T A N D E M
SYSTEMS REVIEW

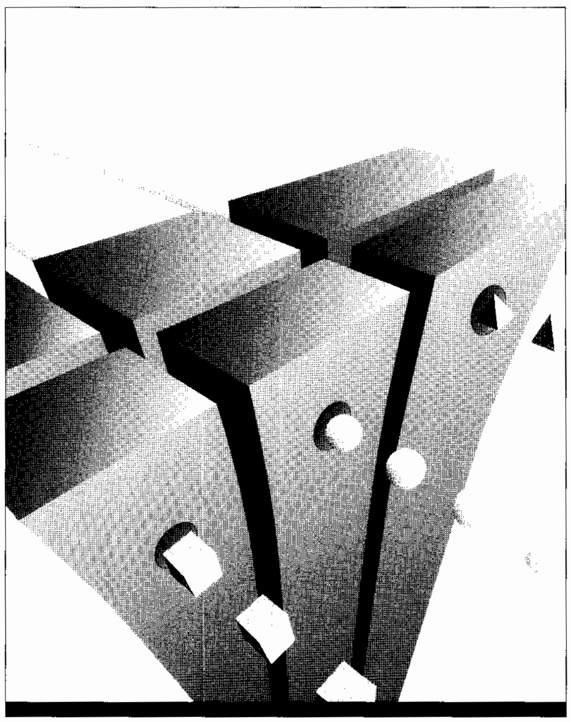

Online Information Processing

DAL Server • RESPOND

Technical Information and Education

Product Update

T A N D E M

SYSTEMS REVIEW

VOLUME 9, NUMBER 1

WINTER 1993

TANDEM

Editor's Note

End users are seeing the increasing benefits of accessing host data from their desktops and making immediate decisions that are based on this information. The demand for workstation access to current mainframe data is making dramatic changes in the requirements of business computing. The first article, "Online Information Processing," is an overview that discusses the benefits of desktop user analysis, possible information processing architectures, and how Tandem addresses issues concerning the efficiency of information processing.

The second article in this issue describes an important tool that provides information access on Tandem systems. "The DAL Server: Client/Server Access to Tandem Databases" describes Tandem's DAL Server and how it is used with Tandem's SQL relational database system and workstation programs that operate on the Apple Macintosh and Microsoft Windows.

Immediate access to online information is also an issue in manufacturing. Tandem's RESPOND system provides continuous access to current data, allowing users to manage events as they occur on the factory floor. The third article in this issue, "The RESPOND OLTP Business Management System for Manufacturing," describes the software and tells why Tandem system architecture makes the transition from batch to OLTP manufacturing applications possible. *—SWT*

Tandem Systems Review is published quarterly by Tandem Computers Incorporated. All correspondence and subscriptions should be addressed to *Tandem Systems Review*, 18922 Forge Drive, Loc 216-05, Cupertino, CA 95014.

Subscriptions: $75.00 per year; single copies are $20.00. Detailed subscription information is provided on the subscription order form at the end of this book.

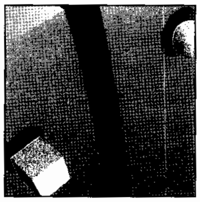

INFORMATION ACCESS

FEATURE

DEPARTMENTS

Guardian 90 Based Software

CLX/R Transaction Servers

November 1992

The CLX/R Transaction Servers are new hardware-and-software packages for client/server computing. They are based on the CLX/R G1100- and G1200-series processors and are offered in four models: GS110, GS112, GS120, and GS122. The GS110 contains two G1100-series CISC processors, each with 16 megabytes of memory, and 600 megabytes of disk storage. The GS112 has four G1100-series CISC processors and 600 megabytes of disk storage. The GS120 contains two G1200-series RISC processors, each with 32 megabytes of memory, and 1.2 gigabytes of disk storage. The GS122 has four G1200-series RISC processors and 2.4 gigabytes of disk storage. All models include a 5126 tape drive and a 3613-1 Ethernet LAN controller.

In addition to the base system hardware, each CLX/R Transaction Server is packaged with the following software: Guardian 90XL operating system, Remote Server Call (host component), Pathway Open Environment Toolkit (POET), Tandem SQL Server Gateway, Data Access Language (DAL) Server, TandemTalk, COBOL85 run-time, and Tandem LAN Access Method (TLAM).

NonStop SQL, Release 2.1

December 1992

Release 2.1 of NonStop SQL provides significant improvements in the areas of database availability, operability, and manageability, as well as in online query performance. These improvements enhance the supportability of NonStop SQL and boost its performance, in many areas by as much as 50 percent. In addition, Release 2.1 reduces the burden of scheduled outages.

Pathway Open Environment Toolkit

December 1992

Tandem's Pathway Open Environment Toolkit (POET) assists programmers in generating client programs that communicate with Pathway servers on Guardian 90 systems. POET, which requires Tandem's Remote Server Call (RSC) to gain access to Pathway servers, currently supports construction of Windows 3.x clients.

POET allows an open choice of low-level client development tools, such as CASE:W, or high-level tools, such as VisualBasic. Windows applications that call a Dynamic Link Library (DLL), and understand and import C data structures, can be integrated with Tandem transaction services more easily through the use of POET.

OSI/AS and OSI/TS, Release 3

November 1992

Release 3 of Tandem's OSI/AS (Application Services) and OSI/TS (Transport Services) provides improved diagnostics, expanded interoperability, and manageability enhancements through Tandem's Subsystem Control Facility (SCF). This release supports Tandem's RISC systems and the C-series and D-series of the Guardian 90 operating system.

NonStop Virtual Hometerm Subsystem

November 1992

The NonStop Virtual Hometerm Subsystem (VHS) acts as a virtual home terminal for Tandem NonStop system applications. NonStop VHS software receives and logs prompts and other messages normally displayed on a home-terminal screen. It thus eliminates the need to use terminals, printers, or other hardware devices as home terminals.

Communications and Networking Products

3650 Communications Subsystem on CLX/R Systems

November 1992

The 3650 Communications Subsystem is now available on all CLX/R systems, including CLX/R Transaction Servers. The 3650 improves the communications connectivity of the CLX/R by providing for a greater number of communications lines. The 3650 communications interface unit (CIU) takes up two of the three I/O controller slots in a single CLX/R cabinet and leaves the third slot for a single-board controller. This configuration provides a 50-percent increase in the number of synchronous and asynchronous lines that can be accommodated, as compared with the use of single-board controllers only.

Workstation and Terminal Products

Tandem Terminal Emulators for Windows

November 1992

Tandem Terminal Emulators (TTE) for Windows comprises three terminal emulators that make it possible to connect a Tandem PSX workstation, or compatible personal computer, to a variety of host computers. The Tandem emulator in TTE supports 65-series connections to Guardian 90 hosts and IXF file transfers; the VT-compatible emulator supports connections to UNIX hosts; and the ASCII emulator supports connections to other hosts. Different emulation sessions can run simultaneously, allowing connection to multiple hosts and facilitating communication throughout an enterprise.

TTE for Windows runs from within Microsoft Windows and features a graphical user interface that provides pull-down menus, resizable windows, font selection, and display color selection. TTE for Windows integrates with standard Windows applications and the multitasking environment lets users move information between host and PC applications with simple copy-and-paste tools.

PSX AP486dE and AP386E Workstations

October 1992

The PSX AP486/50dE and AP386/33E workstations are high-performance desktop computers with Extended Industry Standard Architecture (EISA) I/O bus standard. The AP486/50dE uses the Intel 50-MHz 486DX2 clock-doubler chip to process information internally at twice the clock speed for I/O operations. The AP386/33E uses the 386DX processor with 33-MHz clock speed for both internal and I/O operations.

Both models also incorporate all the standard features and functions of the PSX AP-series workstations. These include an upgradable processor card, integrated 8-kilobyte memory cache, expandable memory (to 16 megabytes on the processor card and to a total of 80 megabytes for the system), integrated graphics support, and the AST FlashBIOS for simple system BIOS upgrades from a diskette.

H26529 Ergonomic Terminal

October 1992

The 6529 Ergonomic Terminal, available only in Europe, is a new, high-quality terminal designed to enhance the comfort level and productivity of the terminal operator. The 6529 meets the highest European ergonomic and safety standards, including Sweden's MPRII 1990:8, and complies with the 1993 EEC Directive.

The 6529 terminal features an adjustable, high-resolution monitor with a 78-Hertz refresh rate, large character-cell display, and antistatic screen. Other features include a detached, two-position keyboard; two RS-232C serial ports and one Centronics-compatible parallel printer port; black text displayed on paper-white background; 10 pages of screen memory; and a built-in menu with a full range of configuration options. In addition, the 6529 supports 6530 and ANSI 3.64 data streams, supports 12 languages, and uses flash EPROMs for easy firmware upgrades.

Electronic Funds Transfer Products

BASE24 Open Network Control Facility

August 1992

The BASE24 Open Network Control Facility (ONCF) is a new product designed to facilitate an automated operational environment within BASE24 networks and to support an open network architecture. ONCF consists of enhancements to the BASE24 Network Control Supervisor (NCS) subsystem and the delivery of two new processes, the Network Control Point Communications (NCPCOM) and Network Control Point Gateway (NCPGATE). ONCF provides a command set commonly used for network control operations of lines, stations, and processes using Tandem's Command Language Standard (CLS) format.

BASE24 users can achieve automated operations by employing ONCF in conjunction with Tandem's Event Management Service (EMS) and NonStop NET/MASTER or Programmatic Network Administrator (PNA). ONCF also works with custom developed Distributed System Management Applications (DSMA) designed to the ONCF command set.

Storage Products

5410 Object Storage Facility

October 1992

The 5410 Object Storage Facility is a high-capacity, data-storage device that provides online, random-access storage for any Tandem Guardian 90 system. The 5410 uses write once, read many (WORM) technology to store and retrieve data on optical disks. The write-once feature ensures that data or documents written to the 5410 cannot be altered. Once the unit is online, the stored data is accessible at any node in a Tandem network.

The 5410 contains two read/write optical disk drives and handles up to 50 disk cartridges for a total storage capacity of 328 gigabytes. The option of constant angular velocity (CAV) or constant linear velocity (CLV) media allows users to optimize operations for fast access time or large capacity. Applications written for the 5200 Optical Storage Facility can run on the 5410 with no modifications.

5200 Optical Storage Facility on Cyclone/R

September 1992

The 5200 Optical Storage Facility (OSF) is now available on NonStop Cyclone/R systems. The 5200 is a write once, read many (WORM) data storage device that provides up to 83.9 gigabytes of online storage on optical disks. To use the 5200 on their Cyclone/R systems, users must have the C30.08, or later, release of the Guardian 90 operating system.

Online Information Processing

Modern technology has progressed to a point where it can provide MIS managers with more than enough data and preformatted reports to run their companies. The problem today is finding a way to turn this wealth of data into information that users can access at their desktops to make day-to-day, and sometimes instant, business decisions.

User analysis of both current and historical data to make daily decisions at user desktops is called *online information processing*. MIS managers are concerned about affecting the efficiency and response time of their operational systems if they add on the processing power to do this type of processing. If they allow more users to have access to the data, managers feel that the response time in their operational systems will decrease.

This article describes the general benefits of online information processing, reasons for using it, and some information processing architectures. It also shows how the mixed workload enhancement technology, provided by Tandem™, and parallel query execution in Tandem's NonStop™ SQL relational database management system solve management's perceived concerns about degrading their current operational systems.

Defining Operational and Information Processing

Operational processing consists of the computing and data collection activities required to run day-to-day business operations. Tandem's online transaction processing (OLTP) is a subset of operational processing.

Information processing also involves computing, but it is more involved with data analysis, often of historical data. When a user requests the system to summarize historical payment data into a past due report so that selected customers' payment histories can be analyzed, the user is doing information processing. The distinction between the two types of processing is small but quite different in terms of the kinds of data each uses, the volumes of data involved, and the types of processing power required.

Operational Systems for Core Computer Processing

Nearly all organizations begin by building operational data processing systems. They are the heart or core of the computer processing. The following are examples of operational systems and their transactions:

- Banking demand-deposit systems manage checks cleared and money withdrawn and deposited.
- Retail systems process incoming inventory and items sold.
- Manufacturing systems generate transactions for raw materials ordered and received and finished goods shipped.
- Utility systems process energy consumed, raw product delivered, and service connected or disconnected.

Operational systems are relatively simple to design and implement in terms of the data input, the processes required on the data, and the data output. They are also easy to cost-justify in terms of reduction of clerical tasks, ability to handle rapidly growing volumes, and efficiencies achieved. For example, posting transactions for 10 million checking accounts cannot be done manually. This requires an operational computer system.

From a data perspective, one of the key attributes of an operational system is that the only data it requires is data for the current business cycle. Although an operational system may handle large volumes of data input during a cycle, an equal amount of data is removed or discarded from the system at the end of the cycle. The cycle may be as long as 60 or 90 days or as short as a few hours. An example is order entry.

Information Processing for Extracting Additional Value

At the end of a cycle, an operational system may retain some summary information for quarterly or year-end processing, but it will archive or purge most detail information. System managers may make archived data accessible to users for further analysis. Users can examine the saved data to gain a better understanding of how the business interacts with its customers, suppliers, and competitors. By doing so, users can find ways to become more efficient and more competitive.

Information processing in most MIS organizations usually begins quite innocently as an extension to operational systems. In fact, application designers almost always define some information processing as part of an operational application system design. During systems analysis, the designer will receive user requests for information processing by-products. For example, a user might request the generation of an accounts past due report at the same time that the system produces monthly invoices. Printing the invoices is operational processing, while printing a report showing a 12-month trailing past due account for trend analysis is an information report for measuring the effectiveness of an accounts receivable unit.

The first information request usually starts an influx of demands on the data processing systems that far exceed the data and processing power needed merely to run the business. At some point, the MIS organization decides to consciously reserve additional data storage and processing capability to meet growing informational needs.

Information service personnel find that managing the transition of data from the operational systems to the information systems is quite complex. They understand the operational systems in terms of the output required, the inputs available, and the processes required to go from input to output. However, they soon find that they can define information systems only by the next question that a user might ask.

Figure 1

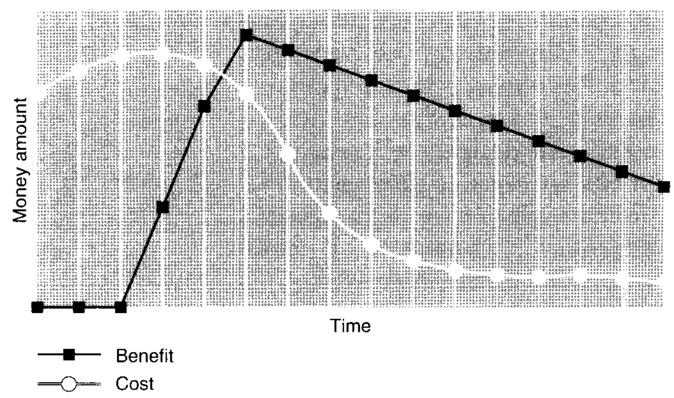

Figure 1.
An operational system may not provide a company a net benefit until a year or more after it has been in operation.

Benefits of Information Processing

Most MIS managers understand at least some of the benefits of analyzing and summarizing operational data to extract additional information. For example, when a company uses information processing to better understand customer needs and buying habits, it can better serve its customers and improve customer satisfaction. Companies can also use online information processing technology to offer advanced services to their customers. An example is direct access to order entry or customer accounts information. In addition, by studying how a company does business and by identifying bottlenecks, managers can determine ways to make a company run more efficiently and become more responsive to both customers and vendors.

If managers analyze business operations and lines of business, they can help identify potential cost savings, eliminate less profitable ventures, dedicate resources to more profitable lines of business, and suggest direction for future investment in products and services. Then, using business data to better predict market direction and customer demands, they can help a company to stay ahead of the competition.

The payback on investment is more rapid for an information processing system than it is for an operational processing system. When building a new operational system, a company must make heavy initial investments to either acquire and customize software and hardware or build an entire system from the ground up. Financial benefits begin to accrue only after the system has operated for some time. Figure 1 shows the typical relationship of investment (cost) to return (benefit) for an operational system. Depending on its complexity, a system may not provide a company a net benefit until a year or more after it has been in operation.

On the other hand, when a company builds a new information processing system, it often realizes nearly immediate payback. Figure 2 shows the typical relationship between costs and benefits for information processing systems.

One reason for the quicker payback is that the data required for information processing already exists in an operational system. Therefore, the company does not have to build or shop for applications that gather the data. Another reason is that the effort a company must invest in building an informational system can be spread over a period of time. The company need not solve the entire information processing problem immediately. In fact, system designers may have difficulty identifying all the potential information processing requirements until after some of the most obvious ones have been satisfied. They will find that the answers to first questions often lead to further questions, which are, in reality, requirements.

A company spends most of its initial investment on three tasks. First, MIS professionals must identify the subset or summary of operational data required to solve a particular problem. Then they must add system storage to keep available data that would normally be archived. Last, they must purchase readily available tools to allow access to and analysis of this data.

The instant a user can gain access to even a small amount of information, the company begins to accrue benefits. Ongoing benefits from expanding information processing generally far exceed investment in management tools and additional hardware.

Reasons Companies Go Online

Online information processing means providing decision support or question-answering capabilities to company management, front-line personnel, and perhaps even customers. Users of online information processing often require access to the most current data available, usually from operational system files.

Companies who have online access to both current data and summarized historical data can solve two major competitive business problems. First, they can make the best decisions about company business problems that require not only historical analysis but access to information about the company's most current status. Also, by offering customers access to both current information (the last ten transactions) and historical information (a summary of the last quarter's transactions), companies can help their customers to solve their own problems.

Responding to a Recent Event

Companies gain the highest benefit when they push information processing back down into operational systems. For example, if order takers can access customer order history while entering a customer order, they can enhance customer service and increase sales. When they notice that customers place orders for a particular product at the same time each year, they can advise them when the company has a sale on the product. The most successful companies in the 1990s will be those that can figure out how to provide their personnel with desktop information that will help them to close a sale, improve service, save time, or save money.

Good distribution systems use historical sales data to determine how to allocate goods to retail outlets. The best systems use up-to-date inventory information to precisely shift allocations where they are most needed.

