the sign bit of the most significant word set to one.

FLOATING-POINT OPERANDS: 23-bit single-precision and 55-bit double-precision operands, in normalized form; all 9-bit exponents are expressed in excess-256 notation.

INSTRUCTIONS: All 43 floating-point instructions are one word long and can be divided into 17 categories, each with a unique format which varies only for immediate-operand instructions. The instruction word is generally divided into five 3-bit fields and one 1-bit field. Fields are sometimes combined, for example, to represent mode and displacement. Likewise, fields may be subdivided to represent a one-bit indicator (such as indexing or no indexing) and a two-bit register identifier. Immediate-operand instructions have bits 7-15 reserved for the immediate operand.

INTERNAL CODE: Binary.

MAIN STORAGE

TYPE: N-channel dynamic MOS semiconductor.

CYCLE TIME: MOS—500 nanoseconds refresh time with an access time of 500 nanoseconds on the NonStop, 400 nanoseconds refresh time with an access time of 400 nanoseconds on the NonStop II.

Tandem NonStop Systems

PERIPHERALS/TERMINALS

DEVICE	DESCRIPTION & SPEED	MANUFACTURER
MAGNETIC TAPE EQUIPMENT		
T16/5104	Transport; 9-track, 800 bpi (NRZI)/1600 bpi (PE), 125 ips, 10.5-inch reels, tension arms, read-after-write and power-fail/auto-restart electronics	Kennedy
T16/5103	Transport; 9-track, 800 bpi (NRZI)/1600 bpi (PE), 45 ips, 10.5-inch reels, tension arms, read-after-write and power-fail/auto-restart electronics	Kennedy
PRINTERS		
T16/5508	Serial matrix printer; 132 positions, 96 ASCII character set, 10 characters per inch, 15-inch paper (width), 6 lines per inch, electronic VFU; 200 cps	Diablo Series 2300
T16/5503	Drum, impact printer; 132 positions; 64 ASCII character set (96 optional), 10 characters per inch, 4 to 16.75-inch paper (width), 6 or 8 lines per inch, 12-channel VFU optional; 600 lpm (436 lpm with 96-character set)	Dataproducts 2260
T16/5504	Drum, impact printer; 132 positions; 64 ASCII character set (96 optional), 10-characters per inch, 4 to 16.75-inch paper (width), 6 or 8 lines per inch, 12-channel VFU optional; 900 lpm (436 lpm with 96-character set)	Dataproducts 2290
T16/5510	Line printer, 132 columns, 1350 LPM band printer, 64 character ASCII set, 12 channel VFU, powered paper stacker, 1350 lpm	Documation
PUNCHED CARD EQUIPMENT		
T16/5301	Reader; 80-column, 1000-card input hopper and output stacker; 600 cpm (for NonStop systems only)	Documation M6002
TERMINALS		
T16/6603	Hard-Copy Terminal; 132 positions, 63 or 95-character set, 10 characters per inch, 3 to 14.9-inch paper (width), 6 lines per inch, standard typewriter keyboard with 10-key numeric pad, 30-cps printing, half duplex transmission, RS-232 interface; 110-1200 bps	DEC DECwriter II
T16/6604	Same as T16/6603 but 20-mA current loop interface; 110-1200 bps	DEC DECwriter II
T16/6520	Multi-page display terminal. In block mode the T16/6520 can be transmitting from or receiving into any one of three 1920 character memory pages. In conversational mode 144 lines x 80 characters. One parity bit per byte. Operational modes include certain combinations of asynchronous, synchronous, character or block RS-232 or 20 mA current loop; point to point or multipoint. Full complement of video and data attributes, local editing, program function keys and 25th status display line. 2000 characters; 25 lines by 80 characters, 7 x 9 or 14 x 9 dot matrix, reverse video, blinking, underscore, non-display, half and full duplex transmission, RS-232 interface, 110 to 19.1 bps.	
T16/6524	Same as T16/6520 except an additional RS-232 communication port is provided for output to serial printer	

Tandem continues to manufacture and support the original NonStop system, and offers a future upgrade plan allowing customers to purchase the current NonStop system for applications development. As the applications come on line, customers can then move up to the NonStop II and take advantage of exchange credits against the list price. A basic NonStop II configuration with dual processors, 512K-bytes of memory per processor, a magnetic tape drive and controller, and an Operations and Service Processor costs $144,475.

The processing modules for the NonStop II have 16-bit architecture including three microcoded processing units, one for CPU and Dynabus (interprocessor bus) control, one for I/O control and a third for Diagnostic Data Transceiver (DDT) control; 512K bytes of 400 nanosecond semiconductor (MOS) memory with EDAC; maximum virtual memory addressing of 1-billion bytes, 8K of 32-bit words of loadable control store; 1K of 32-bit words of read only storage; string manipulation and double word arithmetic; complete DMA-only I/O sys-

CAPACITY: 393,216 to 2,097,152 bytes per processor in increments of 393,216 on the NonStop, 512K-bytes to 2M-bytes on the NonStop II in 512K increments.

CHECKING: MOS memory employs 6 check bits per word for error checking and correcting (ECC).

ECC detects and corrects all single-bit main memory errors, and detects all double-bit errors and most multiple-bit errors. ECC generates a 6-bit check field for each 16-bit data word as it is written, and recomputes the field when the word is read. If the check bits do not match, the erroneous bit is corrected before data is transmitted to the processor. Correction and virtual memory address translation is included in the stated access times of 500 nanoseconds for the NonStop system, and 400 nanoseconds for the NonStop II. ECC helps to ensure uninterrupted operation and is transparent to the user. A Hamming code is used by the hardware encoder on each memory board to construct the check field.

STORAGE PROTECTION: Via memory mapping provisions.

RESERVED STORAGE: Up to 1280 bytes of memory are reserved in the system data area for three specific tables. The

Tandem NonStop Systems

tem; power fail/auto restart; bootstrap loader; interval timer; control panel; Dynabus controller and interface; interfacing for 32 I/O controllers; a power supply and one battery pack.

Both the NonStop and NonStop II systems support Tandem's Guardian operating system. The standard Guardian operating system includes the Enscribe data base access method. Another version of the Guardian operating system includes SPOOLER, a means of storing application output for later retrieval, and XRAY, a systems performance monitor.

Packaged systems include twin processing modules with MOS memories, system cabinets, and a magnetic tape module. Systems may be configured to include up to 16 processing modules.

Tandem's Dynabus hardware provides two autonomous data paths for high-speed transmission of data among processors without tying up normal communications between the processors and peripheral devices. The Dynabus is also used for periodic interprocessor checks, which permit isolation of any processor failure and also indicate at which point in the sequence of program steps the remaining operational system should begin in order to complete an interrupted assignment. Similarly, Tandem's peripheral device controllers are dual-ported. This construction provides communications paths from each controller to two different modules in the system.

Tandem offers a variety of peripherals, including 64- or 240-megabyte pack removable disc drives and 64-megabyte non-removable disc drives; magnetic tape drives; a 600-cpm card reader; 600-, 900-, or 1350-lpm line printers; character or page-mode CRT terminals; and synchronous or asynchronous controllers.

Tandem software includes the Guardian operating system, ENCOMPASS distributed data base management system, Expand data communications networking system, Exchange remote batch workstation, AXCESS (AM3270 and TR3271 Access Methods), Envoy data communications manager, TAL (Tandem's system level language), COBOL, FORTRAN, TIL (Tandem to IBM Link), and THL (Tandem Hyper Link). All of these software products are unbundled.

Tandem systems are marketed for both OEM and end-user sales, although primary emphasis is on the end user. Arrangements for third-party leasing are primarily the responsibility of the user, although Tandem will provide assistance in making such arrangements. Employing its own sales organization, Tandem markets its products through sales offices in Atlanta; Boston; Chicago; Cincinnati; Cleveland; Columbus; Dallas; Denver; Detroit; Fort Worth; Greensboro; Houston; Indianapolis; Kansas City; Los Angeles; Memphis; Milwaukee; Minneapolis; Hasbrouck Heights, N.J.; New York; Omaha; Philadelphia; Phoenix; Pittsburgh; Richmond; Rochester; San Diego; San Francisco; San Mateo; Santa Clara; Seattle; St. Louis; Tampa; Tulsa; and Washington, D.C.

System Interrupt Vector table consists of 128 bytes divided into 16 four-word entries. Each of these entries defines the executing environment for one of the 16 operating system interrupt handlers. The Bus Receive Table consists of 128 bytes divided into 16 four-word entries. Each of the 16 entries corresponds to one of the 16 possible processors. The I/O Control Table consists of up to 1024 bytes divided into up to 256 entries corresponding to the 32 possible controllers. Each controller can handle up to eight units connected to an I/O channel. An entry describes the number of bytes to be transferred and the system data location where the data transfer takes place.

CENTRAL PROCESSOR

Each Tandem processor (T16/1412-1 or T16/1420) is built around a pipelined, microprogrammed central processor with a cycle time of 100 nanoseconds. The processor is operated by 32-bit micro-instructions, each of which consists of combinations of up to 7 of the 168 micro-operations available.

The Tandem processors feature 16-bit data paths and memory addressing; memory mapping; an instruction set that includes string manipulation and double-word arithmetic; a floating-point option; a decimal arithmetic option; stack architecture; non-modifiable code areas of memory; hardware power fail/auto restart; hardware multiply/divide; a block-multiplexed I/O channel; and dual high-speed interprocessor buses.

Tandem employs a true virtual memory system in the sense that besides the software designed for memory management, a number of features are incorporated into the processor module's hardware to aid in reducing the amount of swapping that occurs:

- A "dirty" bit is associated with each data map entry; only if a data page has been modified is it ever swapped out.
- A reference bit is associated with each map entry to record access to a particular page. The memory manager also maintains a list of maps (each process has a separate map) active in a processor module. When memory space is needed for an overlay, the memory manager selects the physical page that has been the least recently used for overlaying.
- Paging hardware is provided in the form of memory maps.

Each Tandem processor module is integrated into a Tandem NonStop system, which consists of from 2 to 16 processor modules. Processor modules are interfaced to one another by means of two interprocessor buses, which are a standard feature of each processor module. A processor module may interface to I/O devices by means of its input/output channel.

Each processor module is capable of operating independently of, and simultaneously with, all other processor modules in the system. The Guardian-controlled interprocessor buses (Dynabuses) are used to transfer data between the memories of the processor modules. Each bus offers a transfer rate of 13 megabytes per second, block transfers of 1 to 32,767 bytes, and packet-multiplexed transfers between any number of processor modules. The dual interprocessor buses contained in each processor module operate simultaneously. Data is actually sent across a bus in 16-word groups (including one check-sum word).

Each NonStop II processing module includes an instruction processing unit (IPU), main memory, a Dynabus interface unit, an I/O processor, and a Diagnostic Data Transceiver (DDT) processor.

Tandem NonStop Systems

Tandem markets internationally in Canada, Hong Kong, Mexico, England, France, Germany, Italy, Japan, the Netherlands, Sweden and Switzerland through subsidiaries; and in Finland, Greece, Australia, the Philippines, Korea, Venezuela, and Taiwan through distributors.

Typically, the company ships its systems to customers within 90 days after receipt of orders, and systems are sold with a 90-day warranty for end-users and 120 days for OEM's. Delivery for the NonStop II is 120 days after receipt of order.

Field engineering is provided on an 8 a.m. to 5 p.m. prime-shift basis. Other hours of service may be arranged by contract with Tandem.

Training is offered at Tandem's corporate headquarters and regional offices in COBOL, FORTRAN, TAL, ENCOMPASS, MUMPS, and hardware. Software and hardware courses are priced at $500 per week.

Tandem considers its prime competitors to be the computer vendors that have the capability to offer dual computer systems on a customized basis. These include Burroughs, Data General, Digital Equipment, Hewlett-Packard, Honeywell, and IBM.

USER REACTION

Ten Tandem users responded to Datapro's 1981 survey of computer users, representing a total of 19 installed systems. Eight users had purchased their systems and two were leasing from third party leasing companies. All systems were installed for an average of 16 months.

One-third of these users said they were using their systems to run accounting/billing applications. Other principal applications mentioned were: banking-check processing, loans, savings; education scheduling/administration, engineering/scientific, health care/medical, order processing/inventory control, and payroll/personnel. Five users were employing in-house personnel to write all application programs, three were using a combination of in-house personnel programming, "ready-made" programs from Tandem, contract programming and proprietary software; one used a combination of contract programming and proprietary software; and one used proprietary software only. The principal programming language being used was COBOL (six users); FORTRAN and TAL each had one user, while other languages had two users.

Memory capacity of the installed systems ranged from 256K bytes to over 8192K bytes, with an average capacity of about 512K bytes to less than 1024K bytes. Disk storage capacity ranged from 256K bytes to over 1200 megabytes, and the number of workstations attached ranged from one to five, to over sixty.

The basic set of 242 machine instructions provides 16-, 32-, and 64-bit integer arithmetic, and byte-oriented functions such as scanning and comparing strings. The NonStop II basic instruction set also supports 32-bit extended addressing and contains the complete instruction set from the original NonStop system. This allows the NonStop II to run programs compiled on the NonStop system, although programs that contain privileged code may require modification. Both the Tandem NonStop and NonStop II use the Guardian operating system, allowing both systems to communicate via the Guardian/Expand networks without software modifications. Files can be exchanged between the systems by using the network or the Backup and Restore utilities.

The NonStop II optional floating-point instruction set provides 43 instructions for high-speed scientific calculations. The instruction sets are implemented in microcode located in a high-speed control store. The control store has 8K 32-bit words of loadable storage, and 1K words of read-only storage. The loadable portion of the control store is initialized when the operating system is loaded. Before loading the control store, the system performs a set of diagnostic routines to be sure the processor is operating correctly. The processor's internal data paths and registers are parity checked.

The Diagnostic Data Transceiver (DDT) is a separate processor included as part of each NonStop II processor module. The DDT provides two distinct functions. First, the DDT allows communications between a processor module and the Operations and Service Processor (OSP), which supports both operational and maintenance functions such as running diagnostics. Secondly, the DDT monitors the status of the instruction processing unit, Dynabus interface, memory, and the I/O processor, and reports any errors to the OSP. The OSP uses the operator's terminal for system diagnostics and maintenance, and includes a built-in modem for remote use. Initial connection to the modem in the OSP requires on-site intervention to protect the user's system against unauthorized access.

CONTROL STORAGE: Up to 3,072 32-bit pre-programmed words of PROM microprogram storage are available to implement the decimal and floating-point instruction options. The basic instruction set utilizes 1024 words. Approximately half of these are used to implement privileged operating system functions. Cycle time is 100 nanoseconds.

The Non Stop II system uses 8K of 32-bit words of RAM, loaded from disk, and 1K of 32-bit words of PROM.

ADDRESSING: Programs in the execution environment are broken into a code area and a data area. The code area consists of instruction codes and program constants. The code area may not be written into and therefore cannot be modified. Code areas are re-entrant and sharable by numerous programs. Each code area can consist of up to 131,072 bytes; the NonStop II provides an additional 131,072 bytes for shared libraries.

The data area consists of up to 131,072 bytes logically separated into global (addressable by any instruction in the program), local, and top-of-stack areas. On NonStop II, users may allocate up to 128 million bytes of extended data segments.

MEMORY MAPPING: Memory maps translate the 16-bit logical address used in a program to the appropriate number of bits needed to address physical memory: 20 bits for the NonStop and 23 bits for the NonStop II. The physical memory addressing limit is 2 million bytes per processor for the NonStop and 16 million bytes per processor for the NonStop II. Memory is arranged in 2048-byte pages. The

Tandem NonStop Systems

provides 64MB of low cost mass storage. The second volume provides a read/write head for every data track. By contrast, the mas storage volume has two read/write heads for each surface of the disc. If the desired data is in a track other than the two currently located under the heads, the mass storage volume must perform a seek to reposition the heads.

Providing a read/write head for each data track eliminates seek time. This gives the fast access volume the potential to perform a read or write operation for every revolution of the disc. The fast access volume of the T16/4109 improves system throughput by providing faster access to critical files that can limit system performance.

Performance specifications include a transfer rate of 1.2MB per second. Average latency is 8.3 milliseconds; maximum is 17.3 milliseconds. Recording technique is modified frequency modulation; the density is 6417 bits per inch (inner track). Full seek time is 55 milliseconds maximum; average is 30 milliseconds; single track is 7 milliseconds. Start/stop time is 30 seconds maximum.

INPUT/OUTPUT UNITS

See Peripherals/Terminals table.

COMMUNICATIONS CONTROL

T16/3401 UNIVERSAL INTERFACE (UI): This interface provides the ability to interface custom equipment to a NonStop System. The UI is capable of connecting two devices that have 8- or 16-line parallel data interfaces to the Tandem system. The interface provides a device data path that is buffered (16 words deep), bi-directional, and capable of operating in half-duplex mode at a sustained data tranfer rate of up to 4 megabytes per second, depending on the channel configuration. It interfaces to one device over positive or ground true TTL lines for distances up to 25 feet and to the second device over differential lines for longer distances up to 500 feet. The data path between either or both of the two devices and the UI can be either one byte (8 bits) or one word (16 bits) wide. Configuration of the UI is accomplished by software and by jumpers.

Odd parity is generated and checked for each data word that is transferred between the channel and the UI. The parity that exists between the UI and each device is defined by configuration jumpers in the connector hood. The jumpers select odd, even, or no parity.

T16/6303/6304 ASYNCHRONOUS CONTROLLER: Handles from 2 to 32 asynchronous communications lines per controller with line speeds ranging from 50 to 19.2K bits per second. The controller is sufficiently fast to support 32 lines all running at 19.2K bps. The T16/6303 provides for the first two lines, while the T16/6304 Asynchronous Extension Board provides for an additional 15 lines. Up to two T16/6304 boards are allowed per controller. The T16/6303 provides DMA access to main memory on all I/O transfers and modem support for Bell-type 103 and 202 modems (including reverse channel). Each line is individually programmable with respect to: bps rate; character size; parity generation and checking; enable/disable checking for signal characters; half-duplex modem turn-around character(s); read completion on ETX (end of text) character; default transfer mode; conversational mode line termination character, backspace type, carriage return/line feed delay, forms control delay, and automatic line feed on input; page mode termination character; and pseudo-polling trigger character.

For point-to-point applications, the standard Guardian I/O Subsystem will support any RS-233 or current loop terminal by merely configuring the line in SYSGEN. Standard I/O calls (READ, WRITE, WRITEREAD) provide access to the terminal. For multipoint applications, the Envoy data communication manager provides the ability to interface polling terminals such as the T16/6552 CRT Polling Terminal.

T16/6202 BYTE SYNCHRONOUS COMMUNICATIONS CONTROLLER: This controller, utilizing microprogrammed technology, supports up to four lines. Each line can be configured dynamically for translate enable, transparent text capability, full or half duplex operation, polling address, and selection address. Speeds up to 56K bps per line are supported by the controller.

The T16/6202 also features automatic generation and detection of block check characters with support for VRC, LRC, and CRC; 16 modes of operation; automatic code translation to ASCII and EBCDIC; DMA access to main memory; autoinsertion of DLE and SYN characters; and support for Bell-type 201, 203, 208, and 209 modems.

The polling of multipoint stations is, for the most part, handled by the controller. The Envoy data communication manager formats a polling list (on behalf of the application process) for the controller to use to poll the multipoint tributary stations. Envoy then commands the controller to begin polling. CPU processing is interrupted only when a polled station responds.

The controller has the capability to recognize if a line is being polled or selected. For each line, the controller stores the first byte of the station's polling address and the first byte of the station's selection address. Only when the line is polled or selected and the corresponding poll or select byte matches is CPU processing interrupted.

T16/6203 BIT SYNCHRONOUS COMMUNICATIONS CONTROLLER: This controller is dual-channel connected and may be powered from either connected power supply. The T16/6203 controls from one to four bit synchronous communication lines, either point-to-point or multipoint. Each of the four lines may run at baud rates to 56K bits per second. Variable length and offset translation of EBCDIC to ASCII codes, automatic polling of multipoint lines, automatic control character generation and recognition, and automatic zero bit insertion and deletion are supported by the T16/6203.

TIL (TANDEM TO IBM LINK): This device is a hardware and software communications subsystem that provides a high-speed data path from at least one, and possibly a network, of Tandem systems to an IBM compatible system running an OS or OS/VS operating system. Data is transferred across the link by two cooperating application programs, one running in a NonStop system, and the other running in an IBM 370 system.

Physically, the computer systems are linked via an IBM 3803 model 1 or model 2 magnetic tape drive controller. Therefore, the TIL device can be used with any IBM 370-compatible system equipped with a 3803 controller running under IBM's OS/VS or OS operating system. To the IBM 370-compatible system, the TIL device appears to be an IBM 3420 Model 7 tape drive.

The TIL device transfers data at tape speeds between an IBM and a Tandem system, which provides a burst mode transfer rate of approximately 300 kilobytes per second. Data block sizes can range from 32 to 4096 bytes. Because the TIL device utilizes as much existing IBM hardware as possible, these data transfer rates are achieved at a reasonable cost and with a minimum of system overhead.

THL (TANDEM HYPER LINK): THL allows a NonStop system to be connected to the Network Systems Corporation's (NSC) HYPERchannel network, a high speed (up tc 50 megabits per second) local network to which many

Tandem NonStop Systems

different vendor systems may be attached. THL features Tandem designed hardware and software for connection to the HYPERchannel network. The Tandem system is connected to NSC equipment with a dual-ported controller. Users access THL through Tandem's Guardian file system. Transfer of data is accomplished at 300K bytes per second maximum burst rate over the Tandem I/O channel. By utilizing the Guardian no-wait facility, full-duplex THL interface can be achieved. Support is provided by NSC for connection to a HYPERchannel network to systems from: CDC, CRAY, UNIVAC, IBM, and DEC.

T16/6810 PERIPHERAL LINE ADAPTER (PLA): The T16/6810 Peripheral Line Adapter provides a link between a NonStop system and up to eight strings of Tandem T16/6520 or T16/6524 terminals and/or two serial printers. A maximum of 64 devices may be connected to a single line.

Using a single PLA instead of several multipoint modems reduces line costs, simplifies system configuration, and facilitates peripheral additions/changes. Unlike RS-232 daisy chains that have a 50-foot (15 meter) limitation between terminals, the PLA does not impose severe distance restrictions; and a disabled terminal on one port does not hinder or affect terminal/printer operations on the other ports.