Figure 2

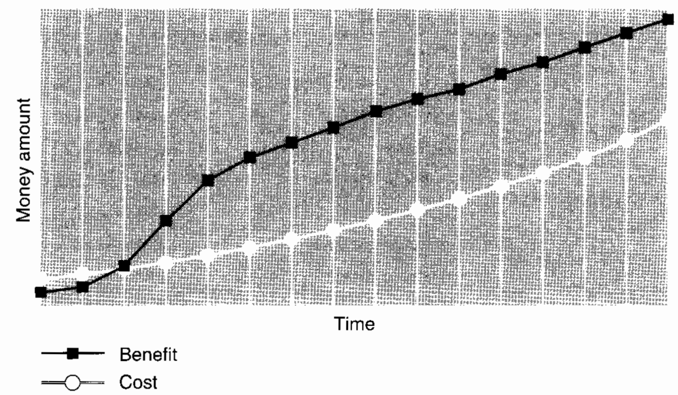

Figure 2.
An information system may provide a company a net benefit soon after it begins operation.

Systems dealing with credit or investment problems can especially benefit from a mix of historical and current transaction data. For example, a person might have an excellent credit history, but information about recent transactions or other credit applications may affect a credit authorization decision. Another example is the investment broker who may need accurate historical and current data to make a decision about taking or liquidating an investment position.

In the field of medicine, a physician may need access to extensive patient history as well as latest test results to perform a correct diagnosis and render treatment. This is especially true in emergencies. Also, a pharmacist may provide a warning or recommend against filling a prescription based on other recently prescribed drugs that could interact negatively.

Table 1.
Differences between operational and informational data characteristics.

Requirement	Operational	Informational
Level of detail	Detailed	Summarized
Performance	Quick response (seconds)	Casual response (hours)
Access technique	Highly structured	Unstructured
Availability	Continuous	Periodic
User type	Customer, clerical	Management, customer
Currentness	Current values	Historical values
Usage	Continuous update	Static, read-only
Definition	Relatively static	Changing
Volume	Limited	Potentially enormous
Processing algorithms	Stable	Dynamic, heuristic
Integrity	Discrete transactions	Some errors tolerated
Data management	All detail	Summarized subsets
Content	Application-based	Subject-based

Offering Information as a Product

Companies have discovered that the information they collect can enhance an existing product or provide a new service to their customers. For example, banking institutions can gain a competitive edge when they provide both retail and wholesale customers with transaction history in addition to allowing them direct entry of debit and credit transactions. Banks can offer these services at ATMs, over phone lines, through personal computer access, or through an online information service such as Compuserve or the PRODIGY service.

In the travel industry, most major airlines can provide special services to travelers who have personal computers and modems. These services include online access to historical information such as on-time records and frequent flier activity. In addition, they may allow customers to book reservations on future flights, choose meals, and select seats.

Hotels and car-rental agencies are also going online. Many major hotels today provide guests with online access to their latest bills. Hotel guests can also order room service and check out through a keypad attached to the television sets located in their rooms.

Several online information services provide users with extensive historical information on securities. In addition, they give them a way to manage a personal portfolio through a home PC.

Some Information Processing Architectures

Designers can use a number of ways to build a business computing system that will support information processing tasks. Some information processing tasks normally are designed as an integral part of an operational system, and, therefore, they run on the same system as the operational tasks. They also share the same data or databases.

As information processing needs grow, many companies begin extracting and summarizing data onto a separate system that is dedicated to information processing. Sometimes this information resides in different databases. When a company brings information processing online, it must implement an architecture that allows users efficient access to both historical and operational data no matter where it resides.

Online Information Processing on Operational Data

MIS professionals can satisfy information processing requests using operational data and operational system resources as long as the users' questions are well-defined and the extra load placed on the operational system is predictable. However, the result of one question may spawn dozens of additional questions. Many questions require the system to summarize the raw data so the user can draw conclusions. When users are allowed to generate additional requests as needed, they can quickly consume resources on operational systems. As a result, the operational system slows down in its response to the user.

To answer the next information processing question, system implementers must anticipate the amount of raw data to collect from an operational system. The operational system may require, at most, 30 or 60 days of data, but a subsequent information system question may need 1, 3, or 5 years of data. From this perspective, the data requirements of operational systems are quite different from those of information systems. Table 1 summarizes the differences in data characteristics.

The characteristics of the processing load are also different in the two types of systems. Systems analysts design operational systems so they are balanced for efficient use of all resources and delivery of consistent response times. The load profile of a well-tuned operational system looks like the one in Figure 3.

On the other hand, information systems have many peaks and valleys in the processing load, as shown in Figure 4. The information system may be idle for extended periods of time, yet a single request can quickly consume all available resources.

Because of these differences, at some point, information system managers isolate information processing data from operational data. To keep information processing from impacting day-to-day business, managers eventually may install an entirely separate system.

Information Processing on a Separate System

When information system managers separate information processing onto its own database or onto an entirely separate system, they minimize the impact on operational systems. In fact, whenever they find a high-volume online transaction processing requirement, many companies go one step further and also implement online systems separate from batch processing systems.

Figure 3

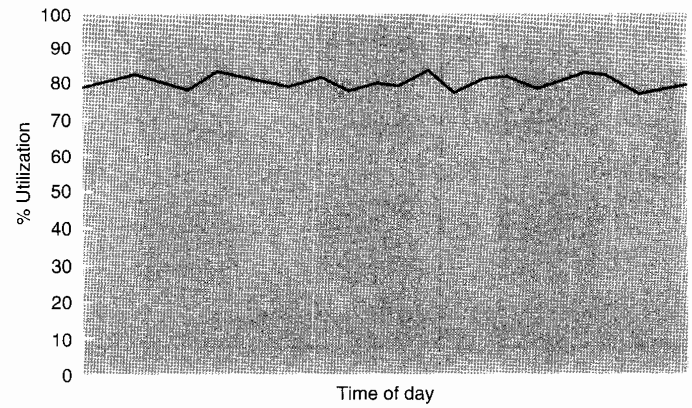

Figure 4

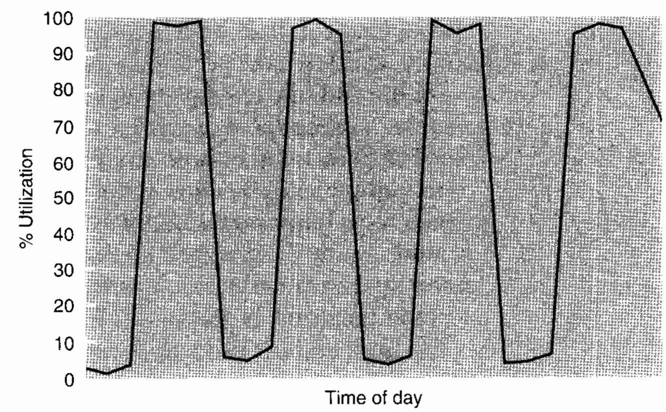

Figure 3.
Typical operational system workload.

Figure 4.
Typical information system workload.

Figure 5

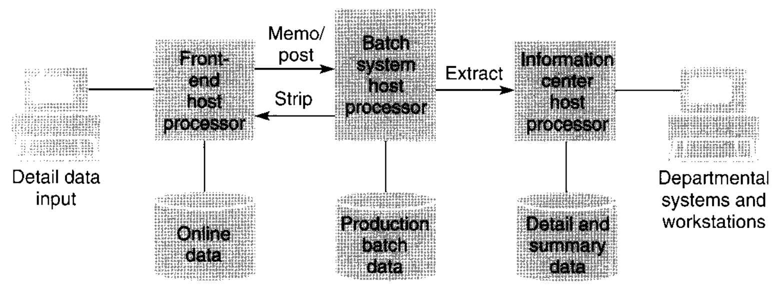

Figure 5.

Typical information processing architecture. The information center data consists of detail plus weekly, monthly, or other summary data.

Figure 5 shows the typical MIS environment that many companies have today. High-volume OLTP applications reside isolated in a front-end processing system, shown on the left of the figure. The most current data resides here. Most users employ Tandem technology to this front end so they can gain high availability, expandability, and low cost per transaction per second.

At the end of each day, or at some point in a cycle, a memo/post operation updates the batch files on another host processor. The data of record typically resides in the batch system where most corporate accounting takes place. A system may be designed so that it strips summary information, such as end-of-day account balances. It places this information in the front end for users to access for online validation in the next day's processing. If the operational system no longer needs this summary data for the current accounting cycle, it may purge the data from its system.

To satisfy information processing needs, companies extract and summarize data on a weekly or monthly basis into a third host processing system, which is dedicated to decision support. From this third system, a complex network of departmental systems and workstations may provide users information on their desktops.

A number of problems arise as an MIS department evolves toward the typical environment. By the time data arrives on the information processing system, it is too old to solve the most difficult problems. For queries that require the most current data, MIS professionals must train users to run them immediately after a periodic update. Otherwise, they must design the system so that users can access data directly from the host.

Another problem is that multiple copies of corporate data may exist in the information center. Users may download any one of these copies to their departmental systems or workstations and have difficulty in determining which copy of the data is most current and correct. Decisions that users make in this environment become questionable.

At times when operational systems can benefit from decision support processing, the operational system may have difficulty gaining access to the informational data. A classic example is a bank that wants to offer instant credit approval. The customer credit history data that is necessary to make good decisions may be isolated on an information processing system that does not have a direct network link to the operational system.

The typical information processing system can provide companies a reasonable return on investment when the system is used to assist in the strategic decision-making process. On the other hand, it cannot help solve user problems of an immediate nature in a cost-effective way. To do this, information processing needs to be more than just online. It must be online to the operational systems.

Obstacles to Online Information Processing

Although many companies recognize the benefits of moving information processing online, most see too many obstacles in implementing these systems. Companies complain that information processing is too expensive to implement online. On the contrary, if MIS managers plan carefully, online information processing systems are no more expensive to implement or maintain than any other kind of automated system. When a manager chooses to solve problems whose solutions lend themselves to online information processing, this type of processing provides a better return on investment than any batch, decision support, or OLTP system.

Some feel that connectivity is a difficult and expensive problem. Today's information systems vendors cannot be successful without providing good connectivity to other host platforms, departmental systems, and workstations. Managers will find technologies and systems integration expertise that provide simple and cost-effective solutions to most connectivity challenges.

Some managers are reluctant to trust the reliability of data in an online information processing system. When a company's information processing systems use the typical architecture in Figure 5, even with good management procedures, MIS managers find it impossible to always guarantee the accuracy and timeliness of information. However, alternative information processing architectures can more easily provide these guarantees.

Others think that they have to pay too high a price for a system that has continuous availability to users. On the contrary, systems with availability high enough to meet online processing requirements are not necessarily more expensive than any other type of information processing system.

Many companies feel they need large monolithic systems to meet the demands of information processing requests. In actuality, parallel processing technology is available to help companies build systems that are well-matched to resource and response-time requirements. These parallel systems can grow to the required size in small, affordable steps. Some grow even larger than the largest monolithic systems.

Some managers are concerned that they will not be able to mix online transaction processing with online information processing. They cannot visualize a system that combines many small, response-critical tasks with large, unpredictable tasks that can consume all resources for extended periods of time.

One of the strengths of today's parallel systems is that systems designers can spread many small pieces of work or break up several large requests across multiple processors. A designer can balance priorities to ensure that transaction-oriented work will run with no interference from information requests. Simple planning and tuning of both the OLTP applications and information processing tasks across a parallel system can ensure adequate service levels for both kinds of tasks on the same system.

Management presents another obstacle when it feels the company does not need online information processing. MIS professionals can overcome this resistance by suggesting that management become aware of what their competitors are doing. If the competition is not already implementing online information processing systems, they probably will be doing so soon.

Implementing Online Information Processing

No company should attempt to convert to online information processing overnight. MIS professionals must first understand the company's overall data architecture and needs. They must then identify those tasks or systems that will provide the best payback once they become part of online processing. Finally, professionals need to understand available technologies and use them.

Building a Data Architecture

MIS professionals agree that they must develop a data model and architecture to successfully implement information processing systems. They must minimize the time it takes to do so, for businesses change quickly over time. If designers spend too much time, the business may change so significantly that they will have to start over again.

Even if the company has never considered modeling the MIS environment, systems analysts can probably complete the identification of major systems and their interactions in three to six months. During the process, they most likely will identify emerging requirements that current systems have not yet met.

Identifying Online Information Processing Opportunities

Once MIS professionals have gained a high-level understanding of the processes and data that drive the company, they can begin to look for opportunities to enhance profitability, productivity, and competitiveness through online information processing. They may find that they will have to make major revisions to existing systems to get accurate and timely data. If these systems are ten or more years old, they may be near the end of their useful investment life cycle anyway. A major online information processing opportunity may provide sufficient justification to begin a complete retooling of these old applications. In addition, MIS professionals may also find opportunities to implement entirely new competitive-edge systems.

In the process of developing data and a process model, systems analysts must interview the current users of the company's information processing systems. They must listen for statements that reveal facts about how a job could be done faster or cheaper and what it would take. They must also note user problems such as poor access to data, old data, and unreliable data. Systems with these problems become prime candidates for an online information processing solution.

Defining an Information Processing Architecture

A *data warehouse* is a storage area for read-only data that has been gathered from multiple sources. MIS professionals recognize W. H. Inmon as a consultant and author who is often credited as the father of the data warehouse concept.

Inmon advocates distinct storage structures for operational and information data, as shown in Figure 6 (Inmon, 1989). The left side of the diagram represents the data required to support operational systems. The support information often resides in the online transaction processing portion of operational systems and is the raw source for the data of record for batch systems.

Once the data of record becomes older than the normal operational processing cycle, it moves from a state of being current to a historical state, and it then resides in the data warehouse. The data warehouse forms the foundation for historically based information processing.

As detail data moves into the first, or atomic, level of the data warehouse, systems must usually add timestamps to identify the time frame for this level of information. Inmon uses the word *atomic* to describe the lowest level of detail available, such as individual transactions. Some of this detail data may eventually migrate to departmental systems or individual user desktops.

The seven boxes on the right side of Figure 6 represent the different summarizations of the data as it moves down in the organization. A department may look at the data in only a few ways, but individuals may perform many different extractions and summarizations.

Within the data warehouse, an information processing system must keep a certain amount of detail data to meet short-term, low-level information processing needs. Maintaining the detail of all transactions for extended periods of time can be costly, so data administrators typically implement programs to summarize much of this detailed data, particularly to satisfy long-term data needs at the department or individual level.

Figure 6

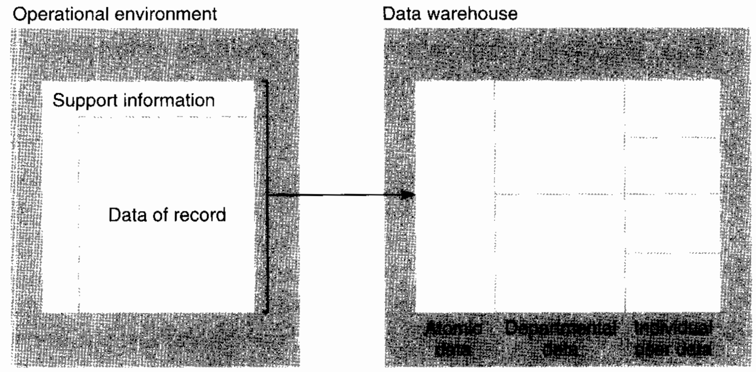

Figure 6.
Building the data warehouse. (Redrawn by permission, from W. H. Inmon, *Building the Data Warehouse,* Fig. 1, ©1992 by QED Information Sciences, Inc.)

Perhaps one of the most important tasks in maintaining a data warehouse is analyzing and predicting future data summarization requirements. A data administrator can accomplish this by creating a matrix for each major subject area.

The matrix shows where the level of summary intersects with the level of user. It usually indicates that personnel in the lower levels of an organization need the most detailed information for the smallest historical time span. Higher in the organization, strategic analysis tasks require a broader historical perspective with highly summarized information. By carefully selecting intersection points, the data administrator can design a data warehouse that stores detailed data for a minimum of time yet with perhaps only three or four summarization levels to handle longer-term needs.

For example, customer service representatives in a particular bank need access to detail customer transactions for 90 days. First-line management needs detail for only the last 30 days, plus monthly summaries for six months. Middle management personnel have no need for detail, but they require monthly summaries by type of account for one year. Senior management is only interested in quarterly summaries, but they need it sorted by type of customer for five years. A data administrator must provide these different levels of detail.

In this scenario, the data administrator might decide to keep detail for 90 days and provide routines to create the monthly and quarterly summaries for upper management from the detail. Once detail data is more than 90 days old, the data administrator can run programs to summarize the data by account type and customer type for up to one year. When the data is more than a year old, the account-type information can be purged.

Using Technology Effectively

Once MIS professionals have defined the online information processing architecture, these professionals can use current technologies to overcome some of the perceived obstacles to implementing the online information systems. Tandem offers two key technologies: mixed workload enhancement and parallel query execution in Tandem's NonStop SQL relational database management system.

Mixed Workload Enhancement. Most OLTP systems must be highly responsive and capable of completing transactions within a short time, measured in seconds or even less. When all tasks on a system require few resources and run quickly, service levels remain stable. This is because no one task can consume more than its share of resources and lock out other tasks for more than a few microseconds.

One of the key reasons why a system cannot run information processing requests simultaneously with operational online transaction processing is because information processing can often require large amounts of both CPU and disk resources for extended periods of time. If an operational OLTP system dispatches a long-running task, the task may lock out key resources from OLTP tasks.

For example, if an information processing task needs to scan large volumes of data, this task can monopolize the services of the routines that manage access to records on disk. To prevent this from happening, the operating system must contain special features.

Recent enhancements in Tandem's Guardian™ 90 Release C30 disk process ensure that no lower-priority task can monopolize access to data on disk. Experience with Release 1 of Tandem's NonStop SQL relational database management system showed that a long-running information processing task can monopolize the disk process for a long time with devastating effects on high-priority OLTP tasks. This can occur even when the information task is dispatched on a processor or system separate from the OLTP work.

The mixed workload enhancement allows the disk process to honor the relative dispatch priority of the requesting tasks and interrupt long-running work, such as an SQL partition scan, whenever it needs to service higher-priority work. It allows operators to assign appropriate dispatching priorities to the OLTP and information processing work so that the OLTP tasks can continue to run at required service levels. At the same time, it allows the system to service simultaneous information requests.

Parallel NonStop SQL Software for Queries. Release 2 of Tandem's NonStop SQL relational database management system gives systems the ability to automatically break long and complex SQL scan requests into smaller units of work that can run in parallel and be spread across multiple processors. For example, before parallel execution was available, a query could run for more than an hour if it needed data from a table that spread across five disk drives. However, with scans running in parallel, the query's execution time can decrease by as much as 80 percent.

For parallel execution to take effect, the SQL compiler must determine that at least part of the query will be solved by a table scan and that the table is spread across three or more partitions. Database administrators can execute the EXPLAIN command in the SQLCI utility to determine how NonStop SQL software will solve a query. The EXPLAIN utility reports whether any portion of the query will be solved by using a scan.

If a table is 40 megabytes or larger and is subject to scanning, the database administrator needs to allocate it across three or more partitions, preferably on disk drives available to separate processors. For example, if three 600-megabyte tables are subject to query scans, the administrator needs to place them each in three 200-megabyte partitions on three separate disk drives, instead of allocating each one to a single partition on its own drive.

Figure 7 shows how to do this. Once the load has been spread in this way, the allocation of disks, processors, and application processes can be balanced as in any other online system.

Building an Online Information Processing Solution

The best way to implement an information processing solution is to start slowly and release the solution in small increments. Build a pilot system first and then determine whether it meets key requirements.

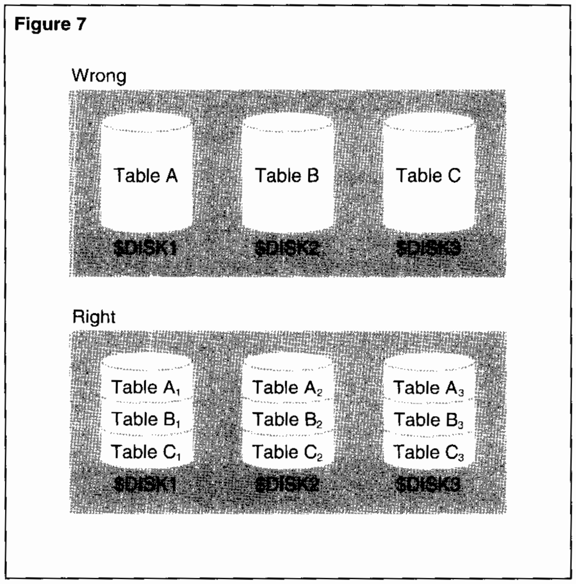

Figure 7.
Allocating tables for parallel processing.

Information processing systems tend to be more loosely defined than most operational systems, and they have requirements that shift over time. MIS professionals who try to implement an entire solution all at once, rather than a piece at a time, find their work obsolete on the day it is released. Relational databases are helpful because they provide the flexibility to implement a partial solution and then enhance it later to meet new requirements. Because information processing systems excel in providing a return on investment in a relatively small amount of time, the sooner one implements a portion of the system, the quicker one will realize a return.

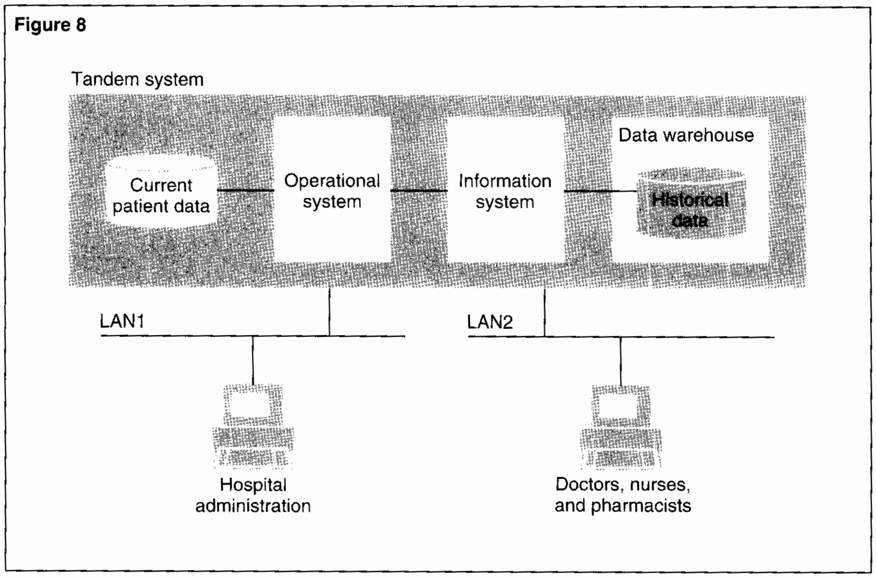

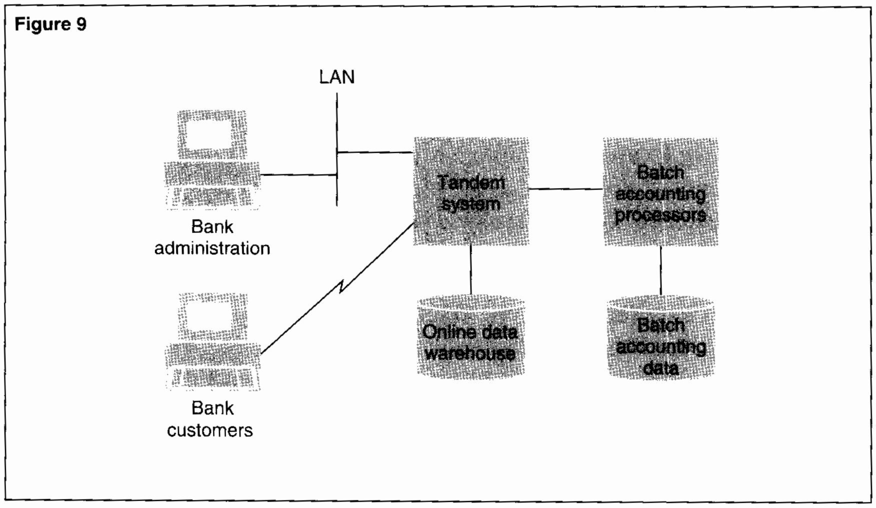

Figure 8.
Running a hospital online.

Figure 9.
Online wholesale banking.

Responding to Recent Events in a Hospital

Hospital care is an example of the kind of system that can gain from combining historical information with recent event data. A major hospital in North America has significantly improved health care professional efficiency by adding online information processing capabilities to its online patient and business management systems. Figure 8 shows the hospital's system configuration.

In an operational database on its Tandem systems, the hospital maintains current patient status, including the results of recent tests. Doctors, nurses, and pharmacists have access to a separate data warehouse that contains comprehensive patient history and is part of the same logical system. Health care workers on the patient floors can request recent test data and history from conveniently located workstations. Nurses record new life-sign readings, and doctors enter diagnoses and new care instructions directly online. To provide optimal online performance to users, system professionals have predetermined and tuned most of the information queries.

Hospital administrators also have online access to current and historical data. They use this information to track costs both by service and by patient-day. Their continual input provides information so they can adjust available services and control costs. Administrators have access to ad hoc query tools to provide added flexibility for their information processing needs.

Offering Information as a Product in a Bank

The ongoing competitive and regulatory challenges in the banking industry make it a prime candidate for online information processing. This industry is an example of the kind of system that can gain from using OLTP systems to provide an improved or a more competitive product.

A major U.S. bank offers enhanced online services to its wholesale customers by adding access to full account history in its dial-in banking services system. The bank implemented this system to increase wholesale customers' access to their account information and provide greater control over collection and disbursement of funds. The bank benefits from a more competitive, reliable offering to customers. This product offering also reduces bank expenses by shifting transaction generation to the customer.

Figure 9 shows that bank customers connect to the Tandem front-end system with terminals or PCs over an X.25 network. From terminals, customers can access previous and current balance information, report on automated clearing house activity, monitor float, and control disbursement of funds. In addition, the bank provides PC customers with custom client software so they can initiate wire transfers and stop payments.

Bank administration also uses this system to monitor customer and account profiles, examine customer activity, control security, and produce extensive reports on wholesale customer use of the system. Every evening, the system posts daily activity to back-end batch accounting machines and verifies the data in the online data warehouse on the Tandem system.

Conclusion

Most Tandem users recognize they can use Tandem systems cost-effectively in implementing online transaction processing solutions, and many have discovered that NonStop SQL software allows them to make use of relational features to support information processing. On a Tandem system, implementing information processing concurrent with operational processing allows Tandem customers to become more competitive, reduce their costs, and improve their productivity.

By providing key technical features coupled with a wide range of information processing tools, Tandem has provided a way for users to add an information processing solution to a Tandem system. These features include the mixed workload enhancement and parallel query execution in NonStop SQL software. To achieve success in implementing information systems, MIS professionals must understand the issues surrounding the design of information processing systems and the use of these key technical features.

Reference

Inmon, W. H. 1989. *Data Architecture: The Information Paradigm.* QED Information Sciences, Inc.

John Viescas is a consulting analyst in Technical Support Services. Over the last four years, he has specialized in recommending relational database solutions for Tandem users worldwide.

The DAL Server: Client/Server Access to Tandem Databases

With the proliferation of off-the-shelf workstation software, end users see the increasing benefits of accessing host data from their workstations and manipulating the data with packaged third-party applications. To provide this capability, Tandem™ offers the Data Access Language (DAL) Server subsystem, client/server software that gives workstation applications access to SQL data stored on Tandem servers.

With the DAL Server, client programs can reside on Apple Macintosh or Microsoft Windows workstations. End users can take advantage of the economy and ease of use of workstation graphical user interfaces (GUIs) and a variety of packaged applications such as spreadsheets, application generators, and database programs. Tandem servers provide the high level of data integrity, large capacity, and data security of the NonStop™ SQL relational database management system.

The DAL Server is designed to support information analysis, decision support, and low-volume online transaction processing (OLTP). Tandem offers complementary client/server products designed for high-volume, fast-response OLTP applications based on Tandem's Pathway transaction processing system and TMF™ (Transaction Monitoring Facility) software. Developers can use Tandem's Remote Server Call (RSC) software together with the Pathway Open Environment Toolkit (POET) product to implement client/server applications for high-volume OLTP. For information about using RSC, see the article by Iem and Kocher in the Fall 1992 issue of the *Tandem Systems Review*.

This article describes how the DAL Server operates in a client/server environment. Next, it describes three DAL Server applications in which users query NonStop SQL databases to perform demographic analysis, corporate support for field marketing personnel, and project planning and control. Readers of this article should have a general knowledge of the concepts of client/server architecture. For readers interested in the details of DAL Server data communications, the article also briefly describes the TandemTalk subsystem, Tandem's implementation of AppleTalk, which provides communication between Tandem systems and Macintosh workstations.

DAL Server Architecture

DAL is an SQL-based data connectivity and manipulation language developed by Apple Computer. It provides access to database information in the client/server environment. Figure 1 shows how DAL supports the client/server model of computing.

Tandem's DAL Server subsystem provides the workstation user with transparent access to Tandem NonStop SQL databases from third-party, DAL-compatible client applications. With these tools, users with little or no SQL programming ability can create decision-support and low-volume OLTP systems.

Figure 2 shows how the DAL Server connects workstations to a NonStop SQL database in a Guardian™ 90 operating system. Macintosh workstations connect through TandemTalk software, and PC workstations connect through Transmission Control Protocol/Internet Protocol (TCP/IP). Tandem licensed the DAL Server and the AppleTalk protocols from Apple Computer and programmed them to run on Tandem NonStop systems. Tandem's Expand™ data communications network connects Tandem systems to one another, thereby allowing the DAL Server to work with remote or distributed databases.

The DAL Server converts DAL statements from client applications into NonStop SQL statements for compilation and query execution. Data selected by the query can be retained in the DAL Server for temporary storage or returned immediately to the client, and the client application can control the flow of data. The DAL Server can also reformat SQL data to meet the requirements of the workstation application. Using the capabilities of the Macintosh or Windows workstation, users can see host data in easy-to-understand graphic formats and quickly manipulate that data.

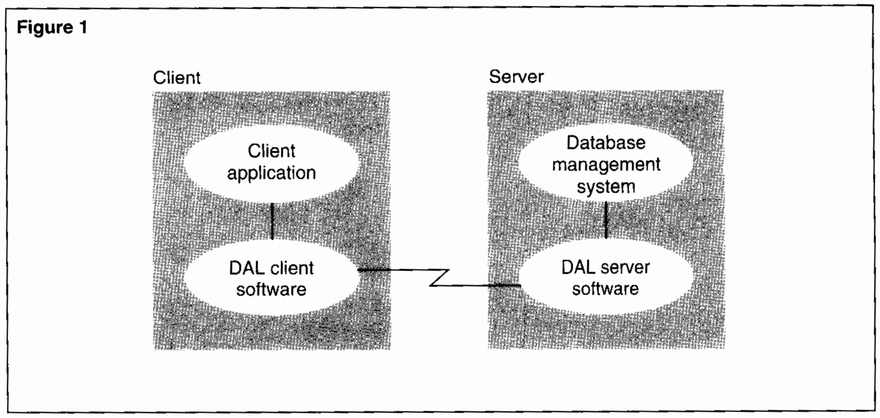

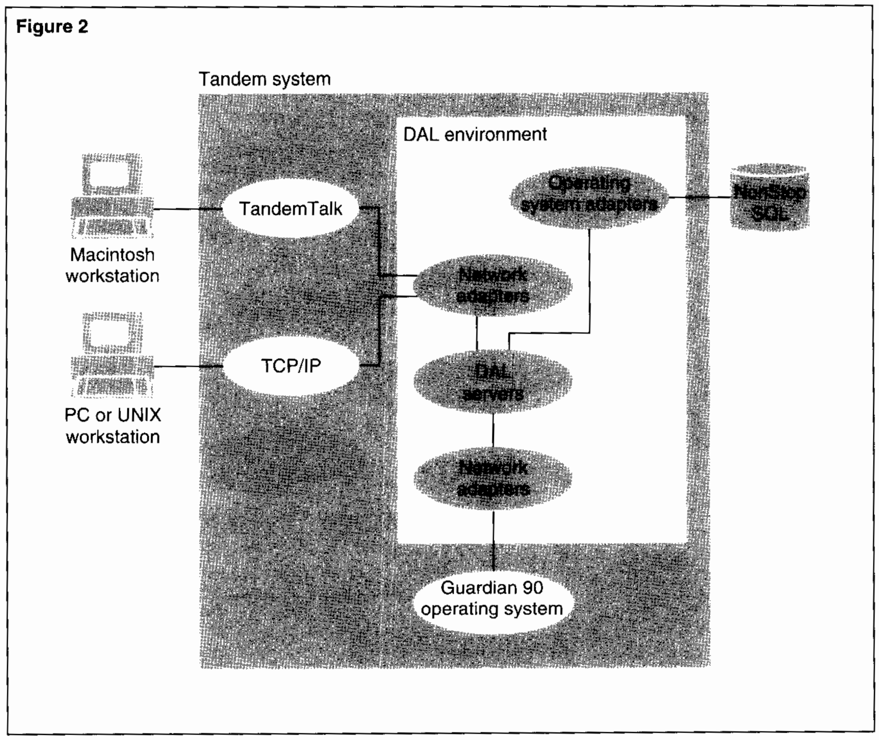

Figure 1.
Client/server architecture.

Figure 2.
DAL Server environment.

Figure 3.
The DAL Server in the Tandem environment.

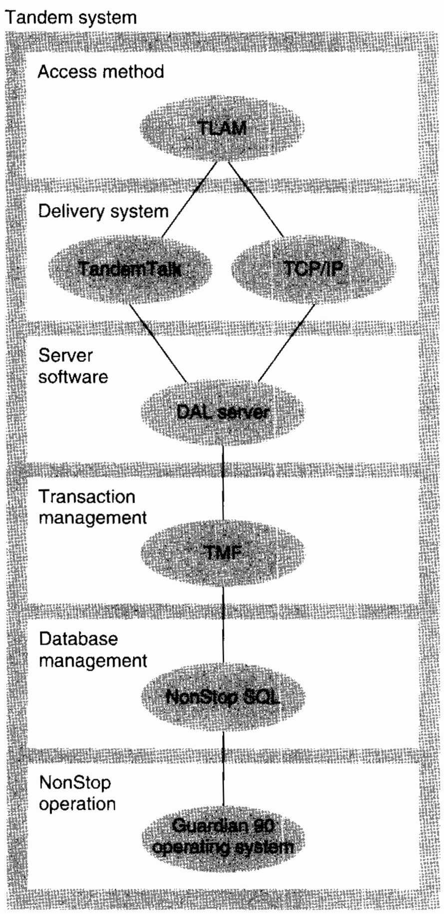

A Typical Inquiry

Using Figure 2 as an example, the following sequence shows how the DAL Server handles a typical inquiry:

1. The workstation (Windows or Macintosh) user requests information.
2. The client application formats the request as DAL statements and sends the request to the DAL Server.
3. The DAL Server converts the DAL statements to SQL statements and sends them to the NonStop SQL database management system (DBMS).
4. The DBMS processes the statements and returns the result to the DAL Server.
5. The DAL Server reformats the result of the inquiry and returns the information to the client application.
6. The requested information is now available to the user through the client application.
7. The user can display, manipulate, and, if the client application allows, transfer the information to another application.

The DAL Server in the Tandem Environment

As shown in Figure 3, the Guardian 90 operating system, file system, and disk process provide NonStop operation of the DAL Server subsystem. This software provides features to maintain data integrity, including continuous availability, structural integrity of files, protection of storage media, and data-sharing protection. The Tandem LAN Access Method (TLAM) communications subsystem, which provides local area network (LAN) connection methods to Tandem systems, is discussed later in this article.

Tandem's NonStop SQL, a fault-tolerant implementation of the industry-standard SQL, performs database management, and TMF software ensures the integrity of the database by rolling back transactions that cannot be completed and by restoring damaged host databases. NonStop SQL databases are fully protected by TMF whenever a DAL client performs an update. Data security is provided by protection mechanisms at several levels, including the system, the database, and specific database tables.

Stored Procedures

In addition to supporting dynamic SQL queries, the DAL Server supports stored procedures, edit files that contain groups of DAL statements. These procedures usually execute standardized or commonly used SQL queries. Either the user or the client application can execute these procedures by specifying the procedure name and optionally including one or more parameters.

Users also can employ stored DAL procedures together with the PATHSEND feature of Pathway, enabling the DAL Server to perform a transaction server function (relaying messages) as well as its usual database server function (relaying SQL statements). PATHSEND allows any program running on a NonStop system to invoke Pathway servers. Previously, Pathway servers could be invoked only by Pathway client programs.