Features of the T16/6180 PLA include sixteen data transmission speeds from 50 bps to 19,200 bps; eight expansion ports for attaching Tandem T16/6520 or T16/6524 terminals or up to eight terminal strings; two printer ports for hard copy output; operation up to 1500 feet (454 meters) or farther. Expansion is easy by cascading T16/6810s to increase the number of ports available or increase the distance to remote terminal clusters. Modem-elimination capability reduces user costs. Parity checking preserves data integrity, while internal self test ensures reliable operation.

SOFTWARE

OPERATING SYSTEM: The *Guardian operating system* is the only operating system currently offered by Tandem.

The Guardian operating system provides the multiprocessing (parallel processing in separate processor modules), multiprogramming (interleaved processing in one processor module), and NonStop capabilities of the NonStop and NonStop II systems. The NonStop character of the Tandem systems is preserved by keeping master copies of Guardian, configured for the specific application, in a system area on mirrored (duplicate) disc volumes. Each mirrored volume is meant to be mounted on a separate disc device and contains Guardian, virtual swapping memory space, system and application programs, and application data. This mirroring is transparent to the user. Other transparent-to-user functions include allocation of buffer space and control blocks, fault handling, the preparation of a program for execution in virtual memory when a request is made to run a program, the capability for processes to communicate with each other regardless of the processor modules in which they are executing, the scheduling of processor module time among multiple executing programs according to their application-assigned priorities, and provision of the virtual memory function by automatically bringing absent memory pages in from disc when needed.

Guardian isolates the application from physical memory constraints by providing a virtual memory management system. Paging hardware on the NonStop system is provided in the form of four memory maps (user code, user data, system code, and system data); paging hardware on the NonStop II is provided in the form of 16 memory maps. All code is both sharable by multiple programs and non-modifiable, two features which reduce overlay and swapping overhead. Code pages are never swapped out to disc. In addition, hardware is used to record the frequency of access to all memory pages and modification of data pages, thus providing a low-overhead method of determining the correct page to be replaced.

In a typical system, master copies of the Guardian operating system, configured for the specific application, are kept in a system area on mirrored disc volumes. Critical and frequently used parts of Guardian are resident in each processor module's memory. Thus, the system's capabilities are maintained even if a processor module, I/O channel, or disc drive fails. Non-critical or less frequently used parts of Guardian are virtual and are brought into a processor module's memory from disc only when needed.

Likewise, a mirror principle permits data to be written simultaneously onto two disc devices. If a failuare occurs in one of the drives, data is read from the other device. Upon restoration of the failed disc drive, the system automatically updates the stored device while continuing to process transactions. The mirroring is transparent to both the application programmer and the computer operator.

To prevent data degradation resulting from a processor failure, an active processor signals an alternate processor (which contains a backup program) at the start of a file update. If the update is not completed correctly, the operating system alerts the other processor to perform the update.

Concurrently with application program execution, Guardian continually checks the integrity of the system. This is accomplised by having each processor module, at a predefined interval, transmit an "I'm alive" message to Guardian in every other processor module. (This interval is typically one second). If Guardian in one processor module finds that an "I'm alive" message has not been received from another processor module, it first verifies that it can transmit a message to its own processor module; if it can, it assumes that the non-transmitting processor module is inoperative; if it can't, it takes action to ensure that its own module does not impair the operation of other processor modules. In either case, Guardian then informs the system processes and interested application processes of the failure.

Guardian provides multiple system services that can be requested by programs or that affect application program design. These include the program control functions of run, suspend, and stop; I/O file management; system messages; utility procedures; a checkpointing facility; traps and trap handling; a command interpreter program for communicating run-time information to an application program; an interactive debug facility; and security for controlling access to data files.

TYPICAL STORAGE REQUIREMENTS OF TANDEM SOFTWARE

	Main Memory (bytes)	Disk Storage (bytes)
Guardian	94K	488K
TAL	60K	208K
COBOL	90K	64K
FORTRAN	56K	80K
Sort	48K	50K
Envoy	23K	—
Expand	25K	—

The executing environment of a given program is a single processor module (specified at run time). A program's environment consists of a code area, containing instruction codes and program constants, and a separate data area, containing variables and hardware environment information.

Tandem NonStop Systems

A given code area is shared by all processes that are executing the same program file. This is permissible because information within the code area cannot be modified. Each program, however, has its own separate, private data area.

It is possible for a properly coded application process to recover from any type of hardware failure except one—a failure of the processor module in which it is executing. Because of this, NonStop applications consist of one or more primary/backup "process-pairs." A process-pair consists of two copies of the same program and data. Logic in the program indicates whether the process is executing in the primary mode, to perform the designated work, or in the backup mode, to monitor the operability of the primary.

Procedures provided by Guardian to perform utility operations include a debug facility, provision of the global address of the last word in the application's data area, conversion of the ASCII representation of a number into its binary equivalent, conversion of the internal machine representation of a number to its ASCII equivalent, and provision of the current date and time.

The checkpointing facility provides the capability for writing application programs that can recover from a processor module failure. To use the checkpointing facility, an application program must be executing as a process-pair. The checkpointing facility is used by the primary process of a process-pair to "checkpoint" pertinent data to its backup process. It is used by the backup process to receive the checkpoint data, to monitor the primary, and to begin executing at the point indicated by the latest checkpoint message.

Certain critical error conditions occurring during process execution prevent the normal execution of a process. The errors, which are for the most part unrecoverable, cause traps to operating system trap handlers. The conditions include illegal address reference, instruction failure, arithmetic overflow, stack overflow, process loop timer time-out, memory manager disc read error, no memory available, uncorrectable memory error, and map parity error.

The Guardian operating system's security capability is designed to prevent inadvertent destruction of files through purging or overwriting and to prevent unauthorized access to sensitive data files by programmers or operations personnel.

Security is enforced by assigning a group name, user name, and (optionally) a password to individuals that are to access the system. For each file, file access at each level may be restricted to reading, writing, executing, and/or purging.

Using the debug facility, a program location can be designated a "breakpoint." The subsequent execution of the breakpoint location during otherwise-normal processing causes the process to enter the debug state. While the process is in the debug state, the programmer can interactively display and modify the contents of the process's variables, display and modify the contents of the process's registers, and designate other breakpoints.

The file management system provides access to, and operations on, "unstructured" disc files. An unstructured disc file contains information placed in it by an application program. There is no logical record size imposed by the system, and there is no implicit difference between sequential and random access to an unstructured disc file. Data base access is an extension of Guardian that provides access to and operations on three additional file structures; relative files, in which records are stored relative to the beginning of the file; entry-sequenced files, in which records are appended to a file in the order presented to the system; and key-sequenced files, in which records are placed in a file in ascending sequence according to the value of a key field in the record.

Guardian program development tools include the command interpreter (Comint), the text editor (Edit), the TAL compiler, the cross-reference program (Xref), the object file editor (Update), and the interactive debugging facility (Debug).

Comint is a program that is used interactively to run programs, check system status, create and delete disc files, and alter system hardware states. An important feature of Comint is its ability to pass user-specified parameter information to a program at run time.

Edit is used to prepare source programs written in TAL. It is an interactive program that allows the programmer to enter and make changes to the source program through an interactive terminal. Text entered through the text editor is stored in a file on disc under a name given by the programmer. This name also specifies the source program to the TAL compiler and is used later if the programmer wishes to "edit" the file.

LANGUAGES: Tandem currently offers four high-level languages: TAL, a language developed exclusively for use on Tandem systems; COBOL; FORTRAN; and MUMPS, a high-level interpretive language.

TAL is a high-level, block-structured, procedure-oriented language designed, according to Tandem, for ease of programming and efficient use of the architectural features of the Tandem systems. Similarities exist between TAL and ALGOL, PL/1, and COBOL in that they all offer machine-independent procedure blocks and high-level constructs. Other characteristics of TAL include free-form structure; machine-dependent statements; programmer-assigned identifiers; single- and double-word integer, byte string, and 18-digit fixed-point data types; multiple-element block operation such as move, compare, or scan block; bit operations such as bit deposit, bit extraction, and bit shift; and procedures, recursive procedures, and subprocedures.

TAL's machine dependence makes it possible for the compiler to generate optimized code that takes advantage of the Tandem systems' hardware characteristics. It also enables the programmer to code at an assembly-language level through the use of the COBOL statement and to operate on hardware registers through the use of STACK and STORE statements.

Program elements such as constants, variables, labels, and procedures are identified throughout a source program by the use of symbolic, programmer-assigned identifiers. This eliminates the need for a programmer to keep track of specific memory addresses. An identifier can contain up to 31 alphanumeric characters.

A procedure is a block of machine instructions that exists only once in a program but can be called into execution from any point in the program. Procedures, as implemented in the TAL language, have special properties. A program has a global data area that is accessible only by statements within that program; a procedure has its own (local) data area that is accessible only by statements within that procedure. Unlike the program's global data area, however, a procedure's local data area is allocated and initialized only when the procedure is executed. Because a procedure has its own local data area which is initialized each time the procedure is entered, a procedure can call itself (recursive procedure). Subprocedures are similar to procedures in that they can have their own variables and can be called recursively. However, a subprocedure is a part of a procedure and therefore can be called only from the procedure in which it resides.

The TAL compiler program reads source statements from one or more files and compiles the statements into a ready-to-run object program. Like the source program file, the object program file is given a name. The object program name is used to run or modify the program. As a by-product of the

Tandem NonStop Systems

compilation, a completely annotated listing of the source program is provided. Certain listing options provide the machine instruction code generated and a map of all the identifiers used in the program.

Tandem COBOL conforms to the ANSI COBOL-1974 language specifications and runs in a mixed-language environment with TAL. Features of Guardian are supported through extensions to the language, including NonStop operations, shared and re-entrant code, virtual memory, geographic independence of I/O devices, and checkpoint/check-monitor facilities.

STARTBACKUP and CHECKPOINT are the verbs that make a program nonstop. STARTBACKUP is normally called once at the beginning of the program to set the NonStop mode. Thereafter, the CHECKPOINT verb is used to pass information to the backup process at critical points in the processing. In a nonstop program, checkpoints will also occur automatically upon any OPEN or CLOSE executed after the backup has been established. Both of these verbs will set the special register, PROGRAM-STATUS, to indicate the outcome of the checkpointing operation.

Geographic independence of I/O devices is provided by several Tandem extensions to COBOL. The verbs LOCKFILE, UNLOCKFILE, and UNLOCKRECORD allow the use of the corresponding system file and record locking routines. This addition allows separate processors to share a common data base.

The READ and REWRITE verbs are extended to allow the specification of a LOCK or UNLOCK operation. The OPEN syntax is extended to specify the file access (EXCLUSIVE, SHARED, or PROTECTED), and to permit the SYNCDEPTH for files opened in the OUTPUT, I-O, or EXTEND mode.

Guardian data base access features included in Tandem COBOL include key-sequenced, entry-sequenced, and relative file structures; logical file sizes of up to four billion bytes; one primary and up to 31 alternate keys; and optional mirror data base recording.

The level of support provided by Tandem COBOL for the various ANSI modules includes: Nucleus, level 2; Table Handling, level 1; Sequential I/O, level 2; Relative I/O, level 2; Indexed I/O, level 2; Sort-Merge, level 2; Library, level 1; Debug, level 1; and Interprogram Communication, level 1. Report-Writer, Communication, and Segmentation are not supported. Segmentation, however, is supported automatically via Guardian.

Tandem FORTRAN conforms to the full language specifications of ANSI FORTRAN-1977. All features of Guardian are supported through extensions to the language, including NonStop operation, shared and re-entrant code, virtual memory, checkpoint/checkmonitor facilities, data base facilities, and interprocess communications.

STARTBACKUP and CHECKPOINT functions allow a FORTRAN program to utilize the NonStop capabilities of Guardian. STARTBACKUP is called once at the beginning of a program to establish the NonStop mode. Thereafter, CHECKPOINT is used to pass critical information to the backup process. Checkpoints will occur automatically upon any OPEN or CLOSE after the backup has been created.

Extensions to the FORTRAN READ and WRITE statements permit the full use of the data base facilities. Thus, it is possible with FORTRAN statements to access key-sequenced, relative, and entry-sequenced files by primary keys or up to 255 alternate keys. Provision has been made to allow exact, approximate, or generic positioning into a file structure using FORTRAN. Concurrent record access is supported with LOCK mechanisms at either the record or file level.

FORTRAN processes can communicate with one another or with processes written in other languages through the standard FORTRAN READ and WRITE statements. Communication with other processes is implemented using the interprocess communication facilities of the Guardian. In addition, FORTRAN is enhanced so that the Tandem Data Definition Language can also be used to transcribe a schema into FORTRAN record structures.

ENCOMPASS DISTRIBUTED DATA BASE MANAGEMENT SYSTEM: This package is designed to simplify the complexities of on-line transaction processing. It was designed from scratch to provide a high performance relational data base and a very high level of data integrity for a single system or a network of systems. It fulfills the needs of applications to process large numbers of transactions on-line without fail.

A transaction processing software package supplies all the procedures, programs, and application structures necessary to allow users to write single-threaded application program modules. User written application modules are written without concern for terminal characteristics and communications protocols. ENCOMPASS transaction processing software includes a terminal control process, a COBOL-like screen language, an application monitor, and an interactive screen definition facility.

All terminal oriented functions are isolated into terminal control processes (TCPs). Each TCP interacts with one or more terminals. The TCP performs four major application functions: terminal interface, field validation, data mapping, and transaction control. The handling of each terminal is defined by a high-level language known as Screen COBOL. The Screen COBOL compiler produces an intermediate code file, which is executed by the TCP. The screen builder facility enables the user to design screen formats directly on the terminal screen. The screen builder program then generates the appropriate Screen COBOL source statements that describe the screen format. The new screen description may then be added to an existing Screen COBOL program. The screen builder will also take an existing Screen COBOL screen description, display the screen, and allow the user to make modifications directly at the terminal.

Overall control of the ENCOMPASS transaction processing software is carried out by the Application Monitor. This program is used to supervise and control all the working processes in the transaction processing system. It performs on-line addition, modification, or deletion of transaction types, screen formats, and terminals under user control. The application monitor also controls the start-up of all other working processes of the system.

The second software package included in ENCOMPASS is a data definition language which allows centralized administration of a data base to accommodate any number of application programs. It is used to facilitate the definition of the data base schema. The data base is defined in a Schema Definition File using the Edit program. The Schema Definition File is used as the input to the Schema program. This program generates a TAL library file that, when compiled along with the application program, produces an object application program that is tailored to the particular data base.

Record types are accessed through the Guardian operating system. Guardian provides high-level access to, and manipulation of, records in data bases. It offers three disc file structures: key-sequenced, with variable record lengths; relative, with fixed record lengths; and entry-sequenced, with

Tandem NonStop Systems

variable record lengths. Records can be accessed randomly by specifying a key or keys, or sequentially by consecutively accessing the records in the collated order of an index. Since multiple key fields can be defined, multiple indices allow an indexed file to appear as sequential. Moreover, Guardian provides three indexing options: exact key match, approximate key match, and generic key match. Approximate match means that the record key may be equal to or greater than the search key; this allows a user program to access records without knowing the exact key. Generic key match means that only the initial portion (partial key) of a key need be specified (such as the prefix to a part number in a vendor's record).

Other features of the operating system's access methods include multi-key access to records (up to 255 alternate key fields), automatic maintenance of all keys, data compression for key-sequenced files, index compression, record locking, multiple-volume files, a cache buffering scheme, and a file utility program to create data base files.

The third major capability of ENCOMPASS is a system-level capability which simplifies the task of designing transaction oriented data management applications on a Tandem system by providing a number of facilities to maintain data base consistency during concurrent transaction processing. The data base can reside on a single Tandem system or can be distributed over multiple nodes of an Expand network. In either case, this facility ensures that the data base remains consistent in the event of a program failure, a single component failure or even the total loss of communications between nodes.

Data base consistency is protected through a network-wide system identification of all data base activity initiated by each transaction. ENCOMPASS monitors the data base activity by every application program which processes the transaction on any of the nodes in the network. This is implemented in a distributed manner, with no single point of control.

If a failure such as a program abort, CPU or other module failure affects the processing of the transaction, the system automatically backs out the data base changes made by that transaction up to the time of the failure. Processing of the transaction is then restarted. Only the individual failed transaction is backed out. The backout occurs without impacting other transactions in progress. Transaction backout occurs on-line and is transparent to the application program. No restart of the system occurs.

ENCOMPASS uses multiple disc audit trails for high transaction throughput. The Guardian operating system automatically provides "before" and "after" images of data base updates by application programs to the audit processes that write the audit trails. Auditing of the data base is totally transparent to user-written application programs. For transactions that span data bases on multiple nodes of an Expand network, all audit images for records residing on a particular node are contained in audit trails on that node.

Data base consistency is ensured both within a single node and across nodes by treating all data base updates made by a single transaction as a group. If a failure occurs, all the changes are backed out as a group, returning the data base to its original state. This takes place even when the changes were distributed over multiple nodes in the Expand network.

This data base consistency facility also provides: transaction completion via a "two phase commit" principle which resolves the problem of loss of network communications during a distributed transaction; a concurrency control mechanism that protects individual transactions from seeing uncommitted data produced by other concurrent transactions; on-line backup to tape of portions of the data base while full data base update activity continues; and recovery from catastrophic failures by reload of prior on-line backup dumps, and application of the data base activity audit trails.

ENCOMPASS also includes the Enform relational query/report writing language, a high-level query specification language. Enform enables programmers as well as non-programmers to retrieve information easily from a relational data base and generate professional looking reports by using a few short English-like statements. At the same time, Enform provides the experienced programmer with a powerful, efficient means of accessing, sorting, and formatting data so that application development time is reduced.

Enform can be used interactively or programmatically by COBOL, FORTRAN or TAL (Tandem's system level language) programs. The versatility of the language lets the user adapt Enform to a wide spectrum of applications, from simple ad hoc inquiries to sophisticated software development. Output can be formatted into clear, intelligible reports with centered titles, subtitles, column headings, underlining, totals, subtotals, averages and percentages.

Enform features include: shared access to data base and dictionary; efficient information retrieval; transparent optimization of data base searching; comprehensive report formatting; easy customization of query/report applications; use interactively or through a host language; decreased amount of data transmitted over a network.

TANDEM SORT: This program reorders a set of records according to the values of sort key fields defined within the records in either ascending or descending order. It can be driven by a set of commands entered conversationally, by a text file containing the commands, or by a user's application program.

Records can be passed to Sort from a file, or sent one by one through procedure calls from a user's program. Similarly, the sorted set of records can be written to a file, or the user's program can call a procedure to retrieve the records, one per call.

Actual sorting runs as a separate process from the host program. Standard interface procedures, which are present in the sort command interpreter program or called from the user's program, handle process creation, control, and communication.

In most large sorts, the memory is insufficient to sort all the data at once, so Sort splits the input data into sorted pieces (runs) it can handle and puts them into a scratch file. Sort then uses the replacement/selection method of merging runs together to produce the final product.

SPOOLER: The Tandem SPOOLER provides a means of storing application output in holding areas for later retrieval. Output may be passed to other processes or printed on one or more devices. The SPOOLER is actually several processes working in unison to provide spooling facilities. These processes include the SPOOLER supervisor, collectors, print processes, and SPOOLCOM.

The SPOOLER supervisor functions as the SPOOLER monitor and communicates with the other SPOOLER processes to determine which tasks to perform or schedule. It interfaces with the other processes as well as applications calling SPOOLER procedures. SPOOLER collectors accept output from application processes and store it on disk. There can be one or ore more collectors. Print processes retrieve spooled data and print it. The Tandem supplied print processes are capable of handling multiple jobs and devices. Users may supply their own print processes. SPOOLCOM is an operator/user interface with the SPOOLER subsystem. It can be run interactively on a terminal or can be passed commands from an application process. SPOOLCOM

Tandem NonStop Systems

performs such functions as downing a device or ordering extra copies of a report. An application process can open the SPOOLER one or more times to perform spooled output. The standard file management procedures WRITE, CONTROL, and SETMODE are used. The application may also use the SPOOLER library procedures to implement more advanced features of the SPOOLER.

The Exchange remote job entry subsystem enables a NonStop system to emulate two types of batch remote job entry: a multileaving HASP workstation or an IBM 2780/3780 data transmission terminal. Input and output can be from/to any media supported by the NonStop system including disc, magnetic tape, terminals, card readers, line printers and other processes.

General capabilities of Exchange include transmitting and receiving in ASCII or EBCDIC, accepting horizontal tab codes, accepting vertical forms control codes, transmitting or receiving EBCDIC transparent data, short record truncation, blank field compression, transmitting and receiving block data link messages and generation of WACK and TTD control codes when temporarily unable to transmit or receive.

Files can be sent or received by entering commands from a terminal. For the 2780/3780 terminal, an Exchange subsystem on one Tandem system can perform remote file transfers to another Tandem system which is also running the Exchange subsystem. If Exchange is emulating an IBM HASP workstation, files can be sent and received simultaneously.

HASP multileaving workstation console support is provided via the Exchange command interpreter or programmatic interface.

ENVOY DATA COMMUNICATIONS MANAGER: Envoy provides an interface between transaction oriented applications and a data communications network. It comprises a set of communications processes that run as a part of the Guardian operating system. These system processes control one or more data communication lines and process requests received from application processes through file management procedure calls. The communications processes secure the system buffer space needed to complete requests, monitor the state of communication lines, and return completion status and data to the application process.

The protocol procedure determines the message format for a request, the sequence of line operations to be performed, and the corrective action to be taken when a line error is detected. The driver procedure schedules I/O operations to the communications controller. Envoy supports the following protocols: bisync, ADM-2, TINET, Burroughs, Full Duplex, Asynchronous, SDLC, and HDLC.

Through use of Envoy, the NonStop system can act as: a station in a point-to-point network; a supervisory station in a multipoint network; and a tributary station in a multipoint network.

Envoy also allows NonStop systems to connect to the international electronic funds transfer network known as the Society for Worldwide Interbank Financial Telecommunication (S.W.I.F.T.). Users can plan and design totally integrated domestic/international funds transfer systems utilizing standard Tandem supported interfaces. The software provides a protocol level (byte synchronous) interface to the S.W.I.F.T. network. The application level of the S.W.I.F.T. interface must be provided by the user.