The stored procedures, instead of containing files of SQL statements, can contain compiled object code of any language supported on Tandem systems. Figure 4 shows how the DAL Server can exercise any of these options.

DAL Clients

The DAL Server lets users work with SQL databases on Tandem NonStop systems while benefiting from the easy-to-use features of Macintosh or Windows-based personal computers. The AppleTalk protocols used by the DAL Server are part of the Macintosh operating system and are ready for use. Other personal computers can become AppleTalk-compatible by using plug-in adapter cards and third-party software. DOS (Windows) workstations can use TCP/IP protocols as the communications link to the Tandem system.

Workstation Software

Workstation-based application software can provide multiple session interfaces to the server processes on Tandem NonStop systems. One instance of the DAL Server supports one workstation. Communications functions are transparent to both the workstation-based client and the host-based server. Table 1 lists some off-the-shelf, DAL-compatible client applications.

Workstation software uses the AppleTalk protocols to carry EtherTalk (or LocalTalk) frames from Macintosh workstations connected to LANs. Users working at Macintosh workstations on a LAN can log on to the Tandem NonStop system. Applications at the Macintosh workstation use standard commands. For example, DAL-compatible applications can access NonStop SQL databases by standard DAL procedure calls and transfer data directly into spreadsheets, databases, or HyperCard stacks on the Macintosh.

Figure 4

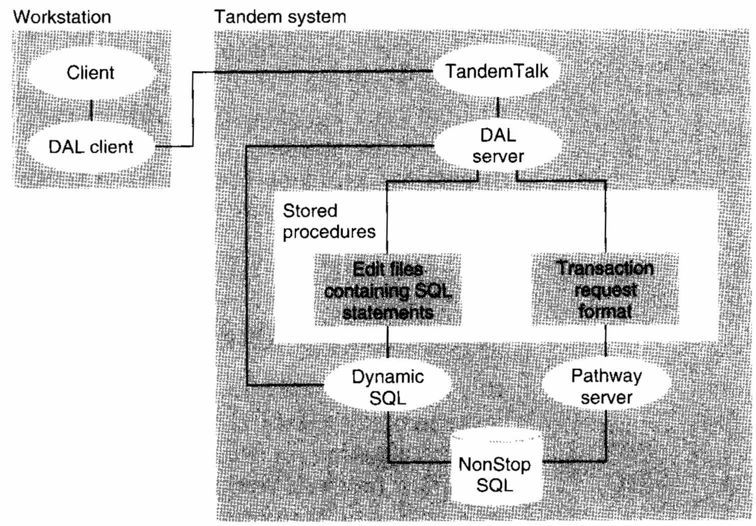

Figure 4.
The DAL Server can provide stored procedures as well as dynamic SQL queries.

Table 1.
A partial list of off-the-shelf client applications compatible with the DAL Server.

Application type	Application name
Query and reporting tools	
	ClearAccess
	DataPrism
	Graphical Query Language (GQL)
Spreadsheets	
	Microsoft Excel
	Wingz
Client/server application generators	
	4th Dimension
	HyperCard
	Omnis 7
	SuperCard
Geographical analysis tools	
	GeoQuery
	Tactician
Expert systems	
	NEXPERT Object

With AppleTalk Remote Access software, users can establish dial-up links to other Macintosh workstations or AppleTalk networks. Users at remote sites can access Tandem NonStop systems and Tandem host-resident applications like the DAL Server, the MacOffender product, or TACL™ (Tandem Advanced Command Language) software using standard modems and telephone lines.

Using the DAL Server

The following examples show how users have implemented the DAL Server to collect database information for analysis. In the first two examples, the DAL Server works with off-the-shelf applications to present database information in graphical format to aid analysts in supporting marketing and sales efforts. In the third example, the DAL Server works with a custom-written application to enable managers to define, analyze, and control projects, and to update database information.

Demographic Analysis

A large direct-mail marketing organization annually mails over five billion advertising inserts from its distribution centers. One of the key elements in securing maximum return on direct-mail advertising is getting the material into the hands of those persons most likely to buy the advertised product. The basic problem is threefold:

- Determining the characteristics of potential buyers.
- Finding where these buyers are located.
- Getting the advertising material to the buyers.

Detailed demographic data is readily available from both government and commercial sources. One can develop profiles of the characteristics of potential buyers by choosing from information recorded on these multiple databases.

Before the availability of the DAL Server, this organization had to maintain an extensive SQL programming effort to cross-reference the demographic profile of a potential buyer to the databases to determine the location of the most likely customers. This company has implemented the DAL Server to connect its Tandem NonStop SQL databases with standard off-the-shelf Macintosh client applications.

One such application is Tactician, a DAL-compatible desktop mapping package from Tactics International Ltd. Tactician displays the relative population density of a specified profile in an area, enabling the direct-mail advertiser to select areas with the best potential return. This user can now offer its direct-mail advertisers a wide choice of demographic selections to target households with similar characteristics (such as income, age, education).

This company also uses the DAL Server as a code generator. Its data processing department develops the SQL code for sets of standard transactions that are used to maintain the database and process a wide variety of common queries and transactions. After the SQL code has been developed and tested on the DAL Server, it is then transported to a number of different environments, including Macintosh and PC workstations, as well as 6530 terminals. Users can perform these transactions and queries from any type of terminal or workstation in the system.

Corporate Marketing Support

The marketing support department of a large manufacturing company supports field personnel by providing worldwide information about historical product shipments and customer applications for those products. Data from the company's marketing, financial, and manufacturing OLTP systems resides in NonStop SQL databases, and is routinely extracted to decision-support databases on non-OLTP systems. Moving data from multiple sources to a read-only storage area is called data warehousing. For a detailed discussion of the data warehouse concept, see the Viescas article in this issue of the *Tandem Systems Review*.

Analysts query the tables in the extract databases, as well their own departmental databases, to produce cross-functional views of where the company's products have been shipped and how they are being used. This information constitutes a critical marketing support resource.

Before the DAL Server was installed, analysts had to request support from the MIS department either to write an application to extract data or to produce ad hoc reports using the NonStop SQL Command Interface (SQLCI). Now, they can autonomously collect cross-functional database information and can use the GUI available on the Macintosh.

The analysts cooperatively query their databases using the Graphical Query Language (GQL) product on Apple Macintosh workstations. GQL connects to the DAL Server on the Tandem NonStop system through an Ungermann-Bass™ Access/One™ LAN, which permits open data connectivity among many environments. The analysts call GQL functions that extract data from the databases and put that data in a spreadsheet, where it can be flexibly and efficiently manipulated. They can then forward that information, in printed or electronic form, to the field personnel making the original request. They can also track the information requested by the field and the information they provided.

Project Planning and Control

The research and development division of an electronics manufacturing firm uses a custom-written 4th Dimension application to plan, schedule, and track the progress of 1,200 company projects. Using the DAL Server and the GUI provided by the application, managers load information about projects and personnel resources into a NonStop SQL database. They then schedule, prioritize, and assign resources to the projects, and they can also request predetermined or ad hoc graphical or text reports.

If resources are changed or moved among projects, the managers can analyze the impact of those changes on project schedules. They can then forecast the effects of the changes and reassign resources or reschedule milestones as necessary. After altering this information, the managers upload the changes to the NonStop system. Thus, all 100 managers using the system have access to the latest versions of schedules and resource allocations for all company projects.

The DAL Server supports both off-the-shelf and custom-written applications.

The client application functions are optimized to run on the NonStop SQL database, thereby maximizing the efficient use of host resources. Because the workstation software performs analyses and generates reports, those workloads are removed from the host.

The application uses stored DAL procedures, which consist of edit files that contain DAL statements. These edit files are stored on the Tandem NonStop system and are optimized for the specific function to be performed.

As client usage patterns have evolved, the application programmers have further optimized the database queries. They changed the SQL statements in the stored procedures without changing the client application or the steps users take to invoke the function.

For example, if managers select an inquiry by using a combination of keystrokes, the programmers can provide additional views to optimize the inquiry simply by changing the SQL statements in the DAL procedures. The managers do not change the way they perform the inquiry and do not have to reload the client software on the workstation.

Before the DAL Server was installed, the division managers had to choose between having a centralized Tandem database and having the simplicity and flexibility of a workstation tool. They decided that a GUI was essential and selected an off-the-shelf Macintosh application to perform project planning and control. The managers approved of the GUI, but the lack of a centralized database limited their capabilities and led to discrepancies when they compared information from different workstations.

With the DAL Server and the custom project-planning application, the managers can take advantage of both the NonStop SQL database and the workstation GUI. The company's unified project control system is now a significant factor in the efficient use of resources.

The TandemTalk Subsystem

The TandemTalk subsystem provides an interface that allows many different types of host-based servers to communicate with client applications running on Apple Macintosh computers connected to an attached Ethernet LAN. The TandemTalk subsystem includes the following components:

- The TandemTalk process provides an application programmatic interface (API) to the AppleTalk Data Stream Protocol (ADSP), as implemented in the Tandem environment.
- The Macintosh Access Method (MacAM) provides a standard, high-level Guardian 90 file-system interface that runs on top of the AppleTalk protocols.
- The Listener Process is a remote server creator that listens for connection requests from clients on the LAN and automatically initiates requested server processes on the Tandem host.
- The Tandem LAN Access Method (TLAM) communications subsystem interfaces are a host-resident, stand-alone software component that supports a number of LAN connection methods, such as TandemTalk, TCP/IP, and OSI (Open Systems Interconnection).

Users at workstations on a LAN can log on to the Tandem NonStop system through a TLAM communications subsystem. Several Tandem products, including the DAL Server, require access to many of the AppleTalk data communication protocols, which are provided by the TandemTalk data communications subsystem running on top of TLAM. Several Tandem products, including the DAL Server, use ADSP to support process-to-process communications between Macintosh and Guardian 90 applications.

A Tandem 5600 or 3613 Ethernet/IEEE 802.3 controller provides the I/O interface from the NonStop system to an Ethernet network using standard 10-megabit-per-second Ethernet connections. Ungermann-Bass Access/One intelligent wiring systems are recommended for bringing LANs and individual workstations into the Ethernet network. If users add a MaxTalk Interface Module to an Access/One system, they can use existing twisted-pair wiring along with a Farallon PhoneNET connector that plugs into a Macintosh LocalTalk connector. To connect the workstation directly to the Ethernet network, users can add an Apple EtherTalk adapter card to the Macintosh.

Conclusion

The Tandem DAL Server gives workstation users direct access to mainframe databases for decision-support and low-volume OLTP functions. It simplifies the task of viewing and integrating data from multiple hosts and databases and insulates workstation applications from network and database complexities. The DAL Server helps users combine the simplicity, versatility, and economy of workstation software with the fault tolerance, large capacity, and data security of Tandem NonStop systems.

References

Iem, M. and Kocher, T. 1992. Implementing Client/Server Applications Using Remote Server Call. *Tandem Systems Review*. Vol. 8, No. 3. Tandem Computers Incorporated. Part no. 89803.

Viescas, J. 1993. Online Transaction Processing. *Tandem Systems Review*. Vol. 9, No. 1. Tandem Computers Incorporated. Part no. 89804.

Acknowledgments

The authors wish to thank David Cooper, Dick Deranleau, Tom Hammer, Ed Hasenkamp, Steve Moriarty, Bob Unanski, and John Viescas for their assistance.

William Schlansky has been a technical writing consultant for 15 years. Previously he was a systems and telecommunications programmer for a computer manufacturer. At Tandem, he has edited a publication contrasting IMS with Pathway and TMF.

Joe Schrengohst is a free-lance technical writer with over 20 years experience in data communications. The projects he has worked on at Tandem include Multilan, TLAM, TCP/IP, RBS, MacOffender, and the DDE Gateway.

The RESPOND OLTP Business Management System for Manufacturing

Although business practices and computer technologies have advanced significantly in the past decade, manufacturing management applications have not kept pace, nor have they provided the level of availability required to support complex around-the-clock operations. Manufacturing systems typically have been batch-oriented and have supplied information that, at best, was obsolete. This has resulted in the growth of informal systems and high-cost, paper-driven expediting practices.

If manufacturers are to remain competitive, they need a new generation of manufacturing management applications. These applications must be based on today's practices and technologies and must replace the traditional batch systems used for material planning and factory management.

The online transaction processing (OLTP) implementation of the RESPOND business management system from Tandem™ provides users continuous access to current data. As a result, RESPOND users find they can move from a reactive environment to an exception-based environment that manages events as they occur. They can also benefit from higher end-user service levels, lower production costs, higher margins, and reduced time to market.

The first part of this article describes the general evolution of manufacturing planning systems; the second part explains the details of how RESPOND software works and tells why Tandem architecture makes the transition from batch to OLTP manufacturing applications both possible and desirable. It also shows the impact this transition makes on the operating environment and the information systems professionals who support it.

Traditional Planning Systems

Material requirements planning (MRP) is a way of developing a time-phased schedule for the procurement or manufacture of materials used in building an end product based on the master schedule, current supply, and demand. MRP provides a way to delay inventory and resource investments until they are absolutely required, and, therefore, has become a critical tool for managing inventory levels in investment-conscious environments.

Figure 1

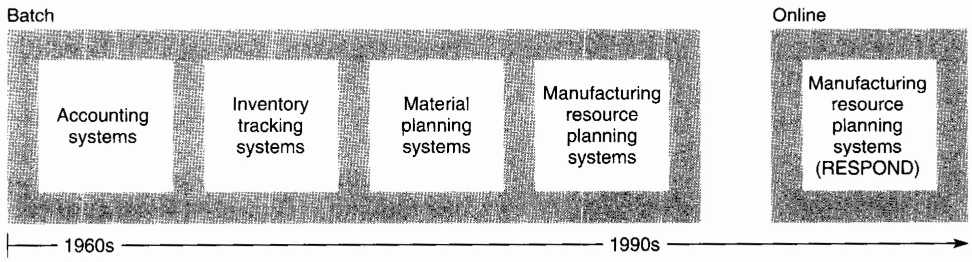

Figure 1.
The evolution of the RESPOND system from accounting systems.

Traditional MRP systems contain batch information processing. Figure 1 shows that these systems were initially based on accounting systems.

The information generated by these systems eventually influenced the management of manufacturing. Information used to track the value of inventory and to project cash flow for both purchased and manufactured items was fundamentally the same as that used to track and plan material. As a result, managers decided to extend the use of information from the accounting systems to inventory tracking and later to material planning. The techniques used by these material planning systems were formalized into a discipline called MRP I.

Since the constraints of batch processing required the complete rebuilding of all records, this new use of MRP universally replanned all materials, discarding all calculations that had been performed in previous MRP runs. As a consequence, processing was lengthy and could only be performed when a significant number of system resources were available. MRP I was the first step in the evolution of the RESPOND system.

As MRP I increased the users' access to data, management and operational personnel began using it to control and track costs. They used information that was generated by one system as input or feedback for other systems. These uses of MRP I lead to the development of manufacturing resource planning (MRP II), a set of practices to synchronize and maintain manufacturing and resource information.

At this point, they could extract accounting information from the data they used to track and maintain the status of the manufacturing process. Figure 2 shows this process.

Figure 2

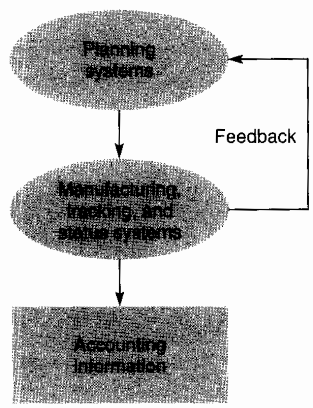

Figure 2.
MRP II systems provide accounting information as a byproduct of processing.

Soon increasing business requirements for data currency outstripped technology's ability to provide it. The main reason for this failing was the large amount of data that users needed when they were developing a plan to determine the number of items they needed to acquire or build. The complete redevelopment of this plan on a periodic basis, called *regenerative planning*, took more processing time than was available.

Figure 3.

As time since processing increased, data accuracy declined, and the risk of using inaccurate data increased. Knowing this, users relied less on the automated system.

Figure 3

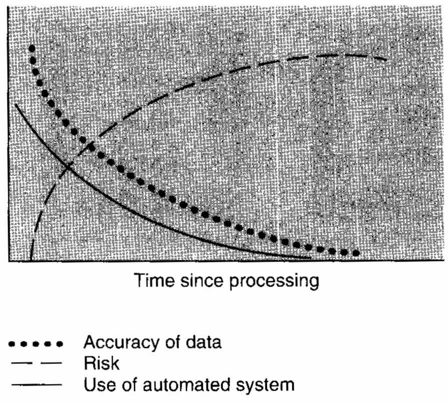

Figure 4

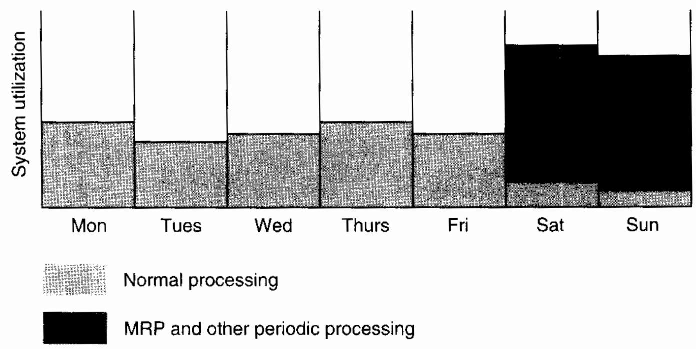

Figure 4.

MRP and other periodic processes caused an increase in system utilization.

Regenerative Planning

Before the advent of MRP II, users found that developing a material plan required a great deal of processing. To determine how many parts were needed within a specified time frame, material planners had to break down every finished item listed in the master production schedule.

Once material planners learned which assemblies, subassemblies, manufactured parts, and purchased parts were required to build a product, they used the material planning system to compute the status of every inventory item, including those items that were in various stages of procurement or manufacture. The system then compared the requirements with the status of each item and generated stacks of unwieldy reports. Material planning personnel used these reports to create, change, or expedite orders for the required items. This processing had a large impact on system availability.

Most early-generation systems prevented users from accessing the system when it was completely dedicated to MRP processing. As the company's products became more complex, the processing time spanned most of a weekend.

When users coupled MRP processing with periodic accounting processing, an interruption in computer operation could have a devastating effect on business. Furthermore, a loss of the system or data during an MRP run sometimes caused purchasing agents to make erroneous large-scale investments in inventory. They ordered quantities and items based on unreliable data. Expediters also wasted time and money by concentrating on the wrong items. In addition, the risk of making poor decisions increased dramatically with time. Figure 3 shows the impact of this risk on the accuracy of data.

The data processing also had an impact on system sizing. Figure 4 shows that the capacity planner had to configure hardware to provide enough cycles to complete this massive processing. MRP, combined with accounting and other periodic processes, had to run in a reasonable amount of time, such as a weekend. Printers had to accommodate large volumes of output, typically tens of thousands of pages in a week, for a moderately sized manufacturer.

Net Change Processing

The two factors of limited system capacity and high business risk led to the evolution of net change processing and MRP II systems. In a net change environment, found in the RESPOND system and many recent MRP II systems, one need only replan items that have experienced changes in supply or demand. This type of processing results in lower levels of periodic processing requirements but does not necessarily reduce the number of generated reports.

As systems became more powerful, planning applications could run daily in a net change mode. In many instances, this still meant that the system was unavailable to users for much of each night. Therefore, the system still contained outdated information, and users had to maintain informal, manual systems to make business decisions and sustain work flow. The accuracy and risk curves shown in Figure 3 remained the same except that they became daily instead of weekly.

Since the normal workload was now spread over every night of the week, a systems capacity planner could configure smaller systems to handle it. Figure 5 shows this environment. At this point, traditional MRP II processing and practices stopped evolving.

Manufacturing as a Transaction Environment

Any event in a manufacturing environment that has an impact on material or labor has the potential to affect the data in the system. The impact can be on supply, demand, availability, or status. The constantly changing nature of a manufacturing environment requires systems that can handle these changes online and in small units or transactions. The information in the systems can then closely match the requirements of the environment. Therefore, a state-of-the-art manufacturing business management system must be event-driven and can only be adequately designed around an OLTP system. The RESPOND business management system fills this need.

Figure 5

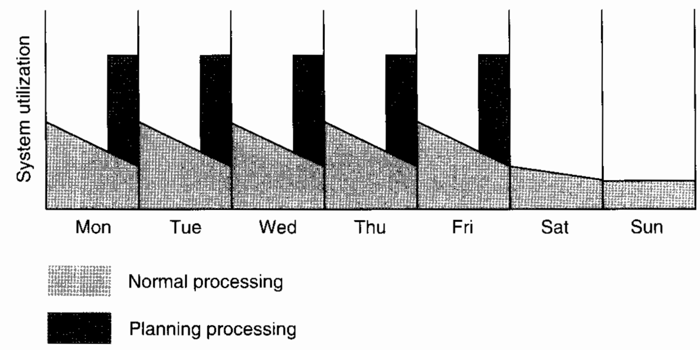

Figure 5.

In a daily net change environment, processing is spread over every night, and capacity planners can configure smaller systems.

The RESPOND System in the Tandem Environment

Today's manufacturers are decentralizing and giving their employees more direct control. Management is shifting this control to local plants, and, in some cases, to departments or work groups within those plants. To support the decentralized environment, users must be able to access information no matter where it is stored.

The RESPOND system uses Tandem's Enscribe database and remote partitioning to provide this distributed database environment. File partitioning is also used when multiple NonStop™ systems are joined in Tandem's Expand™ data communications network software.

Figure 6.
RESPOND is a modular application.

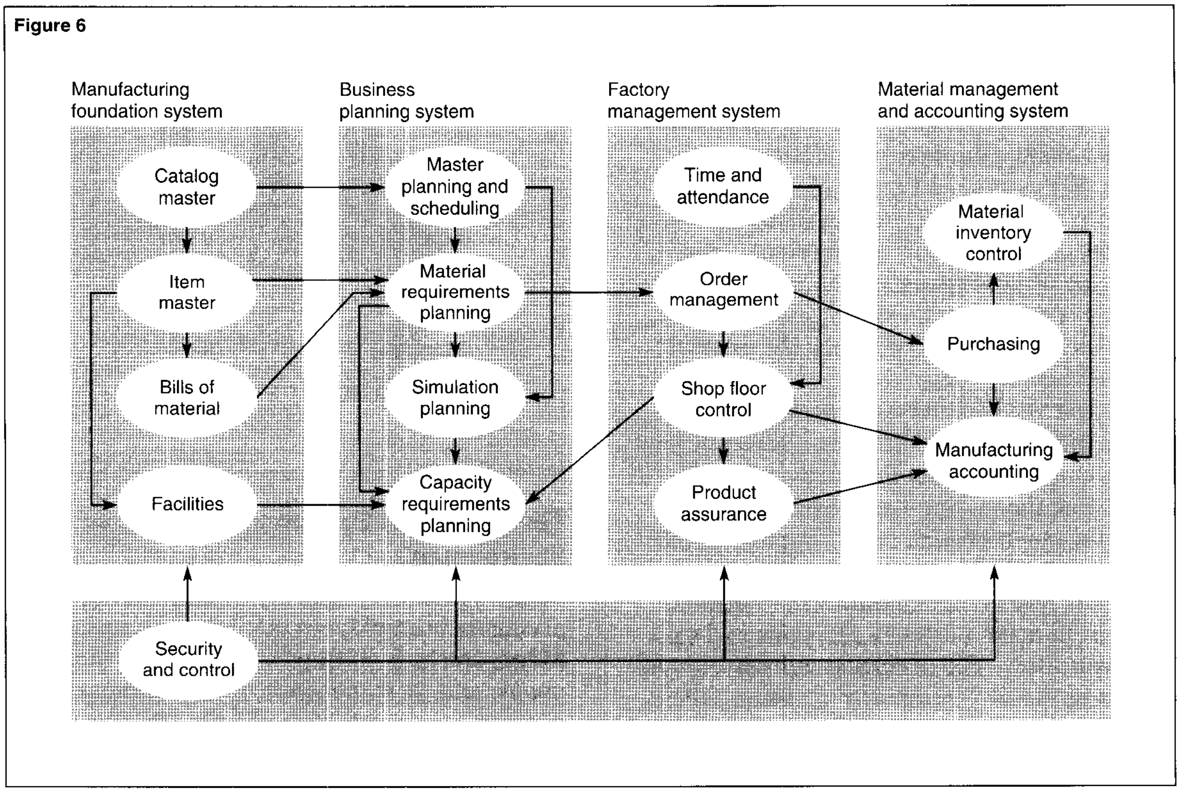

Tandem's Guardian™ 90 operating system supports the RESPOND application's fault tolerance as a NonStop system. This fault tolerance is necessary to maintain optimal performance with no downtime while system configurations are changing. Fault tolerance is also important to manufacturers who must ensure that the location, status, and process instructions in an MRP II system are always current.

In addition to networking capabilities and fault tolerance, manufacturers require a relational database to manage the large volume of data in their environments. For this reason, the initial versions of the RESPOND system incorporated databases created by Tandem's Enscribe database record manager. A newer version of the RESPOND system will incorporate databases created by Tandem's NonStop SQL relational database management system.

How the RESPOND System Differs from Typical Business Management Applications

While other business management applications rely on some level of periodic batch processing, Tandem's RESPOND system is a completely online application. It uses Tandem's OLTP features as well as other design concepts to provide an online, event-driven, net change operating environment.

Figure 7

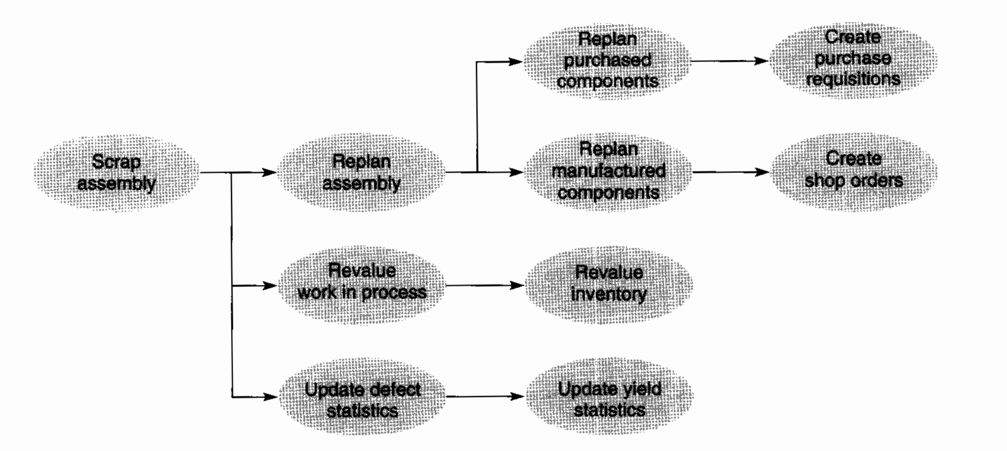

Figure 7.
One manufacturing transaction initiates multiple transactions.

The RESPOND application uses the Tandem Pathway transaction processing system. With Pathway, users can reconfigure requesters, servers, and terminals without stopping online service to end users.

Typical systems may contain modules that are not fully integrated. Users may have to input the same data more than once or may have to manipulate the data and pass it on to another user who inputs it again.

Figure 6 shows that the RESPOND system is a modular application that provides a complete *closed loop* requirements planning system. This means that any event that is reported to one module is automatically available to all of the other RESPOND modules. As a result, the RESPOND system requires no additional user input.

Typical systems may contain huge programs that perform several large concurrent tasks. The RESPOND system differs. It separates logical business processes into small sets of highly efficient transactions. Developers created the RESPOND application specifically for the Tandem environment. They designed it so that the system combines the requester-server architecture of Pathway with the advantages and strengths of Tandem's Guardian 90 operating system. This combination allows processes to communicate with each other by means of system-managed messages and provides users a more responsive application.

Manufacturing transactions often have a ripple effect in which one transaction generates several others. Figure 7 shows an example of an assembly that has been scrapped in manufacturing. The scrapped assembly causes the RESPOND system to generate a transaction that replans the assembly. This transaction, in turn, causes replanning of both manufactured and purchased components of the assembly. In addition, it may result in further transactions that initiate purchasing and manufacturing activities. At the same time, the system needs to calculate the value of work in process and update statistics that pertain to quality.

Table 1.
The primary consumers that RESPOND uses for continuous net change planning.

Module	Consumer (secondary transactions)
Master planning and scheduling	
	MPS net change processing
	MPS out-of-phase processing
Material requirements planning	
	Application interface processing
	MRP net change processing
	MRP out-of-phase processing
	Shop order and production schedule creation
Capacity requirements planning	
	CRP planned order generation
	CRP load generation
	CRP I/O generation

The user need not wait while the RESPOND system performs these secondary transactions. The application contains a special type of process called a *consumer*, which performs these processes online without user intervention.

Typical systems that use sequential batch processes usually prevent the user from accessing the data during batch processing. This is not true with the RESPOND system. Users can continue their work while the consumers perform secondary transactions.

Consumers are not Pathway processes, so they do not require terminal input. Consumers function continuously from *work queues*, which are files created by RESPOND transactions. These files contain pointers to information that requires processing by consumers. Work queues differ from batch inputs used in typical systems; instead of containing images of transactions, they contain pointers. The work queues allow continual updating in the system with no need for sequential batch processing of transactions.

Consumers and Net Change Planning

The RESPOND system's architecture differs most from that of traditional systems in the area of net change planning. In the past, this area has been served exclusively by batch processing. Whereas typical batch systems have a large sequential process for performing net change planning and reporting, the RESPOND system uses Tandem architecture and a number of consumers that may run in parallel. Each consumer splits the manufacturing system environment into logical business functions and performs a logical unit of the required work.

Table 1 shows the primary consumers that the RESPOND system uses for continuous net change planning. These consumers are *context-free*, which means that, in order to process a transaction, they do not need to remember data from a previous transaction.

The consumers only process one particular type of transaction. Traditional architectures operate differently; they provide a large program to process many kinds of transactions.

RESPOND processes are also *single-threaded*, processing only one transaction at a time. Thus, the system avoids processing conflicting information. The RESPOND application manages demand for consumer processing with a set of work queues.

Quiet States and Secondary Work Queues

The level of reported activity on the manufacturing floor dictates whether planning processes need to use system resources. Therefore, if activity is low and there is no work for these processes, the consumers remain in a quiet state. This means that the system consumes resources only when changes occur in the factory environment.

RESPOND consumers follow a hierarchy of planning processes by setting work queues, as required, for initial and subsequent processes. For example, suppose a consumer detects a condition that requires the generation of a message to material planning personnel. If the message relates to a condition that exists when supply and demand are not properly aligned in time or quantity, the consumer will write a record to a specific work queue. In this example, the queue is for a consumer that checks for out-of-phase conditions. The RESPOND system calls this particular consumer only when a potential exists for an out-of-phase condition.

Data Contention

The Tandem OLTP environment prevents multiple sources from updating the same data concurrently. If, for example, a user creates a material order at the same time that another user records a scrap transaction for the same item, the system must prevent the second transaction from overlaying the quantity-on-order data that it received from the first transaction. To prevent an inaccurate update to the database, the RESPOND system timestamps all data. Before it applies data to the database, it uses the timestamp to check for intervening transactions.

The RESPOND application uses Tandem's TMF™ (Transaction Monitoring Facility) software to protect databases from the effects of transaction failures. If the RESPOND system attempts a conflicting update, TMF simply aborts the transaction.

If the transaction occurs as a result of user interaction with the system, the RESPOND system displays a message to inform the user of values that have changed since the last display. The user can either redisplay or retry the update with confidence that the system will use current data.

Older technology systems lack this feature and must resort to fallback capabilities such as allowing inventory and order balances that are negative. These negative values are compensations for a system that is not completely online. In a RESPOND environment, an item with a quantity less than zero never exists.

Database Design and Data Currency

One shortcoming of traditional systems is that the same data resides in the system in different states of currency. These systems may contain similar copies of data in different parts of the application. For instance, planning reports may show an item as *required but not on order* while open purchase order and inventory statuses may show it as *on order* or *received and in inspection*. These discrepancies can occur when the system processes the data at different times or the user reads the information from outdated printed reports instead of from online data.

The RESPOND system uses the Tandem architecture to ensure an environment in which only one current view of the data exists. TMF verifies that this view is also internally consistent.

Figure 8

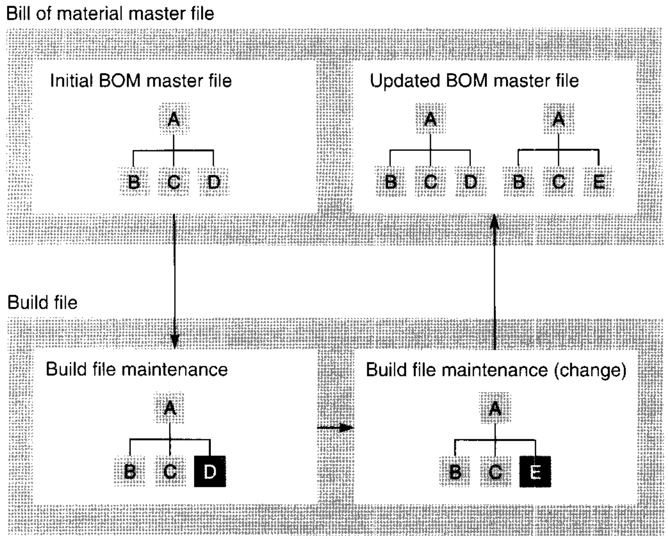

Figure 8.
The build file provides control of bill of material master file updates.

The Role of TMF

Transactions in a manufacturing environment tend to be extremely complex. For example, a purchased item receipt requires the system to check for shop floor shortages; update the purchase order, purchase order line item, and purchase order delivery status; post a record for accounts payable; change the value of inventory; and initiate a receiving document print process. The complex nature of this transaction requires a means for ensuring database consistency in the event of a process or data failure.

Systems with traditional batch architecture depend on error logs and transaction journals to manage exceptions. More recent systems contain code that simulates single logical transactions. The RESPOND system, on the other hand, takes advantage of TMF to provide a mechanism for grouping several transactions into a logical whole, and, therefore, guarantee data consistency.

Build Files

In a traditional manufacturing system environment, the batch planning processes deny users access to critical data while the processes are running. During potentially long computing cycles, users cannot make changes to product definitions and product process specifications. Therefore, these definitions and processes do not reflect up-to-date production methods or changing policies that tell how the item should be ordered or stored. This lack of access is unacceptable to users in an environment where planning systems have to run continuously.

The RESPOND system uses a *build file*, or temporary staging area, in which it keeps copies of master data while it is being updated by the user. The first diagram in Figure 8 shows the initial condition of the data. The arrows in the figure show the sequence of action. The second diagram depicts the data being copied to the build file. The third shows the changed data, and the fourth represents the production file after the change has been released to the build file. At this point, the production file contains both versions of data.

This four-change process starts when a user makes a copy of the production data and assigns it to the employee responsible for maintenance. The RESPOND system allows one employee at a time to change the same piece of data.

After the user completes the changes, a release transaction updates the changes on the production data either by overwriting the original data or by adding a new row to the table. The latter condition takes place when a user makes an engineering change that will be effective at a future date.

Moving the data from the build file to the production file, including deletion of the data in the build file, is a single online TMF transaction. The RESPOND system ensures that a release from the build file is, in itself, cause to evaluate the material plan for the item. If one of the planning consumers is using the data when a user or consumer retrieves it for a change, the time-stamping process prevents an incorrect update.

This approach also ensures that the RESPOND system simultaneously applies all data related to each change. For instance, if a user modifies a product's structure by adding one component and deleting another, the two changes that make up that change appear together and are applied to the production data as a single logical unit by a single TMF transaction.

This differs from some recent systems that perform these updates as two separate transactions without benefit of TMF. As a result, these systems may not consistently apply the entire change.

Manufacturing business operations benefit from the build file as well. When managers implement a controlled process in which users can make changes to product and process descriptions, they provide an opportunity to establish procedures for formal review by those designated to approve changes. Separate RESPOND transactions for each change give supervisors the ability to review and approve changes before they update the production data.

Time Keeping

Typically, time-keeping systems have been either batch or nonintegrated interactive systems whose sole purpose is to perform after-the-fact status reporting and to feed payroll systems. The RESPOND online system performs real-time checks of employee attendance records and validates the order condition before it accepts transactions from an employee working on an order. This prevents users from inputting spurious data that generates invalid orders.

The online implementation can also reduce lines at check-in stations and time clocks. The RESPOND system reads a job-start transaction that allows for a grace period and proceeds on the assumption that employees are on time. The RESPOND application also assumes that employees leave at their scheduled times, unless they log on to a job during overtime.