EXPAND NETWORK: In October 1978, Tandem announced Expand, a communications network system. An extension of the Guardian operating system, Expand can accommodate up to 255 geographically dispersed systems, each with a maximum of 16 central processors, all operating in Tandem's fault-tolerant mode.

Expand provides for automatic routing, and rerouting as necessary, of communications among network nodes. Any system located within the network can communicate with any other system without traditional point-to-point connections between all the systems. Also, systems capable of network communications may be added to the network by simply being connected to any one of the existing systems.

Single or multiple-line paths between adjacent nodes are available in order to provide increased bandwidth and recoverability. Lines may be added to or deleted from a multiple-line path between adjacent nodes to adjust bandwidth in economic increments. Transmission capacity may be matched to growing requirements without forcing the use of expensive high speed communications facilities.

For increased flexibility in communications interface, each of the nodes in an Expand network may be connected by either dedicated private lines, or they may be linked via a public X.25 packet carrier. The availability of an X.25 interface may reduce the data communication costs of network users.

Computing resources and data bases in an Expand network are location-transparent to operators and programmers, and application programs are identical for single or multiple systems. A program resident in any system in a Tandem network may access data (subject to user-imposed security restrictions) resident in any other system. The program will be executed in the system on which the program file resides, or the user may request that the program be run remotely.

Automatic routing in the Expand network is accomplished by the Network Control Process (NCP) which resides in each system. Through an algorithm, the NCP assures that data is communicated via the best path in terms of minimum travel time. In the event of a line failure, the NCP automatically reroutes communications in accordance with the next best travel time between communicating systems.

Keys to the automatic routing process are NETMAP, the Network Routing Table (NRT) data arrays, and NETCHANGE, a message service, all managed by the NCP. If the NCP is advised by NETCHANGE of a change in status of a line or system node, it determines via its algorithm if there is a path of communications shorter than the one currently being utilized and accordingly alters NETMAP. All systems are then notified by NETCHANGE, and the NRT is updated with this latest information and forwarded data packets are directed to the proper systems. Residing within the Network Control Process is a tie-breaker algorithm for use in the event that two or more alternate routes are identical in terms of travel time between two communicating systems.

X.25AM is a software package supporting the X.25 packet-switched network protocol, and is designed to allow Tandem computer systems to communicate directly over public packet-switched networks using standard protocols defined for use with X.25. Included with the package are protocols to allow access to virtually any kind of computer or terminal. The interactive terminal interface (ITI) provides for remote terminal access to the host Tandem system from anywhere in the country, or from around the world via international gateways. Remote terminals look like local terminals, with no special programming necessary. In addition, X.25AM allows an application within the Tandem system to communicate with any other Data Terminal Equipment (DTE) connected to the packet network using any agreed upon protocol.

An additional higher level protocol is used to allow distributed Tandem systems to be linked together through a public packet network using Expand.

Tandem NonStop Systems

The AM3270 AXCESS METHOD gives application programs the ability to access individual 3270 compatible devices connected to a non-switched bisynchronous, multipoint communications line. It provides the ability for Tandem programs to access individual 3270-type devices on the line without concern for communication-related requirements. In addition, a pass-through mode is provided to allow terminals connected to a Tandem system to directly access IBM-compatible hosts.

The TR3271 AXCESS METHOD provides a simplified means to connect a NonStop system to an IBM-compatible host system by emulating one or more 3271-type cluster controllers and associated terminals and/or printer devices. It provides interface by allowing a Tandem application program to appear like a terminal to the IBM host. In conjunction with the AM3270 pass-through protocol, TR3271 allows real 3270 devices to be connected to Tandem, yet access application programs running on the IBM-compatible system.

XRAY: A performance monitor for the Tandem systems whose applications include:

- **Mix balancing—the distribution of applications across system hardware to eliminate bottlenecks.**
- **Growth management—the long term appraisal of system components usage for planning, budgeting, and control purposes.**
- **On-line monitoring—a provision for immediate detection of performance difficulties and continuous feedback on system usage.**
- **Application tuning—a feature showing where the application program should be restructured to increase transaction throughput.**

XRAY consists of two programs, XRAYCOM and XRAYSCAN. XRAYCOM is used to control measurements. Data is collected in a disc file and system performance is analyzed by running the program XRAYSCAN against the collected data for data reduction and analysis. On-line performance monitoring is achieved by running XRAYSCAN against the currently active disc file. Any item or set of items in the meaurement can be plotted on a terminal, as they are observed.

APPLICATIONS: Tandem does not provide application software at the present time, but from time to time it may assist customers with particular programs.

PRICING

POLICY: The manufacturer offers the Tandem systems on a purchase-only basis, with separately priced software. Maintenance is also separately priced. Tandem can assist the customer in obtaining third-party leasing.

System deliveries are F.O.B. Tandem's plant. The warranty period extends for 90 days after initial delivery to end users only, and for 30 days to OEM's. Tandem provides on-site installation for fixed fees stated in the equipment price list of the purchaser's location within the contiguous United States. For other locations, consult Tandem. Installation facilities, including electrical power and connector requirements, are the responsibility of the user and must be completed prior to installation.

Tandem provides training classes on a fee basis in software design and programming (COBOL, FORTRAN, TAL, ENCOMPASS, and MUMPS) and hardware maintenance. Classes range in duration from one to three weeks. Software and hardware coures are priced at $500 per week per student.

Tandem provides post-warranty maintenance service under a standard maintenance contract.■

EQUIPMENT PRICES

PACKAGED SYSTEMS		Purchase Price	Monthly Maint.	Installation
NonStop				
The NonStop packaged system consists of the following modules: two processors, each containing 384K bytes of 500 nanosecond semiconductor (MOS) memory, power fail/auto restart, memory mapping, bootstrap loader, interval timer, and DMA for all I/O. Each processor may be expanded to 2M bytes of memory. NonStop systems also include dual Dynabus redundant interprocessor links rated at 13M bytes per second each, two block multiplexed I/O channels rated at 4M bytes per second each, 13 unassigned I/O slots for system expansion, a system cabinet, one hard copy console, one dual channel connected asynchronous controller, one terminal patch panel, one magnetic tape controller, one magnetic tape drive, and two battery packs.				
T16/244-3	NonStop System; includes two T16/1412-1 processors, each with 384K bytes of main memory expandable to 2M bytes, and 13 unassigned slots for system expansion. Also included in the package are a T16/7104 system cabinet, T16/3202 magnetic tape controller, T16/5103 magnetic tape drive, T16/6603 hard copy console, and two battery packs	$94,975	$726	—
NonStop II				
The NonStop II packaged system consists of the following modules: two processors, each containing 512K bytes of 400 nanosecond semiconductor (MOS) memory, power fail/auto restart, memory mapping, support for 16-bit and 32-bit addressing, bootstrap loader, interval timer, and DMA for all I/O. Each processor may be expanded to 2M bytes of memory. NonStop II systems also include dual Dynabus redundant interprocessor links rated at 13M bytes per second each, two block multiplexed I/O channels rated at 5M bytes per second each, 23 unassigned I/O slots for system expansion, one system cabinet, one magnetic tape controller, one magnetic tape drive, two battery packs, and one Operations and Service Processor (OSP).				
T16/250	NonStop II System; includes two T16/1420 processors, each with 512K bytes of 400 nanosecond semiconductor (MOS) memory expandable to 2M bytes, and 23 unassigned slots for system expansion. Also included in the package are a T16/7120 system cabinet, T16/3202 magnetic tape controller, T16/5103 magnetic tape drive, one T16/3910 Operations and Service Processor (OSP), and two battery packs	144,475	883	—

Tandem NonStop Systems

EQUIPMENT PRICES

		Purchase Price	Monthly Maint.	Installation
PROCESSORS				
T16/1412-1	NonStop processors include two microcoded processing units, one for CPU and Dynabus control and one for I/O control (each with 100 nanosecond cycle time); 384K bytes of 500 nanosecond semiconductor (MOS) memory with EDAC; complete DMA only I/O system; virtual memory control; memory mapping protection for up to 2M bytes of memory; bootstrap loader; interval timer; control panel; Dynabus controller and interface; interfacing for 32 I/O controllers; power fail/auto restart; power supply; and one battery pack	33,500	243	600
T16/1420	NonStop II processors include three microcoded processing units, one for CPU and Dynabus control, one for I/O control (each with 100 nanosecond cycle time), and one for Diagnostic Data Transceiver Control (DDT); 512K bytes of 400 nanosecond semiconductor (MOS) memory with EDAC; complete DMA only I/O systems; virtual memory addressability—to one billion bytes; 8K of 32-bit words of loadable control store, 1K of 32-bit words of read only storage; power fail/auto restart; bootstrap loader; interval timer; control panel; Dynabus controller and interface; interfacing for 32 I/O controllers; string manipulation and double word arithmetic; power supply; and one battery pack	50,500	237	600
PROCESSOR OPTIONS				
	Decimal Arithmetic Package	2,000	22	100
	Floating Point Arithmetic Package; adds additional instructions to basic set	2,000	22	100
	Power Module; supplies power for I/O controllers in large configurations	3,500	40	600
	Battery Pack Backup for MOS Memory	1,500	28	60
T16/3910	Operations and Service Processor (OSP, for NonStop II system only)	14,875	163	300
MEMORY				
T16/2412	384K bytes Error Correcting MOS Memory Module; 500-nanosecond cycle time (for NonStop only)	9,600	132	100
T16/2420	512K bytes Error Correcting MOS Memory Module; 400 nanosecond cycle time (for NonStop II only)	12,800	47	113
MASS STORAGE				
T16/3106	Disc Controller; dual channel connected; can control 1 to 8 drives; uses 2 I/O slots	10,500	58	200
T16/4104	Disc; moving head; 240MB	26,500	208	450
T16/4105	Disc; moving head; 64MB	15,500	174	300
T16/4106	Disc; moving head; 64MB; drawer mounted	15,200	174	413
T16/4109	Disc; fixed & moving head; 64MB	18,270	123	450
T16/7202	Two-high disc cabinet	1,500	—	—
T16/7504	Disc Patch Panel—STD	775	—	80
T16/7603	Cable; disc patch panel	—	—	—
T16/7606	Cable set, backup controller; Drive 0	—	—	—
		—	—	—
MAGNETIC TAPE EQUIPMENT				
T16/3202	Magnetic tape controller for up to two drives; 800/1600 bpi controller; can be powered from either processor	4,800	23	188
T16/3203	Magnetic tape controller; 200 or 800 bpi; 7 track	4,800	23	188
T16/5103	Magnetic Tape Drive; 45 ips, 9-track, 800 bpi NRZI and 1600 bpi P/E, 36 or 72K BS; requires T16/3202	8,000	57	300
T16/5104	Magnetic Tape Drive; 125 ips, 9-track, 800 bpi NRZI and 1600 bpi P/E, 100 or 200K BS; required T16/3201 or T16/3202	14,500	110	300
PRINTERS				
T16/3302	Line Printer Controller	2,800	20	188
T16/3305	Card Reader/Line Printer Controller for one printer and one card reader; dual channel connected; can be powered from either processor	2,800	20	188
T16/5508	Serial Printer; 200 cps, 96 character set	4,500	55	150
T16/5503	Line Printer; 600 lpm, 64 character set	14,000	178	225
T16/5504	Line Printer; 900 lpm, 64 character set	21,000	202	225
T16/5510	Line Printer; 1350 lpm; 132 columns; 67 character ASCII set	45,500	507	225
PUNCHED CARD EQUIPMENT				
T16/5301	Card reader; 80 columns, 600 cpm; requires T16/3305 (for NonStop only)	5,800	46	175
TERMINALS				
T16/6603	Hard-Copy Terminal; 30 cps, 132 columns; RS-232 interface	2,900	48	150
T16/6604	With 20-ma current loop interface	3,200	44	150
T16/6520	CRT, multi-page; 3 or 6 displayable memory pages; RS-232 or 20-ma current loop	3,150	24	150
T16/6524	Same as 6520 except with printer port	3,150	24	150
T16/6810	Peripheral line adapter; 1 input port and 8 output ports	1,975	15	113

Tandem NonStop Systems

EQUIPMENT PRICES

		Purchase Price	Monthly Maint.	Installation
COMMUNICATIONS EQUIPMENT				
T16/7501	Terminal Patch Panel; provides connection between asynchronous controllers and up to 17 terminal ports	775	6	75
T16/7502	Synchronous Patch Panel; provides connection between synchronous controllers and up to 12 synchronous lines (3 controllers)	775	6	75
T16/6202	Byte Synchronous Controller; dual channel connected; may be powered from either processor; controls up to four synchronous communications lines; either point-to-point or multidrop	5,800	29	100
T16/6203	Bit Synchronous Controller, dual channel connected	5,800	39	100
T16/6303	Asynchronous Controller; dual channel connected; may be powered from either processor; controls up to two terminal lines, either hard-wired or modem-connected	3,600	18	125
T16/6304	Asynchronous Extension Board; provides additional control for 15 asynchronous lines	4,300	22	188
T16/3401	Universal Interface; dual channel connected; can be powered from either processor; controls any two devices having 16-line parallel interface	2,800	20	188
T16/3810	Tandem to IBM Link (TIL); high speed communications link to an IBM system	24,500	125	600
T16/3812	Tandem to IBM Link; additional link, required for each additional link or site	12,900	91	375
T16/3830	Tandem HyperLink (THL); connection between a Tandem system and other systems using Network System Corporation's HYPERchannel	14,900	151	300
HARDWARE				
All system cabinets include Dynabus redundant interprocessor priority resolution circuitry and data paths, provide I/O controller power switching, and provide processor I/O channel wiring.				
T16/7104	Large System Cabinet; provides space for up to 4 processors and 16 I/O controllers, each processor having 6 memory slots; can accommodate up to 512K bytes of semiconductor memory per processor; can accommodate up to 2 additional I/O expansion boards (T16/7801) (for NonStop system only)	8,600	88	400
T16-7107	Expansion Cabinet; for systems requiring more than 2 system cabinets and/or more than 8 processors	2,500	N/C	400
T16/7801	I/O Expansion Board; provides eight additional I/O slots for systems cabinets; maximum of two per cabinet (for NonStop system only)	970	—	600
T16/7109	Cabinet, I/O (for NonStop system only)	8,600	88	400
T16/7110	Cabinet, peripheral	2,500	N/C	400
T16/7120	System Cabinet; provides space for up to 4 processors and 24 I/O controllers; each processor having 4 memory slots; can accommodate up to 2M bytes of semiconductor memory per processor (for NonStop II system only)	15,800	166	600
T16/7121	Cabinet, I/O (for NonStop II system only)	15,800	104	600

SOFTWARE PRICES

The following software packages are available for the NonStop and the NonStop II. Some of the software packages require microcode for each processor and/or a bit or byte synchronous controller. Microcode prices range from $500 to $7,000 per processor; each controller costs $5,800.

	One-Time License Fee	Monthly Maint.
Base Guardian—standard operating system; includes Enscribe data base access method	$11,000	$110
Exchange—Remote Job Entry and HASP	2,000	20
XRAY software subsystem performance monitor**	2,500	25
Expand Networking System	10,000	100
X25AM-X.25 Access Method	2,000	20
Guardian operating system with SPOOLER, XRAY, and Base Guardian	13,000	130
AM3270 Axcess Method	2,000	20
TR3271 Axcess Method	4,000	40
AM 6520 Axcess Method	2,000	20
Transaction Monitoring Facility*	10,000	100
SPOOLER**	2,000	20
Enform Query Report Writer*	7,000	70
Pathway Transaction Processing System*	8,500	85
ENCOMPASS distributed data base management system; includes Pathway, Enform, Transaction Monitoring Facility, and Data Definition Language (DDL)	22,000	220
DDL—Data Definition Language*	3,000	30
COBOL—ANSI 74	7,000	70
FORTRAN—ANSI 77	6,000	60
MUMPS—ANSI 77	7,000	70

*Can be licensed as part of the ENCOMPASS distributed data base management package.
**Can be licensed as part of the Guardian operating system.

DATAQUEST INCORPORATED
A Subsidiary of A.C. Nielsen Co.
RESEARCH NEWSLETTER

SCIS Code: Tandem Computers Inc.

THE NONSTOP II
TANDEM'S FIRST MAJOR PRODUCT FOR THE EIGHTIES

SUMMARY

On April 14, 1981, Tandem Computers Incorporated announced the Tandem NonStop II, its second computer system. According to David R. Mackie, vice president of Tandem's Headquarters Marketing Operations, the NonStop II is intended to "carry Tandem forward into the eighties," and toward its stated goal of achieving revenues of $1 billion by the close of fiscal 1985.

The NonStop II provides all the features of the original fault-tolerant NonStop system plus the following new features and enhancements:

- A new 32-bit data access architecture, which is the key to implementing the system's extended addressing capability. Each processor can address up to 16 Mbytes of physical memory without architectural changes as denser memory devices become cost effective. The maximum main memory size is currently 2 Mbytes per processor. Virtual memory addressing capability is now one billion bytes per processor.

- Software compatibility with current NonStop systems at the object code level.

- Enhanced input/output capability that will significantly improve the terminal throughput and the number of terminals and communications circuits available.

- Processor microcode resident in 32 Kbytes of loadable control storage, which allows the processor's instruction set to be easily modified.

- A new Operations and Service Processor (OSP) that provides both local and remote detailed system status information and diagnostic facilities to assist in reducing repair times.

- An exchange program that allows current customers to upgrade their systems by returning their original NonStop systems to Tandem for credit against the list price of a NonStop II.

19055 Pruneridge Ave. / Cupertino, CA 95014 / (408) 725-1200 / TWX (910) 338-7695 / DATAQUEST CPTO

FEATURES

32-Bit Data Access Architecture

A major feature of NonStop II is an extended 32-bit addressing mode that allows access to the system's entire virtual memory space. A fully configured 16-processor NonStop II system can address up to 16 billion bytes of virtual memory (one billion bytes per processor), giving users "access to almost unlimited data space," according to Mr. Mackie. The extended address ability provides hardware and software support for very large applications with hundreds of terminals, communications lines (including X.25 virtual circuits), and large, distributed data bases.

Each NonStop II processor currently has a main memory capacity of 512 Kbytes to 2 Mbytes, expandable in increments of 512 Kbytes. However, each processor is capable of addressing up to 16 Mbytes of physical memory. As denser memory devices become cost effective, the current maximum memory of the processor can be increased without architectural changes. This strategy recognizes potential user needs for larger memories in the future and should allow Tandem to meet these needs in a cost-effective manner that is advantageous to its customers.

Software Compatibility

NonStop II is fully compatible with the original Tandem NonStop 16 system at both application code and network levels. "Our users . . . will be able to retain virtually all their original investment in software and most hardware if they choose to upgrade to the NonStop II system to take advantage of its new capabilities," stated Mr. Mackie. Futhermore, NonStop II systems can be used in a common data communications network with the current NonStop systems without software modification. However, NonStop II processors cannot be combined with older NonStop processors in a single system.

Enhanced Input/Output Capability

According to James G. Treybig, Tandem's president, the NonStop II is a direct response to the needs of the Company's customers with large, on-line transaction processing needs. The 32-bit data access architecture can support very large I/O intensive applications needed by organizations having critical business operations on-line. The NonStop II expanded architecture provides more capability in the amount of data that can be transferred to and from disks, terminals, and other peripherals, as there is a 16-fold increase in both I/O buffer space available per processor and in the maximum I/O transfer size. Each processor can now support 1 Mbyte of I/O buffer space. I/O channel speed can reach up to 5 Mbytes/second in a burst mode.

Loadable Control Storage

The processor microcode resides in 32-Kbytes of loadable control storage. The loadable control store allows the processor's instruction set to be modified for new features or performance enhancements as part of Tandem's standard software update procedure.

are not really comparable. NonStop II systems are in production and six systems are already in place. Tandem is promising deliveries 120 days after receipt of an order.

Comments

The NonStop II clearly emphasizes two aspects of Tandem's strategy and philosophy.

By bringing out a system that contains significant enhancements to its original product, the Company has reaffirmed its intention to keep its place in the forefront of the on-line transaction processing systems market and its unique position as the only manufacturer of a fully redundant, fail-safe system. Furthermore, the features of the new system show a high regard for both the user's convenience and pocketbook. In our opinion, the planned memory expandability of the NonStop II virtually guarantees it a long duration of usefulness, and the complete software compatibility and serviceability features show a real concern for the customer's needs and welfare.

Ellen K. Clements
Richard J. Matlack

Prudential-Bache
Securities

Industry Outlook
Technology

Donald H. Brown
Marianne Bye

Research

DIGITAL SYSTEMS -- THE MARKET MONITOR November 22, 1982

Preliminary analysis of survey contacts with 96 third-party marketing organizations suggests that there is little change in the outlook in comparison to the summer months. Demand continues to bump along with no signs yet appearing for a significant upturn.

- Revenue growth expectations are at 18% for the three-month period ending in December versus the three-month period ending in September. This compares with a 23% expected growth rate for the June/September survey period, and 36% for the March/December period early this year.

- There is still no sign of a sustained upturn in demand for larger systems as expectations there continue to be weak; low-end systems under $32,000 show continued strength.

- Expected employment gains of 8% are predicted for the December quarter. Earlier surveys had shown expected increases of 4% for the September versus June quarter, and 15% for the first quarter of this year.

- Anticipated improvement in employee productivity implies that profitability for third-party marketing organizations will rise.

- Feedback on the DEC personal computer confirms a high level of interest, and indiates that volume production is likely to be established by February, 1983.

- Hewlett-Packard may not be getting much marketing leverage in third-party distribution channels for the new 32-bit desktop product which will be announced shortly.

REVENUE GROWTH

Mail contacts with 71 small third-party marketing organizations and telephone surveys of 25 larger companies indicate that revenue growth expectations of 18% for the three months ending in December are down slightly from an expected 23% growth for the June/September period, and

26% for the March/December quarter. Our hopes are fading for a significant upturn in demand for small computer systems this quarter.

Results indicate that demand for smaller systems selling for under $32,000, while down slightly from the September/June quarter, continues to lead in terms of growth expectations over the higher-priced system classes. Large systems still exhibit vulnerability in the marketplace, and continue to be weak in terms of projected revenue growth. This is a plus for personal computer and word processing vendors, but a short-term negative to the systems suppliers such as DEC, HP and Prime.

Historically, "minicomputer" suppliers and their customers have "powered through" a recession. Overall, total market expansion has averaged 36% annually for the 1975 to 1981 period using the combined revenues for Data General, Digital Equipment, and Prime as an index. In truth, the expansion was highly volatile around the average. Growth for the group, for example, dropped briefly from 58% in early 1974 to 19% during the 1974-1975 recession. Demand quickly rebounded coming out of the last recession with expansion of 40%-50% registered in 1976. In short, growth might slow -- but only temporarily. With the first turn in industrial production, a concurrent business indicator, orders picked up. (For a discussion of the historic economic relationship of demand to economic trends, see Digital Systems, No. 19, April 7, 1982.)

In such an environment, the name of the game was to match employment expansion levels to short-term business fluctuations, which were derived principally from the economic cycle. For those that missed by overexpanding going into the downturn, margins were hammered. Even so, there was continued pressure to maintain position for a subsequent upturn, which would be only around the corner. The real priority was to keep the technical work force intact in the face of earnings pressures.

This may no longer be the case for larger systems in the $128,000-$516,000 price class. These systems represent a sharply higher proportion of revenues for DEC and Prime in particular than in the prior economic downturn. The success of the "minicomputer" manufacturers in pushing up the product line means that a rising component of revenues and earnings for small computer systems now falls in the $128,000-$516,000 price class, which is treated as a "capital spending investment". An upturn in this class of equipment could be both delayed and modest in comparison to historic patterns.

In our summer survey, we specifically focused on shifts in spending patterns in the scientific/engineering applications. We found that Department of Defense budget delays contributed to intense pressures through the summer months in scientific/engineering applications. Further, these pressures are likely to be reversed going into calendar 1983. In the current survey, we did not recheck the earlier findings, but focused attention in our current effort on DEC's personal computer. (See Outlook for Small Computer Systems in Scientific/Engineering Applications in our DIGITAL SYSTEMS UPDATE of July 21, 1982.)

Table I
REVENUE GROWTH - EXPECTED % CHANGE QUARTER TO QUARTER

	Low-End Systems <$32K	Mid-Range Systems $32-$128K	Larger Systems $128-$516K	Total
Current Fall Survey (Dec./Sept.)				
Total Revenue ($ thous.)	$14,174	$17,021**	$24,020	$61,117*
Exp. Chg. Q/Q	+36%	+9%	+4%	+18%
No. of Respondents	45	27	9	96
Summer Survey (Sept./June)				
Total Revenue ($ thous.)	$4,101	$19,667	$22,103	$98,934*
Exp. Chg. Q/Q	+39%	+10%	+6%	+23%
No. of Respondents	32	20	5	83
Winter Survey (Mar. 1982/Dec. 1981)				
Total Revenue ($ thous.)	$17,867	$19,072	$16,210	$53,149
Exp. Chg. Q/Q	+77%	+9%	+10%	+26%
No. of Respondents	117	59	15	191

* Includes companies not classified by system price class.
** Excludes one very large company reporting 75% expected revenue growth for the December period from $4 million for the September quarter.

EMPLOYMENT GAINS

Projected employment gains for the September/December period have aligned with revenue growth expectations fairly evenly across the board. Employment at the end of the December quarter is expected to rise 8% overall, a slight uptick from the 4% gain anticipated for the September quarter, but well under the 15% expectation of January, when most respondents were still ebullient.

Table II
PROJECTED EMPLOYMENT GAINS - % CHANGE QUARTER TO QUARTER

	Low-End Systems <$32K	Mid-Range Systems $32-$128K	Larger Systems $128-$516K	Total
Current Fall Survey (Dec./Sept.)				
Sept. Employment	539	1,010	354	1,903
Exp. Chg. Q/Q	+15%	+5%	+5%	+8%
No. of Respondents	45	27	9	81
Summer Survey (Sept./June)				
Exp. Chg. Q/Q	+7%	+6%	+2%	+4%
Winter Survey (Mar. 1982/Dec. 1981)				
Exp. Chg. Q/Q	+23%	+10%	+16%	+15%

EMPLOYEE PRODUCTIVITY RISING

Employee productivity is likely to improve through the fourth quarter, particularly in very small systems in the $1,000-$8,000 price class. Although third-party marketing firms will rarely comment on profitability, the revenue per man ratio provides some indication of trends. Overall, survey respondents expect the revenue/man ratio to rise 9% in the December quarter, which implies a significant improvement in profitability. While the top-end, very large systems, may achieve the highest profitability (it is currently $67,850 in revenue per man), survey results show that little further improvement is anticipated here. Surprisingly, very small systems show a respectable ratio of $34,900 in revenues per man. The strongest improvement is expected here. These companies have been successful in selling enough units at the lower price to maintain reasonable prospects for good profits. They've achieved the volume to compensate for lower unit prices.

Table III
EMPLOYEE PRODUCTIVITY SHIFTS

	Average System Price ($ Thous.)				
	Very Small $1-$8	Low-End $8-$32	Mid-Range $32-$128	Larger Systems $128-$516	Total
Fall Survey Revenues Per Man: (Sept. Qtr.)	$34.91	$16.34	$20.81	$67.85	$32.12
Expected Gain in Revenues Per Man: (Dec. Qtr.)	+18%	+15%	+15%	-1%	+9%

DEC'S PERSONAL COMPUTER CAPTURES STRONG ACCEPTANCE

Of specific interest, momentum for the DEC personal computer appears to be rising in the third-party distribution channels. Of 18 DEC authorized distributors contacted, two have ordered systems for February, with one of the two already experiencing an order cut. Thirteen are seriously evaluating the product. Only three have decided not to carry the line. The timing of the contacts was particularly appropriate since the company just recently held a meeting for the "Commercial OEM Product Line" with these customers in Florida.

None of the distributors, however, will be receiving units before February and March of next year. At first, the finding appeared to suggest delays in the production build-up of DEC's new pc's, and to directly conflict with the recent statement by the company that "there are about 1,000 personal computers installed at test sites." Further discussions clarified the difference in DEC's parlance between the "availability" of a new product, and "volume deliveries." The first term means that shipments have been initiated to targeted accounts. Indeed, installations have begun to companies such as AT&T and TRW, as well as important software organizations heavily involved with the development and testing of the product. Only later, with "volume deliveries" would the product begin shipping in high quantities to mass distribution channels such as authorized DEC DEALERS (which concentrate on the terminal product lines) and authorized DEC DISTRIBUTORS (which work with systems.) Thus, when the 15 DISTRIBUTORS report that they will not be getting the product before February or March, that represents the point at which volume deliveries will be rolling off the production line.

DEC will be relying, as usual, on multiple -- and often conflicting -- marketing channels. Large "mass merchandisers" include Computerland and Hamilton/Avnet. DEC's authorized DISTRIBUTORS and DEALERS can qualify, provided they sign a special agreement outlining their plans which will generally require an emphasis on vertical marketing to gain DEC's support. In part, the company may want to discourage system houses, which have primarily concentrated on larger systems in the $16,000 to $64,000 price class, from being drawn inadvertently into the $1,000 to $16,000 sector which requires a totally different approach to selling and supporting systems. Thus, DEC's requirement that its established DISTRIBUTORS file a business plan to qualify for the personal computers does not represent the papering over of program delays, as expressed by one such organization several weeks ago. Rather, DEC appears to be pushing these organizations hard to ensure that they have thoroughly thought through the pitfalls and challenges of supporting the new product as well as the opportunities. Finally, DEC will be selling and supporting large corporate accounts with its own direct sales force, and through DEC's stores.

HP'S DISTRIBUTION LEVERAGE FOR THE NEW 32-BIT TECHNICAL DESKTOP

Fifteen HP System Houses were contacted in the latest effort, of whom none had decided that they would definitely carry the new HP 32-bit, technical desktop computer to be announced this month. Only two indicated it could be of interest to them, and the remainder were both unaware and uninterested in the product announcement. In comparison to the strong acceptance within the existing DEC distribution channel for DEC's new low-end offering, HP would seem to be obtaining little marketing leverage in System Houses for the new desk top. In large part, HP's third-party marketing organizations appear to distribute a relatively narrow range of "point" products from only one of HP'S divisions. Even so, DEC's product has been announced and marketed for many months, whereas HP's had not even been announced. As a result, it remains premature to assess timing and magnitude of the contribution of the new products from HP through System House marketing channels.

Prices as of 11/17/82: Digital Equipment (DEC-102 3/8), Hewlett-Packard (HWP-68 1/4), Prime Computer (PRM-35 1/4).

Additional information is available upon request. 82-569

Donald H. Brown
Marianne Bye
(212) 791-2946

Bache
Bache Halsey Stuart Shields Incorporated

Institutional Research

Digital Systems

No. 20

Donald H. Brown
Marianne Bye
(212) 791-2946
September 16, 1982

Mini-MicroSystems
CAHNERS PUBLISHING COMPANY

1982 BACHE/MINI-MICRO ANNUAL MARKET SURVEY

COMING IN DIGITAL SYSTEMS No. 21: "DEC's Return to Basics" -- The success of the VAX product lines created a new class of "super mini's," but left the company vulnerable to inroads at the low end of the product line. Our next issue concentrates on the strategic shifts taking place at DEC over the latest year which have contributed to a new focus on the low end.

LISTING OF CHARTS AND TABLES

SECTION I

SECTION II

TOP 10 Rankings

Application Analysis

Price Class Analysis

SUMMARY AND CONCLUSION: The Year of the "Micro"

With this issue of Digital Systems, Bache Halsey Stuart Shields Incorporated formally releases the first highlights of the 1982 BACHE/MINI-MICRO ANNUAL MARKET SURVEY. Questionnaires were mailed late last year to 71,500 subscribers of Mini-Micro Systems, the leading trade publication covering small computer systems. The study included 11,202 responses that were returned early this year of which 10,801 represented usable responses from separate buying sites. These responses included 28,498 separate entries covering individual computer models, representing 139,206 minicomputers or microcomputers with an estimated end-user sales value of $2.7 billion purchased in 1981. The survey represents one of the few reliable sources of detailed data available covering mini/micro users and provides a necessary tool of analysis in a sector replete with inflated claims from individual suppliers. The final report will be released in full in late September.

All surveys reflect sample bias. This effort concentrates on the technically sophisticated, trained engineer, who has traditionally defined the core of the mini/micro market. The results demonstrate the unique aspects of the subsector, which contrasts starkly with other surveys narrowly focusing on commercial applications -- an area dominated by IBM. Further, the survey would tend to be biased against the hobbyist personal computer user -- which makes the positive results for Apple and Tandy all the more surprising.

A comparison of the 1982 results with the 1981 findings focuses on critical shifts. Last year's survey was generally neutral to favorable for the traditional "mini" manufacturers including Digital Equipment, Data General, and Hewlett-Packard, which have supplied complete computer systems in the past. With the 1982 study, the traditional minicomputer manufacturers (led by DEC) retained their strong leadership position in the industry in terms of estimated shipments because of the successful drive up the product line with larger systems in the $124,000-$512,000 price class. Over the last five years, this new price class in minicomputers has fueled revenue expansion. Last year, the larger systems accounted for slightly over a third of total shipments, and fully half of DEC's shipments.

Even so, this year's results suggest that the traditional "mini" suppliers lost substantial market position on the basis of purchasing sites, whereas the "micro's" blossomed. "Micro" vendors included Intel from the semiconductor sector and Apple and Tandy selling personal computers in the $1,000-$8,000 price class. Intel grabbed first place in unit deliveries for third-party marketing firms; on the basis of units, Apple scored as No.-1 supplier to end-users. Moreover, these two coupled with Motorola and Tandy made deep inroads in the sector in terms of their breadth of penetration as measured by the number of sites purchasing their proucts. The broadening of the "micro" customer base with high unit volumes secures the position of these companies in the sector and provides them with the means for future growth. Once a customer base is established, the potential for profitable add-on and upgrade business remains.

The results of the survey once again conclusively demonstrate that the mini/micro markets are not the domain of IBM, despite the company's dominant position in traditional commercial markets. The company has fallen far short of its long-term goals to achieve leadership in the sector. Nonetheless, IBM scored strong gains in share overall in the latest year, with particular success in business data processing applications -- its home market. For the first time, IBM appeared as a TOP 10 participant in engineering/scientific applications as well. Part of

IBM's success in the sector may be attributed to the company's increased willingness to match the discount terms required to deal with third-party marketing firms and, evidently, to accept the conflicts internal to IBM that result from whatever competition may arise between third-party marketing firms and IBM's own direct selling efforts.

The continuing absence of an effective presence for the other traditional mainframe computer vendors represents a severe long-term negative, which confirms the findings of earlier surveys. Similarly, the "intelligent terminal" subsector, including Datapoint, Four-Phase, and Mohawk Data, does not fare well in the results -- another trend that was apparent in last year's results.

While 1981 shaped up as the "Year of the Micro," the direct challenge to the established minicomputer manufacturers still remains largely in the future. The new low-end products represent a high proportion of shipments and purchasing sites, but a small fraction of shipment value. Even so, the low end has traditionally fueled the growth of mini/micro markets by continuously introducing new customers and applications to computers. The established suppliers have recognized the challenge, and 1982 may be shaping up as the "Year of the Response." The next issue of Digital Systems will focus on "DEC's Return to Basics," with particular emphasis on pricing trends.

TOP INDUSTRY STANDINGS: Diverging Strategies Between "Mini" and "Micro" Vendors

Results from the BACHE/MINI-MICRO ANNUAL MARKET SURVEY confirm that the "traditional" minicomputer suppliers, led by DEC, continue to dominate the small computer systems market in terms of shipment value -- the criterion which counts most. DEC leads the pack with a substantial 36% share of the $2.7 billion estimated as value of shipments for 1981 from the survey responses. Hewlett-Packard, IBM, and Data General follow up as near ties for second through fourth positions. For the most part, these four suppliers have maintained their leadership position with a successful drive up the product line with ever larger and more expensive systems. Their recent growth is derived largely from their ability to meet the expanding requirements for established customers and applications. In terms of shipment value, there remains only the hint of change with Intel and Motorola scoring high at fourth and fifth positions, respectively, in the third-party sector, while Apple makes the TOP 10 SHIPMENTS standings in terms of 1981 shipment value for end-users. These "micro" vendors, primarily selling products in the $1,000-$8,000 price class, represent new entrants in the industry, as well as new competition to the "mini" manufacturers, which have generally concentrated on systems. Overall, these "micro" vendors captured a small 8% of total shipment value compared with 58% for the four leading "mini" suppliers. Table I (on page 6) provides summary overall data; Tables 1A and 1B in the "Market Share Analysis" Section provide a breakdown between third-party and end-user responses.

In terms of their breadth of penetration, however, the "micro" vendors led by Apple and Tandy have registered dramatic success in the latest year. These two, along with Intel and Motorola, account for a highly respectable 20% share overall of sites, or buying locations, with all four making the TOP 10 SITES standing for third parties and three making it on the end-user listing. DEC's share drops from 36% for shipment value to 22% for sites. Apple with 8% of the sites and Tandy and Intel both at 5%, all challenge the position of Data General (6%), Hewlett-Packard (9%), and IBM (8%). Table II (page 7) summarizes the overall data while Tables 3A and 3B in the second section provide additional detail.

In terms of unit shipments (Tables 4A and 4B), Apple's victory as the No.-1 supplier, edging out DEC with end-users, and Intel's top position with third parties represent significant accomplishments. High unit volumes represent a first strategic objective for low-end products. Unit competition provides one key to understanding the mini/micro markets. Historically, DEC's success as the industry leader was partly predicated on high-volume manufacturing for large numbers of users. Once a customer base is established, the potential for profitable add-on and upgrade business remains.

Even these large gains in terms of units could easily be scored by shipping unusually high numbers of units to only a few clients. If this were true, then the success could be short-lived, since competitors could come in on the next round of technology and pricing to dislodge the leaders. Their drive, however, has been more broadly based. The new low end has cracked open new markets and substantially extended the customer base in terms of numbers of applications. In short, the high unit volumes have broadened the market.

Many customers brought into the fold at the low end subsequently upgrade to larger processors or indirectly contribute to demand for larger processors by tying into computer networks. Thus, the "micro" vendors have secured their position with a breadth of penetration which will be difficult to dislodge. The breadth of penetration provides a degree of security and staying power, and establishes a strong base for potential growth.

TABLE I

TOTAL 1981 TOP STANDING BY SHIPMENT VALUE

Supplier	End-User			Third Party			Total		
	Rank	$ Mil.	% Total	Rank	$ Mil.	% Total	Rank	$ Mil.	% Total
Digital Equipment	1	$ 388	34%	1	$ 582	37%	1	$ 970	36%
Hewlett-Packard	3	108	9	3	119	8	2.5	227	8
IBM	2	174	15	6	49	3	2.5	223	8
Data General	5	41	4	2	155	10	4	196	7
Intel	NA	16	1	4	112	7	5	128	5
Honeywell	6	40	4	10	38	2	6	78	3
Systems Engineer.	NA	17	1	7	48	3	9	65	2
Perkin-Elmer	8.5	22	2	8.5	41	3	9	63	2
Motorola	NA	11	1	5	52	3	9	63	2
Texas Instruments	NA	18	2	8.5	41	3	9	59	2
Tandem	4	47	4	NA	11	1	9	58	2
Prime	7	35	3	NA	14	1	NA	49	2
Apple	10	21	2	NA	12	1	NA	33	1
Wang Labs	8.5	22	2	NA	7	0	NA	29	1
Tandy	NA	2	0	NA	8	1	NA	10	0
Subtotal		$ 962	84%		$1,289	83%		$2,251	83%
Other		$ 187	16%		$ 267	17%		$ 454	17%
Total		$1,149	100%		$1,556	100%		$2,705.	100%

TABLE II

TOTAL TOP STANDINGS BY SITES

Supplier	End-User			Third Party			Total		
	Rank	No. Mentions	% Total	Rank	No. Mentions	% Total	Rank	No. Mentions	% Total
Digital Equipment	1	2,908	22%	1	1,331	20%	1	4,239	22%
Hewlett-Packard	2	1,260	10	5	417	6	2	1,677	9
IBM	3	1,151	9	3	453	7	3	1,604	8
Apple	4	1,129	9	6	388	6	4	1,517	8
Data General	5	746	6	4	429	7	5	1,175	6
Intel	7	457	4	2	486	7	6.5	943	5
Tandy	6	665	5	7	276	4	6.5	941	5
Motorola	NA	203	2	8	187	3	8	390	2
Texas Instruments	10	222	2	9	166	3	8	388	2
Wang Labs	8	280	2	NA	107	2	8	387	2
Perkin-Elmer	NA	207	2	10	113	2	NA	320	2
Prime	9	233	2	NA	73	1	NA	306	2
Subtotal		9,461	73%		4,426	68%		13,887	71%
Other		3,578	27%		2,087	32%		5,665	29%
Total		13,039	100%		6,513	100%		19,552	100%

CHART 1

"MICRO'S" AND IBM GAIN PENETRATION, "MINI'S" LOSE

Supplier	Composite Market Share*	1981 Gain or Loss
Digital Equipment	34.75%	-7.25%
Hewlett-Packard	11.75	-1.00
Data General	8.50	-3.25
"Mini's" Subtotal	55.00%	-11.50%
IBM	11.30%	3.75%
Intel	7.75%	-1.50%
Motorola	2.00	1.50
Apple	10.75	7.00
Tandy	6.50	3.75
"Micro's" Subtotal	27.00%	10.75%

* Average of share for ranking by primary manufacturer designation and by sites for third parties and end-users. Detailed analyses are included in Tables 2A-3B in Section II.

The comparisons which can be drawn relating the year-to-year trend in market share for the TOP 10 vendors further dramatize the gains achieved by the personal computer manufacturers, Apple and Tandy. IBM scored the only increase in share of those supplying larger systems. Chart 1 presents a composite average for the TOP 10 which combines the site and the primary vendor rankings (discussed below) in the third-party and end-user market sectors. On average, DEC suffered a substantial decline of 7.25 percentage points in position within the TOP 10 vendors from 1980 to 1981. The three traditional "mini" vendors -- Data General, DEC, and Hewlett-Packard -- suffered a drop of 11.50 percentage points taken together. With the loss, DEC accounted for 34.75% on average of the sites and primary vendor mentions for the TOP 10 suppliers in 1981 while the three traditional "mini" vendors as a group accounted for 55.00% of the total for the TOP 10. IBM increased its share by 3.75 percentage points, on average, to 11.30% of the total. In "micro's," the semiconductor manufacturers held roughly even while the personal computer suppliers registered strong gains. Apple garnered 7.00 additional percentage points to reach the level of 10.75% of the total; Tandy gained 3.75 to hit 6.50% overall.

Because of the importance of the rankings for "primary" vendors, this ranking has been combined with the site data to arrive at a composite average share of the market as an indication of the breadth of penetration. While the data covering site standings indicate the breadth of penetration of individual vendors, the listing for "primary" vendors provides an added indication of loyalty. That is, whenever a survey respondent lists an individual supplier as its "primary" source of computers in 1981, then that user generally stays with that supplier and does not second-source its requirements. For example, DEC furnished over 98% of the units purchased by respondents listing DEC as their primary supplier; Intel furnished 88% of the demand. The dearth of purchasing from multiple vendors represents a particular surprise in the case of the "micro" suppliers, which currently offer a narrow product line concentrated in systems with an average selling prices of less than $8,000.

In terms of shipment values, however, market shares have probably recorded little, if any, shift -- with the exception of a likely gain for IBM. Since data for average systems value delivered to end-users were not collected for 1980, year-to-year comparisons in terms of shipments cannot be made directly. Even so, "micro" shipments in the $1,000-$8,000 price class represent a very high proportion of unit shipments, but a small percentage of the value of shipments. As a result, the market share "erosion" that may have been incurred by the traditional "mini" manufacturers represents a negligible shift in terms of shipments. On the other hand, the data do suggest that IBM's gain in terms of share of shipment value probably rose about in line with the company's expanded penetration in terms of sites. On balance, IBM has fallen far short of the long-term strategic objectives of attaining leadership status in small systems which were targeted by the company in the mid-1970s with the initial introduction of the Series 1 and followed up with additional proliferating products. Over the short term of the latest year, the company has made measurable progress and clearly achieved the status of a second-tier competitor.