This implementation uses an efficient set of transactions, including several consumer processes to reconcile records and provide easy interfaces to payroll systems. As with all RESPOND consumers, the processes perform logical units of work and function independently through the use of work queues. This means that the RESPOND system only invokes those specific processes that are required for a transaction.

Online Purchasing

Purchasing departments typically make heavy use of printed reports. The nature of the work requires that users have quick access to large amounts of information. This information may include vendor histories, vendor qualifications, quotations, and price performance. As the purchasing department's need for information has expanded, cycle time reduction initiatives (such as Just-in-Time manufacturing) have reduced the amount of time for purchasing agents to react to requirements and to procure materials in the most cost-effective way.

The RESPOND system provides a complete online implementation of purchasing functions, including vendor review and selection. Purchasing agents can review quality, delivery, and qualification history. In the context of online planning, these purchasing activities immediately update the material plan and provide necessary communications back to material planners.

Online Cost Accounting

Cost accounting functions were one of the initial reasons for the creation of MRP systems. Traditionally, these functions have been processor-intensive, but the RESPOND system's completely online design provides these functions in a new way.

Online accounting functions eliminate month-end processing that makes strong demands on schedules and processing. By performing updates periodically throughout the accounting period, the RESPOND system removes the need for a steamroller approach to bring accounts to a current status at the end of the period. As a result, functions such as plant ledger updates to the general ledger and inventory balances no longer create processing bottlenecks.

The RESPOND system's inventory valuation, work-in-process valuation, standard cost roll up, and other functions are fully online. Typical systems require users to correlate data, make decisions, and re-enter information. The RESPOND system provides a set of functions that are driven by user-defined parameters. These functions replace the need for intermediate user processing. They are linked with a group of consumer processes that provide the required periodic close functions and also interface to external general ledger systems.

Online Business Simulation

The RESPOND system's online processes help the user to perform trial-fit calculations for schedule changes and other events. An example is a calculation that changes the master schedule, material availability, or product structure.

When traditional system users perform simulations, they often have to wait overnight for results. In many cases, they have only old data to use as a basis for comparison to the results of previous simulations. The RESPOND system's online architecture allows users to perform multiple simulations at the same time, using current data. Moreover, they can perform these simulations in a short time, such as during a meeting, with no effect on production data.

Traditional systems may require users to change production data so they can perform simulations. Users must then recreate the production data to implement the changes. The RESPOND system replaces this tedious process by providing users the capability to implement a new production plan with a single transaction.

Online Audit Trails

Users are sometimes uncomfortable with paperless systems because these systems do not provide a history of each transaction performed, the time the transaction took place, and the person responsible for it. The RESPOND system addresses this concern with online audit trails for all items that affect an item's cost or location. These audit trails give users all the information they need to track the source of costs and material changes.

The audit trails also provide users the required information to conform to some government contract regulations. In addition, TMF transactions give systems personnel the ability to recover the database if a catastrophic hardware failure occurs.

System Installation Considerations

With a traditional batch MRP system, the user can load the system, schedule batch jobs, and let the system execute. Little system monitoring is required.

Because access to the RESPOND system is online, systems professionals must become more involved in ensuring a successful implementation. This involvement extends from system capacity planning during pre-implementation activities to maintaining the mix of processes on each processor and maximizing resource availability throughout the application's life.

This additional involvement provides a better understanding between users and information services personnel. The results of this improved relationship include greater systems acceptance by users, better definition of required changes, and more cost-effective operations in information services.

Online Planning Evolution

In most RESPOND system implementations, the users do not immediately begin using planning functions in an online, continuous net change way. Instead, they start in a daily, weekly, or less frequent net change mode and evolve toward continuous replanning over a period of time. Whether one initially implements continuous replanning or less frequent net change planning has a large effect on how one configures the system and how one prioritizes processes.

In a normal manufacturing operation, where work usually takes place during daylight hours on weekdays, the system load may look like the one in Figure 9. When planning loads are superimposed over a weekend, capacity usually does not present a problem. This is more apparent in the RESPOND system than it is in a traditional batch environment because RESPOND planning processes are more efficient. They process only those items that require attention on an exception basis.

Figure 9

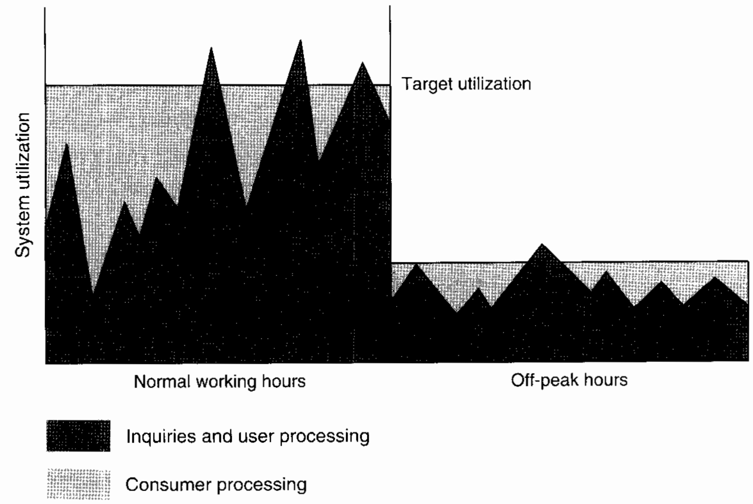

Conventional wisdom suggests an increase in processor size to compensate for the increase in net change processing frequency. However, the RESPOND system does not require additional processing power. Instead, it provides a way for system administrators to individually adjust frequency and priority of planning processes so the system can use the slack time between online transactions. This results in a daily system load profile like the weekly profile shown in Figure 9.

Figure 9.
RESPOND consumers use spare CPU cycles without impacting user response time.

Multiplant Capabilities and Tandem Distributed Processing

The RESPOND system is a multiplant application. A single logical copy of the RESPOND application can control multiple physical or logical manufacturing facilities. In addition, multiple plants may share a common master production schedule, while each logical plant controls inventory and demand.

If a system is to yield optimal performance, it must have personnel who are willing to invest time in planning complex manufacturing applications before they install the system. The RESPOND application's architecture makes it possible for a system manager to partition data by logical plant and to place data and system resources where they can provide the most efficient solution for the users. In some complex environments, the RESPOND system makes departmental computing a desirable and obtainable option.

The RESPOND application uses the Tandem Pathway transaction processing system to provide distributed processing capabilities. The Tandem Expand data communications network software may be used to link a single logical copy of the RESPOND software across multiple company locations. Each location maintains local control over its own data. On the other hand, the system runs in accordance with a central master production schedule and reports cost accounting results to a headquarters location.

Exception-Based Planning Environment

The RESPOND system environment is driven by exception rather than by routine. Those involved in daily manufacturing operations must approach the operating environment in a way different from their approach to a traditional batch environment. In turn, information services personnel must support the application differently.

Planners in traditional planning environments expect periodic reports that they can analyze for required actions. The planners can then generate transactions that update the system before the next run of the report. Even though this mode of operation may require significant resources to support routine operations, it is a predictable mode and it can be scheduled.

The RESPOND system, on the other hand, is driven by exceptions. Therefore, the workload of planners is reduced to responding to errors and inconsistencies. Planners no longer need large numbers of reports to do lengthy analysis. Instead, they plan daily or even more frequently. Their work cycles include checking for online action messages, using online transactions to perform analysis, and then performing corrective action, again using online transactions. They do not need printed reports or information services support.

Changing Job Descriptions

With an online system, job descriptions change significantly. Higher-level analysis and decision-making activities replace lower-level activities such as report analysis, expediting, and part search activities. This means that a smaller core of higher-level individuals is making decisions. Therefore, the RESPOND application requires less system support than traditional systems.

Impact on Information Systems Organization

An online manufacturing system requires different levels of information systems support and planning than those required in a batch environment. Most important, online reports require minimal print management and distribution. In many environments, systems personnel implement the RESPOND application so that it generates little or no paper.

On the other hand, those who support the RESPOND system must have basic manufacturing knowledge and receive training on the details of the RESPOND system's architecture. These requirements provide them the background to develop custom queries and assist users in formulating ad hoc reports.

Operations Management Considerations

The RESPOND system and the Tandem Pathway transaction processing system work together to give systems personnel the ability to set operational parameters so that operators are no longer needed to run the system. Instead, these operators can concentrate on making the application perform as efficiently as possible in the hardware environment.

Users as well as operators can view the status of the RESPOND system's material requirements planning consumer work queues online within the application. This immediate feedback provides them a quick gauge of how responsive the application is at any point in time.

In addition, operators can use Tandem's Measure™ system performance measurement tool. This tool provides a record of system activity that can be analyzed, and the results can be used to optimize applications on the Tandem system. Operators can also use other tools, available from Tandem and third parties, to optimize application performance.

Conclusion

Traditional manufacturing management applications isolate the manufacturing planning function from the daily operations of a company. This is a natural result of the time lag, inherent in batch systems, in delivering critical information to users. The RESPOND business management system integrates planning into ongoing operations and control. It has thus transformed the reactive, paper-driven environment into one that is exception-based, online, and always current.

The impact of the RESPOND system on the way information services supports manufacturing personnel is a positive one. The RESPOND application requires fewer low-level support personnel than traditional batch systems. Instead, relatively few high-level decision makers can support the application. On the other hand, information services personnel must become more involved in planning and implementation, acquire a basic manufacturing knowledge, and learn the details of the RESPOND system's architecture.

Acknowledgments
The authors gratefully acknowledge the technical review and input provided by Peter Oswald.

William R. Bronson is one of the original architects of the RESPOND system. He has been associated with the product since its inception and is currently guiding the product's future direction. Prior to joining Tandem, William was the vice president of product development for a computer services and software company, specializing in business management systems for the manufacturing industry.

Howard I. Bolling, CPIM, has been integral to the RESPOND system's implementation and support since 1985. He has 20 years of manufacturing and manufacturing systems experience in the fields of mechanical engineering, manufacturing engineering, and production.

Technical Information & Education

CD Read Documentation

Tandem CD Read provides a complete set of Guardian 90 operating system software manuals on a single CD-ROM disc. Periodically, CD Read subscribers receive a new disc containing the latest set of Guardian software manuals.

CD Read, Version C30_08_02

January 1993

Version C30_08_02 of CD Read, now available, includes a number of enhancements to the viewing software. The enhancements include printing capabilities for A4 size paper, Macintosh System 7 compatibility, default to manual title instead of part number in global searches, and display of both absolute page numbers and chapter-page numbers.

Tandem Education

The following paragraphs provide highlights of the latest education courses offered by Tandem. To sign up for a class or to order an independent study program (ISP), users should call 1-800-621-9198. Full descriptions of all available courses and ISPs appear in the *Tandem Education Course Catalog* and on InfoWay.

Managing Operations in a Tandem Environment

This three-day course offers a solid foundation in management responsibilities for Tandem systems in homogeneous, heterogeneous, and distributed systems environments. The main objective of the course is to train students in establishing processes and procedures, service-level agreements, tool utilization, and organizational teamwork for effective management of the Tandem system environment.

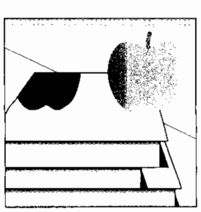

The Technical Information and Education department is an annotated list of new Tandem education courses, consulting and information services, and books Tandem is offering its users.

Migrating to the D10 Release of the Guardian Operating System

This four-day course provides an introduction to Guardian 90 system concepts unique to, or changed by, the D-series software releases. The course enables students to determine what modifications, if any, their systems may require to migrate to the D-series software.

Object-Oriented Analysis and Design

In this five-day course students gain a solid foundation in object-oriented analysis and design. They also learn the benefits of the object-oriented analysis and design paradigm and how the paradigm differs from traditional software design methods.

Object-Oriented Programming with C++

This five-day course teaches students how to use C++ to implement object-oriented programs. Students also learn the elements of object-oriented analysis and design as applied to C++ programming.

Pathway Application Programming I

In this combination classroom-and-lab course students learn how to develop Pathway applications quickly and efficiently. The five-day course provides hands-on practice on the Tandem system and thus enables students to demonstrate practical application programming skills in the Pathway environment. Students work with either Enscribe or NonStop SQL databases, or both.

Pathway Application Programming II

This combination classroom-and-lab course teaches students how to develop Pathway applications using advanced features such as intelligent device support, PATHSEND, and unsolicited-message processing. The four-day course provides hands-on practice on the Tandem system and thus enables students to demonstrate advanced application programming skills in the Pathway environment. Students work with either Enscribe or NonStop SQL databases, or both.

Programming Client/Server Applications with POET

This four-day workshop teaches how to use Tandem's Pathway Open Environment Toolkit (POET) product to help generate client/server applications. Students learn to identify each POET component, its function, and its relationship to other components. Using POET and CASE:W, students also create a Windows program that communicates with an existing Pathway server.

Windows 3.1 for End Users

This course provides a solid foundation for new computer users who plan to use Microsoft Windows software and applications based on Microsoft Windows software. The course lasts one full day.

Tandem Professional Services

Tandem Professional Services is a program in which trained Tandem experts deliver technical consulting services at the user site. Tandem now offers a number of standardized services and will announce additional ones in this department of the *Tandem Systems Review* as they become available. Following are brief descriptions of new professional services offered by Tandem. For more information, users should contact their local Tandem representative.

Capacity Planning Service

This service produces a capacity plan that predicts the number of processors and logical disk volumes needed in a well-balanced system to satisfy the user's future business requirements. The Tandem service provider works with the user to gather information about the user's application, to obtain Measure data from the system to be modeled, and to identify peak operating periods. On the basis of this information, and using the Tandem Capacity Model (TCM), Tandem designs and validates a capacity model for the system. Tandem then presents capacity forecasts and recommendations to the user.

Performance Review and Analysis Service

This service provides an in-depth review and analysis of the user's system performance. The service includes two levels of performance analysis. The first provides the user with recommendations for improving system balance and resource utilization. The second analyzes key subsystems for performance problems that are keeping the system from achieving its optimum performance level. The result of the analysis is a report of the findings with recommendations for further action.

Application Design

This service is aimed at users who are designing applications to run under Guardian 90 systems. The service assists the user in designing a high-quality, high-performing application that meets the user's business requirements.

In the course of the service, Tandem identifies the user's main business transactions. For each of these transactions, Tandem completes a functional script, interprocess message definition, performance objective, and recommendation for implementing the transaction. Additional areas covered during delivery of the service include requester structuring, Transaction Monitoring Facility (TMF) boundary definition, message design, and server packaging.

NonStop NET/MASTER Implementation Service

This service assists Tandem users in configuring and implementing NonStop NET/MASTER software on their NonStop production systems. Security issues and network considerations in the Tandem environment are discussed and implemented. Operator training specific to the newly installed environment is included.

The service begins with an evaluation of the user's installation, configuration, and customization objectives. The Tandem service provider then assists the user in installing and configuring a customized version of NonStop NET/MASTER software based on the user's requirements. After guiding the user's staff through the program's operational procedures, the service concludes with a review of NonStop NET/MASTER maintenance and management procedures.

Books

The following are brief descriptions of three books on C++ programming that may be of interest to Tandem users.

C++ Primer
2nd edition
Stanley B. Lippman
Addison-Wesley, 1991

This book covers the complete C++ language, including templates, with an abbreviated section on exception handling. The book has two chapters on object-oriented design and includes an introduction to the C++ I/O streams library, as well as an appendix that contrasts C++ with ANSI C. The current (3.0) level of C++ is described throughout the book and the changes from the 2.0 level are specified in an appendix.

The C++ Programming Language
2nd edition
Bjarne Stroustrup
Addison-Wesley, 1991

This is an updated reference guide for the C++ language. It covers the language in its entirety and thoroughly describes both templates and exception handling. The book has a chapter on designing C++ libraries, a section on run-time type of information, and a detailed description of C++ I/O streams. The C++ reference manual is provided as an appendix.

Advanced C++, Programming Styles and Idioms
James O. Coplien
Addison-Wesley, 1992

This book is intended for readers who already have a working knowledge of C++. It presents the material in three phases: (1) creating safe, efficient C++ data types; (2) object-oriented design and reuse; and (3) programming using meta-objects. The book includes an appendix that contrasts C++ with C; most often, however, the comparison is with pre-ANSI C.

TandemSystemsReview*Index*

The *Tandem Journal* became the *Tandem Systems Review* in February 1985. Four issues of the *Tandem Journal* were published:

Volume 1, No. 1	Fall 1983
Volume 2, No. 1	Winter 1984
Volume 2, No. 2	Spring 1984
Volume 2, No. 3	Summer 1984

As of this issue, 20 issues of the *Tandem Systems Review* have been published:

Volume 1, No. 1	Feb. 1985	Volume 6, No. 1	March 1990
Volume 1, No. 2	June 1985	Volume 6, No. 2	Oct. 1990
Volume 2, No. 1	Feb. 1986	Volume 7, No. 1	April 1990
Volume 2, No. 2	June 1986	Volume 7, No. 2	Oct. 1991
Volume 2, No. 3	Dec. 1986	Volume 8, No. 1	Spring 1992
Volume 3, No. 1	March 1987	Volume 8, No. 2	Summer 1992
Volume 3, No. 2	Aug. 1987	Volume 8, No. 3	Fall 1992
Volume 4, No. 1	Feb. 1988	Volume 9, No. 1	Winter 1993
Volume 4, No. 2	July 1988		
Volume 4, No. 3	Oct. 1988		
Volume 5, No. 1	April 1989		
Volume 5, No. 2	Sept. 1989		

The articles published in all 24 issues are arranged by subject below. (*Tandem Journal* is abbreviated as TJ and *Tandem Systems Review* as TSR.) A second index, arranged by product, is also provided.