Ironically, "micro" vendors have adopted the same strategies that contributed to the extended growth of the traditional "mini" suppliers throughout the latest two decades. Historically, DEC birthed the small computer systems sector with the pdp-8 product family in the mid-1960s, followed by the pdp-11 family in the 1970s. These minimal performance products offered stripped-down functionality and aimed at new applications and new markets which heretofore could not afford the

expensive processors of the traditional suppliers. The constant influx of new customers with expanding applications requirements fueled the growth of minicomputer manufacturers throughout the last decade. Ironically, the "micro" group has followed DEC's lead in capturing new users by coming underneath the older products to secure a broadly based, high-unit penetration. These new customers, as their needs expand, provide the means for future growth for the "micro" vendors as they push up the product line to offer larger, more powerful systems in their own turn. The momentum of expanding the customer base clearly passed from the "mini" to the "micro" vendors in 1981, which drives future growth. The breadth of that base suggests that the "micro" vendors will be difficult to dislodge as well.

TRENDS IN THE $1,000-$8,000 PRICE CLASS: "Micro" Vendors Define a New Low End

Semiconductor manufacturers and personal computer suppliers have created a new product class with systems selling in the $1,000-$8,000 price class underneath the established business from traditional minicomputer manufacturers. This price class accounted for roughly half of the total units included in the survey, around a fourth of the sites, but only an eighth of the value of shipments. As summarized in detailed charts in Section II, Intel, Motorola, Apple, and Tandy accounted for almost three fourths of shipments for the third-party and end-user sectors in this price class. Only one of the traditional "mini" suppliers, Hewlett-Packard, makes it to the TOP 10 standings in the price class. Data General, DEC, and IBM make their presence most felt by their absence.

The semiconductor houses have effectively used their expertise in component manufacturing as a springboard into the systems business. In the move, they fundamentally benefited from the technology itself which has generally offered a 15%-20% improvement annually in price/performance in small computer systems for the last twenty years. Technology has enabled these suppliers to integrate dramatically greater function on individual component chips -- their mainstay products.

In effect, it has been easier and more palatable for the component suppliers to move their product line upward than for the traditional "mini" manufacturers to stake out new territory underneath their established systems business. The trend represents a normal shift up the product line. Moves down the product line historically have presented designers of low-cost, minimal systems with the dilemma of stripping out functions. Severe internal disciplines must be established to implement minimum-performance design strategies because dramatic performance gains are so readily achievable for what appears to be only a nominal increase in system costs. With small increases in cost, performance escalates. Grosch's law, for example, postulates that performance is an exponential function of the cost of systems. As a result, designers almost always tend to look up the product line to larger systems for the "greener grass."

The real success of the traditional component suppliers relates to the range of options they have developed for their customers. To be able to move customers down to minimum-performance systems at the component level, the "micro" vendors first had to expand upwards with full-fledged systems. Large-volume, third-party customers can rapidly develop prototype products using fully supported computer systems with a broad range of software to cut product development time and cost. A user can start with an integrated system such as the System 86/330 from Intel, which has operating system software, disks, power supply, and packaging. As one's volume grows, one may integrate backwards to the board and then semiconductor-chip level. Even more, Intel has extended its effective range with a "foundry"

service. That is, many customers desire a product uniqueness which is not available from standard catalog chips. Intel now offers its customers assistance in designing and processing these custom requirements. The semiconductor manufacturers have established the product range which offers users the option of integrating their designs back to the component level.

At the other end of the scale, Apple and Tandy have scored in the latest year through the combined emphasis on new software packaging and distribution strategies. At $1,000-$8,000 per system, no supplier can afford the high cost of a technically sophisticated salesmen to support marketing. The mass-merchandizing approaches developed by both companies represent a leading-edge wave in the market. Equally important has been the software approach. Visicalc, the astonishingly successful electronic worksheet, represented a new extreme in simplicity -- or stripped-down functionality. With only thirteen basic commands, the limited capability represented a new standard in stripped-down functionality. At the same time, the approach benefited from providing a new class of users immediate access to a useful computer-based tool of analysis which did not require training.

Given a broad base and the likelihood of further moves up the product line to offer upgrades to their customers as applications requirements rise, "micro" suppliers do not represent a current challenge as much as a future threat to the established "mini" manufacturers. While the low end of the product line ranging from $1,000-$8,000 in average system price represents almost half of the units in the survey, the value accounts for less than an eighth of the total. The rising challenge derives from the potential for futher moves up the product line where direct competition between the two subsectors will increase. Already, Intel and Motorola have established a TOP 10 position in the $8,000-$32,000 range of systems, in part through the sale of development systems.

TRENDS IN THE NEW $124,000-$512,000 PRICE CLASS: "Mini's" Solidify Position

Contrary to its origins in birthing the minicomputer industry with the minimum-performance-and-cost pdp-8 and pdp-11 product lines, DEC has led its competitors in driving up the product line with ever larger systems and breakthrough performance. With the pdp 11/45, the pdp 11/70, and the VAX 11/780, DEC extended the product line upward with a dramatic performance increase accompanied by rising system prices. The company in effect created a new price class and new performance range for minicomputers which contributed to roughly half of DEC's shipments, and slightly over a third of all shipments in the survey. Yet, these systems accounted for only 13% of third-party sites, and 17% of end-user sites. The fundamental strategic objectives related to the top-end range of the product line represent near opposites to the minimum performance/cost strategies originally employed by the minicomputer manufacturers to develop their markets.

Yet, the trend was a natural evolution driven by strong factors both within the companies and among customers. As discussed above, a move down the product line presents designers with a dilemma of stripping out function. Further, the concentration on improving performance serves to develop existing markets to the extent that the user can effectively capitalize on the technology to register productivity gains with personnel. With the bulk of data processing expenditures now concentrated in software and programming support, cheaper computers may have little real impact on overall user budgets. When the end-user can leverage the technology to improve the productivity of personnel, the established customers register the greatest gains in efficiency overall. The objective requires more function, not less.

By way of contrast, the minimum performance/cost strategies of the "micro" vendors provide the means to cost-justify new applications with lower software and programming overhead costs. The approach generally requires a new organization and new applications in terms of the user community, and serves as a vehicle to enter new markets on the part of the suppliers.

STRATEGIC PRICING OPTIONS IN TECHNOLOGY

Digital systems have offered customers dramatic improvements in capability at continuously lower prices. The Apple III and the IBM personal computers currently provide processing capability at least equivalent to the IBM System/360, Model 30, which was the most popular commercial data processing system of the late 1960s. Yet personal computers cost less than ten thousand dollars, whereas the cost of the IBM Model 30 ran into the hundreds of thousands of dollars fifteen years ago.

The industry discusses the trend in terms of "price/performance." That is, the technology provides the means to drive prices down while holding performance constant or, as another option, to push perfomance and functionality up while maintaining prices constant. The choice is to minimize price or to maximize performance. Generally, a rate of 20% has been a reasonable price/performance trend for many of the products in the sector. At 20% annually, the price of today's $50,000 system can be cut to a tenth, or to $5,000, while maintaining constant performance in only ten years. Following the first alternative of minimizing price, the "micro" vendors have benefited from the underlying technology curve and created a new product class in terms of price which now incorporates on individual semiconductor chips the processing power of full computer systems. The second alternative of maximizing performance suggests that functional capability can be increased ten times in ten years while holding the price constant. This trend represents the major thrust for the established minicomputer suppliers since the introduction of the DEC VAX series in 1977.

The price/performance carve severely tests distribution and marketing. In terms of minimizing price, the salesman trying to earn a living pushing these products may have to sell ten times as many units, possibly by covering ten times as many clients. Because this is generally not an attractive option to existing marketing forces, distribution channels shift with the technology price/performance curve. The more common and attractive alternative for marketing is to master the technologial challenge of a product that becomes ten times as complex in only one decade and to continue to sell roughly the same number of units to a similar number of clients.

Competitive strategic alternatives between the two extremes discussed above determine the middle ground. In midrange systems, DEC and Hewlett-Packard effectively apply technology to increase performance while maintaining roughly constant prices. Admittedly, DEC has been slow to develop product alternatives in the new minimum-price range of $1,000-$8,000 for systems and has applied most of its energy to capitalizing on the potential high-volume business available with large systems, where the company accounts for over half of shipment and site totals in the survey. Nonetheless, DEC has applied technology aggressively in midrange systems, as well, to maintain its position of leadership. In both the $8,000-$32,000 and $32,000-$64,000 prices classes, DEC enjoys the top position with an average 40% share of shipments and sites for both third-party and end-user sectors. In the $64,000-$124,000 range, however, DEC's share drops to an average 19%. Here, Hewlett-Packard's concentration with an average 31% share edges out DEC as the TOP 10 supplier in terms of shipment value and site penetration in the latest survey.

Overall, three strategic alternatives are readily apparent in the survey:

1. The "micro" vendors have established a major presence in the market for systems by emphasizing minimum performance/cost products which have created new users and applications.

2. Traditional "mini" suppliers have applied technology to improve performance while maintaining constant prices in midrange systems.

3. DEC, in particular, has primarily relied on breakthrough performance at higher prices to fully capitalize on its established market position.

In this context, average systems prices represent one of the most important strategic variables of planning to those in the industry. Further, the approaches used in integrating technological advances into the existing product line represents a critical determinant in the success of individual companies over the longer term. Competitive strategies must be defined as to which price class a company will concentrate on at any point in time, and as to the shifts that can constantly be anticipated from technologically driven gains in price/performance. Companies must prepare to migrate existing levels of performance down and to absorb greater performance into the current product offering.

THE COLLECTION OF MINI/MICRO MARKETS IN TERMS OF APPLICATIONS

Small computers are general-purpose machines which can be programmed for a wide variety of tasks in nearly all functions of every industry. The systems have been inexpensive enough to be purchased by individuals and departments without requiring approval or participation by a monolithic, centralized data processing department. By comparison, the traditional commercial mainframe business dominated by IBM has tended to concentrate on relatively narrow, financial applications managed by a small and dominant central staff.

Vendors confront each other in specific submarkets by application and by industry. As the mini/micro sector continues to experience rapid growth, and suffuses throughout the economy, niche-competitive strategies will become increasingly important as a determinant of success.

The BACHE/MINI-MICRO ANNUAL MARKET SURVEY collected data on thirteen different market sectors. As summarized in Chart 2 (page 14), opposite, one company stands in a class by itself in the analysis -- Digital Equipment. Across all thirteen applications sectors, DEC was first in eleven in terms of breadth of penetration on the basis of sites and first in terms of shipment value in ten. Below DEC, four manufacturers vie for the second-place position in a close race: Hewlett-Packard, Data General, IBM, and Intel. (Apple places surprisingly high in site penetration, but does not represent as significant a factor in terms of the value of shipments.) Each of the second tier tends to be a leader in a few areas, offset by weakness in other selected applications sectors. Below the second tier, individual companies penetrate selected niches.

DEC's leadership across the full applications spectrum of the market testifies to the strength of the organization. The proliferation of products, the array of customers with applications requirements most notable for their diversity, an industry environment of accelerated innovation with overlapping and conflicting

CHART 2

DEC'S LEADERSHIP EXTENDS ACROSS ALL APPLICATIONS

	Top 10 Ranking*			
	By Sites		By Shipment Value	
	No. of Sectors in Top 10	Average Rank	No. of Sectors in Top 10	Average Rank
Digital Equipment	13	1.2	13	1.2
Hewlett-Packard	13	4.5	9	3.3
Data General	12	4.3	13	4.1
IBM	12	5.3	10	4.9
Apple	11	4.2	2	5.0
Intel	10	4.0	10	6.2
Tandy	9	6.8	2	5.5
Texas Instruments	7	8.6	4	6.8
Motorola	7	6.9	9	6.8
Perkin-Elmer	6	7.9	7	7.0
Tektronix	4	7.3	NA	NA
Prime	3	8.7	5	6.2
Modular Computer	2	7.3	3	8.7
Commodore	2	6.0	1	9.0
Rockwell	2	8.0	NA	NA
Wang Labs	2	5.0	2	5.5
Honeywell	1	7.0	3	5.0
Sperry	1	8.5	4	6.5
Systems Engineer.	1	5.5	6	5.8
Tandem	1	8.0	5	8.2
Zilog	1	10.5	2	3.0

* Includes only those companies making the TOP 10 ranking overall. The average rank included only those applications where a company makes the TOP 10 list. Additional detail is provided in Section II.

distribution channels -- these factors exert extraordianary internal pressures on DEC's management. In response to these challenges, DEC has evolved a partial matrix management. The approach formally acknowledges the interdependencies of individuals at a low level in the corporation. Authority is pushed down so that people define and recognize joint responsibilities and requirements. Theoretically, operational issues and conflicts are resolved at a low level in the organization instead of traveling up the line of authority which is the case in a more functional organization. In short, DEC has developed an organization which is at least as distinctive as its products.

APPLICATIONS SECTORS

IBM and Personal Computer Manufacturers Gain Share in Business DP

Three primary applications -- business data processing, engineering/scientific data processing, and industrial automation -- account for roughly half of the value of shipments and sites in the survey in 1981, while ten other applications account for the remainder of the market.

Business DP (data processing) now represents the largest market for small computer systems. IBM's traditional dominance of commercial markets shows up fully in its 35% share of the shipments to end-users -- the only applications area where IBM beat DEC (13% of shipments). With third-party marketing organizations, however, IBM came in fourth at 8% of shipments compared with DEC's placing first with 24%, Data General second with 12%, and Hewlett-Packard third with 9%. These results confirm the high strategic priority that IBM of necessity will accord to marketing efforts through third-party organizations; this has been a traditional competitive weakness in business DP applications for small computer systems to date. Indeed, IBM signed contracts with eight master distributors for the Series/1 with reported discounts of up to 40%. Under the more general contract terms with independent marketing organizations, the maximum discount has ranged up to 30%. More recently, some trade sources contend that additional products will be included in the program such as the 4300, 8100, Displaywriter, and personal computer. Such a move would represent an aggressive increase in IBM's commitment to third-party marketing organizations. Even so, IBM temporarily suspended further expansion of the number of master distributors, in part to evaluate the effects of the program already in place. In any case, IBM will clearly be expanding its effort with third-party marketing channels, although the timing of specific moves is uncertain.

Of particular note, IBM's share of sites in business DP for third parties rose from 12% in 1980 to 20% in 1981; IBM's share of sites in the end-user sector rose from 25% to 31%. The traditional "mini" manufacturers suffered in general a substantial loss of position; the personal computer manufacturers captured substantial market share.

Engineering/Scientific DP: DEC Maintains Dominance

DEC captured over 50% of the value of shipments in engineering/scientific applications -- riding on the continuing success of the high-end VAX product line. Indeed, DEC's success in pushing price/performance has historically found its warmest reception in the technically oriented markets. Given its leadership in the complementary instrumentation markets, Hewlett-Packard records a surprisingly low 6% of shipments with third-party organizations and 11% of end-users.

In terms of sites, gains for Intel and Motorola represent a natural expansion whereby the products should find a ready reception among technically sophisticated users. The strong progress for Apple and Tandy, however, represents a more surprising trend.

TABLE III

MARKET SHARE SHIFTS FOR MAJOR APPLICATIONS*

	Business DP		Engineering/Scientific		Industrial Automation	
	1981 Share*	Gain or Loss	1981 Share*	Gain or Loss	1981 Share*	Gain or Loss
"Mini's"						
Digital Equipment	21.5%	-5.5%	44.0%	-0.5%	31.0%	-9.0%
Data General	9.0	-5.5	9.0	-4.5	9.0	-3.0
Hewlett-Packard	10.0	-2.0	15.5	-3.5	8.0	-6.0
Subtotal	40.5%	-13.0%	68.5%	-8.5%	48.0%	-18.0%
IBM	25.5%	7.0%	4.5%	3.0%**	NA	NA
"Micro's"						
Intel	NA	NA	7.5%	6.0%**	20.0%	7.0%
Motorola	NA	NA	2.0	2.0**	6.0	4.0
Subtotal	NA	NA	9.5%	8.0%	26.0	11.0%
Apple	10.5%	7.0%	6.5%	3.5%	11.0%	5.0%
Tandy	9.0	4.5	3.5	3.5**	NA	NA
Subtotal	19.5%	11.5%	10.0%	7.0%	11.0%	5.0%

* Average share of TOP 10 sites for third-party and end-user sectors except industrial automation, which is based solely on third-party activity.

** IBM, Motorola, and Tandy made the TOP 10 in engineering for the first time in 1981; Intel appeared with end-users in engineering/scientific in 1981.

Industrial Automation: Intel's Biggest Success

Last year's survey credited Intel with a dramatic increase in share: from 2% to 7% in terms of units for industrial automation. Intel introduced its first single-board microcomputer in 1976, the iSBC/80/10. The continuing popularity of the 8-bit "micro" products relates to their simplicity and economy in use. It appears difficult, however, to differentiate clearly between 8080/8086 chip set deliveries and board shipments, since users may lump their purchases under the processor designation, even when they are purchasing the processor chip set mounted on a board. In any case, Intel's momentum lifted its share of sites for the TOP 10 manufacturers for third-party organizations from 13% in 1980 to 20% in 1981. Still well behind DEC's 29% share of shipments, Intel nonetheless placed second at 16% -- comfortably ahead of Data General and Hewlett-Packard, both having come in with 11%.

MARKET SHIFTS FOR MAJOR APPLICATIONS

Table III summarizes the shift in market share by applications sector in terms of sites from 1980 to 1981; detailed tables are included in Section II. The trends by applications sector generally follow the totals for the survey: Tthe traditional minicomputer manufacturers lost significant market share in terms of site data, whereas IBM and the "micro" suppliers have gained. Again, a comparison in terms of estimated shipment value is not available since data on the value of system shipments was not collected for 1980. Any shift in terms of value would probably not be significant given the low average systems value of the "micro" products. IBM's gain in market share in terms of sites, however, would be fully reflected in terms of the value of shipments as well.

82-430

"MINI'S" SECURE; INTEL AND MOTOROLA CRACK THE TOP 10 IN SHIPMENTS

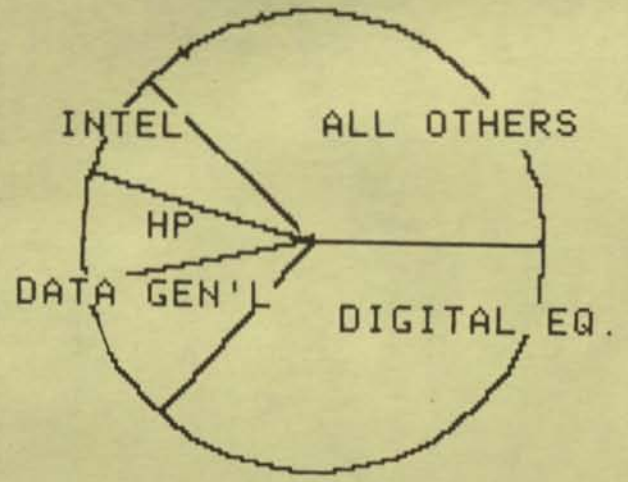

SHARE OF 1981 SHIPMENT VALUE

TABLE 1A

TOP 10 BY 1981 SHIPMENT VALUE: THIRD PARTY

Rank	Supplier	$ Mil.	% Total
1	Digital Equipment	582	37
2	Data General	155	10
3	Hewlett-Packard	119	8
4	Intel	112	7
5	Motorola	52	3
6	IBM	49	3
7	Systems Engineer.	48	3
8.5	Texas Instruments	41	3
8.5	Perkin-Elmer	41	3
10	Honeywell	38	2
	Subtotal	1,237	79
	Other	319	21
	Total	1,556	100

TRADITIONAL "MINI'S" MAINTAIN SHARE OF SHIPMENTS; IBM GAINS

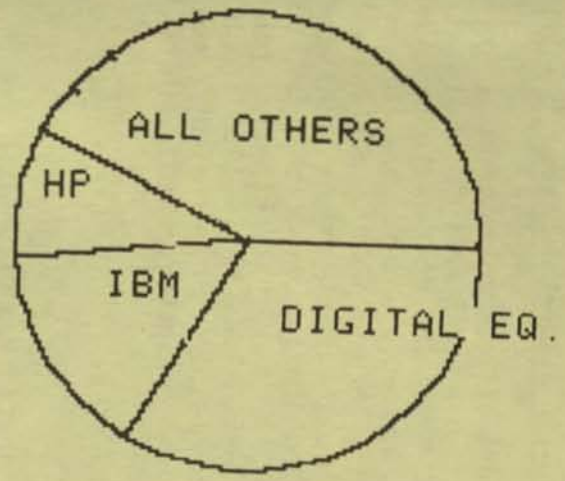

SHARE OF 1981 SHIPMENT VALUE

TABLE 1B

TOP 10 BY 1981 SHIPMENT VALUE: END-USER

Rank	Supplier	$ Mil.	% Total
1	Digital Equipment	388	34
2	IBM	174	15
3	Hewlett-Packard	108	9
4	Tandem	47	4
5	Data General	41	4
6	Honeywell	40	4
7	Prime	35	3
8.5	Wang Labs	22	2
8.5	Perkin-Elmer	22	2
10	Apple	21	2
	Subtotal	898	79
	Other	251	21
	Total	1,149	10

TOP 10 "PRIMARY" SUPPLIERS: THIRD PARTY

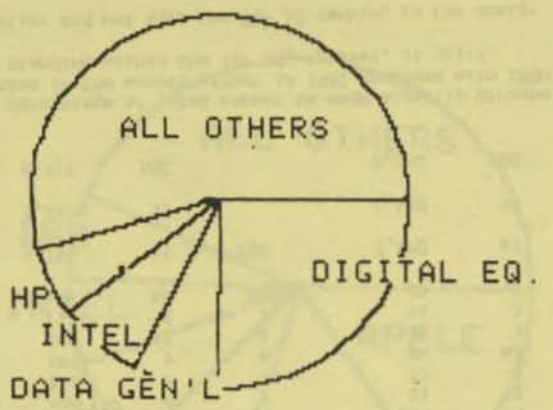

% OF 1981 MENTIONS

TABLE 2A

TOP 10 "PRIMARY" SUPPLIERS: THIRD PARTY

1981 Rank	Supplier	1981 No. Mentions	1981 % Total	1981 % TOP 10	1980 No. Mentions	1980 % Total	1980 % TOP 10
1	Digital Equipment	766	24	38	783	34	46
2	Intel	229	7	11	291	13	17
3	Data General	219	7	11	142	6	8
4	Hewlett-Packard	180	6	9	199	9	12
5.5	Apple	156	5	8	31	1	2
5.5	IBM	155	5	8	43	2	3
7	Tandy	110	4	5	NA	NA	NA
8	Texas Instruments	91	3	4	96	4	6
9	Motorola	72	2	4	38	2	2
10	Perkin-Elmer	54	2	3	54	2	3
NA	Systems Engineer.	NA	NA	NA	41	2	2
	Subtotal	2,032	65	100	1,718	75	100
	Other	1,109	35		560	25	
	Total	3,141	100		2,278	100	
	Unspecified	914			207		

Note: Comparisons for the percentage of total cannot be made directly because of moderate differences in the survey effort in 1981 compared with 1980. A comparison of the standing within the TOP 10, however, is valid.

NA: Not applicable; supplier did not make the TOP 10 ranking in the designated year.

TOP 10 "PRIMARY" SUPPLIERS: END-USER

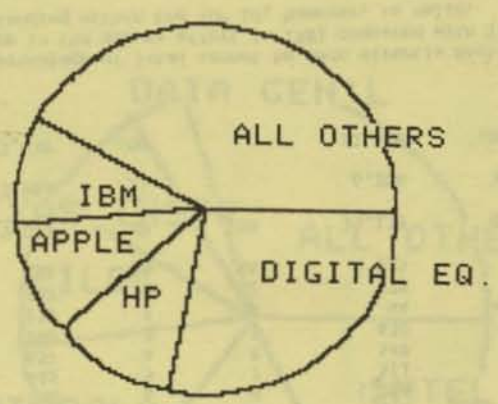

% OF 1981 MENTIONS

TABLE 2B

TOP 10 "PRIMARY" SUPPLIERS: END-USER

1981 Rank	Supplier	1981 No. Mentions	1981 % Total	1981 % TOP 10	1980 No. Mentions	1980 % Total	1980 % TOP 10
1	Digital Equipment	1,458	27	37	2,712	34	46
2	Hewlett-Packard	558	10	14	858	11	15
3	Apple	513	10	13	245	3	4
4	IBM	498	9	13	656	8	11
5	Tandy	249	5	6	161	2	3
6	Data General	207	4	5	441	6	8
7	Intel	143	3	4	280	4	5
8	Wang Labs	125	2	3	190	2	3
9.5	Texas Instruments	88	2	2	NA	NA	NA
9.5	Prime	87	2	2	156	2	3
NA	Datapoint	NA	NA	NA	163	2	3
	Subtotal	3,926	74	100	5,862	73	100
	Other	1,393	26		2,132	27	
	Total	5,319	100		7,994	100	
	Unspecified	1,416			1,720		

Note: Comparisons for the percentage of total cannot be made directly because of moderate differences in the survey effort in 1981 compared with 1980. A comparison of the standing within the TOP 10, however, is valid.

NA: Not applicable; supplier did not make the TOP 10 ranking in the designated year.

PERSONAL COMPUTER MANUFACTURERS AND IBM BROADEN THEIR CUSTOMER BASE AT THE EXPENSE OF TRADITIONAL "MINI'S"

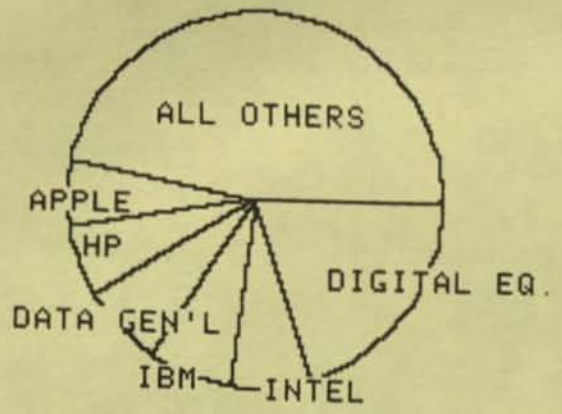

% OF TOTAL SITES ; THIRD PARTY

TABLE 3A

TOP 10 BY SITES: THIRD PARTY

1981 Rank	Supplier	1981 No. Mentions	1981 % Total	1981 % TOP 10	1980 No. Mentions	1980 % Total	1980 % TOP 10
1	Digital Equipment	1,331	20	31	1,005	25	40
2	Intel	486	7	11	237	6	10
3	IBM	453	7	11	129	3	5
4	Data General	429	7	10	463	11	19
5	Hewlett-Packard	417	6	10	199	5	8
6	Apple	388	6	9	81	2	3
7	Tandy	276	4	7	77	2	3
8	Motorola	187	3	4	NA	NA	NA
9	Texas Instruments	166	3	4	136	3	5
10	Perkin-Elmer	113	2	3	94	2	4
NA	Systems Engineer.	NA	NA	NA	69	2	3
	Subtotal	4,246	65	100	2,490	62	100
	Other	2,267	35		1,550	38	
	Total	6,513	100		4,040	100	

Note: Comparisons for the percentage of total cannot be made directly because of moderate differences in the survey effort in 1981 compared with 1980. A comparison of the standing within the TOP 10, however, is valid.

NA: Not applicable; supplier did not make the TOP 10 ranking in the designated year.

APPLE AND TANDY CRACK THE TOP 10 WITH BROAD SITE PENETRATION

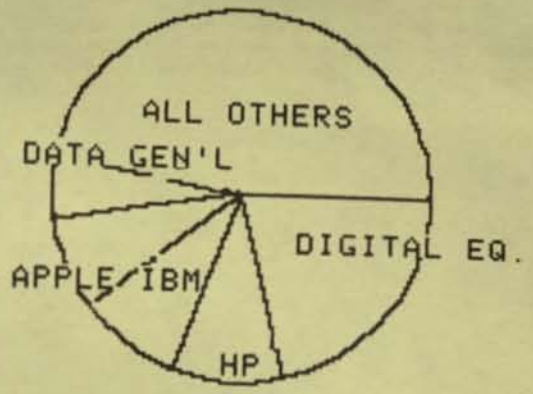

% OF TOTAL SITES; END-USER

TABLE 3B

TOP 10 BY SITES: END-USER

1981 Rank	Supplier	1981 No. Mentions	1981 % Total	1981 % TOP 10	1980 No. Mentions	1980 % Total	1980 % TOP 10
1	Digital Equipment	2,908	22	32	3,967	23	36
2	Hewlett-Packard	1,260	10	14	1,786	10	16
3	IBM	1,151	9	13	1,248	7	11
4	Apple	1,129	9	12	651	4	6
5	Data General	746	6	8	1,286	7	12
6	Tandy	665	5	7	511	3	5
7	Intel	457	4	5	549	3	5
8	Wang Labs	280	2	3	430	2	4
9	Prime	233	2	3	NA	NA	NA
10	Texas Instruments	222	2	2	375	2	3
NA	Perkin-Elmer	NA	NA	NA	335	2	3
	Subtotal	9,051	69	100	11,138	64	100
	Other	3,988	31		6,206	36	
	Total	13,039	100		17,344	100	

Note: Comparisons for the percentage of total cannot be made directly because of moderate differences in the survey effort in 1981 compared with 1980. A comparison of the standing within the TOP 10, however, is valid.

NA: Not applicable; supplier did not make the TOP 10 ranking in the designated year.

APPLE EDGES OUT DEC IN UNITS

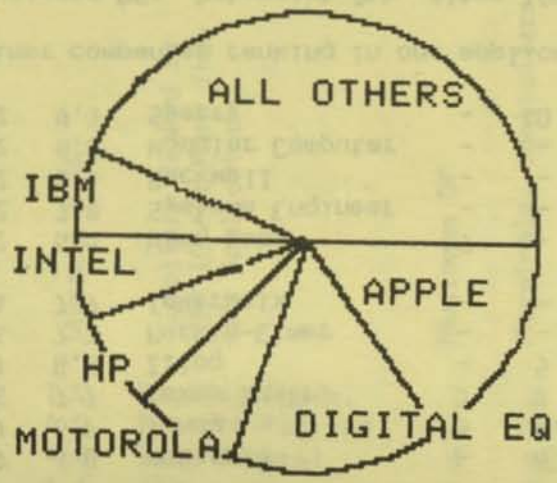

% OF TOTAL UNITS; END-USER

TABLE 4B

TOP 10 BY UNITS ACQUIRED: END-USER

1981 Rank	Supplier	1981 No. Units	1981 % Total	1981 % TOP 10
1	Apple	5,391	17	23
2	Digital Equipment	5,351	17	22
3	Motorola	2,538	8	11
4	Hewlett-Packard	2,347	7	10
5	Intel	2,321	7	10
6	IBM	2,162	7	9
7	Tandy	1,657	5	7
8	Texas Instruments	809	3	3
9	Wang Labs	656	2	3
10	Data General	651	2	3
	Subtotal	23,883	75	100
	Other	7,848	25	
	Total	31,731	100	

INTEL DISPLACES DEC

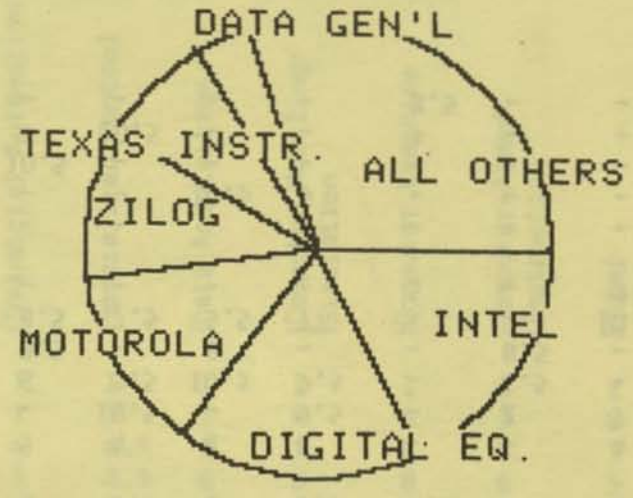

% OF TOTAL UNITS; THIRD PARTY

TABLE 4A

TOP 10 BY UNITS ACQUIRED: THIRD PARTY

1981 Rank	Supplier	1981 No. Units	1981 % Total	1981 % TOP 10
1	Intel	17,985	18	22
2	Digital Equipment	17,021	17	21
3	Motorola	12,261	12	15
4	Zilog	11,240	11	14
5	Texas Instruments	8,075	8	10
6	Data General	3,934	4	5
7	Apple	3,213	3	4
8	Hewlett-Packard	2,858	3	4
9	Tandy	2,674	3	3
10	Ohio Scientific	926	1	1
	Subtotal	80,187	78	100
	Other	22,446	22	
	Total	102,633	100	

TABLE 5A
TOP 10 RANKING COMPANIES

BY APPLICATION, THIRD-PARTY MARKETERS, SITES

"DEC's Leadership Extends Across Market Spectrum of Applications; Second Tier Now Includes Apple and Intel"

No. of Sectors	Average Rank	Suppliers Ranked in TOP 10 for More Than One Application	Business DP	Terminals, Controllers	Communications	WP/Office Automation	Engineering/ Scientific DP	Laboratory Control	Industrial Automation	Specialized Control	Data Acquisition	Computer Modeling/ Simulation	Computer Graphics	Computer-Aided Instruction	Other
13	1.2	Digital Equip.	1	1	1	1	1	1	1	1	1	1	1	4	1
12	3.8	Intel	-	2	2	8	3	3	2	2	3	3.5	7	8.5	2
11	3.6	Data General	3	3	3	-	4	4	3	3	4	6	2	-	5
11	4.7	Apple	4	7.5	-	2.5	5	5	-	7.5	8	2	3.5	1	6
11	5.7	IBM	2	4	4	5.5	6	10.5	-	9	5.5	-	10	6.5	3
10	4.2	Hewlett-Packard	6	-	5.5	-	2	2	4	5.5	2	3.5	3.5	-	8
9	6.0	Motorola	-	6	5.5	-	7	6.5	5	4	8	8	-	-	4
9	6.7	Tandy	5	7.5	-	2.5	8	8.5	-	10.5	-	9.5	-	2	7
8	7.7	Texas Instr.	7	9	7.5	-	-	-	6	5.5	10.5	9.5	-	6.5	-
7	8.6	Zilog	-	5	9.5	-	10	10.5	8.5	-	8	-	-	8.5	-
6	7.7	Perkin-Elmer	-	-	7.5	-	-	-	8.5	7.5	5.5	7	-	-	10
4	7.9	Tektronix	-	-	-	-	9	8.5	-	-	-	-	5	-	9
2	6.5	Wang Labs	9	-	-	4	-	-	-	-	-	-	-	-	-
2	7.8	Systems Engineer	-	-	-	-	-	-	-	-	10.5	5	-	-	-
2	8.5	Rockwell	-	-	-	-	-	6.5	10.5	-	-	-	-	-	-
2	8.8	Modular Computer	-	-	-	-	-	-	7	10.5	-	-	-	-	-
2	9.3	Sperry	-	10	-	-	-	-	-	-	-	-	8.5	-	-

Other companies ranking in one applications area include:

Business DP: Datapoint 8th, Altos 10th.
Communications: Tandem 9.5.
WP/Office Automation: Xerox 5.5, North Star 7th, Lanier 9th, Zenith 10th.
Industrial Automation: Cromenco 10.5.
Computer Graphics: Computervision 6th, Intel. Sy[illegible] 8.5.
Computer-Aided Instruction: Commodore 3rd, At[illegible]n.

TAB
TOP 10 RANKIN[illegible]OMPANIES

BY APPLICATION, END-USERS, SITES

No. of Sectors	Average Rank	Suppliers Ranked in TOP 10 for More Than One Application	Business DP	Terminals, Controllers	Communications	WP/Office Automation	Engineering/ Scientific DP	Laboratory Control	Industrial Automation	Specialized Control	Data Acquisition	Computer Modeling/ Simulation	Computer Graphics	Computer-Aided Instruction	Other
13	1.2	Digital Equip.	2	1	1	1	1	1	1	1	1	1	1	2	1
13	4.3	Hewlett-Packard	3	9	4	8	2	2	2	3	2	4	4	5	8
12	4.6	Data General	6	4.5	3	-	3	3	4	4.5	3	3	6	8	7
12	4.6	IBM	1	2	2	4	6	-	5	4.5	6	6	8	6	5
12	4.8	Apple	4	7.5	8.5	5	4	4	-	9	5	2	5	1	2
10	5.1	Intel	-	6	6	-	7	5	3	2	4	8	-	7	3
9	6.0	Tandy	5	-	8.5	3	5	10.5	-	10	-	5	-	3	4
6	8.8	Texas Instr.	-	-	-	-	-	10.5	8	7	8	-	-	9	10
5	7.3	Motorola	-	-	-	-	-	7.5	7	6	10	-	-	-	6
5	8.3	Perkin-Elmer	-	-	5	-	9.5	-	9	-	9	9	-	-	-
4	7.3	Tektronix	-	-	-	-	8	-	-	-	-	10	2	-	9
4	8.1	Honeywell	10	4.5	10	-	-	-	-	8	-	-	-	-	-
3	8.2	Datapoint	8	7.5	-	9	-	-	-	-	-	-	-	-	-
3	8.5	Prime	9	-	-	-	9.5	-	-	-	-	-	7	-	-
2	4.5	Wang Labs	7	-	-	2	-	-	-	-	-	-	-	-	-
2	5.8	Commodore	-	-	-	-	-	7.5	-	-	-	-	-	4	-
2	6.5	Modular Computer	-	-	-	-	-	-	6	-	7	-	-	-	-
2	8.0	Rockwell	-	-	-	-	-	6	10	-	-	-	-	-	-
2	9.8	Sperry	-	10	-	-	-	-	-	-	-	-	9.5	-	-
2	9.8	Terak	-	-	-	-	-	-	-	-	-	-	9.5	10	-

Other companies ranking in one applications area include:

Terminals, Controllers: Northern Telecom 3rd.
Communications: Tandem 7th.
WP/Office Automation: Xerox 6th, CPT 7th, AM Jacquard 10.5, Lanier 10.5.
Laboratory Control: Cromenco 9th.
Computer Modeling/Simulation: Systems Engineering 7th.
Computer Graphics: Computervision 3rd.

BUSINESS DP, INDUSTRIAL AUTOMATION AND ENGINEERING/SCIENTIFIC DP ARE THE LARGEST THIRD-PARTY MARKETERS

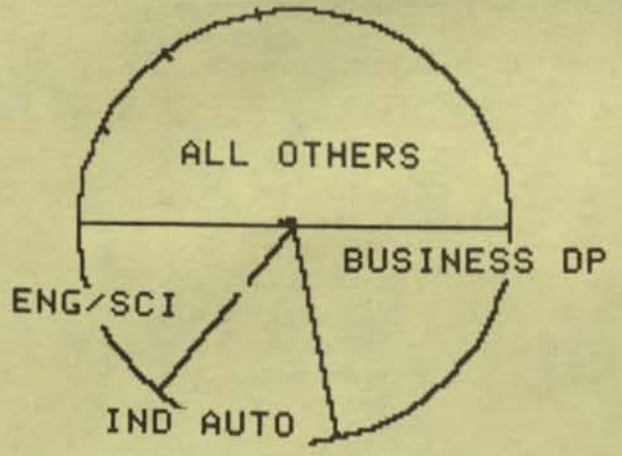

1981 SYSTEM VALUE SHARE BY APPLICATION FOR THIRD-PARTY MARKETERS

TABLE 6A

APPLICATION ANALYSIS: THIRD PARTY

Application	Third Party (% of Total) Sites* 1980	Sites* 1981	Units 1981	Value 1981
Business DP	23	30	16	21
Engineering/Scientific DP	18	17	9	14
Industrial Automation	12	7	16	14
Design, Drafting, Graphics (incl. CAD)	3	3	2	5
Data Acquisition	NA	4	6	8
Word Processing/Electronic Mail/Office Automation	4	8	5	4
Communication/Telephone Systems	5	5	7	8
Laboratory Control	3	3	5	4
Terminal & Peripheral Controllers	4	3	12	4
Computer-Aided Instruction	1	1	1	1
Computer Modeling/Simulation	NA	4	7	4
Specialized Control	4	4	8	7
Other	22	11	5	7
Subtotal	100	100	100	100
Total in Category	4,417	6,645	95,667	$1,489 (mil.)

* Sites include multiple responses for those purchasing either from multiple suppliers and/or purchasing for multiple applications from the same manufacturer.

BUSINESS DP AND SCIENTIFIC ENGINEERING DP ARE THE LARGEST END-USER MARKETS

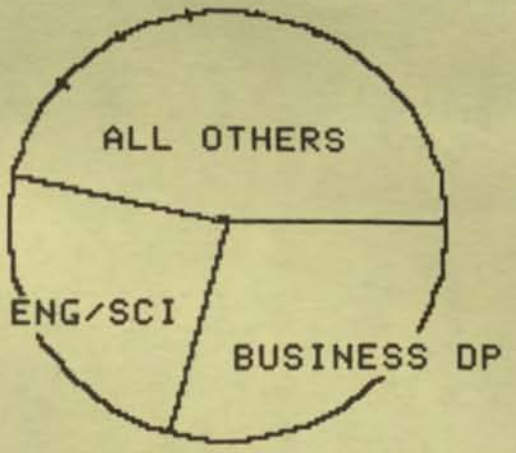

1981 SYSTEM VALUE SHARE BY APPLICATION FOR END-USERS

TABLE 6B

APPLICATION ANALYSIS: END-USER

Application	End-User (% of Total) Sites* 1980	Sites* 1981	Units 1981	Value 1981
Business DP	22	21	16	29
Engineering/Scientific DP	25	23	15	24
Industrial Automation	8	6	12	8
Design, Drafting, Graphics (incl. CAD)	3	4	3	6
Data Acquisition	NA	7	7	6
Word Processing/Electronic Mail/Office Automation	5	7	7	5
Communication/Telephone Systems	5	3	3	4
Laboratory Control	9	9	11	4
Terminal & Peripheral Controllers	5	3	4	3
Computer-Aided Instruction	4	7	14	3
Computer Modeling/Simulation	NA	3	2	3
Specialized Control	2	2	2	1
Other	11	5	4	4
Subtotal	100	100	100	100
Total in Category	18,747	13,867	30,311	$1,102 (mil.)

* Sites include multiple responses for those purchasing either from multiple suppliers and/or purchasing for multiple applications from the same manufacturer.

BUSINESS DP: END-USER

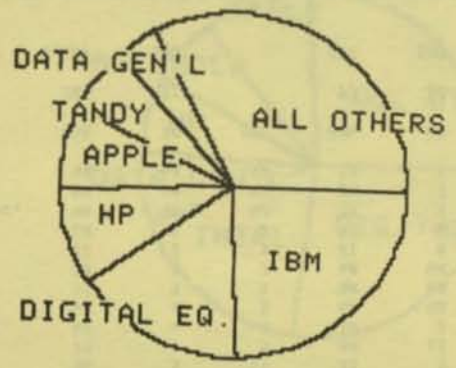

SHARE OF 1981 SITE MENTIONS

TABLE 7B

TOP SUPPLIERS IN BUSINESS DP: END-USER

Sites Rank	Supplier	Sites 1981 Number	Sites 1981 % Total	Sites 1981 % TOP 10	Sites 1980 Number	Sites 1980 % Total	Sites 1980 % TOP 10	$ Shipments 1981 $ Mil.	$ Shipments 1981 % Total
1	IBM	719	24	31	744	18	25	111	35
2	Digital Equipment	463	16	20	739	18	25	43	13
3	Hewlett-Packard	264	9	11	408	10	14	26	8
4	Apple	222	8	10	118	3	4	3	1
5	Tandy	177	6	8	124	3	4	1	
6	Data General	132	4	6	256	6	9	8	2
7.5	Wang Labs	90	3	4	158	4	5	8	2
7.5	Honeywell	90	3	4	101	2	3	32	10
9	Datapoint	89	3	4	185	4	6	3	1
10	Prime	79	3	3	NA	NA	NA	13	4
NA	Burroughs	NA	NA	NA	121	3	4	NA	NA
	Subtotal	2,325	79	100	2,954	70	100	135	43
	Other	623	21		1,258	30		182	57
	Total	2,948	100		4,212	100		317	100

Note: Comparisons for the percentage of total cannot be made directly because of moderate differences in the survey effort in 1981 compared with 1980. A comparison of the standing within the TOP 10, however, is valid.

NA: Not applicable; supplier did not make the TOP 10 ranking in the designated year.