Index by Subject

Article title	Author(s)	Publication	Volume, Issue	Publication date	Part number
APPLICATION DEVELOPMENT AND LANGUAGES *(cont.)*					
State-of-the-Art C Compiler	E. Kit	TSR	2,2	June 1986	83937
TACL, Tandem's New Extensible Command Language	J. Campbell, R. Glascock	TSR	2,1	Feb. 1986	83936
Tandem's New COBOL85	D. Nelson	TSR	2,1	Feb. 1986	83936
The DAL Server: Client/Server Access to Tandem Databases	W. Schlansky, J. Schrengohst	TSR	9,1	Winter 1993	89804
The ENABLE Program Generator for Multifile Applications	B. Chapman, J. Zimmerman	TSR	1,1	Feb. 1985	83934
TMF and the Multi-Threaded Requester	T. Lemberger	TJ	1,1	Fall 1983	83930
Writing a Command Interpreter	D. Wong	TSR	1,2	June 1985	83935
CLIENT/SERVER					
An Overview of Client/Server Computing on Tandem Systems	H. Cooperstein	TSR	8,3	Fall 1992	89803
Designing Client/Server Applications for OLTP on Guardian 90 Systems	W. Culman	TSR	8,3	Fall 1992	89803
Gateways to NonStop SQL	D. Slutz	TSR	6,2	Oct. 1990	46987
Implementing Client/Server Using RSC	M. Iem, T. Kocher	TSR	8,3	Fall 1992	89803
The DAL Server: Client/Server Access to Tandem Databases	W. Schlansky, J. Schrengohst	TSR	9,1	Winter 1993	89804
CUSTOMER SUPPORT					
Customer Information Service	J. Massucco	TSR	3,1	March 1987	83939
Remote Support Strategy	J. Eddy	TSR	3,1	March 1987	83939
Tandem's Software Support Plan	R. Baker, D. McEvoy	TSR	3,1	March 1987	83939
DATA COMMUNICATIONS					
An Overview of SNAX/CDF	M. Turner	TSR	5,2	Sept. 1989	28152
A SNAX Passthrough Tutorial	D. Kirk	TJ	2,2	Spring 1984	83932
Changes in FOX	N. Donde	TSR	1,2	June 1985	83935
Connecting Terminals and Workstations to Guardian 90 Systems	E. Siegel	TSR	8,2	Summer 1992	69848
Introduction to MULTILAN	A. Coyle	TSR	4,1	Feb. 1988	11078
Overview of the MULTILAN Server	A. Rowe	TSR	4,1	Feb. 1988	11078
SNAX/APC: Tandem's New SNA Software for Distributed Processing	B. Grantham	TSR	3,1	March 1987	83939
SNAX/HLS: An Overview	S. Saltwick	TSR	1,2	June 1985	83935
TLAM: A Connectivity Option for Expand	K. MacKenzie	TSR	7,1	April 1991	46988
Using the MULTILAN Application Interfaces	M. Berg, A. Rowe	TSR	4,1	Feb. 1988	11078

Article title	Author(s)	Publication	Volume, Issue	Publication date	Part number
DATA MANAGEMENT					
A Comparison of the B00 DP1 and DP2 Disc Processes	T. Schachter	TSR	1,2	June 1985	83935
An Overview of NonStop SQL Release 2	M. Pong	TSR	6,2	Oct. 1990	46987
Batch Processing in Online Enterprise Computing	T. Keefauver	TSR	6,2	Oct. 1990	46987
Concurrency Control Aspects of Transaction Design	W. Senf	TSR	6,1	March 1990	32968
Converting Database Files from ENSCRIBE to NonStop SQL	W. Weikel	TSR	6,1	March 1990	32986
DP1-DP2 File Conversion: An Overview	J. Tate	TSR	2,1	Feb. 1986	83936
Determining FCP Conversion Time	J. Tate	TSR	2,1	Feb. 1986	83936
DP2's Efficient Use of Cache	T. Schachter	TSR	1,2	June 1985	83935
DP2 Highlights	K. Carlyle, L. McGowan	TSR	1,2	June 1985	83935
DP2 Key-sequenced Files	T. Schachter	TSR	1,2	June 1985	83935
Gateways to NonStop SQL	D. Slutz	TSR	6,2	Oct. 1990	46987
High-Performance SQL Through Low-Level System Integration	A. Borr	TSR	4,2	July 1988	13693
Improvements in TMF	T. Lemberger	TSR	1,2	June 1985	83935
Online Information Processing	J. Viescas	TSR	9,1	Winter 1993	89804
Online Reorganization of Key-Sequenced Tables and Files	G. Smith	TSR	6,2	Oct. 1990	46987
Optimizing Batch Performance	T. Keefauver	TSR	5,2	Sept. 1989	28152
Overview of NonStop SQL	H. Cohen	TSR	4,2	July 1988	13693
Parallelism in NonStop SQL Release 2	M. Moore, A. Sodhi	TSR	6,2	Oct. 1990	46987
NetBatch: Managing Batch Processing on Tandem Systems	D. Wakashige	TSR	5,1	April 1989	18662
NetBatch-Plus: Structuring the Batch Environment	G. Earle, D. Wakashige	TSR	6,1	March 1990	32986
NonStop SQL: The Single Database Solution	J. Cassidy, T. Kocher	TSR	5,2	Sept. 1989	28152
NonStop SQL Data Dictionary	R. Holbrook, D. Tsou	TSR	4,2	July 1988	13693
NonStop SQL Optimizer: Basic Concepts	M. Pong	TSR	4,2	July 1988	13693
NonStop SQL Optimizer: Query Optimization and User Influence	M. Pong	TSR	4,2	July 1988	13693
NonStop SQL Reliability	C. Fenner	TSR	4,2	July 1988	13693
The NonStop SQL Release 2 Benchmark	S. Englert, J. Gray, T. Kocher, P. Shah	TSR	6,2	Oct. 1990	46987
The Outer Join in NonStop SQL	J. Vaishnav	TSR	6,2	Oct. 1990	46987
The Relational Data Base Management Solution	G. Ow	TJ	2,1	Winter 1984	83931
Tandem's NonStop SQL Benchmark	Tandem Performance Group	TSR	4,1	Feb. 1988	11078
The TRANSFER Delivery System for Distributed Applications	S. Van Pelt	TJ	2,2	Spring 1984	83932
TMF Autorollback: A New Recovery Feature	M. Pong	TSR	1,1	Feb. 1985	83934
MANUALS/COURSES					
B00 Software Manuals	S. Olds	TSR	1,2	June 1985	83935
C00 Software Manuals	E. Levi	TSR	4,1	Feb. 1988	11078
New Software Courses	M. Janow	TSR	1,2	June 1985	83935
New Software Courses	J. Limper	TSR	4,1	Feb. 1988	11078
Subscription Policy for Software Manuals	T. McSweeney	TSR	2,1	Feb. 1986	83936
Tandem's New Products	C. Robinson	TSR	2,1	Feb. 1986	83936
Tandem's New Products	C. Robinson	TSR	2,2	June 1986	83937

Article title	Author(s)	Publication	Volume, Issue	Publication date	Part number
OPERATING SYSTEMS					
Highlights of the B00 Software Release	K. Coughlin, R. Montevaldo	TSR	1,2	June 1985	83935
Increased Code Space	A. Jordan	TSR	1,2	June 1985	83935
Managing System Time Under GUARDIAN 90	E. Nellen	TSR	2,1	Feb. 1986	83936
New GUARDIAN 90 Time-keeping Facilities	E. Nellen	TSR	1,2	June 1985	83935
New Process-timing Features	S. Sharma	TSR	1,2	June 1985	83935
NonStop II Memory Organization and Extended Addressing	D. Thomas	TJ	1,1	Fall 1983	83930
Overview of the C00 Release	L. Marks	TSR	4,1	Feb. 1988	11078
Overview of the NonStop-UX Operating System for the Integrity S2	P. Norwood	TSR	7,1	April 1991	46988
Robustness to Crash in a Distributed Data Base: A Nonshared-memory Approach	A. Borr	TSR	1,2	June 1985	83935
The GUARDIAN Message System and How to Design for It	M. Chandra	TSR	1,1	Feb. 1985	83935
The Tandem Global Update Protocol	R. Carr	TSR	1,2	June 1985	83935
PERFORMANCE AND CAPACITY PLANNING					
A Performance Retrospective	P. Oleinick, P. Shah	TSR	2,3	Dec. 1986	83938
Buffering for Better Application Performance	R. Mattran	TSR	2,1	Feb. 1986	83936
Capacity Planning Concepts	R. Evans	TSR	2,3	Dec. 1986	83938
Capacity Planning With TCM	W. Highleyman	TSR	7,2	Oct. 1991	65248
C00 TMDS Performance	J. Mead	TSR	4,1	Feb. 1988	11078
Credit-authorization Benchmark for High Performance and Linear Growth	T. Chmiel, T. Houy	TSR	2,1	Feb. 1986	83936
Debugging Accelerated Programs on TNS/R Systems	D. Cressler	TSR	8,1	Spring 1992	65250
DP2 Performance	J. Enright	TSR	1,2	June 1985	83935
Estimating Host Response Time in a Tandem System	H. Horwitz	TSR	4,3	Oct. 1988	15748
FASTSORT: An External Sort Using Parallel Processing	J. Gray, M. Stewart, A. Tsukerman, S. Uren, B. Vaughan	TSR	2,3	Dec. 1986	83938
Getting Optimum Performance from Tandem Tape Systems	A. Khatri	TSR	2,3	Dec. 1986	83938
How to Set Up a Performance Data Base with MEASURE and ENFORM	M. King	TSR	2,3	Dec. 1986	83938
Improved Performance for BACKUP2 and RESTORE2	A. Khatri, M. McCline	TSR	1,2	June 1985	83935
Improving Performance on TNS/R Systems With the Accelerator	M. Blanchet	TSR	8,1	Spring 1992	65250
MEASURE: Tandem's New Performance Measurement Tool	D. Dennison	TSR	2,3	Dec. 1986	83938
Measuring DSM Event Management Performance	M. Stockton	TSR	8,1	Spring 1992	65250
Message System Performance Enhancements	D. Kinkade	TSR	2,3	Dec. 1986	83938
Message System Performance Tests	S. Uren	TSR	2,3	Dec. 1986	83938
Network Design Considerations	J. Evjen	TSR	5,2	Sept. 1989	28152
NonStop VLX Performance	J. Enright	TSR	2,3	Dec. 1986	83938
Optimizing Sequential Processing on the Tandem System	R. Welsh	TJ	2,3	Summer 1984	83933
Pathway TCP Enhancements for Application Run-Time Support	R. Vannucci	TSR	7,1	April 1991	46988

Article title	Author(s)	Publication	Volume, Issue	Publication date	Part number
PERFORMANCE AND CAPACITY PLANNING *(cont.)*					
Performance Benefits of Parallel Query Execution and Mixed Workload Support in NonStop SQL Release 2	S. Englert, J. Gray	TSR	6,2	Oct. 1990	46987
Performance Considerations for Application Processes	R. Glasstone	TSR	2,3	Dec. 1986	83938
Performance Measurements of an ATM Network Application	N. Cabell, D. Mackie	TSR	2,3	Dec. 1986	83938
Predicting Response Time in On-line Transaction Processing Systems	A. Khatri	TSR	2,2	June 1986	83937
The 6600 and TCC6820 Communications Controllers: A Performance Comparison	P. Beadles	TSR	2,3	Dec. 1986	83938
The ENCORE Stress Test Generator for On-line Transaction Processing Applications	S. Kosinski	TJ	2,1	Winter 1984	83931
The PATHWAY TCP: Performance and Tuning	J. Vatz	TSR	1,1	Feb. 1985	83934
The Performance Characteristics of Tandem NonStop Systems	J. Day	TJ	1,1	Fall 1983	83930
Sizing Cache for Applications that Use B-series DP1 and TMF	P. Shah	TSR	2,2	June 1986	83937
Sizing the Spooler Collector Data File	H. Norman	TSR	4,1	Feb. 1988	11978
Tandem's 5200 Optical Storage Facility: Performance and Optimization Considerations	S. Coleman	TSR	5,1	April 1989	18662
Tandem's Approach to Fault Tolerance	B. Ball, W. Bartlett, S. Thompson	TSR	4,1	Feb. 1988	11078
Understanding PATHWAY Statistics	R. Wong	TJ	2,2	Spring 1984	83932
PERIPHERALS					
5120 Tape Subsystem Recording Technology	W. Phillips	TSR	3,2	Aug. 1987	83940
An Introduction to DYNAMITE Workstation Host Integration	S. Kosinski	TSR	1,2	June 1985	83935
Data-Encoding Technology Used in the XL8 Storage Facility	D. S. Ng	TSR	2,2	June 1986	83937
Data-Window Phase-Margin Analysis	A. Painter, H. Pham, H. Thomas	TSR	2,2	June 1986	83937
Introducing the 3207 Tape Controller	S. Chandran	TSR	1,2	June 1985	83935
Peripheral Device Interfaces	J. Blakkan	TSR	3,2	Aug. 1987	83940
Plated Media Technology Used in the XL8 Storage Facility	D.S. Ng	TSR	2,2	June 1986	83937
Streaming Tape Drives	J. Blakkan	TSR	3,2	Aug. 1987	83940
Terminal Selection	E. Siegel	TSR	8,2	Summer 1992	69848
The 5200 Optical Storage Facility: A Hardware Perspective	A. Patel	TSR	5,1	April 1989	18662
The 6100 Communications Subsystem: A New Architecture	R. Smith	TJ	2,1	Winter 1984	83931
The 6600 and TCC6820 Communications Controllers: A Performance Comparison	P. Beadles	TSR	2,3	Dec. 1986	83938
The DYNAMITE Workstation: An Overview	G. Smith	TSR	1,2	June 1985	83935
The Model 6VI Voice Input Option: Its Design and Implementation	B. Huggett	TJ	2,3	Summer 1984	83933
The Role of Optical Storage in Information Processing	L. Sabaroff	TSR	3,2	Aug. 1987	83940
The V8 Disc Storage Facility: Setting a New Standard for On-line Disc Storage	M. Whiteman	TSR	1,2	June 1985	83935

Article title	Author(s)	Publication	Volume, Issue	Publication date	Part number
PROCESSORS					
Fault Tolerance in the NonStop Cyclone System	S. Chan, R. Jardine	TSR	7,1	April 1991	46988
NonStop CLX: Optimized for Distributed On-Line Transaction Processing	D. Lenoski	TSR	5,1	April 1989	18662
NonStop VLX Hardware Design	M. Brown	TSR	2,3	Dec. 1986	83938
Overview of Tandem NonStop Series/RISC Systems	L. Faby, R. Mateosian	TSR	8,1	Spring 1992	65250
The High-Performance NonStop TXP Processor Transaction Processing	W. Bartlett, T. Houy, D. Meyer	TJ	2,1	Winter 1984	83931
The NonStop TXP Processor: A Powerful Design for On-line Transaction Processing	P. Oleinick	TJ	2,3	Summer 1984	83933
The VLX: A Design for Serviceability	J. Allen, R. Boyle	TSR	3,1	March 1987	83939
SECURITY					
Dial-In Security Considerations	P. Grainger	TSR	7,2	Oct. 1991	65248
Distributed Protection with SAFEGUARD	T. Chou	TSR	2,2	June 1986	83937
Enhancing System Security With Safeguard	C. Gaydos	TSR	7,1	April 1991	46988
SYSTEM CONNECTIVITY					
Building Open Systems Interconnection with OSI/AS and OSI/TS	R. Smith	TSR	6,1	March 1990	32986
Connecting Terminals and Workstations to Guardian 90 Systems	E. Siegel	TSR	8,2	Summer 1992	69848
Implementing Client/Server Using RSC	M. Iem, T. Kocher	TSR	8,3	Fall 1992	89803
Network Design Considerations	J. Evjen	TSR	5,2	Sept. 1989	28152
Terminal Connection Alternatives for Tandem Systems	J. Simonds	TSR	5,1	April 1989	18662
Terminal Selection	E. Siegel	TSR	8,2	Summer 1992	69848
The OSI Model: Overview, Status, and Current Issues	A. Dunn	TSR	5,1	April 1989	18662
SYSTEM MANAGEMENT					
Configuring Tandem Disk Subsystems	S. Sitler	TSR	2,3	Dec. 1986	83938
Data Replication in Tandem's Distributed Name Service	T. Eastep	TSR	4,3	Oct. 1988	15748
Enhancements to TMDS	L. White	TSR	3,2	Aug. 1987	83940
Event Management Service Design and Implementation	H. Jordan, R. McKee, R. Schuet	TSR	4,3	Oct. 1988	15748
Introducing TMDS, Tandem's New On-line Diagnostic System	J. Troisi	TSR	1,2	June 1985	83935
Instrumenting Applications for Effective Event Management	J. Dagenais	TSR	7,2	Oct. 1991	65248
Measuring DSM Event Management Performance	M. Stockton	TSR	8,1	Spring 1992	65250
Network Statistics System	M. Miller	TSR	4,3	Oct. 1988	15748
Overview of DSM	P. Homan, B. Malizia, E. Reisner	TSR	4,3	Oct. 1988	15748
SCP and SCF: A General Purpose Implementation of the Subsystem Programmatic Interface	T. Lawson	TSR	4,3	Oct. 1988	15748
RDF: An Overview	J. Guerrero	TSR	7,2	Oct. 1991	65248
RDF Synchronization	F. Jongma, W. Senf	TSR	8,2	Summer 1992	69848

Article title	Author(s)	Publication	Volume, Issue	Publication date	Part number
SYSTEM MANAGEMENT *(cont.)*					
Tandem's Subsystem Programmatic Interface	G. Tom	TSR	4,3	Oct. 1988	15748
Using FOX to Move a Fault-tolerant Application	C. Breighner	TSR	1,1	Feb. 1985	83934
Using the Subsystem Programmatic Interface and Event Management Services	K. Stobie	TSR	4,3	Oct. 1988	15748
VIEWPOINT Operations Console Facility	R. Hansen, G. Stewart	TSR	4,3	Oct. 1988	15748
VIEWSYS: An On-line System-resource Monitor	D. Montgomery	TSR	1,2	June 1985	83935
Writing Rules for Automated Operations	J. Collins	TSR	7,2	Oct. 1991	65248)
UTILITIES					
Enhancements to PS MAIL	R. Funk	TSR	3,1	March 1987	839393