BUSINESS DP: THIRD PARTY

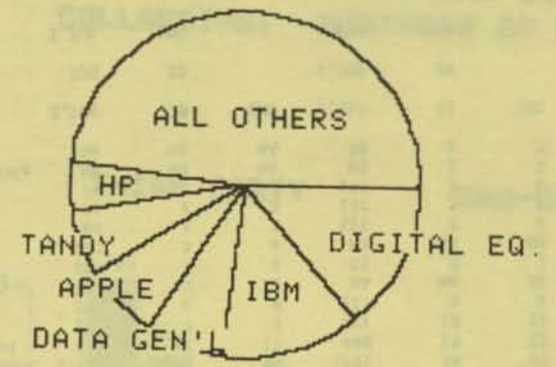

SHARE OF 1981 SITE MENTIONS

TABLE 7A

TOP SUPPLIERS IN BUSINESS DP: THIRD PARTY

Sites Rank	Supplier	Sites 1981 Number	Sites 1981 % Total	Sites 1981 % TOP 10	Sites 1980 Number	Sites 1980 % Total	Sites 1980 % TOP 10	$ Shipments 1981 $ Mil.	$ Shipments 1981 % Total
1	Digital Equipment	286	14	23	203	20	29	78	24
2	IBM	250	13	20	83	8	12	25	8
3	Data General	148	7	12	142	14	20	38	12
4	Apple	141	7	11	22	2	3	7	2
5	Tandy	120	6	10	32	3	5	4	
6	Hewlett-Packard	106	5	9	67	7	10	28	9
7	Texas Instruments	67	3	5	61	6	9	20	6
8	Wang Labs	47	2	4	32	3	5	5	2
9	Altos	43	2	3	NA	NA	NA	1	0
10	Prime	39	2	3	NA	NA	NA	8	3
NA	Datapoint	NA	NA	NA	33	3	5	NA	NA
NA	Microdata	NA	NA	NA	30	3	4	NA	NA
	Subtotal	1,247	63	100	705	68	100	215	67
	Other	732	37		325	32		105	33
	Total	1,979	100		1,030	100		320	100

Note: Comparisons for the percentage of total cannot be made directly because of moderate differences in the survey effort in 1981 compared with 1980. A comparison of the standing within the TOP 10, however, is valid.

NA: Not applicable; supplier did not make the TOP 10 ranking in the designated year.

ENGINEERING/SCIENTIFIC·THIRD PARTY

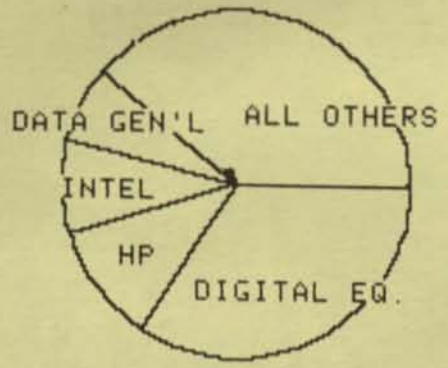

SHARE OF 1981 SITE MENTIONS

TABLE 8A

TOP SUPPLIERS IN ENGINEERING/SCIENTIFIC DP: THIRD PARTY

Sites Rank	Supplier	Sites 1981 Number	Sites 1981 % Total	Sites 1981 % TOP 10	Sites 1980 Number	Sites 1980 % Total	Sites 1980 % TOP 10	$ Shipments 1981 $ Mil.	$ Shipments 1981 % Total
1	Digital Equipment	373	34	44	278	36	41	108	53
2	Hewlett-Packard	119	11	14	114	15	17	12	6
3	Intel	94	9	11	23	3	3	11	5
4	Data General	79	7	9	88	11	13	26	13
5	Apple	46	4	5	14	2	2	1	0
6	IBM	40	4	5	NA	NA	NA	4	2
7	Motorola	37	3	4	NA	NA	NA	12	6
8	Tandy	28	3	3	NA	NA	NA	1	
9	Tektronix	22	2	3	61	8	9	1	0
10	Ziloq	18	2	2	32	4	5	2	1
NA	Systems Engineer.	NA	NA	NA	39	5	6	NA	NA
NA	Perkin-Elmer	NA	NA	NA	23	3	3	NA	NA
NA	Sperry	NA	NA	NA	12	2	2	NA	NA
	Subtotal	856	78	100	684	88	100	177	86
	Other	245	22		93	12		29	14
	Total	1,101	100		777	100		206	100

Note: Comparisons for the percentage of total cannot be made directly because of moderate differences in the survey effort in 1981 compared with 1980. A comparison of the standing within the TOP 10, however, is valid.

NA: Not applicable; supplier did not make the TOP 10 ranking in the designated year.

ENGINEERING/SCIENTIFIC·END-USER

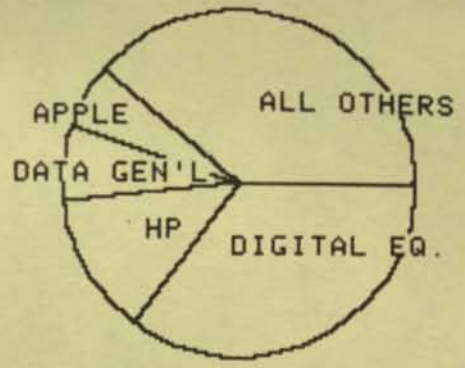

SHARE OF 1981 SITE MENTIONS

TABLE 8B

TOP SUPPLIERS IN ENGINEERING/SCIENTIFIC DP: END-USER

Sites Rank	Supplier	Sites 1981 Number	Sites 1981 % Total	Sites 1981 % TOP 10	Sites 1980 Number	Sites 1980 % Total	Sites 1980 % TOP 10	$ Shipments 1981 $ Mil.	$ Shipments 1981 % Total
1	Digital Equipment	1,115	35	44	1591	34	48	151	56
2	Hewlett-Packard	419	13	17	690	15	21	28	11
3	Data General	217	7	9	451	10	14	9	3
4	Apple	203	6	8	123	3	4	2	1
5	Tandy	111	3	4	NA	NA	NA	1	0
6	IBM	105	3	4	85	2	3	11	4
7	Intel	100	3	4	NA	NA	NA	2	1
8	Tektronix	84	3	3	121	3	4	1	0
9.5	Prime	76	2	3	122	3	4	7	3
9.5	Perkin-Elmer	76	2	3	134	3	4	4	1
NA	Texas Instruments	NA	NA	NA	88	2	3	NA	NA
NA	Wang Labs	NA	NA	NA	88	2	3	NA	NA
	Subtotal	2,506	78	100	3,317	71	100	216	81
	Other	700	22		1,359	29		52	19
	Total	3,206	100		4,676	100		268	100

Note: Comparisons for the percentage of total cannot be made directly because of moderate differences in the survey effort in 1981 compared with 1980. A comparison of the standing within the TOP 10, however, is valid.

NA: Not applicable; supplier did not make the TOP 10 ranking in the designated year.

INDUSTRIAL AUTOMATION: THIRD PARTY

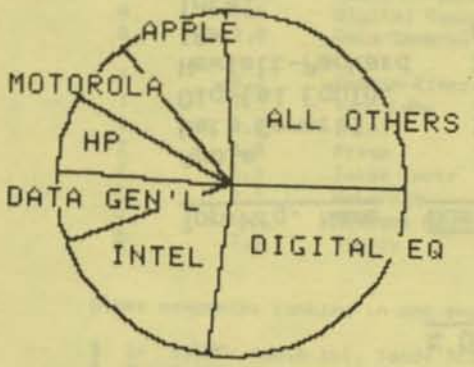

SHARE OF 1981 SITE MENTIONS

TABLE 9

TOP SUPPLIERS IN INDUSTRIAL AUTOMATION: THIRD PARTY

Sites Rank	Supplier	Sites 1981 Number	Sites 1981 % Total	Sites 1981 % TOP 10	Sites 1980 Number	Sites 1980 % Total	Sites 1980 % TOP 10	$ Shipments 1981 $ Mil.	$ Shipments 1981 % Total
1	Digital Equipment	126	27	31	163	32	40	59	29
2	Intel	81	17	20	54	11	13	33	16
3	Apple	46	10	11	23	5	6	0	0
4	Data General	35	7	9	47	9	12	21	11
5	Hewlett-Packard	33	7	8	59	12	14	22	11
6	Motorola	26	6	6	8	2	2	6	3
7	Texas Instruments	18	4	4	17	3	4	3	2
8	Modular Computer	13	3	3	13	3	3	9	5
9.5	Perkin-Elmer	12	3	3	12	2	3	7	3
9.5	Zilog	12	3	3	NA	NA	NA	2	1
NA	Comp. Automation	NA	NA	NA	11	2	3	NA	NA
	Subtotal	402	86	100	407	80	100	162	80
	Other	65	14		104	20		41	20
	Total	467	100		511	100		203	100

Note: Comparisons for the percentage of total cannot be made directly because of moderate differences in the survey effort in 1981 compared with 1980. A comparison of the standing within the TOP 10, however, is valid.

NA: Not applicable; supplier did not make the TOP 10 ranking in the designated year.

CHART 10

SCIENTIFIC/ENGINEERING SECTOR COLLAPSING; BUSINESS DP FIRM

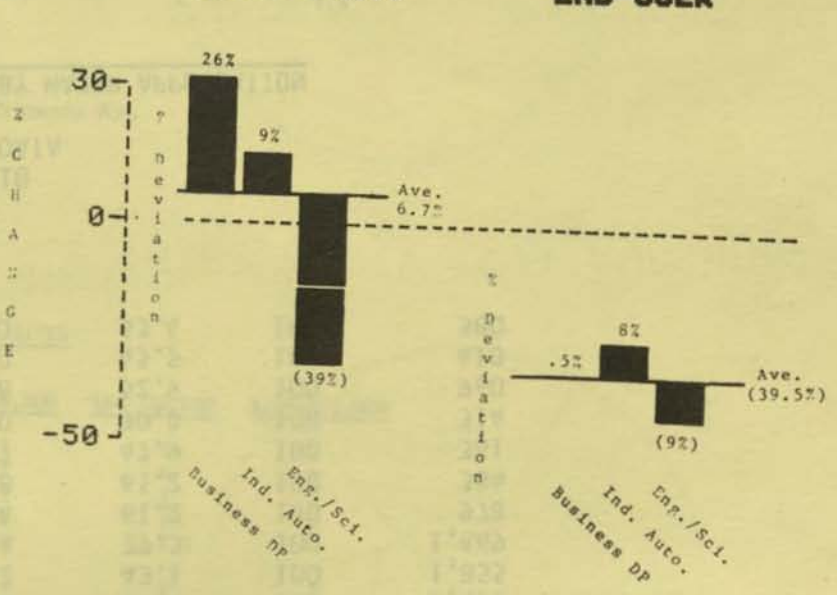

See Digital Systems Update, July 7, 1982; Digital Systems Update, "Outlook for Small Computer Systems in Scientific/Engineering Applications," July 21, 1982; and Small Computer Systems, August 24, 1982.

* Relative growth rates by major application; expected 1982 purchases versus 1981 purchases.

TABLE 11A
COMPANY DATA

% OF MENTIONED SITES BY MAJOR APPLICATION

Supplier	Percentages					Total No. Mentions
	Bus. DP	Sci./Eng.	Ind. Auto.	Other	Total	
Apple	24.8	17.3	1.2	56.7	100	1,488
Data General	22.2	24.3	7.4	46.1	100	1,290
Digital Equip.	14.3	28.6	7.2	49.9	100	5,477
Hewlett-Packard	20.9	29.8	6.2	43.1	100	1,855
IBM	59.3	9.0	2.4	29.3	100	1,669
Intel	1.9	20.5	16.4	61.2	100	978
Motorola	2.6	21.4	14.8	61.2	100	384
Perkin-Elmer	15.8	26.6	9.7	47.9	100	361
Prime	39.5	28.7	1.0	30.8	100	314
Tandy	31.6	14.7	0.8	52.9	100	960
Texas Instr.	27.6	15.9	11.0	45.5	100	410
Wang Labs	37.1	9.5	0.0	53.4	100	380

TABLE 11B
COMPANY DATA

% OF 1981 SHIPMENT VALUE BY MAJOR APPLICATION

Top Mfg. Name	Percentages					$ Total
	Bus. DP	Sci./Eng.	Ind. Auto.	Other	Total	
Apple	31.3	8.2	1.0	59.5	100	32,668
Data General	24.1	17.9	13.1	44.9	100	198,140
Digital Equip.	12.7	27.6	9.1	50.6	100	984,569
Hewlett-Packard	25.9	17.8	13.7	42.6	100	234,186
IBM	65.0	7.4	5.2	22.4	100	215,204
Intel	3.0	10.8	31.3	54.9	100	121,700
Motorola	0.8	20.9	17.0	61.3	100	64,588
Perkin-Elmer	15.7	11.2	16.8	56.3	100	62,624
Prime	45.3	16.2	1.1	37.4	100	53,869
Tandy	38.4	7.2	0.2	54.2	100	13,497
Texas Instr.	47.5	1.7	12.6	38.2	100	60,064
Wang Labs	44.7	3.4	0.0	51.9	100	28,893

TABLE 12A
TOP 10 RANKING COMPANIES

BY AVERAGE PRICE CLASS, THIRD-PARTY MARKETERS, SITES

Number of Classes	Average Rank	Supplier	$1-8,000	$8-32,000	$32-64,000	$64,128,000	$128-512,000
5	4.6	IBM	5	5	2	9	2
5	4.9	Hewlett-Packard	8	4	4	1	7.5
4	1.5	Digital Equip.	-	1	1	3	1
4	2.8	Data General	-	3	3	2	3
3	6.3	Perkin-Elmer	-	-	6	4	9
3	7.3	Wang Labs	-	8	8	6	-
2	2.0	Intel	2	2	-	-	-
2	4.8	Prime	-	-	-	5	4.5
2	5.5	Texas Instr.	-	6	5	-	-
2	5.5	Motorola	4	7	-	-	-
2	6.3	Systems Engineer.	-	-	-	8	4.5
2	7.3	Sperry	-	-	7	-	7.5

Other companies ranking in one average price class include:

$ 1- 8,000: Apple 1st, Tandy 3rd, Zilog 6th, North Star 7th, Cromenco 9th.
$ 8- 32,000: Tektronix 9th, Altos 10th.
$ 32- 64,000: Datapoint 9.5, Microdata 9.5.
$ 64-128,000: Modular Computer 7th, Honeywell 10th.
$128-516,000: Tandem 6th, Harris 10th.

TABLE 12B
TOP 10 RANKING COMPANIES

BY AVERAGE PRICE CLASS, END-USERS, SITES

Number of Classes	Average Rank	Supplier	$1-8,000	$8-32,000	$32-64,000	$64-128,000	$128-512,000
5	2.8	Hewlett-Packard	4	2	3	1	4
5	4.2	IBM	6	5	2	6	2
4	1.8	Digital Equip.	-	1	1	3	1
4	3.5	Data General	-	3	4	2	5
3	6.8	Wang Labs	-	7	6	7.5	-
2	4.0	Prime	-	-	-	5	3
2	4.5	Intel	3	6	-	-	-
2	4.5	Perkin-Elmer	-	-	5	4	-
2	6.8	Sperry	-	-	7.5	-	6
2	8.3	Honeywell	-	-	7.5	9	-
2	8.5	Motorola	7	10	-	-	-
2	8.8	Texas Instr.	-	8	9.5	-	-
2	9.2	Datapoint	-	9	9.5	-	-
2	10.0	Harris	-	-	-	10	10

Other companies ranking in one average price class include:

$ 1- 8,000: Apple 1st, Tandy 2nd, Commodore 5th, Cromenco 8th, Rockwell 9th, North Star 10th.
$ 8- 32,000: Tektronix 4th.
$ 64-128,000: Modular Computer 7.5.
$128-512,000: Tandem 7th, Systems Engineering 8th, Computervision 9th.

END-USER CONCENTRATION AT HIGH END

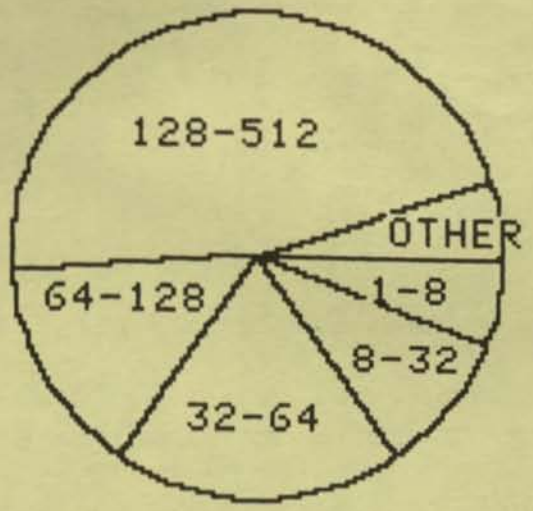

END-USER 1981 PURCHASES
BY VALUE

TABLE 13B

ANALYSIS OF BUSINESS BY PRICE CLASS: END-USER

Price Class	% of Total Sites*	Units	Value
$ 1- 8,000	25	49	6
$ 8- 32,000	28	22	10
$ 32- 64,000	19	14	19
$ 64-128,000	10	6	14
$128-512,000	17	8	47
Other	1	1	5
Total	100	100	100

* Sites include multiple entries for respondents purchasing from several suppliers and/or respondents purchasing computers in different price classes from the same supplier.

THIRD-PARTY BUSINESS SHIFTING TO LOW END

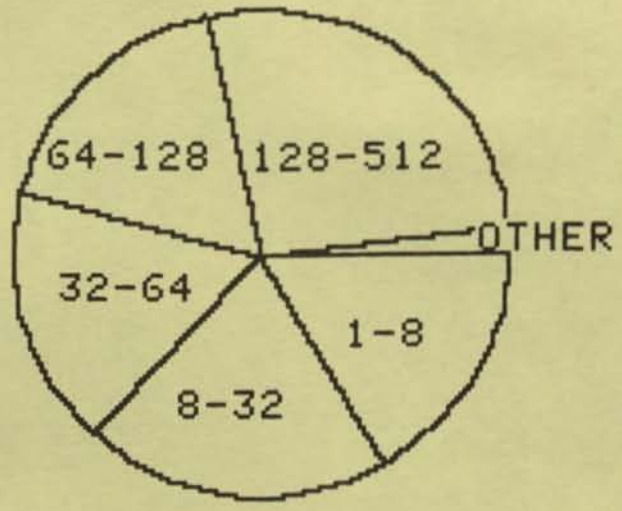

THIRD-PARTY 1981 PURCHASES
BY VALUE

TABLE 13A

ANALYSIS OF BUSINESS BY PRICE CLASS: THIRD PARTY

Price Class	% of Total Sites*	Units	Value
$ 1- 8,000	30	53	16
$ 8- 32,000	30	21	22
$ 32- 64,000	16	6	19
$ 64-128,000	9	3	15
$128-512,000	13	2	27
Other	1	16	1
Total	100	100	100

* Sites include multiple entries for respondents purchasing from several suppliers and/or respondents purchasing computers in different price classes from the same supplier.

TOP 10 IN $1-8,000 SYSTEMS: END-USER

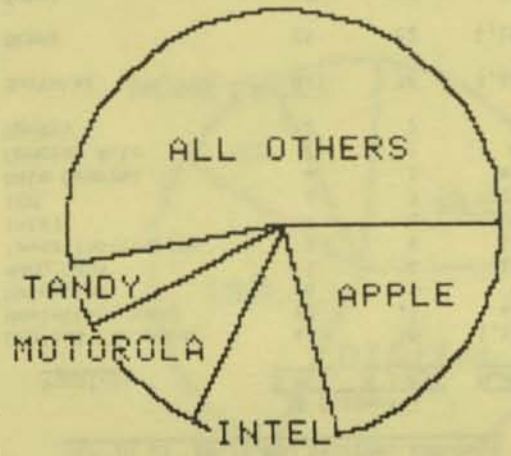

SHARE OF 1981 TOTAL SHIPMENT VALUE

TABLE 14B

TOP 10 IN $1-8,000 SYSTEMS: END-USER

Rank	Supplier	$ Shipments $ Mil.	$ Shipments % Total	Sites Number	Sites % Total
1	Apple	21	33	1,129	28
2	Intel	11	17	299	7
3	Motorola	10	15	146	4
4	Tandy	5	8	665	17
5	Hewlett-Packard	3	5	243	6
6	Zilog	2	3	45	1
8.5	RCA	1	2	16	0
8.5	North Star	1	1	98	2
8.5	Xerox	1	1	32	1
8.5	Commodore	1	1	191	5
	Subtotal	57	88	2,864	71
	Other	8	12	1,142	29
	Total	65	100	4,006	100

TOP 10 IN $1-8,000 SYSTEMS: THIRD PARTY

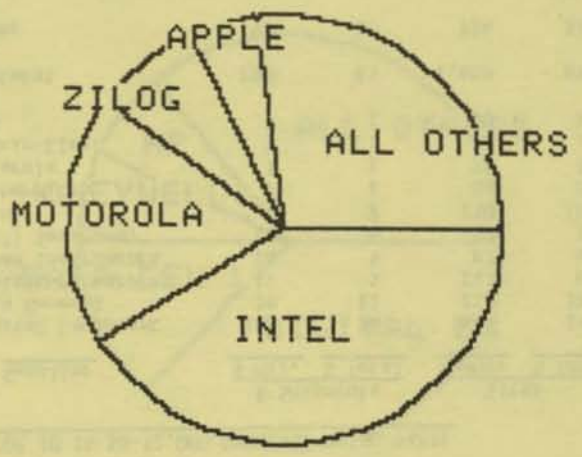

SHARE OF 1981 TOTAL SHIPMENT VALUE

TABLE 14A

TOP 10 IN $1-8,000 SYSTEMS: THIRD PARTY

Rank	Supplier	$ Shipments $ Mil.	$ Shipments % Total	Sites Number	Sites % Total
1	Intel	105	41	308	13
2	Motorola	49	19	123	5
3	Zilog	20	8	81	3
4	Apple	12	5	388	17
5	Tandy	8	3	276	12
6	RCA	7	3	7	0
7	Advanced Micro.	6	2	5	0
8	Ohio Scientific	5	2	27	1
9.5	North Star	3	1	78	3
9.5	Texas Instruments	3	1	37	2
	Subtotal	219	85	1,330	57
	Other	37	15	1,003	43
	Total	256	100	2,333	100

TOP 10 IN $8-32,000 SYSTEMS: END-USER

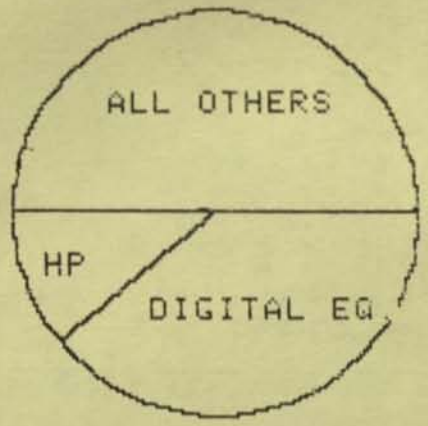

SHARE OF 1981 TOTAL SHIPMENT VALUE

TABLE 15B

TOP 10 IN $8-32,000 SYSTEMS: END-USER

Rank	Supplier	$ Shipments		Sites	
		$ Mil.	% Total	Number	% Total
1	Digital Equipment	44	39	1,528	35
2	Hewlett-Packard	12	11	416	10
3	Datapoint	6	5	92	2
4.5	Wang Labs	5	4	150	3
4.5	Texas Instruments	5	4	132	3
7	Intel	4	4	190	4
7	IBM	4	3	233	5
7	Data General	4	3	412	9
9.5	General Auto	2	2	46	1
9.5	Sperry	2	2	11	0
	Subtotal	87	78	3,210	74
	Other	25	22	1,155	26
	Total	112	100	4,365	100

TOP 10 IN $8-32,000 SYSTEMS: THIRD PARTY

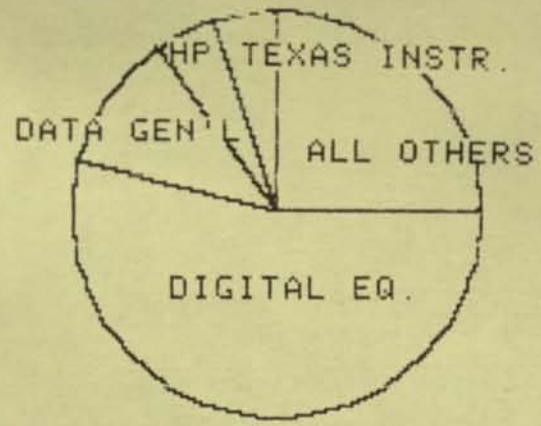

SHARE OF 1981 TOTAL SHIPMENT VALUE

TABLE 15A

TOP 10 IN $8-32,000 SYSTEMS: THIRD PARTY

Rank	Supplier	$ Shipments		Sites	
		$ Mil.	% Total	Number	% Total
1	Digital Equipment	181	54	684	29
2	Data General	38	11	223	10
3	Hewlett-Packard	17	5	143	6
4	Texas Instruments	16	5	87	4
5	Nat'l Semicond.	14	4	19	1
6	Intel	7	2	230	10
7	Alpha Micro	5	1	41	2
9	Motorola	3	1	76	3
9	Perkin-Elmer	3	1	26	1
9	IBM	3	1	90	4
	Subtotal	288	85	1,619	69
	Other	50	15	724	31
	Total	338	100	2,343	100

TOP 10 IN $32-64,000 SYSTEMS: END-USER

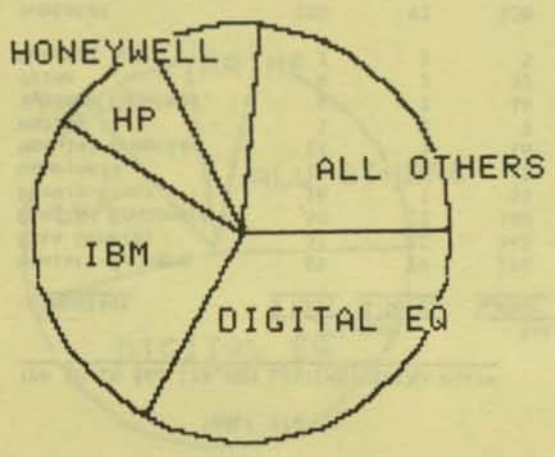

SHARE OF 1981 TOTAL SHIPMENT VALUE

TABLE 16B

TOP 10 IN $32-64,000 SYSTEMS: END-USER

Rank	Supplier	$ Shipments $ Mil.	$ Shipments % Total	Sites Number	Sites % Total
1	Digital Equipment	73	33	1,398	47
2	IBM	57	26	434	15
3	Hewlett-Packard	20	9	247	8
4	Honeywell	19	9	73	2
5	Texas Instruments	12	5	52	2
6.5	NCR	7	3	17	1
6.5	Data General	7	3	153	5
8.5	Sperry	4	2	73	2
8.5	Wang Labs	4	2	78	3
10	Perkin-Elmer	3	1	86	3
	Subtotal	206	93	2,611	88
	Other	15	7	349	12
	Total	221	100	2,960	100

TOP 10 IN $32-64,000 SYSTEMS: THIRD PARTY

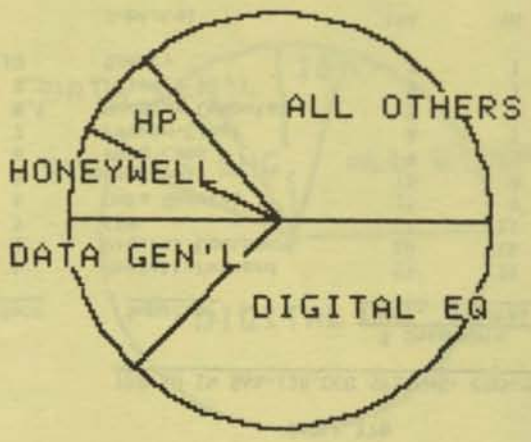

SHARE OF 1981 TOTAL SHIPMENT VALUE

TABLE 16A

TOP 10 IN $32-64,000 SYSTEMS: THIRD PARTY

Rank	Supplier	$ Shipments $ Mil.	$ Shipments % Total	Sites Number	Sites % Total
1	Digital Equipment	107	37	580	46
2	Data General	36	13	121	10
3	Honeywell	20	7	18	1
4	Hewlett-Packard	19	7	71	6
5	IBM	18	6	158	12
6	Texas Instruments	15	5	45	4
7	Sperry	14	5	42	3
8	Perkin-Elmer	12	4	44	3
9	Microdata	9	3	23	2
10	NEC	4	2	3	0
	Subtotal	255	89	1,105	87
	Other	33	11	168	13
	Total	288	100	1,273	100

TOP 10 IN $64-128,000 SYSTEMS: THIRD PARTY

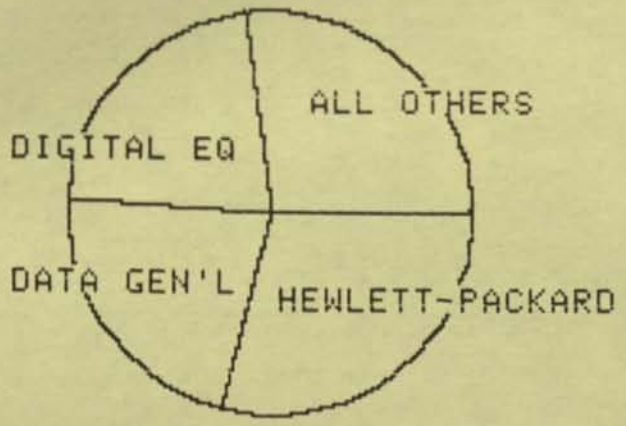

SHARE OF 1981 TOTAL SHIPMENT VALUE

TABLE 17A

TOP 10 IN $64-128,000 SYSTEMS: THIRD PARTY

Rank	Supplier	$ Shipments $ Mil.	$ Shipments % Total	Sites Number	Sites % Total
1	Hewlett-Packard	69	29	196	27
2	Data General	51	22	162	23
3	Digital Equipment	50	22	160	22
4	Perkin-Elmer	16	7	54	8
5	Honeywell	15	6	10	1
6	Modular Computer	11	5	16	2
7	Harris	5	2	9	1
8.5	Systems Engineer.	4	2	14	2
8.5	Prime	4	2	32	4
10	General Auto	3	1	5	1
	Subtotal	226	97	658	91
	Other	8	3	62	9
	Total	234	100	720	100

TOP 10 IN $64-128,000 SYSTEMS: END-USER

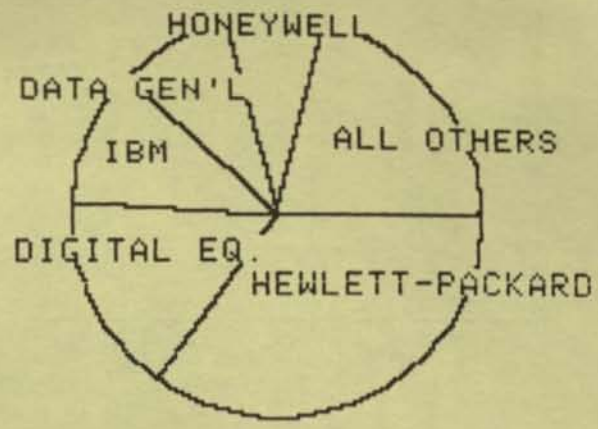

SHARE OF 1981 TOTAL SHIPMENT VALUE

TABLE 17B

TOP 10 IN $64-128,000 SYSTEMS: END-USER

Rank	Supplier	$ Shipments $ Mil.	$ Shipments % Total	Sites Number	Sites % Total
1	Hewlett-Packard	55	35	558	34
2	Digital Equipment	26	16	245	15
3	IBM	17	11	85	5
4	Data General	15	9	267	16
5	Honeywell	13	8	36	2
6	Wang Labs	9	6	64	4
7	Perkin-Elmer	8	5	99	6
8.5	Modular Computer	5	3	64	4
8.5	Prime	5	3	89	5
10	Sperry	2	1	16	1
	Subtotal	154	96	1,523	94
	Other	6	4	102	6
	Total	160	100	1,625	100

TOP 10 IN $128-512,000 SYSTEMS: END-USER

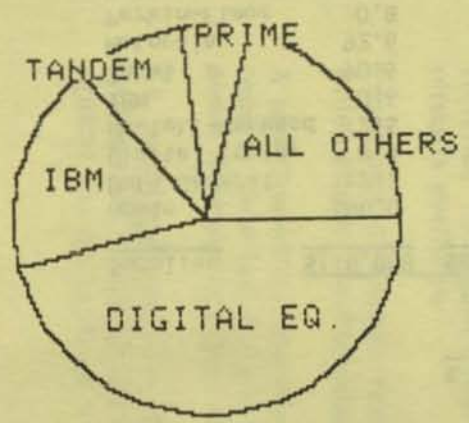

SHARE OF 1981 TOTAL SHIPMENT VALUE

TABLE 18B

TOP 10 IN $128-512,000 SYSTEMS: END-USER

Rank	Supplier	$ Shipments $ Mil.	$ Shipments % Total	Sites Number	Sites % Total
1	Digital Equipment	246	46	1,279	49
2	IBM	94	18	525	20
3	Tandem	47	9	74	3
4	Prime	30	6	160	6
5	Hewlett-Packard	17	3	111	4
6	Data General	16	3	102	4
7	Systems Engineer.	15	3	57	2
8	Computervision	12	2	47	2
9	Perkin-Elmer	11	2	32	1
10	Sperry	9	2	80	3
	Subtotal	496	92	2,467	94
	Other	41	8	170	6
	Total	537	100	2,637	100

TOP 10 IN $128-512,000 SYSTEMS: THIRD PARTY

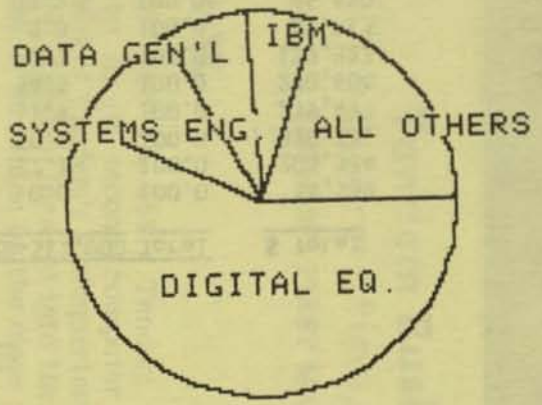

SHARE OF 1981 TOTAL SHIPMENT VALUE

TABLE 18A

TOP 10 IN $128-512,000 SYSTEMS: THIRD PARTY

Rank	Supplier	$ Shipments $ Mil.	$ Shipments % Total	Sites Number	Sites % Total
1	Digital Equipment	240	57	584	55
2	Systems Engineer.	43	10	43	4
3	Data General	28	7	53	5
4	IBM	25	6	166	16
5.5	Tandem	11	3	27	3
5.5	Hewlett-Packard	11	3	23	2
7.5	Harris	10	2	12	1
7.5	Prime	10	2	43	4
9.5	Modular Computer	9	2	11	1
9.5	Sperry	9	2	23	2
	Subtotal	397	94	985	93
	Other	25	6	69	7
	Total	422	100	1,054	100

TABLE 19A
COMPANY DATA

% OF MENTIONED SITES BY PRICE CLASS

Supplier	$1-8,000	$8-32,000	$32-64,000	$64-128,000	$128-512,000	Total	Total No. Mentions
Apple	100.0	0.0	0.0	0.0	0.0	100.0	1,562
Data General	2.6	41.5	17.8	27.8	10.2	100.0	1,592
Digital Equip.	0.4	33.8	30.5	6.4	28.9	100.0	6,788
Hewlett-Packard	15.1	26.5	15.3	36.6	6.5	100.0	2,140
IBM	12.0	16.6	30.6	5.0	35.8	100.0	1,992
Intel	56.1	39.3	0.3	0.0	0.5	100.0*	1,105
Motorola	65.0	35.0	0.0	0.0	0.0	100.0	429
Perkin-Elmer	4.0	15.8	31.0	36.4	12.7	100.0	448
Prime	0.0	0.0	4.8	34.8	60.4	100.0	356
Tandy	100.0	0.0	0.0	0.0	0.0	100.0	962
Texas Instr.	16.9	47.6	21.8	3.0	0.0	100.0*	473
Wang Labs	0.0	51.1	25.3	19.9	3.7	100.0	427

* Includes 3.9% of mentions for Intel and 10.8% of mentions for Texas Instruments in an average price class of less than $1,000.

TABLE 19B
COMPANY DATA

% OF 1981 SHIPMENT VALUE BY PRICE CLASS

Supplier	$1-8,000	$8-32,000	$32-64,000	$64-128,000	$128-512,000	Total	$ Total
Apple	100.0	0.0	0.0	0.0	0.0	100.0	34,532
Data General	2.2	21.4	21.5	32.8	22.1	100.0	204,334
Digital Equip.	0.3	23.2	18.3	7.9	50.4	100.0	1,028,751
Hewlett-Packard	2.6	13.3	16.6	55.6	11.9	100.0	239,477
IBM	0.6	3.2	33.2	8.7	54.3	100.0	230,604
Intel	90.5	8.8	0.2	0.0	0.3	100.0*	129,827
Motorola	92.6	7.4	0.0	0.0	0.0	100.0	67,913
Perkin-Elmer	0.8	5.8	23.4	37.8	32.2	100.0	65,550
Prime	0.0	0.0	1.0	15.3	83.6	100.0	55,585
Tandy	100.0	0.0	0.0	0.0	0.0	100.0	13,816
Texas Instr.	6.5	35.2	46.5	3.8	0.0	100.0*	61,475
Wang Labs	0.0	24.3	26.1	36.0	13.6	100.0	29,010

* Includes 0.2% of value for Intel and 8.1% of value for Texas Instruments in an average price class of less than $1,0

DATAQUEST INCORPORATED
A Subsidiary of A.C. Nielsen Co.
RESEARCH NEWSLETTER

SCIS Code: Tandem Computers Inc.

THE NONSTOP II
TANDEM'S FIRST MAJOR PRODUCT FOR THE EIGHTIES

SUMMARY

On April 14, 1981, Tandem Computers Incorporated announced the Tandem NonStop II, its second computer system. According to David R. Mackie, vice president of Tandem's Headquarters Marketing Operations, the NonStop II is intended to "carry Tandem forward into the eighties," and toward its stated goal of achieving revenues of $1 billion by the close of fiscal 1985.

The NonStop II provides all the features of the original fault-tolerant NonStop system plus the following new features and enhancements:

- A new 32-bit data access architecture, which is the key to implementing the system's extended addressing capability. Each processor can address up to 16 Mbytes of physical memory without architectural changes as denser memory devices become cost effective. The maximum main memory size is currently 2 Mbytes per processor. Virtual memory addressing capability is now one billion bytes per processor.

- Software compatibility with current NonStop systems at the object code level.

- Enhanced input/output capability that will significantly improve the terminal throughput and the number of terminals and communications circuits available.

- Processor microcode resident in 32 Kbytes of loadable control storage, which allows the processor's instruction set to be easily modified.

- A new Operations and Service Processor (OSP) that provides both local and remote detailed system status information and diagnostic facilities to assist in reducing repair times.

- An exchange program that allows current customers to upgrade their systems by returning their original NonStop systems to Tandem for credit against the list price of a NonStop II.

19055 Pruneridge Ave. / Cupertino, CA 95014 / (408) 725-1200 / TWX (910) 338-7695 / DATAQUEST CPTO

FEATURES

32-Bit Data Access Architecture

A major feature of NonStop II is an extended 32-bit addressing mode that allows access to the system's entire virtual memory space. A fully configured 16-processor NonStop II system can address up to 16 billion bytes of virtual memory (one billion bytes per processor), giving users "access to almost unlimited data space," according to Mr. Mackie. The extended address ability provides hardware and software support for very large applications with hundreds of terminals, communications lines (including X.25 virtual circuits), and large, distributed data bases.

Each NonStop II processor currently has a main memory capacity of 512 Kbytes to 2 Mbytes, expandable in increments of 512 Kbytes. However, each processor is capable of addressing up to 16 Mbytes of physical memory. As denser memory devices become cost effective, the current maximum memory of the processor can be increased without architectural changes. This strategy recognizes potential user needs for larger memories in the future and should allow Tandem to meet these needs in a cost-effective manner that is advantageous to its customers.

Software Compatibility

NonStop II is fully compatible with the original Tandem NonStop 16 system at both application code and network levels. "Our users . . . will be able to retain virtually all their original investment in software and most hardware if they choose to upgrade to the NonStop II system to take advantage of its new capabilities," stated Mr. Mackie. Futhermore, NonStop II systems can be used in a common data communications network with the current NonStop systems without software modification. However, NonStop II processors cannot be combined with older NonStop processors in a single system.

Enhanced Input/Output Capability

According to James G. Treybig, Tandem's president, the NonStop II is a direct response to the needs of the Company's customers with large, on-line transaction processing needs. The 32-bit data access architecture can support very large I/O intensive applications needed by organizations having critical business operations on-line. The NonStop II expanded architecture provides more capability in the amount of data that can be transferred to and from disks, terminals, and other peripherals, as there is a 16-fold increase in both I/O buffer space available per processor and in the maximum I/O transfer size. Each processor can now support 1 Mbyte of I/O buffer space. I/O channel speed can reach up to 5 Mbytes/second in a burst mode.

Loadable Control Storage

The processor microcode resides in 32-Kbytes of loadable control storage. The loadable control store allows the processor's instruction set to be modified for new features or performance enhancements as part of Tandem's standard software update procedure.

Operations and Service Processor

Tandem estimates that during an assumed seven-year life of a system purchased in 1982-83, service will cost a user about three times the original price of the hardware; i.e., for every $100 spent on hardware, $288 will be spent on hardware repair and maintenance. To keep these costs from overwhelming the user, Tandem has introduced the Operations and Service Processor (OSP) designed to lower service costs through early fault detection and isolation, and through reducing mean time to repair. Each NonStop II system is supplied with an OSP, which functions both as an operator terminal and as a system diagnostic and maintenance tool that communicates with the Diagnostic Data Transceiver (DDT) that is included in each NonStop II processor module. The operator can thus diagnose software and hardware problems through the operator console. The OSP includes a built-in modem and can be connected to a remote terminal or another OSP, allowing an operator or Tandem service personnel to diagnose, and perhaps correct, problems from a remote site. Furthermore, a remote terminal connected to the OSP can be used as the operator's console, allowing users at the remote terminal to load and run an unattended system. However, initial connection to the OSP modem must be done at the system site.

Exchange Program

Current Tandem users wishing to migrate to the NonStop II may participate in an exchange program that permits them to upgrade their systems as their application needs expand. This plan provides for on-site upgrade of the customer's NonStop system to a NonStop II system, with the customer's original equipment being returned to Tandem for credit against the list price of the upgrade. The returned equipment will be refurbished and resold by Tandem.

Under a similar program, customers can purchase the original NonStop system for application development and subsequently migrate to a NonStop II system as applications come on-line and the application support features of the NonStop II system are needed.

Tandem will continue to manufacture and support the original NonStop system.

PRICING AND DELIVERY

A basic NonStop II configuration consisting of two 512-Kbyte processors, a magnetic tape drive and controller, and an OSP, is priced at $144,475. A basic configuration of the original NonStop system with two 384-Kbyte processors, a magnetic tape drive and controller, and an operator console sells for $94,975. A medium-sized system with six 2-Mbyte processors, four 300-Mbyte disk drives, a magnetic tape drive and controller, an OSP, and software consisting of the Guardian operating system, database management, COBOL, Network/Expand, and access method for 6520 terminals, is priced at $750,000. A roughly similar configuration of the original NonStop system sells for about $700,000. In general, NonStop II systems are priced 15 to 20 percent higher than their earlier counterparts. However, because of the increased memory and the addition of the OSP, the two systems

are not really comparable. NonStop II systems are in production and six systems are already in place. Tandem is promising deliveries 120 days after receipt of an order.

Comments

The NonStop II clearly emphasizes two aspects of Tandem's strategy and philosophy.

By bringing out a system that contains significant enhancements to its original product, the Company has reaffirmed its intention to keep its place in the forefront of the on-line transaction processing systems market and its unique position as the only manufacturer of a fully redundant, fail-safe system. Furthermore, the features of the new system show a high regard for both the user's convenience and pocketbook. In our opinion, the planned memory expandability of the NonStop II virtually guarantees it a long duration of usefulness, and the complete software compatibility and serviceability features show a real concern for the customer's needs and welfare.

Ellen K. Clements
Richard J. Matlack