Index by Product

Article title	Author(s)	Publication	Volume, Issue	Publication date	Part number
DP1 AND DP2					
A Comparison of the B00 DP1 and DP2 Disc Processes	T. Schachter	TSR	1,2	June 1985	83935
Determining FCP Conversion Time	J. Tate	TSR	2,1	Feb. 1986	83936
DP1-DP2 File Conversion: An Overview	J. Tate	TSR	2,1	Feb. 1986	83936
DP2 Highlights	K. Carlyle, L. McGowan	TSR	1,2	June 1985	83935
DP2 Key-sequenced Files	T. Schachter	TSR	1,2	June 1985	83935
DP2 Performance	J. Enright	TSR	1,2	June 1985	83935
DP2's Efficient Use of Cache	T. Schachter	TSR	1,2	June 1985	83935
Sizing Cache for Applications that Use B-series DP1 and TMF	P. Shah	TSR	2,2	June 1986	83937
DSM					
Data Replication in Tandem's Distributed Name Service	T. Eastep	TSR	4,3	Oct. 1988	15748
Event Management Service Design and Implementation	H. Jordan, R. McKee, R. Schuet	TSR	4,3	Oct. 1988	15748
Instrumenting Applications for Effective Event Management	J. Dagenais	TSR	7,2	Oct. 1991	65248
Measuring DSM Event Management Performance	M. Stockton	TSR	8,1	Spring 1992	65250
Network Statistics System	M. Miller	TSR	4,3	Oct. 1988	15748
Overview of DSM	P. Homan, B. Malizia, E. Reisner	TSR	4,3	Oct. 1988	15748
SCP and SCF: A General Purpose Implementation of the Subsystem Programmatic Interface	T. Lawson	TSR	4,3	Oct. 1988	15748
Tandem's Subsystem Programmatic Interface	G. Tom	TSR	4,3	Oct. 1988	15748
Using the Subsystem Programmatic Interface and Event Management Services	K. Stobie	TSR	4,3	Oct. 1988	15748
VIEWPOINT Operations Console Facility	R. Hansen, G. Stewart	TSR	4,3	Oct. 1988	15748
Writing Rules for Automated Operations	J. Collins	TSR	7,2	Oct. 1991	65248
DYNAMITE					
An Introduction to DYNAMITE Workstation Host Integration	S. Kosinski	TSR	1,2	June 1985	83935
The DYNAMITE Workstation: An Overview	G. Smith	TSR	1,2	June 1985	83935
ENABLE					
The ENABLE Program Generator for Multifile Applications	B. Chapman, J. Zimmerman	TSR	1,1	Feb. 1985	83934
ENCOMPASS					
The Relational Data Base Management Solution	G. Ow	TJ	2,1	Winter 1984	83931
ENCORE					
The ENCORE Stress Test Generator for On-line Transaction 83931Processing Applications	S. Kosinski	TJ	2,1	Winter 1984	
ENSCRIBE					
Converting Database Files from ENSCRIBE to NonStop SQL	W. Weikel	TSR	6,1	March 1990	32986
FASTSORT					
FASTSORT: An External Sort Using Parallel Processing	J. Gray, M. Stewart, A. Tsukerman, S. Uren, B. Vaughan	TSR	2,3	Dec. 1986	83938

Article title	Author(s)	Publication	Volume, Issue	Publication date	Part number
FOX					
Changes in FOX	N. Donde	TSR	1,2	June 1985	83935
Using FOX to Move a Fault-tolerant Application	C. Breighner	TSR	1,1	Feb. 1985	83934
FUP					
Online Reorganization of Key-Sequenced Tables and Files	G. Smith	TSR	6,2	Oct. 1990	46987
GUARDIAN 90					
B00 Software Manuals	S. Olds	TSR	1,2	June 1985	83935
C00 Software Manuals	E. Levi	TSR	4,1	Feb. 1988	11078
Highlights of the B00 Software Release	K. Coughlin, R. Montevaldo	TSR	1,2	June 1985	83935
Improved Performance for BACKUP2 and RESTORE2	A. Khatri, M. McCline	TSR	1,2	June 1985	83935
Increased Code Space	A. Jordan	TSR	1,2	June 1985	83935
Managing System Time Under GUARDIAN 90	E. Nellen	TSR	2,1	Feb. 1986	83936
Message System Performance Enhancements	D. Kinkade	TSR	2,3	Dec. 1986	83938
Message System Performance Tests	S. Uren	TSR	2,3	Dec. 1986	83938
New GUARDIAN 90 Time-keeping Facilities	E. Nellen	TSR	1,2	June 1985	83935
New Process-timing Features	S. Sharma	TSR	1,2	June 1985	83935
NonStop II Memory Organization and Extended Addressing	D. Thomas	TJ	1,1	Fall 1983	83930
Overview of the C00 Release	L. Marks	TSR	4,1	Feb. 1988	11078
Robustness to Crash in a Distributed Data Base: A Nonshared-memory Multiprocessor Approach	A. Borr	TSR	1,2	June 1985	83935
Tandem's Approach to Fault Tolerance	B. Ball, W. Bartlett, S. Thompson	TSR	4,1	Feb. 1988	11078
The GUARDIAN Message System and How to Design for It	M. Chandra	TSR	1,1	Feb. 1985	83934
The Tandem Global Update Protocol	R. Carr	TSR	1,2	June 1985	83935
INTEGRITY S2					
Overview of the NonStop-UX Operating System for the Integrity S2	P. Norwood	TSR	7,1	April 1991	46988
MEASURE					
How to Set Up a Performance Data Base with MEASURE and ENFORM	M. King	TSR	2,3	Dec. 1986	83938
MEASURE: Tandem's New Performance Measurement Tool	D. Dennison	TSR	2,3	Dec. 1986	83938
MULTILAN					
Introduction to MULTILAN	A. Coyle	TSR	4,1	Feb. 1988	11078
Overview of the MULTILAN Server	A. Rowe	TSR	4,1	Feb. 1988	11078
Using the MULTILAN Application Interfaces	M. Berg, A. Rowe	TSR	4,1	Feb. 1988	11078

Article title	Author(s)	Publication	Volume, Issue	Publication date	Part number
NETBATCH-PLUS					
NetBatch: Managing Batch Processing on Tandem Systems	D. Wakashige	TSR	5,1	April 1989	18662
NetBatch-Plus: Structuring the Batch Environment	G. Earle, D. Wakashige	TSR	6,1	March 1990	32986
NONSTOP SQL					
An Overview of NonStop SQL Release 2	M. Pong	TSR	6,2	Oct. 1990	46987
Concurrency Control Aspects of Transaction Design	W. Senf	TSR	6,1	March 1990	32986
Converting Database Files from ENSCRIBE to NonStop SQL	W. Weikel	TSR	6,1	March 1990	32986
Gateways to NonStop SQL	D. Slutz	TSR	6,2	Oct. 1990	46987
High-Performance SQL Through Low-Level System Integration	A. Borr	TSR	4,2	July 1988	13693
NonStop SQL Data Dictionary	R. Holbrook, D. Tsou	TSR	4,2	July 1988	13693
NonStop SQL: The Single Database Solution	J. Cassidy, T. Kocher	TSR	5,2	Sept. 1989	28152
NonStop SQL Optimizer: Basic Concepts	M. Pong	TSR	4,2	July 1988	13693
NonStop SQL Optimizer: Query Optimization and User Influence	M. Pong	TSR	4,2	July 1988	13693
NonStop SQL Reliability	C. Fenner	TSR	4,2	July 1988	13693
Overview of NonStop SQL	H. Cohen	TSR	4,2	July 1988	13693
Parallelism in NonStop SQL Release 2	M. Moore, A. Sodhi	TSR	6,2	Oct. 1990	46987
Performance Benefits of Parallel Query Execution and Mixed Workload Support in NonStop SQL Release 2	S. Englert, J. Gray	TSR	6,2	Oct. 1990	46987
Tandem's NonStop SQL Benchmark	Tandem Performance Group	TSR	4,1	Feb. 1988	11078
The NonStop SQL Release 2 Benchmark	S. Englert, J. Gray, T. Kocher, P. Shah	TSR	6,2	Oct. 1990	46987
The Outer Join in NonStop SQL	J. Vaishnav	TSR	6,2	Oct. 1990	46987
OSI					
Building Open Systems Interconnection with OSI/AS and OSI/TS	R. Smith	TSR	6,1	March 1990	32986
The OSI Model: Overview, Status, and Current Issues	A. Dunn	TSR	5,1	April 1989	18662
PATHFINDER					
PATHFINDER–An Aid for Application Development	S. Benett	TJ	1,1	Fall 1983	83930
PATHWAY					
A New Design for the PATHWAY TCP	R. Wong	TJ	2,2	Spring 1984	83932
PATHWAY IDS: A Message-level Interface to Devices and Processes	M. Anderton, M. Noonan	TSR	2,2	June 1986	83937
Pathway TCP Enhancements for Application Run-Time Support	R. Vannucci	TSR	7,1	April 1991	46988
The PATHWAY TCP: Performance and Tuning	J. Vatz	TSR	1,1	Feb. 1985	83934
Understanding PATHWAY Statistics	R. Wong	TJ	2,2	Spring 1984	83932
POET					
Designing Client/Server Applications for OLTP on Guardian 90 Systems	W. Culman	TSR	8,3	Fall	89803
PS MAIL					
Enhancements to PS MAIL	R. Funk	TSR	3,1	March 1987	83939

Article title	Author(s)	Publication	Volume, Issue	Publication date	Part number
RDF					
RDF: An Overview	J. Guerrero	TSR	7,2	Oct. 1991	65248
RDF Synchronization	F. Jongma, W. Senf	TSR	8,2	Summer 1992	69848
RESPOND					
The RESPOND OLTP Business Management System for Manufacturing	H. Bolling, W. Bronson	TSR	9,1	Winter 1993	89804
RSC					
Implementing Client/Server Using RSC	M. Iem, T. Kocher	TSR	8,3	Fall 1992	89803
SAFEGUARD					
Dial-In Security Considerations	P. Grainger	TSR	7,2	Oct. 1991	65248
Distributed Protection with SAFEGUARD	T. Chou	TSR	2,2	June 1986	83937
Enhancing System Security With Safeguard	C. Gaydos	TSR	7,1	April 1991	46988
SNAX					
An Overview of SNAX/CDF	M. Turner	TSR	5,2	Sept. 1989	28152
A SNAX Passthrough Tutorial	D. Kirk	TJ	2,2	Spring 1984	83932
SNAX/APC: Tandem's New SNA Software for Distributed Processing	B. Grantham	TSR	3,1	March 1987	83939
SNAX/HLS: An Overview	S. Saltwick	TSR	1,2	June 1985	83935
SPOOLER					
Sizing the Spooler Collector Data File	H. Norman	TSR	4,1	Feb. 1988	11078
TACL					
Debugging TACL Code	L. Palmer	TSR	4,2	July 1988	13693
TACL, Tandem's New Extensible Command Language	J. Campbell, R. Glascock	TSR	2,1	Feb. 1986	83936
TAL					
New TAL Features	C. Lu, J. Murayama	TSR	2,2	June 1986	83837
TCM					
Capacity Planning With TCM	W. Highleyman	TSR	7,2	Oct. 1991	65248
TLAM					
TLAM: A Connectivity Option for Expand	K. MacKenzie	TSR	7,1	April 1991	46988
TMDS					
C00 TMDS Performance	J. Mead	TSR	4,1	Feb. 1988	11078
Enhancements to TMDS	L. White	TSR	3,2	Aug. 1987	83940
Introducing TMDS, Tandem's New On-line Diagnostic System	J. Troisi	TSR	1,2	June 1985	83935
TMF					
Improvements in TMF	T. Lemberger	TSR	1,2	June 1985	83935
TMF and the Multi-Threaded Requester	T. Lemberger	TJ	1,1	Fall 1983	83930
TMF Autorollback: A New Recovery Feature	M. Pong	TSR	1,1	Feb. 1985	83934
TNS/R					
Debugging Accelerated Programs on TNS/R Systems	D. Cressler	TSR	8,1	Spring 1992	65250
Improving Performance on TNS/R Systems With the Accelerator	M. Blanchet	TSR	8,1	Spring 1992	65250
Overview of Tandem NonStop Series/RISC Systems	L. Faby, R. Mateosian	TSR	8,1	Spring 1992	65250

Article title	Author(s)	Publication	Volume, Issue	Publication date	Part number
TRANSFER					
The TRANSFER Delivery System for Distributed Applications	S. Van Pelt	TJ	2,2	Spring 1984	83932
TXP					
The High-Performance NonStop TXP Processor	W. Bartlett, T. Houy, D. Meyer	TJ	2,1	Winter 1984	83931
The NonStop TXP Processor: A Powerful Design for On-line Transaction Processing	P. Oleinick	TJ	2,3	Summer 1984	83933
V8					
The V8 Disc Storage Facility: Setting a New Standard for On-line Disc Storage	M. Whiteman	TSR	1,2	June 1985	83935
VIEWSYS					
VIEWSYS: An On-line System-resource Monitor	D. Montgomery	TSR	1,2	June 1985	83935
VLX					
NonStop VLX Hardware Design	M. Brown	TSR	2,3	Dec. 1986	83938
NonStop VLX Performance	J. Enright	TSR	2,3	Dec. 1986	83938
The VLX: A Design for Serviceability	J. Allen, R. Boyle	TSR	3,1	March 1987	83939
XL8					
Data-encoding Technology Used in the XL8 Storage Facility	D. S. Ng	TSR	2,2	June 1986	83937
Plated Media Technology Used in the XL8 Storage Facility	D. S. Ng	TSR	2,2	June 1986	83937

TandemSystemsReview*OrderForm*

Use this form to order new subscriptions, change subscription information, and order back issues.

☐ I am a Tandem customer. My Tandem sales representative is ______________________.

☐ I am not a Tandem customer and am enclosing a check or money order for the requests indicated on this form. (Subscriptions are $75 per year and each back issue is $20. Make checks payable to Tandem Computers Incorporated.)

Subscription Information

☐ New subscription

☐ Update to subscription information
Subscription number: ______________

Your subscription number is in the upper right corner of the mailing label.

____
COMPANY

____
NAME

____
JOB TITLE

____
DIVISION

____
ADDRESS

____

____
COUNTRY

____
TELEPHONE NUMBER (include all codes for U.S. dialing)

Title or position:

☐ President/CEO
☐ Director/VP information services
☐ MIS/DP manager
☐ Software development manager
☐ Programmer/analyst
☐ System operator
☐ End user
☐ Other: ______________________

Your association with Tandem:

☐ Tandem customer
☐ Third-party vendor
☐ Consultant
☐ Other: ______________________

Back Issue Requests

Number of copies ***Tandem Systems Review***

___ Vol. 1, No. 1, Feb. 1985
___ Vol. 1, No. 2, June 1985
___ Vol. 2, No. 1, Feb. 1986
___ Vol. 2, No. 2, June 1986
___ Vol. 2, No. 3, Dec. 1986
___ Vol. 3, No. 1, March 1987
___ Vol. 3, No. 2, Aug. 1987
___ Vol. 4, No. 1, Feb. 1988
___ Vol. 4, No. 2, July 1988
___ Vol. 4, No. 3, Oct. 1988
___ Vol. 5, No. 1, April 1989
___ Vol. 5, No. 2, Sept. 1989
___ Vol. 6, No. 1, March 1990
___ Vol. 6, No. 2, Oct. 1990
___ Vol. 7, No. 1, April 1991
___ Vol. 7, No. 2, Oct. 1991
___ Vol. 8, No. 1, Spring 1992
___ Vol. 8, No. 2, Summer 1992
___ Vol. 8, No. 3, Fall 1992
___ Vol. 9, No. 1, Winter 1993

Tandem Journal

___ Vol. 1, No. 1, Fall 1983
___ Vol. 2, No. 1, Winter 1984
___ Vol. 2, No. 2, Spring 1984
___ Vol. 2, No. 3, Summer 1984

For questions or ordering information, call 800-473-5868 in the U.S. and Canada or +1-408-285-8150 in other countries.

Send this form to:
Tandem Computers Incorporated
Tandem Systems Review, Loc 216-05
18922 Forge Drive
Cupertino, CA 95014-0701
FAX: +1-408-285-8369

Tandem employees must order their subscriptions and back issues through Courier.

Menu sequence: Marketing Information→ Literature Orders→ Technical Marketing Pubs (TSR)

▲ FOLD

▲ FOLD

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 482 CUPERTINO, CA U.S.A.

POSTAGE WILL BE PAID BY ADDRESSEE

TANDEM SYSTEMS REVIEW
LOC 216-05
TANDEM COMPUTERS INCORPORATED
19333 VALLCO PARKWAY
CUPERTINO, CA 95015–9862

▼ FOLD

▼ FOLD

TandemSystemsReview*ReaderSurvey*

The purpose of this questionnaire is to help the *Tandem Systems Review* staff select topics for publication. Postage is prepaid when mailed in the United States. Readers outside the U.S. should send their replies to their nearest Tandem sales office.

1. How useful is each article in this issue?

Product Update
01 ☐ Indispensable 02 ☐ Very 03 ☐ Somewhat 04 ☐ Not at all

Online Information Processing
05 ☐ Indispensable 06 ☐ Very 07 ☐ Somewhat 08 ☐ Not at all

The DAL Server: Client/Server Access to Tandem Databases
13 ☐ Indispensable 14 ☐ Very 15 ☐ Somewhat 16 ☐ Not at all

The RESPOND OLTP Business Management System for Manufacturing
09 ☐ Indispensable 10 ☐ Very 11 ☐ Somewhat 12 ☐ Not at all

Technical Information and Education
17 ☐ Indispensable 18 ☐ Very 19 ☐ Somewhat 20 ☐ Not at all

2. I specifically would like to see more articles on (select one):

21 ☐ Overview discussions of new products and enhancements.
22 ☐ Performance and tuning information.
23 ☐ High-level overviews on Tandem's approach to solutions.
24 ☐ Application design and customer profiles.
25 ☐ Technical discussions of product internals.
26 ☐ Other____________________

3. Your title or position:

27 ☐ President, VP, Director 28 ☐ Systems analyst 29 ☐ System operator
30 ☐ MIS manager 31 ☐ Software developer 32 ☐ End user
33 ☐ Other ____________________

4. Your association with Tandem:

34 ☐ Tandem customer 35 ☐ Tandem employee 36 ☐ Third-party vendor 37 ☐ Consultant
38 ☐ Other ____________________

5. Comments

____

____

____

____
NAME

____
COMPANY NAME

____
ADDRESS

____

▲ FOLD

▲ FOLD

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 482 CUPERTINO, CA U.S.A.

POSTAGE WILL BE PAID BY ADDRESSEE

TANDEM SYSTEMS REVIEW
LOC 216-05
TANDEM COMPUTERS INCORPORATED
19333 VALLCO PARKWAY
CUPERTINO, CA 95015–9862

▼ FOLD

▼ FOLD

TANDEM

Tandem Computers Incorporated
19333 Vallco Parkway
Cupertino, CA 95014-2599